THE CYBER TOOLCHAIN

NO. 083

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 23, 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 327 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 57 tools

▾ // THE BRIEF

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

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

Three unrelated things stand out. Tailscale's PAM beta brokers database, RDP, SSH and Kubernetes sessions through the tailnet with

vim ? help

S3-backed recording, retiring standing bastion credentials. Google SecOps drops an autonomous triage agent straight into SOAR playbooks. Tirith 0.4.0 puts a package firewall in front of `pip`, holding dependency installs for approval before code executes on a developer box.

> PROTECT
>
> **Grant a privileged session instead of handing out a standing credential**
>
> Tailscale PAM (beta) fronts databases, RDP, SSH, Kubernetes and HTTPS with connectors, secrets management and Terraform definitions, plus session recording to your own S3 bucket, so the audit trail is yours rather than the vendor's. JumpCloud Access Requests routes the grant through a JIRA webhook and logs it in Directory Insights, and Bytebase's `bytebase recovery` CLI gives a break-glass path that does not require leaving an admin account permanently enabled.
>
> tailscale · JumpCloud · bytebase

> PROTECT
>
> **Stop a malicious dependency before the install script runs**
>
> Tirith's `tirith package inspect` and `pkg approve` / `install` put a human gate between resolving a package and executing its install hooks, which is the exact window npm and PyPI attacks use; the `tirith capsule run --preset untrusted-project` sandbox fails closed for code you have not vetted. TruffleHog's pre-receive hook moves secret scanning server-side so a leaked key is rejected at push rather than found after it is already in history, and ARC Harden Runner v2.20.0 pulls its API key from AWS Secrets Manager via IRSA so the CI hardening agent itself stops being a place you store a long-lived secret.
>
> tirith · trufflehog · StepSecurity

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Exploitation & C2

Bugcrowd Researcher Platform

Feedback ⓘ

changelog · Sep 23, 2026 CHANGELOG ↗ RANK

Bugcrowd Researcher Platform connects security researchers with organizations to find and report vulnerabilities through coordinated bug bounty programs.

> Bugcrowd introduced Submission Collaborator, letting researchers work together on a single submission, and quickly raised the collaborator cap while automating reward allocation.

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

01 **Submission Collaborator for shared submissions** NEW

60

Researchers can now collaborate with others on a single submission via the new Submission Collaborator feature. The maximum number of collaborators per submission was raised from 8 to 20, with reward allocation for all collaborators handled directly on the platform without manual processing delays.

— Names the limit change and reward handling but no UI/API path

Submission Collaborator (2026-09-22) changelog-20260923-f01e5362

> changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE
>
> Bugcrowd Researcher Platform adds Submission Collaborator, enabling researchers to collaborate on submissions.

> Submission Collaborator (2026-09-22) BLOG ↗ STABLE
>
> Bugcrowd raises the submission collaborator limit from 8 to 20, with full platform-native reward allocation.
>
> Raises the maximum number of collaborators per submission from 8 to 20, with reward allocation for all collaborators handled directly on the platform without manual processing delays.

WAS THIS USEFUL? 👍 👎

PhoneSploit-Pro

Feedback ⓘ v2.3 NOTES ↗ RANK

An all-in-one hacking tool to remotely take over Android devices.

> PhoneSploit-Pro v2.3 adds screen and display control options during ADB sessions and reorganizes the tool's actions into a hub-based menu structure.

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

01 **Screen and display control options** NEW 33

Adds screen and display control options, expanding the set of device interactions available during an ADB session.

— Names capability but no specific commands or mechanism. v2.3

02 **Hub-based menu structure** IMPROVED 25

Introduces a hub-based menu structure, reorganizing available exploits and actions into grouped navigation hubs.

— Describes reorganization but no navigation paths given. v2.3

v2.3 NOTES ↗ LATEST STABLE

PhoneSploit-Pro v2.3 adds screen/display controls and a hub-based menu to the ADB exploitation toolkit.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.3
https://github.com/AzeemIdrisi/PhoneSploit-Pro.git
# already have the repo? check out this version:
$ git checkout v2.3
```

Adds screen and display control options, expanding the set of device interactions available during an ADB session.

Introduces a hub-based menu structure, reorganizing available exploits and actions into grouped navigation hubs.

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Veracode

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

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

> Veracode's biggest update this window is a sweeping expansion of DAST — new target configuration options, authentication methods, Internal Scanning Management, a public Target Configuration Service REST API, and a scheduling overhaul that breaks existing scheduling API integrations — alongside a new DryRun Security integration for Static Analysis and an IaC `policy` parameter for the GitHub Workflow Integration.

└──▷ WHAT SHIPPED · 16 FEATURES most completely described first ? what's the number?

01 **New authentication options for web application and API scans**

NEW 80

Web application and API scans now support time-based one-time password (TOTP) seeds for multi-factor authentication (MFA), configurable via the REST API, and custom cookies for authentication. API Scanning adds Scriptable Request Modification (SRM), allowing JavaScript to modify HTTP requests at runtime during authentication, and DAST Enterprise mode adds support for login scripts, scanner variables, client certificates, and Scriptable Request Modifications (SRM) for APIs as additional authentication configuration options.

— Names REST API and several auth mechanisms, no exact field names

DAST updates (2026-09-23)

02 **Public DAST Target Configuration Service REST API** NEW

75

The DAST Target Configuration Service REST API is now publicly available to automate management of DAST targets and trigger analyses.

— Names public API surface, no example endpoint path given DAST updates (2026-09-23)

03 **Target configuration enhancements in DAST** NEW 70

DAST target creation and configuration now support assigning multiple teams to a single target for shared security ownership, selecting a business unit when creating a web application or API specification target or on the DAST configuration page, Directory Restrictions to control scan scope, and custom host-to-IP address mappings for scans when DNS lookup is unavailable or not desired.

— Names config field but no exact API path DAST updates (2026-09-23)

04 **Internal Scanning Management (ISM) support for DAST Enterprise**

NEW

70

The ISM endpoint 24.10.0 introduces a proxy exclusion list allowing specified hosts to bypass the configured proxy settings, and DAST Enterprise mode adds Internal Scanning Management (ISM) support for scanning web applications and APIs behind a corporate firewall.

— Names endpoint version and capability, lacks exact config steps

DAST updates (2026-09-23)

05 **Enhanced DAST scan scheduling with breaking API change**

70

BREAKING

Adds the ability to schedule an analysis pause and resume on specific days of the week, during specific periods, or both. The DAST scheduling API endpoints have changed as part of this enhanced scheduling rollout, so integrations using the REST API to manage schedules must be updated to use the new endpoints per the DAST API specification.

— States breaking change and migration need, no exact endpoint names given

DAST updates (2026-09-23)

06 **IaC policy parameter in GitHub Workflow Integration** NEW

65

Adds the `policy` parameter to GitHub Workflow Integration v2.1.7 for infrastructure-as-code (IaC) scans.

— Names exact flag and integration version but no usage steps changelog-20260923-e1647b29

07 **AI-Assisted Login for DAST Enterprise** NEW

65

DAST Enterprise mode now provides an AI-based authentication method (AI-Assisted Login) to cover complex login flows without recorded login scripts.

— Names feature and mechanism, no configuration steps DAST updates (2026-09-23)

08 **Manual application linking for DAST analysis runs** NEW

60

DAST now supports manual application linking to connect analysis run results to a Veracode Platform application for policy evaluation and combined multi-scan-type reporting.

— Explains purpose and outcome, no UI/API steps DAST updates (2026-09-23)

09 **Organization-level URL scan settings via REST API** NEW

60

Adds the ability to set URL scan settings at the organization level via the Dynamic Analysis REST API.

— Names API surface, no field or endpoint path DAST updates (2026-09-23)

THINNER COVERAGE BELOW

10 **EASM auto-discovery integration with DAST** NEW 55

EASM now automatically creates discovered targets in DAST, which can be mapped or converted to regular DAST targets for scanning.

— Explains mechanism but no UI or API path DAST updates (2026-09-23)

11 **DryRun Security integration for Static Analysis** NEW 50

Integrates Veracode Static Analysis with DryRun Security, surfacing Veracode findings and remediation guidance inside pull request and code review workflows.

— Describes purpose but no setup detail or config changelog-20260923-e1647b29

12 **Postman Collection uploads for API Scanning** NEW 50

API Scanning now supports uploading and scanning Postman Collections in the Veracode Platform.

— Names format supported, no upload steps described DAST updates (2026-09-23)

13 **Downloadable scan reports in DAST Enterprise** NEW 40

DAST Enterprise mode now supports downloadable scan reports providing visibility into scan coverage and scanner behavior.

— States benefit, no format or location named DAST updates (2026-09-23)

14 **Similarity threshold for web application scans** NEW 40

Adds a similarity threshold setting for web application analyses to ignore similar web pages during scanning.

— Explains purpose, no default value or range given DAST updates (2026-09-23)

15 **Open port scanning in Dynamic Analysis and DAST Enterprise** NEW 35

Dynamic Analysis and DAST Enterprise mode now support scanning web applications for open ports.

— Bare capability statement, no scope or limits DAST updates (2026-09-23)

16 **New HTTP header, TLS, and fingerprinting detections** NEW

Adds new detections for HTTP headers, TLS, and fingerprinting in the scan engine.

— Names categories but no specific checks listed DAST updates (2026-09-23)

└──▷ **BREAKING ON UPGRADE**

! The DAST scheduling API endpoints have changed as part of the enhanced scheduling rollout; integrations using the REST API to manage schedules must be updated to use the new endpoints per the DAST API specification.

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Veracode adds DryRun Security integration, IaC policy parameter, Jira 11 support, and a redesigned EASM UI

Adds the `policy` parameter to GitHub Workflow Integration v2.1.7 for infrastructure-as-code (IaC) scans.

Integrates Veracode Static Analysis with DryRun Security, surfacing Veracode findings and remediation guidance inside pull request and code review workflows.

DAST updates (2026-09-23) BLOG ↗ STABLE

Veracode DAST gains multi-team targets, custom host-to-IP mappings, directory restrictions, open-port scanning, and EASM integration.

Supports assigning multiple teams to a single DAST target when creating a web application or API specification, enabling shared security ownership.

Adds the ability to select a business unit when creating a web application or API specification target, or on the DAST configuration page.

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

Introduces custom host-to-IP address mappings for DAST scans when DNS lookup is unavailable or not desired.

Dynamic Analysis and DAST Enterprise mode now support scanning web applications for open ports.

+16 more

└──▷ **BREAKING ON UPGRADE**

! The DAST scheduling API endpoints have changed as part of the enhanced scheduling rollout; integrations using the REST API to manage schedules must be updated to use the new endpoints per the DAST API specification.

WAS THIS USEFUL?  

Snyk

Feedback ⓘ changelog · Sep 23, 2026 DOCS ↗ BLOG ↗ RANK

The official changelog for the Snyk AI Trust Platform. Discover what's new in our security intelligence, AI-powered tools, and latest product enhancements.

> Snyk shipped GA of the Evo MCP Server for agentic AI asset discovery and policy management, alongside a new `snyk fix --agentic` CLI command, DNS-based domain verification, Snyk Studio for AI-assisted dev workflows, and a migration of Snyk for Jira to Atlassian's Forge platform.

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

01 **`snyk fix --agentic` CLI command** NEW 70

Adds the `snyk fix --agentic` CLI command, which applies and verifies vulnerability fixes automatically in an agentic pipeline without manual review steps.

Automatically apply and verify fixes for detected vulnerabilities in an agentic pipeline without manual review steps.

```
$ snyk fix --agentic
```

Apply and verify a vulnerability fix in an agentic CI pipeline without manual review steps.

```
$ snyk fix --agentic
```

— Exact runnable flag and command given 2026-09-23-284c2c21 2026-09-23-480bc122

02 **Evo MCP Server for agentic asset and policy management** NEW 62

Adds the Evo MCP Server (GA), letting any MCP client — Claude, Codex, Cursor — query discovered AI assets across code, developer machines, and pentest targets, inspect asset relationships (e.g. which MCP servers an agent uses), and read or write policies without leaving the agent.

— Names supported clients and capabilities but no setup steps 2026-09-23-284c2c21

THINNER COVERAGE BELOW

03 **Snyk for Jira migrates from Atlassian Connect to Forge** 46

BREAKING

Migrates Snyk for Jira from Atlassian Connect to Forge, Atlassian's current cloud app platform, ahead of Atlassian's December 2026 end-of-support deadline for Connect.

— Explains why and deadline but not migration steps for users `2026-09-23-480bc122`

04 **DNS TXT record domain verification** NEW 45

Adds DNS TXT record verification as a method to confirm domain ownership for Snyk API & Web targets.

— Names the verification method and target area, no exact record format

`2026-09-23-284c2c21`

05 **IDE Plugin Compatibility Matrix** NEW 45

Adds an IDE Plugin Compatibility Matrix documenting compatible Snyk CLI version ranges for each IDE plugin version released in the past 12 months.

— Documents scope and timeframe but no direct link given `2026-09-23-284c2c21`

06 **Snyk Studio for AI-assisted development** NEW 25

Adds Snyk Studio, which embeds security directives into AI-assisted development workflows.

— Bare description with no mechanism or usage detail `2026-09-23-284c2c21`

changelog · Sep 23, 2026 DOCS ↗ BLOG ↗ LATEST STABLE

Snyk: Introduces Evo by Snyk with an MCP Server surface for agentic security use cases.

└──▷ **TRY IT**

Automatically apply and verify fixes for detected vulnerabilities in an agentic pipeline without manual review steps.

```
$ snyk fix --agentic
```

Migrates Snyk for Jira from Atlassian Connect to Forge, Atlassian's current cloud app platform, ahead of Atlassian's December 2026 end-of-support deadline.

changelog · Sep 23, 2026 DOCS ↗ BLOG ↗ STABLE

Snyk's Evo gains an MCP server so Claude, Codex, or Cursor can query your AI asset estate and manage policies in-agent.

└──▷ **TRY IT**

Apply and verify a vulnerability fix in an agentic CI pipeline without manual review steps.

```
$ snyk fix --agentic
```

Adds the Evo MCP Server (GA), letting any MCP client — Claude, Codex, Cursor — query discovered AI assets across code, developer machines, and pentest targets, inspect asset relationships (e.g. which MCP servers an agent uses), and read or write policies without leaving the agent.

Adds `snyk fix --agentic` CLI command, which applies and verifies vulnerability fixes.

Adds DNS TXT record verification as a method to confirm domain ownership for Snyk API & Web targets.

Adds Snyk Studio, which embeds security directives into AI-assisted development workflows.

Adds an IDE Plugin Compatibility Matrix documenting compatible Snyk CLI version ranges for each IDE plugin version released in the past 12 months.

WAS THIS USEFUL? 👍 👎

StackHawk

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

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

> StackHawk shipped a new CLI install script with SHA-256 verification, organization-wide secret resolution for shared scan credentials, expanded session-verification controls for authenticated scans, and OpenAPI base-path overrides, alongside fixes to technology-flag management and replayed-request headers.

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

01 hawk CLI install script NEW

90

Adds support for installing the `hawk` CLI via `curl -sSf https://download.stackhawk.com/hawk/install.sh | sh`, with SHA-256 verification against the published manifest. Supports `HAWK_INSTALL_VERSION` to pin a release and `HAWK_INSTALL_DIR` to control the install destination (default `~/.hawk/bin`).

Pin the hawk CLI to a specific version in CI to keep pipelines reproducible.

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

Pin the HawkScan CLI to a specific version in CI so pipeline installs are reproducible and unaffected by new releases.


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


— Exact command, env vars and default path given snapshot-20260923

02 OpenAPI basePath overrides NEW

90

Adds `basePath` under `app.openApiConf` to prepend a base path to every operation discovered in an OpenAPI specification, taking precedence over both a path on `app.host` and `servers[].url` in the spec. Adds `disableServerBasePath` under `app.openApiConf` to suppress the base path declared in the specification's `servers[].url` when the API is served at the root of `app.host`.

Scan an API served at the root of your host even though the OpenAPI spec declares a base path in `servers[].url` — use `basePath` and `disableServerBasePath` together to

override both the spec and any path on the host URL.

```yaml
app:
  host: http://localhost:8000
  openApiConf:
    basePath: /api/v2
    disableServerBasePath: true
```

— Exact config keys with a runnable override example product docs

03 **Session verification for authenticated scans** NEW 88

The session-verification block under `app.authentication.sessionVerification` in `stackhawk.yml` / `.stackhawk.yml` enables periodic re-authentication polling during scans, with `pollFrequency`, `activeScanPollFrequency`, and `maxUnauthenticatedRequestPercent` controlling check cadence and tolerance. `maxConsecutiveUnauthenticatedChecks` flags suspect authenticated coverage when a burst of consecutive failed session checks occurs, catching gaps a scan-wide average would miss, and `disabled: true` explicitly turns off session verification, falling back to `app.authentication.loggedInIndicator` / `loggedOutIndicator` regex matching. The block is off by default — providing it (even empty) enables polling, while omitting it requires a `loggedInIndicator` or `loggedOutIndicator` instead, and those two fields remain supported (not deprecated) as the only per-request session check when verification is absent or disabled.

Detect short session-loss bursts during an authenticated scan without waiting for the scan-wide unauthenticated percentage to look unhealthy.

```yaml
app:
  authentication:
    sessionVerification:
      maxConsecutiveUnauthenticatedChecks: 3
```

Disable session-verification polling for an app that gives false positives, and fall back to per-request `loggedInIndicator` regex matching instead.

```yaml
app:
  authentication:
    loggedInIndicator: 'Welcome, '
    sessionVerification:
      disabled: true
```

Enable mid-scan session verification for an authenticated scan so HawkScan re-authenticates automatically if it drifts logged out, catching missed findings during active attack phases.

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

Enable mid-scan session verification for an app that uses form login, so HawkScan re-authenticates automatically if it drifts logged-out during a long active scan.

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

Enable mid-scan session re-authentication to catch missed findings caused by the scanner going unauthenticated mid-attack phase.

```yaml
app:
  authentication:
    sessionVerification:
      pollFrequency: 20
      activeScanPollFrequency: 3
      maxUnauthenticatedRequestPercent: 2
```

— Names every config key and fallback behavior with runnable examples `snapshot-20260923`

04 **Fix for hawk op app tech-flags set** IMPROVED 80

Naming one technology flag in `hawk op app tech-flags set` no longer silently re-enables all other 34 flags; the command now reads the current flags, merges the requested change, and writes only the result.

— Names exact command and precise before/after fix `changelog-20260923-a9739752`

05 **Organization-wide secrets in scan config** NEW 70

Scan configuration now resolves `${VARIABLE}` references against organization-wide secrets, so a team can share one copy of a scan credential instead of one per engineer. Personal secrets take precedence over organization secrets, HawkScan warns when both exist for the same name, and it names any unresolvable variable explicitly.

— Names syntax and precedence rules but no example snapshot-20260923

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

StackHawk: Fixes `hawk op app tech-flags set` so naming one technology flag no longer silently re-enables all other 34 flags; the command now reads current flags, merges the change, and writes only the result.

└──▷ **USE IT**

Enable mid-scan session verification for an authenticated scan so HawkScan re-authenticates automatically if it drifts logged out, catching missed findings during active attack phases.

stackhawk.yml

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

Pin the HawkScan CLI to a specific version in CI so pipeline installs are reproducible and unaffected by new releases.

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

Emits the platform's own message — including minimum supported version and upgrade location — and exits non-zero when a scan is declined because the CLI is out of support; logs deprecation advisories once per scan rather than on every platform call.

snapshot-20260923 NOTES ↗ STABLE

StackHawk: Adds support for installing the `hawk` CLI via `curl -sSf https://download.stackhawk.com/hawk/install.sh | sh`, with SHA-256 verification against the published manifest; supports `HAWK_INSTALL_VERSION` to pin a release and `HAWK_INSTALL_DIR` to control the install destination (default `~/.hawk/bin`).

└──▷ **USE IT**

Enable mid-scan session re-authentication to catch missed findings caused by the scanner going unauthenticated mid-attack phase.

stackhawk.yml

```
app:
  authentication:
    sessionVerification:
      pollFrequency: 20
      activeScanPollFrequency: 3
      maxUnauthenticatedRequestPercent: 2
```

Pin the hawk CLI to a specific version in CI to keep pipelines reproducible.

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

Adds resolution of `${VARIABLE}` references in scan configuration against organization-wide secrets, so a team can share one copy of a scan credential; personal secrets take precedence and HawkScan warns when both exist, and names any unresolvable variable explicitly.

WAS THIS USEFUL?  

Aikido Security

Feedback ⓘ

#56 - Cloud Asset Map, Creating fix PRs with Deep Review and Code Coverage on GitLab

NOTES ↗ RANK

Aikido Security is a developer-first security platform that identifies and remediates vulnerabilities in code, dependencies, and infrastructure.

> Aikido's headline addition this window is Cloud Inventory (Cloud Asset Map) for visualizing cloud resources by region and VPC, alongside the ability to generate fix pull requests directly from Deep Review comments, GitLab CI code coverage support, and Altar — Aikido's first open-weight, on-prem pentesting model.

└──▷ **WHAT SHIPPED · 13 FEATURES** most completely described first ? what's the number?

01 **New APIs and webhook events** NEW 73

Adds a `GET` List user logins history API endpoint for auditing authentication activity, a Broker management API for programmatic control of broker configuration, a webhook event for when a new issue ignore/snooze request is accepted or denied, and a webhook event for when a Domain scan finishes.

— Names endpoint and webhook events but no full paths or payloads.

#56 - Cloud Asset Map, Creating fix PRs…

02 **Fix PRs generated from Deep Review comments** NEW 70

Reply to a Deep Review comment requesting a fix and Aikido opens a new pull request with the fix built on top of your current PR.

— Clear trigger action and resulting mechanism described.

#56 - Cloud Asset Map, Creating fix PRs…

03 **Code Coverage upload from GitLab CI** NEW 60

Aikido now supports Code Coverage upload from GitLab CI, with per-file and per-line results visible in Aikido alongside existing GitHub coverage.

— Names the CI source and result granularity, no setup steps.

#56 - Cloud Asset Map, Creating fix PRs…

04 **Altar: on-prem, air-gapped pentesting model** NEW 60

Altar is Aikido's first open-weight security model, built to run frontier-grade pentesting fully on-prem and air-gapped without sending data to third parties.

— Explains purpose and deployment model but no setup instructions.

#56 - Cloud Asset Map, Creating fix PRs…

THINNER COVERAGE BELOW

05 **Cloud Inventory (Cloud Asset Map)** NEW 58

Cloud Inventory (Cloud Asset Map) lets you see every cloud resource laid out by region and VPC, with exposure findings and firewall rules accessible in one click.

— Describes layout and data shown but no exact navigation path.

#56 - Cloud Asset Map, Creating fix PRs…

06 **Zen Firewall agent expanded to more languages and platforms**

IMPROVED 55

New Zen Firewall agent releases ship for Java, PHP, Golang, and Node, and the PHP agent adds Alpine support.

— Lists supported languages and platform but no install detail.

#56 - Cloud Asset Map, Creating fix PRs…

07 **Cloud scanning reachability and checks expanded** IMPROVED

55

Container scanning Reachability now extends to include BigQuery datasets, VM Reachability is now available for GCP in Cloud scanning, and new infrastructure-focused checks have been added to cloud scanning.

— Names extended targets but no configuration or command steps.

#56 - Cloud Asset Map, Creating fix PRs…

08 **JetBrains IDE plugin gains multi-repository scanning** IMPROVED

40

New JetBrains IDE plugin release adds support for multi-repository scanning.

— States the capability without usage detail. #56 - Cloud Asset Map, Creating fix PRs…

09 **New MCP tools for issue and scan data** NEW 40

New MCP tools let you get issue details including pentest information, Deep Review, and other SCA scanning tool data.

— Lists tool coverage but no invocation examples.

#56 - Cloud Asset Map, Creating fix PRs…

10 **Firefox and Homebrew added to Device Protection** NEW

40

Adds Firefox and Homebrew as new ecosystems in Device Protection.

— Names ecosystems but no scope or setup detail. #56 - Cloud Asset Map, Creating fix PRs…

11 **Task Tracker issues for Code Quality findings** NEW

35

Supports creating Task Tracker issues for Code Quality findings.

— Brief mention without workflow detail. #56 - Cloud Asset Map, Creating fix PRs…

12 **Grouped bar charts in Custom Reports** NEW 25

Custom Reports now support grouped bar charts.

— Bare feature name with no further detail. #56 - Cloud Asset Map, Creating fix PRs…

13 **Sensitive Data information in DSPM** NEW 25

Adds Sensitive Data information to DSPM.

— Bare mention, no detail on data types covered. #56 - Cloud Asset Map, Creating fix PRs…

> #56 - Cloud Asset Map, Creating fix PRs with Deep Review and Code Coverage on GitLab
>
> NOTES ↗ LATEST STABLE
>
> Aikido adds Cloud Asset Map, fix PRs from Deep Review comments, GitLab code coverage, on-prem pentest model, and new APIs
>
> Adds `GET` List user logins history API endpoint for auditing authentication activity.
>
> Adds Broker management API for programmatic control of broker configuration.
>
> Adds webhook event for when a new issue ignore/snooze request is accepted or denied.
>
> Adds webhook event for when a Domain scan finishes.
>
> Introduces Cloud Inventory (Cloud Asset Map): see every cloud resource laid out by region and VPC, with exposure findings and firewall rules accessible in one click.
>
> +14 more

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

Get started with Semgrep to help you catch, flag, and fix real vulnerabilities before they ship.

> Semgrep expanded Agentic Workflows with in-agent status reporting from Semgrep Guardian and added CSV export for Agentic Workflows issues.

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

01 **Guardian status reporting inside AI coding agents** NEW 45

Semgrep Guardian now reports status directly inside AI coding agents, including setup confirmation, per-turn scan summaries, sign-in requirements, and server errors. A banner also appears when viewing older Agentic Workflows runs that links to the newest run of the same workflow on that branch.

— Describes behaviors but no concrete UI path or API changelog-20260923-0de72f2b

02 **CSV export for Agentic Workflows issues** NEW 45

Agentic Workflows issues can now be exported as CSV, available both from the Issues page and through the public API.

— Names export format and two access surfaces but no endpoint details snapshot-20260923

> changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE
>
> Semgrep: Semgrep Guardian now reports status inside AI coding agents, including setup confirmation, per-turn scan summaries, sign-in requirements, and server errors.
>
> Adds a banner when viewing older Agentic Workflows runs that links to the newest run of the same workflow on that branch.

> snapshot-20260923 NOTES ↗ STABLE
>
> Semgrep adds CSV export for Agentic Workflows issues via the Issues page and public API.

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ OpenTelemetry in the GitHub Copilot app (2026-09-23)

BLOG ↗ RANK

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

> GitHub shipped enterprise-grade OpenTelemetry observability for the Copilot app, added Claude Opus 5.5 and the new GPT-6 Sol/Luna models to the Copilot model picker, rolled out a major JetBrains Copilot update (assisted tool approvals, Codex plan mode, MCP controls, re-editable messages, org skills), sped up C++ intelligence with whole-codebase indexing in Copilot CLI, and hardened SSH with post-quantum key exchange plus new RSA key-size enforcement and legacy algorithm deprecations.

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

01 **OpenTelemetry support in the Copilot app** NEW 90

Adds a `telemetry` property to `managed-settings.json` for enterprise-managed Copilot settings, enabling OTel export and specifying the receiving endpoint, with prompt/response content excluded by default. Supports sending Copilot agent activity data — including AI model requests, tool calls, and step-by-step execution traces — to any OpenTelemetry-compatible monitoring backend.

Enable OTel export for all developers in your enterprise so agent sessions flow into your existing monitoring stack without per-developer configuration.

json

```
{
  "telemetry": {
    "endpoint": "https://otel-collector.example.com"
  }
}
```

— Names exact config key, file, and includes a working config example.

OpenTelemetry in the GitHub Copilot app…

02 **New GPT-6 Sol and GPT-6 Luna models** NEW 80

Adds GPT-6 Sol, a balanced model for interactive and agentic coding with multistep validation, available on Copilot Pro+, Max, Business, and Enterprise plans, and GPT-6 Luna, a lightweight, cost-efficient model for smaller and faster tasks — the lowest-cost

option in the GPT-6 family — available on Copilot Pro, Pro+, Max, Business, and Enterprise plans.

Switch to GPT-6 Luna in Copilot CLI for fast, low-cost inline suggestions during a code review session.

> 📍 In Copilot CLI, open the model picker and select `GPT-6 Luna` before running your query.

As a Copilot Business admin, restrict which GPT-6 models are available to your organisation before rollout completes.

> 📍 In GitHub, go to your organisation's Copilot settings › model policy, then locate the GPT-6 model entries and disable any models you do not want enabled by default.

— Names both models, plan tiers, and selection/policy steps.

OpenAI's GPT-6 Sol and GPT-6 Luna now a…

---

03 **Minimum RSA SSH key size enforcement** BREAKING 80

Enforces a minimum RSA SSH key size of 3072 bits for all new keys uploaded after October 14, 2026, for both signing and authentication; smaller keys will be rejected.

Generate a new Ed25519 key (recommended) or a compliant 3072-bit RSA key before the October 14 deadline to avoid connection failures.

```
$ ssh-keygen -t ed25519 -C "your_email@example.com"
# or, if RSA is required for compatibility:
ssh-keygen -t rsa -b 3072 -C "your_email@example.com"
```

— Exact deadline plus a runnable keygen command. Security improvements for SSH (2026-09-…

---

04 **Deprecated legacy SSH algorithms** DEPRECATED 75

The `ssh-rsa` signature type (RSA keys using SHA-1, including `ssh-rsa-cert-v01@openssh.com` certificates) will be removed effective January 13, 2026, with brownouts on November 4 and December 9, 2026, and the `diffie-hellman-group-exchange-sha256` key exchange algorithm will be removed on the same schedule; connections relying on either will fail.

— Exact algorithm names and removal timeline given.

Security improvements for SSH (2026-09-…

---

05 **Whole codebase indexing for C++ in Copilot CLI** NEW 70

GitHub Copilot CLI now builds a persistent symbol index across all project files for C++ to speed up code intelligence; use `/lsp logs` at any time to check indexing progress.

Check how far along the initial C++ index build is while working in a large repository.

```
$ /lsp logs
```

— Names an exact command to check progress. Faster C++ code intelligence with whole…

06 **Post-quantum SSH key exchange** NEW 70

Adds `mlkem768x25519-sha256` post-quantum key exchange for SSH sessions on github.com and GitHub Enterprise Cloud with Data Residency (excluding the U.S. region), with GitHub Enterprise Server support arriving in version 3.24.

— Names exact algorithm and platform scope, no user action needed.

Security improvements for SSH (2026-09-…

07 **Plan mode for the Codex agent** NEW 65

Adds plan mode to the Codex agent, letting you review, refine, or approve a plan before the agent begins making changes.

Use Codex plan mode to review and shape an agent's approach before it touches any files — useful for high-impact refactors or unfamiliar codebases.

> 📍 In the Copilot tool window, open a Codex agent session, describe your task, and select Plan mode. Review the generated plan, refine the steps, then approve it to let the agent begin implementation.

— Exact UI steps given for invoking plan mode. New features and improvements in Copilo…

08 **Claude Opus 5.5 available in Copilot** NEW 65

Claude Opus 5.5 is now available in the Copilot model picker across Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, GitHub Copilot app, github.com, GitHub Mobile (iOS and Android), JetBrains IDEs, Xcode, and Eclipse, aimed at agentic coding and long-running tasks with efficient multistep error recovery.

▶ Watch video

Select Claude Opus 5.5 in your IDE to handle long-running agentic coding tasks that benefit from efficient multistep error recovery.

> 📍 In the Copilot model picker (Visual Studio Code, JetBrains, Xcode, Eclipse, or github.com), select Claude Opus 5.5 from the model list.

— Lists all supported clients and exact selection steps.

Claude Opus 5.5 is now available in Git…

09 **MCP server controls in Copilot for JetBrains** NEW 60

Adds a setting to toggle the built-in GitHub MCP Server on or off independently of manually configured MCP servers (enabled by default), plus persistent per-tool controls for MCP servers in Copilot agent sessions, letting you enable or disable individual tools without affecting others.

Disable the built-in GitHub MCP Server to avoid conflicts with a custom MCP configuration while keeping your other manually configured MCP servers active.

> 📍 In the Copilot tool window, open MCP tool settings, locate the labeled 'built-in GitHub MCP Server' entry, and toggle it off. Manually configured MCP servers remain unaffected.

– UI toggle steps given but no config/flag name. New features and improvements in Copilo…

THINNER COVERAGE BELOW

10 **AI-assisted tool approvals in Copilot agent sessions** NEW 50

Introduces AI-assisted tool approvals (assisted approvals) in public preview for Copilot agent sessions, automatically approving low-risk tool calls while still prompting for higher-risk actions.

– Explains behavior but gives no usage steps. New features and improvements in Copilo…

11 **Re-editable agent messages** NEW 45

Enables re-editing of previous user messages in a Copilot agent session, rewinding both the conversation and file changes before the replacement message is sent.

– Mechanism described but no navigation path given.

New features and improvements in Copilo…

12 **Session browsing and chat UI improvements in JetBrains**

IMPROVED 40

Adds a side-by-side chat panel switcher in the session toolbar, letting you keep a chat open in the editor while browsing sessions in the tool window, and adds browsable usage tips above the chat input with shortcuts to commands, customizations, and settings.

– UI additions named but no exact navigation steps.

New features and improvements in Copilo…

13 **Organization and enterprise skills in Copilot sessions** NEW 30

Extends local and Copilot agent sessions to support organization and enterprise skills and organization-managed custom instructions.

— Brief description with no config or steps given.

New features and improvements in Copilo…

14 **Inline chat hidden in remote JetBrains environments** BREAKING

25

Inline chat and its entry points are now hidden in JetBrains Gateway and remote development environments.

— States the change but offers no workaround. New features and improvements in Copilo…

└──▷ **BREAKING ON UPGRADE**

! Inline chat and its entry points are now hidden in JetBrains Gateway and remote development environments.

! The `ssh-rsa` signature type (RSA keys using SHA-1, including `ssh-rsa-cert-v01@openssh.com` certificates) will be removed effective January 13, 2026 (brownouts November 4 and December 9, 2026); connections using it will fail.

! The `diffie-hellman-group-exchange-sha256` key exchange algorithm will be removed on the same schedule; SSH clients relying on it will be unable to connect.

! New RSA SSH keys uploaded after October 14, 2026 must be at least 3072 bits; smaller keys will be rejected.

OpenTelemetry in the GitHub Copilot app (2026-09-23) BLOG ↗ LATEST

STABLE

GitHub Copilot app gains OpenTelemetry support via enterprise `managed-settings.json`, enabling centralized agent observability.

└──▷ **USE IT**

Enable OTel export for all developers in your enterprise so agent sessions flow into your existing monitoring stack without per-developer configuration.

managed-settings.json

```
{
  "telemetry": {
    "endpoint": "https://otel-collector.example.com"
  }
}
```

Adds a `telemetry` property to `managed-settings.json` for enterprise-managed Copilot settings to enable OTel export and specify the receiving endpoint, with prompt/response content excluded by default.

Supports sending Copilot agent activity data — including AI model requests, tool calls, and step-by-step execution traces — to any OpenTelemetry-compatible monitoring backend.

New features and improvements in Copilot for JetBrains (2026-09-23) BLOG ↗

STABLE

GitHub Copilot for JetBrains 1.18.0 adds AI-assisted tool approvals, re-editable agent messages, org-shared skills, and Codex plan mode.

└──▷ **HOW TO FIND IT**

Use Codex plan mode to review and shape an agent's approach before it touches any files — useful for high-impact refactors or unfamiliar codebases.

📍 In the Copilot tool window, open a Codex agent session, describe your task, and select `Plan` mode. Review the generated plan, refine the steps, then approve it to let the agent begin implementation.

Disable the built-in GitHub MCP Server to avoid conflicts with a custom MCP configuration while keeping your other manually configured MCP servers active.

📍 In the Copilot tool window, open MCP tool settings, locate the labeled 'built-in GitHub MCP Server' entry, and toggle it off. Manually configured MCP servers remain unaffected.

Adds a setting to toggle the built-in GitHub MCP Server on or off independently of manually configured MCP servers, which remains enabled by default.

Adds persistent per-tool controls for MCP servers in Copilot agent sessions, letting you enable or disable individual tools without affecting others.

Introduces AI-assisted tool approvals (assisted approvals) in public preview for Copilot agent sessions, automatically approving low-risk tool calls while still prompting for higher-risk actions.

Enables re-editing of previous user messages in a Copilot agent session, rewinding both the conversation and file changes before the replacement message is sent.

Adds plan mode to the Codex agent, letting you review, refine, or approve a plan before the agent begins making changes.

+3 more

└──▷ **BREAKING ON UPGRADE**

! Inline chat and its entry points are now hidden in JetBrains Gateway and remote development environments.

Faster C++ code intelligence with whole codebase indexing (2026-09-22) BLOG ↗

STABLE

GitHub Copilot CLI gains whole codebase indexing for C++, building a persistent symbol index across all project files.

└──▷ **TRY IT**

Check how far along the initial C++ index build is while working in a large repository.

```
$ /lsp logs
```

Adds the `/lsp logs` command to check C++ whole codebase indexing progress at any time.

Claude Opus 5.5 is now available in GitHub Copilot (2026-09-22) BLOG ↗

STABLE

Claude Opus 5.5 is now available in GitHub Copilot for agentic coding and long-running tasks across all major IDEs and clients.

└──▷ **HOW TO FIND IT**

Select Claude Opus 5.5 in your IDE to handle long-running agentic coding tasks that benefit from efficient multistep error recovery.

> 📍 In the Copilot model picker (Visual Studio Code, JetBrains, Xcode, Eclipse, or github.com), select `Claude Opus 5.5` from the model list.

Adds Claude Opus 5.5 to the Copilot model picker in Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot coding agent, GitHub Copilot app, github.com, GitHub Mobile (iOS and Android), JetBrains IDEs, Xcode, and Eclipse.

OpenAI's GPT-6 Sol and GPT-6 Luna now available (2026-09-22) BLOG ↗ STABLE

GitHub Copilot adds GPT-6 Sol and GPT-6 Luna as selectable models across VS Code, CLI, mobile, and IDEs.

└──▷ **HOW TO FIND IT**

Switch to GPT-6 Luna in Copilot CLI for fast, low-cost inline suggestions during a code review session.

> 📍 In Copilot CLI, open the model picker and select `GPT-6 Luna` before running your query.

As a Copilot Business admin, restrict which GPT-6 models are available to your organisation before rollout completes.

> 📍 In GitHub, go to your organisation's Copilot settings › model policy, then locate the GPT-6 model entries and disable any models you do not want enabled by default.

Adds `GPT-6 Sol` to GitHub Copilot as a balanced model for interactive and agentic coding with multistep validation, available on Copilot Pro+, Max, Business, and Enterprise plans.

Adds `GPT-6 Luna` to GitHub Copilot as a lightweight, cost-efficient model for smaller and faster tasks — the lowest-cost option in the GPT-6 family — available on Copilot Pro, Pro+, Max, Business, and Enterprise plans.

`Security improvements for SSH (2026-09-22)` **BLOG ↗** **STABLE**

GitHub SSH gains post-quantum `mlkem768x25519-sha256` key exchange and enforces 3072-bit minimum RSA key sizes.

└──▷ TRY IT

Generate a new Ed25519 key (recommended) or a compliant 3072-bit RSA key before the October 14 deadline to avoid connection failures.

```
$ ssh-keygen -t ed25519 -C "your_email@example.com"
# or, if RSA is required for compatibility:
ssh-keygen -t rsa -b 3072 -C "your_email@example.com"
```

Adds `mlkem768x25519-sha256` post-quantum key exchange for SSH sessions on github.com and GitHub Enterprise Cloud with Data Residency (excluding the U.S. region), with GitHub Enterprise Server support in version 3.24.

Enforces a minimum RSA SSH key size of 3072 bits for all new keys uploaded after October 14, 2026, for both signing and authentication.

└──▷ BREAKING ON UPGRADE

- **!** The `ssh-rsa` signature type (RSA keys using SHA-1, including `ssh-rsa-cert-v01@openssh.com` certificates) will be removed effective January 13, 2026 (brownouts November 4 and December 9, 2026); connections using it will fail.
- **!** The `diffie-hellman-group-exchange-sha256` key exchange algorithm will be removed on the same schedule; SSH clients relying on it will be unable to connect.
- **!** New RSA SSH keys uploaded after October 14, 2026 must be at least 3072 bits; smaller keys will be rejected.

WAS THIS USEFUL?  

StepSecurity

Feedback ⓘ ARC Harden Runner v2.20.0 (2026-09-23)

BLOG ↗ RANK

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

> StepSecurity's ARC Harden Runner v2.20.0 lets the Helm-deployed daemonset pull its API key straight from AWS Secrets Manager via IRSA, alongside new options for injecting custom environment variables and configuring the Tetragon sidecar's security context.

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

01 Keyless API key fetch from AWS Secrets Manager via IRSA NEW

93

New `keyless.enabled`, `keyless.region`, `keyless.secretName`, and `keyless.refreshIntervalHours` Helm values let the Harden Runner daemonset fetch its StepSecurity API key from AWS Secrets Manager using the pod's IAM role (IRSA) instead of a static Helm value, refreshing on a configurable schedule (default `3` hours). The secret must be a JSON object containing an `api_key` field, and the IAM role requires the `secretsmanager:GetSecretValue` permission.

Eliminate the static API key from Helm values by having the daemonset pull it from AWS Secrets Manager via the pod's IAM role — useful in environments where secrets in Helm values are a compliance risk.

yaml

```yaml
serviceAccount:
  annotations:
    eks.amazonaws.com/role-arn: arn:aws:iam::<account-
id>:role/<irsa-role>
env:
  customer: "<your-org>"
  keyless:
    enabled: "true"
    region: "us-east-1"
    secretName: "stepsecurity/orgs/<github-org>/arc-api-key"
    refreshIntervalHours: "3"
```

— Full mechanism, exact Helm keys, IAM permission, and runnable config example.

ARC Harden Runner v2.20.0 (2026-09-23)

02 Custom environment variables for Harden Runner container NEW

67

New `env.customEnvVars` Helm value injects arbitrary environment variables (e.g. `HTTP_PROXY`, `HTTPS_PROXY`, `NO_PROXY`) into the `harden-runner` container, enabling deployments behind a proxy.

Route all Harden Runner agent traffic through a corporate HTTP proxy in air-gapped or proxied EKS clusters.

```yaml
env:
  customEnvVars:
    - name: HTTP_PROXY
      value: "http://proxy.corp.example.com:3128"
    - name: HTTPS_PROXY
      value: "http://proxy.corp.example.com:3128"
    - name: NO_PROXY
      value: "169.254.169.254,10.0.0.0/8"
```

— Names exact Helm key and a runnable proxy config example.

ARC Harden Runner v2.20.0 (2026-09-23)

THINNER COVERAGE BELOW

03 Configurable Tetragon sidecar security context NEW

40

New `tetragon.containerSecurityContext` Helm value allows configuring the security context of the Tetragon sidecar container.

— Names the config key but gives no example or detail on use.

ARC Harden Runner v2.20.0 (2026-09-23)

ARC Harden Runner v2.20.0 (2026-09-23) BLOG ↗ LATEST STABLE

ARC Harden Runner v2.20.0 adds keyless AWS Secrets Manager API key fetch via IRSA, custom agent env vars, and configurable Tetragon security context.

└──▷ USE IT

Eliminate the static API key from Helm values by having the daemonset pull it from AWS Secrets Manager via the pod's IAM role — useful in environments where secrets in Helm values are a compliance risk.

values.yaml

```
serviceAccount:
  annotations:
    eks.amazonaws.com/role-arn: arn:aws:iam::<account-
id>:role/<irsa-role>
env:
  customer: "<your-org>"
  keyless:
    enabled: "true"
    region: "us-east-1"
    secretName: "stepsecurity/orgs/<github-org>/arc-api-key"
    refreshIntervalHours: "3"
```

Route all Harden Runner agent traffic through a corporate HTTP proxy in air-gapped or proxied EKS clusters.

values.yaml

```
env:
  customEnvVars:
    - name: HTTP_PROXY
      value: "http://proxy.corp.example.com:3128"
    - name: HTTPS_PROXY
      value: "http://proxy.corp.example.com:3128"
    - name: NO_PROXY
      value: "169.254.169.254,10.0.0.0/8"
```

Adds `keyless.enabled`, `keyless.region`, `keyless.secretName`, and `keyless.refreshIntervalHours` Helm values to let the Harden Runner daemonset fetch its StepSecurity API key from AWS Secrets Manager via IRSA (pod IAM role) instead of a static Helm value, refreshing on a configurable schedule (default `3` hours); the secret must be a JSON object with an `api_key` field and the IAM role requires `secretsmanager:GetSecretValue`.

Adds `env.customEnvVars` Helm value to inject arbitrary environment variables (e.g. `HTTP_PROXY`, `HTTPS_PROXY`, `NO_PROXY`) into the `harden-runner` container for proxied cluster deployments.

Adds `tetragon.containerSecurityContext` Helm value to configure the Tetragon sidecar's security context.

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube

Feedback ⓘ announced · Sep 22, 2026

BLOG ↗ RANK

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

> SonarQube Cloud adds first-class Rust support in this window, bringing curated Clippy-based rules, complexity metrics, coverage import, duplicate detection, and quality gate enforcement into the same multi-language project view as C/C++, Java, and JavaScript.

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

01 **Rust language support in SonarQube Cloud** NEW 95

SonarQube Cloud now analyzes Rust as a first-class language: it enforces 122 curated Clippy rules through quality profiles across all projects, and adds library- and framework-aware rules covering tokio async handling, serde serialization, and actix-web misuses beyond what Clippy detects. It calculates Cognitive Complexity and Cyclomatic Complexity for Rust code, runs duplicate code detection (CPD) across the codebase, and imports Rust coverage data in LCOV or Cobertura format, correlating it with analysis results in the issue view. Rust projects appear inside the same multi-language project view as C/C++, Java, and JavaScript with no separate dashboard, and support quality gate enforcement with pass/fail thresholds scoped to new code for CI blocking. The Rust analyzer exposes `sonar.rust.cargo.manifestPaths` to point at non-default `Cargo.toml` locations, and `sonar.rust.clippy.enabled` (set to `false` to import an external Clippy JSON report instead of letting SonarQube run `cargo clippy` automatically).

Analyze a Rust workspace where Cargo.toml is not at the repo root — point the scanner at the correct manifest path.

```
$ sonar-scanner -
Dsonar.rust.cargo.manifestPaths=services/myservice/Cargo.toml
```

Use a pre-generated Clippy JSON report (e.g. from a prior CI step) instead of letting SonarQube run cargo clippy itself.

```
$ sonar-scanner -Dsonar.rust.clippy.enabled=false
```

— Names rule count, formats, both config keys, with runnable scanner commands.

launch-20260922-39805485

announced · Sep 22, 2026 BLOG ↗

SonarQube Cloud adds first-class Rust support: 122 Clippy rules, complexity metrics, coverage import, and quality gates.

└──▷ **TRY IT**

Analyze a Rust workspace where Cargo.toml is not at the repo root — point the scanner at the correct manifest path.

```
$ sonar-scanner -
Dsonar.rust.cargo.manifestPaths=services/myservice/Cargo.toml
```

Use a pre-generated Clippy JSON report (e.g. from a prior CI step) instead of letting SonarQube run cargo clippy itself.

```
$ sonar-scanner -Dsonar.rust.clippy.enabled=false
```

Adds `sonar.rust.cargo.manifestPaths` config property to point the Rust analyzer at non-default `Cargo.toml` locations.

Adds `sonar.rust.clippy.enabled` config property (set to `false` to import an external Clippy JSON report instead of running `cargo clippy` automatically).

Supports Rust as a first-class language in SonarQube Cloud, with 122 curated Clippy rules enforced through quality profiles across all projects.

Adds library- and framework-aware Rust rules covering tokio async handling, serde serialization, and actix-web misuses — beyond what Clippy detects.

Adds Cognitive Complexity and Cyclomatic Complexity metric calculation for Rust code.

+4 more

WAS THIS USEFUL? 👍 👎

tirith

Feedback ⓘ v0.4.0 NOTES ↗ RANK

Terminal security for developers and AI agents. Intercepts homograph URLs, pipe-to-shell, ANSI injection, obfuscated payloads, data exfiltration, and malicious AI skills/configs before they execute.

> Tirith 0.4.0 ships a full Python package firewall (`tirith package inspect`, `pkg approve` / `install`, `pkg attest-npm`, and related subcommands), a fail-closed `tirith capsule run --preset untrusted-project` sandbox, Web3 command gating via `web3_guard`, a browser extension integrity auditor, a signed threat-intelligence v2 database, expanded MCP hardening, and AI-agent host support across 19 hosts.

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

01 **Capsule run for untrusted project execution** NEW 93

`tirith capsule run --preset untrusted-project` copies an untrusted project into a held ephemeral directory and executes an exact argv under fail-closed containment on x86_64 Linux with a usable Landlock ABI; all other platforms refuse before copy or spawn.

Run an untrusted open-source project's build script in a fail-closed OS capsule on x86_64 Linux, producing a signed receipt, without risking your host.

```
$ tirith capsule run --preset untrusted-project --project
./untrusted-repo -- python build.py
```

— Full command, platform scope, and a runnable usage example provided. v0.4.0

02 **Package inspect command for wheels and environments** NEW 83

`tirith package inspect` analyzes wheel files, sets of wheels, or installed environments without downloading them, using a bounded streaming wheel reader that rejects path traversal, entry collisions, encrypted entries, conflicting metadata, CRC failures, decompression abuse, and ambiguous identities.

— Names the command and specific rejection mechanisms but gives no example run. v0.4.0

03 **MCP configuration locking** BREAKING 83

MCP configuration locking binds source path, server identity, transport, descriptor set, declared tools, and code/config drift without persisting raw environment values or URL credentials. Existing MCP privacy lockfiles require an explicit re-lock when the

documented schema migration says so, and generated wrappers must be regenerated via `tirith setup <tool> --update-configs` so they contain the current validated absolute interpreter and Tirith paths.

— Names bound fields and the exact migration command to re-lock configs. v0.4.0

04 **AI-agent host integrations expanded to 19 hosts** NEW

80

`tirith setup` now supports 19 named AI-agent hosts, adding Grok Build, OMP, OpenCode, Vercel Labs fx, Prime Agent, Cline, Roo Code, Continue, and OpenHands alongside previously supported hosts. Blocking pre-execution hooks are added for Prime Agent, OMP, Cline, and OpenHands where the host exposes one; Cline receives the platform-native `PreToolUse` wrapper on POSIX and `PreToolUse.ps1` on Windows.

— Names all nine new hosts and hook file names, no exact setup command shown. v0.4.0

05 **Threat intelligence v2 database** IMPROVED

78

Threat intelligence is upgraded to a signed v2 database that adds artifact SHA-256, installed-file hash, malicious-URL, campaign, and behavior-tag indices to package, domain, and IP reputation, with monotonic sequence checks to prevent rollback. An old client must not be pointed directly at a v2 ThreatDB data file; ThreatDB v1 remains available during the v2 cutover.

— Full mechanism and compatibility note, but no command to invoke. v0.4.0

06 **pkg approve and install workflow** NEW

73

`tirith pkg approve` and `tirith pkg install` bind resolution, exact artifact hashes, inspection results, approval state, contained execution, environment verification, and a tamper-evident receipt into one workflow.

— Lists what the workflow binds but gives no example or flag detail. v0.4.0

07 **Deeper wheel and environment analysis: native triage, RECORD verification, startup chains**

NEW

73

Package inspection gains ELF, Mach-O, and PE triage to find native extensions in Python wheels and correlate native imports, runtime launches, loaders, and payloads across distribution boundaries; Wheel RECORD and installed-environment RECORD verification to detect missing, modified, unowned, or multiply owned files; and modeling of Python startup execution through `.pth`, `.start`, `sitecustomize`, and `usercustomize` as an execution chain rather than ordinary data.

— Names three analysis techniques and their targeted files, no direct command example. v0.4.0

08 **Task check for untrusted task envelopes** NEW 73

`tirith task check` and the preview-gated `tirith_check_task` MCP tool (requires `TIRITH_MCP_PREVIEW=1`) diagnostically assess untrusted task envelopes without executing anything.

— Names command, MCP tool, and required env var but no example run. v0.4.0

09 **Browser extension integrity audit** NEW 73

`tirith browser audit` performs a one-shot, read-only integrity audit of Chrome, Chromium, Brave, and Edge extension source trees with signed baselines and typed drift.

— Names command and browsers covered, no example given. v0.4.0

10 **web3_guard policy section for chain operations** NEW 73

The `web3_guard` policy section classifies network and signer policy, covering state changes, signer risk, and network-policy violations for Cast, Forge, Hardhat, Solana, and Anchor operations, without reading chain state or simulating a transaction.

— Names policy section and tool coverage but no config example. v0.4.0

11 **pkg attest-npm for npm provenance binding** NEW 65

`tirith pkg attest-npm` binds npm's own signature/provenance report to the exact lockfile and installed inventory.

— Names command and binding target, no example given. v0.4.0

12 **MCP tool and resource output sanitization** NEW 64

MCP tool and resource output sanitization is on by default, using bounded structured traversal, leaf preservation, cross-leaf detection, strict blob decoding, MIME checks, URI rules, and explicit incomplete-analysis failures.

— Lists mechanism components but no way to configure or invoke it. v0.4.0

13 **CI-focused GitHub Actions security scanning** NEW 64

CI-focused directory scans correlate fork-reachable artifact producers with privileged `workflow_run` consumers and execution, PATH, publish, or deployment sinks. Structural GitHub Actions rules cover mutable action references, privileged trigger hazards, and expression-to-shell injection.

— Names hazard categories and trigger type, no scan command shown. v0.4.0

14 **Byte-classified file analysis** NEW 63

Files are classified by magic bytes read through a retained no-follow handle rather than by name or suffix, with explicit routing for misnamed, extensionless, malformed,

polyglot, archive, native, and unknown binary content.

— Explains mechanism and routing categories, no direct command. v0.4.0

15 **Expanded malicious pattern detections** NEW 63

Detections now include reverse-shell and suspicious inline-interpreter command shapes, OpenAI/Hugging Face/PyPI/PGP credential patterns, hidden PDF instructions, and wallet and sensitive-asset exfiltration through promoted operands (`xargs`, `find -exec`, `parallel`), compressors, encoders, and staged encryption output.

— Names many concrete patterns but no scan invocation shown. v0.4.0

16 **Prime Agent IPython guard** NEW 62

A Prime Agent IPython guard extracts shell escapes, shell magics, and common Python process-spawn forms before delegating to the Tirith engine; unrecognised execution forms fail closed by default.

— Describes mechanism and fail-closed default, no example run. v0.4.0

17 **Command-card schema 2 for Web3 operations** NEW 60

Command-card schema 2 binds ordered Web3 operations, network identity, signer kind, artifact hashes, and policy identity while keeping schema-1 signing bytes compatible.

— Names schema fields but no usage path. v0.4.0

18 **Checkpoint state migration to per-user storage** BREAKING 60

Checkpoints created by older builds under `/tmp/tirith/checkpoints` are not migrated automatically — inspect and remove that directory manually; 0.4.0 uses private per-user state.

— Gives exact path and required manual migration action. v0.4.0

THINNER COVERAGE BELOW

19 **pkg graph, diff, attest, verify-env, and receipt subcommands** NEW 55

`tirith pkg graph`, `diff`, `attest`, `verify-env`, and `receipt` expose ownership, execution, release-delta, provenance, integrity, and evidence workflows for packages.

— Names five subcommands but no mechanism per command. v0.4.0

20 **Global finding and output budgets** NEW 40

Global finding and output budgets preserve Critical, High, and `analysis_incomplete` evidence ahead of lower-value output.

— States the priority rule only, no configuration detail. v0.4.0

└──▷ **BREAKING ON UPGRADE**

! Checkpoints created by older builds under `/tmp/tirith/checkpoints` are not migrated automatically — inspect and remove that directory manually; 0.4.0 uses private per-user state.

! Existing MCP privacy lockfiles require an explicit re-lock when the documented schema migration says so; generated wrappers must be regenerated via `tirith setup <tool> -- update-configs` so they contain the current validated absolute interpreter and Tirith paths.

! An old client must not be pointed directly at a v2 ThreatDB data file; ThreatDB v1 remains available during the v2 cutover.

v0.4.0 NOTES ↗ LATEST STABLE

Tirith 0.4.0 adds a Python package firewall, Web3 command gating, browser extension auditing, threat intelligence v2, and 9 new AI-agent host integrations.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.4.0 https://github.com/sheeki03/tirith.git
# already have the repo? check out this version:
$ git checkout v0.4.0
```

└──▷ **TRY IT**

Run an untrusted open-source project's build script in a fail-closed OS capsule on x86_64 Linux, producing a signed receipt, without risking your host.

```
$ tirith capsule run --preset untrusted-project --project
./untrusted-repo -- python build.py
```

Adds `tirith package inspect` to analyze wheel files, sets of wheels, or installed environments without downloading them, using a bounded streaming wheel reader that rejects traversal, collisions, encrypted entries, conflicting metadata, CRC failures, decompression abuse, and ambiguous identities.

Adds `tirith pkg approve` and `tirith pkg install` to bind resolution, exact artifact hashes, inspection results, approval state, contained execution, environment verification, and a tamper-evident receipt.

Adds `tirith pkg graph`, `diff`, `attest`, `attest-npm`, `verify-env`, and `receipt` subcommands exposing ownership, execution, release-delta, provenance, integrity, and evidence workflows.

Adds `tirith pkg attest-npm` to bind npm's own signature/provenance report to

the exact lockfile and installed inventory.

Adds `tirith capsule run --preset untrusted-project` to copy an untrusted project into a held ephemeral directory and execute an exact argv under fail-closed containment on x86_64 Linux with a usable Landlock ABI; all other platforms refuse before copy or spawn.

+18 more

└──▷ **BREAKING ON UPGRADE**

- ! Checkpoints created by older builds under `/tmp/tirith/checkpoints` are not migrated automatically — inspect and remove that directory manually; 0.4.0 uses private per-user state.
- ! Existing MCP privacy lockfiles require an explicit re-lock when the documented schema migration says so; generated wrappers must be regenerated via `tirith setup <tool> --update-configs` so they contain the current validated absolute interpreter and Tirith paths.
- ! An old client must not be pointed directly at a v2 ThreatDB data file; ThreatDB v1 remains available during the v2 cutover.

WAS THIS USEFUL? 👍 👎

cve-lite-cli

Feedback ⓘ v1.37.0 NOTES ↗ RANK

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

> cve-lite-cli adds an `--output <dir>` flag (and matching GitHub Action input) to route JSON, SARIF, and SBOM report files into a specified directory.

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

01 **Configurable output directory for reports** NEW 88

The new `--output <dir>` flag directs `--json`, `--sarif`, and `--sbom` output files into a specified directory, which is created if it doesn't exist. Passing `--output` without one of those three flags is treated as an error rather than a silent no-op. A matching `output` input was added to the GitHub Action for the same directory-routing behaviour.

Write all structured output to a dedicated reports directory so a CI step can hand a stable path to the next stage without globbing or renaming.

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

— Names exact flag, error behaviour, and GitHub Action input. v1.37.0

v1.37.0 NOTES ↗ LATEST STABLE

cve-lite-cli v1.37.0 adds `--output <dir>` to control where JSON, SARIF, and SBOM files land, with a matching GitHub Action input.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v1.37.0 https://github.com/OWASP/cve-lite-cli.git
# already have the repo? check out this version:
$ git checkout v1.37.0
```

└──▷ **TRY IT**

Write all structured output to a dedicated reports directory so a CI step can hand a stable path to the next stage without globbing or renaming.

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

Adds `--output <dir>` flag to direct `--json`, `--sarif`, and `--sbom` output files into a specified directory (created when missing); passing `--output` without one of those three flags is an error rather than a silent no-op.

Adds a matching `output` input to the GitHub Action for the same directory-routing behaviour.

WAS THIS USEFUL? 👍 👎

sbomify

Feedback ⓘ v26.9.0 NOTES ↗ RANK

sbomify is a product security artifact hub and a trust center.

> sbomify's v26.9.0 release adds a security advisories API, broadens SPDX 3 support across scanning, validation and VEX handling, and rounds out trust center, workspace and compliance features.

└──▷ **WHAT SHIPPED · 13 FEATURES** most completely described first ? what's the number?

01 **SPDX 3 scanning, validation and VEX support** IMPROVED 70

Scans SPDX 3 documents through a derived copy instead of skipping them, sends SPDX documents to Dependency Track through a derived CycloneDX copy, validates SPDX 3.0.1 claims against the vendored schema with clear version-edge messages, and honours VEX statements carried inside an SPDX 3 document about itself.

– Names format versions and integration but not a config surface v26.9.0

02 **Workspace terminology in API and UI** NEW 60

Adds an `is_default_workspace` field alongside `is_default_team` in the API and on-screen copy, and serves a v2 API that uses workspace terminology from the same views.

– Names exact fields and API version but limited mechanism v26.9.0

THINNER COVERAGE BELOW

03 **Pinned CRA audit logger level** IMPROVED 55

Pins the CRA audit logger so a raised `LOG_LEVEL` cannot silence it.

– Names the env var and the fix it addresses v26.9.0

04 **CSAF discovery through security.txt** NEW 50

Exposes public CSAF discovery through `security.txt`.

– Names the discovery mechanism but no further detail v26.9.0

05 **Security advisories API and data model** NEW 45

Adds a security advisories API, including creation and deletion of advisories through the same endpoint, and closes advisory-to-release gaps in the many-to-many model.

– Names API and data-model fix but not exact endpoint path v26.9.0

06 **CISA 2026 minimum elements compliance plugin** NEW 45

Scores CISA 2026 minimum elements and registers a new compliance plugin.

– Names the standard and plugin but no scoring mechanics v26.9.0

07 **Trust center gating and certification badges** NEW 45

Shows gated documents with a request-access gate on the trust center, and adds certification badges on the trust center.

– Names the trust center surface but not exact UI path v26.9.0

08 **Non-expiring OIDC publish tokens** IMPROVED 40

Keeps OIDC publish tokens off the expiry bell so they do not expire.

– Names the token type but no config or flag v26.9.0

09 **Unique document name and version per component** IMPROVED 40

Makes a document's name and version unique per component.

– States the constraint but not the enforcement mechanism v26.9.0

10 **Workspace-wide document toggle from document page** NEW 40

Lets a document be made workspace-wide from its own page.

– Names the exact UI location for the action v26.9.0

11 **Configurable request body upload ceiling** IMPROVED 35

Makes the request body upload ceiling reachable and configurable.

– No config key or limit value given v26.9.0

12 **Workspace plugin management via API** NEW 30

Enables and disables a workspace's plugins over the API.

– Bare one-line capability, no endpoint or flag named v26.9.0

13 **Dedicated vulnerability finding detail row** IMPROVED 25

Gives each vulnerability finding a dedicated detail row.

– Bare description of a UI change v26.9.0

v26.9.0 NOTES ↗ LATEST STABLE

sbomify v26.9.0 adds a security advisories API, SPDX 3 VEX/scanning support, a Kubernetes Helm chart, CSAF discovery via security.txt, and more.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v26.9.0 https://github.com/sbomify/sbomify.git
# already have the repo? check out this version:
$ git checkout v26.9.0
```

Adds a security advisories API, including creation and deletion of advisories through the same endpoint.

Enables and disables a workspace's plugins over the API.

Exposes public CSAF discovery through `security.txt`.

Adds `is_default_workspace` field alongside `is_default_team` in the API and on-screen copy.

Serves a v2 API that uses workspace terminology from the same views.

+14 more

WAS THIS USEFUL? 👍 👎

Cycode Feedback ⓘ

Introducing Cycode Workstation Protection: ADLC Security Now Stops Software Supply Chain

BLOG ↗ RANK

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

> Cycode shipped Workstation Protection, extending its ADLC Security platform upstream to the developer workstation to block malicious and aged-out packages before install.

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

01 **Workstation Protection for package installs** NEW 73

Cycode Workstation Protection intercepts package install commands in real time on developer workstations to enforce cooldown (release-age gating) policies and block known malicious packages—including typosquatting, slopsquatting, and persistent compromised-maintainer releases—before they install, checking every install against a continuously updated threat intelligence feed. It deploys via Mobile Device Management (MDM) with no new console, no new infrastructure, and no change to developer install workflows, and enforces at the device level so gating applies regardless of which package manager version is used or whether a local config file was set up correctly. It unifies this workstation enforcement with existing AI Bill of Materials, prompt guardrails, pipeline integrity, and secret scanning capabilities inside the Context Intelligence Graph.

— Mechanism and scope well described, but no config keys or commands given.

Introducing Cycode Workstation Protecti…

> Introducing Cycode Workstation Protection: ADLC Security Now Stops Software Supply Cha
>
> BLOG ↗ LATEST STABLE
>
> Cycode Workstation Protection adds real-time package install enforcement with cooldown policies and malicious package threat intelligence.
>
> Adds Workstation Protection, a new capability that intercepts package install commands in real time to enforce cooldown (release-age gating) policies and block known malicious packages before they install on developer workstations.
>
> Deploys Workstation Protection via Mobile Device Management (MDM) with no new console, no new infrastructure, and no change to developer install workflows.
>
> Checks every package install against a continuously updated threat intelligence feed to catch malicious packages that have aged past cooldown windows, including

typosquatting, slopsquatting, and persistent compromised maintainer releases.

Enforces cooldown policies at the device level, so gating applies regardless of which package manager version a developer or coding agent is running or whether a local config file was configured correctly.

Extends ADLC Security coverage upstream to the developer workstation, unifying workstation package enforcement with existing AI Bill of Materials, prompt guardrails, pipeline integrity, and secret scanning into the Context Intelligence Graph.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

Truffle Security trufflehog

Feedback ⓘ v3.97.6

NOTES ↗ CODE ↗ RANK

trufflehog discovers, classifies, and validates leaked credentials

> TruffleHog's biggest window additions are on the Enterprise side — a Topology graph for tracing credential blast radius and a Shared Secrets feature for sharing findings outside the platform — alongside open-source additions like custom regex detectors with webhook verification, a `customVerifiers` / `--verifier` mechanism for self-hosted verification endpoints, a local `notifiers` config system, an `analyze` subcommand, and pre-receive hook support for server-side git scanning.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 **Custom regex detectors with webhook verification** NEW

100

Adds a `detectors` field to the config YAML to define custom regex detectors, each with `name`, `keywords`, `regex`, and optional `verify` sections. Verification supports `verify.endpoint`, `verify.unsafe`, and `verify.headers` to validate secrets via a webhook POST, plus `verify.successranges` and `verify.rotatedranges` (HTTP status codes or inclusive ranges like `'200-202'` or `'401'`) to mark secrets live or rotated. `regex` supports named capture groups where every named pattern must match, a `primary regex name` to pick which pattern's line number is reported, `exclude regexes capture` and `exclude regexes match` to strip or suppress matches, an `entropy` parameter to filter low-randomness false positives, and `exclude words` to ignore strings containing listed words.

Define a custom regex detector with webhook verification and HTTP status-based liveness classification — useful when you have internal tokens with a known format and a validation endpoint.

```yaml
detectors:
  - name: hogtokendetector
    keywords:
      - hog
    regex:
      hogid: '\b(hog[0-9a-z]{17})\b'
      hogtoken: '[^a-zA-Z0-9+\/]{0,1}([a-zA-Z0-9+\/]{40})[^a-zA-
Z0-9+\/]{0,1}'
    verify:
      - endpoint: http://localhost:8000/
        unsafe: true
        headers:
          - 'authorization: super secret authorization header'
        successranges:
          - '200'
        rotatedranges:
          - '401'
          - '403'
```

— Every config key named with a complete runnable example product docs

02 **Self-hosted verification endpoints** NEW 90

Adds `customVerifiers` to the root-level config YAML to point secret verification at self-hosted or enterprise endpoints for GitHub Enterprise Server, GitLab, Jira Data Center, Confluence Data Center, Bitbucket Data Center, JFrog Artifactory (access token and reference token variants), and HashiCorp Vault. Adds a `--verifier` CLI flag to the open-source binary to pass custom verifier endpoints at scan time (e.g. `--verifier github=example1.github.com,example2.github.com`) without a config file.

Verify secrets against a private GitHub Enterprise Server and self-hosted GitLab instance by declaring both in the root-level `customVerifiers` block of your config YAML.

```yaml
customVerifiers:
  - name: github
    endpoints:
      - https://github.example.com/api/v3
  - name: gitlab
    endpoints:
      - https://example.gitlab.com
```

Scan the current directory and verify found secrets against two GitHub Enterprise endpoints and one self-hosted GitLab endpoint without editing a config file.

```
$ trufflehog filesystem . --verifier
github=example1.github.com,example2.github.com --verifier
gitlab=example1.gitlab.com
```

— Named config field, CLI flag, and two runnable examples `product docs`

---

03 **Shared Secrets for external finding access** NEW 90

TruffleHog Enterprise adds `Share Secrets`, turning a finding into a secure, access-controlled link where recipients authenticate via Google SSO or SAML SSO against trusted domains and see a read-only view without a platform login. Links support four expiration options — `on read`, `1 hour`, `1 day`, `7 days`, `1 month` — and are managed on a `Shared Secrets` page with `Active Shared Secrets` (revocation) and `Sharing Settings` (trusted domains, default auth) tabs. Trusted domains are set under `Settings > Shared Secrets > Edit Trusted Domains`, and the `sharing authentication` setting picks `Google OAuth` or `SAML SSO` as default. Shared pages show secret type, locator, verification status, first-detected/last-seen timestamps, expiration time, location details, and rotation guidance, without exposing the plaintext credential.

Share a detected credential finding with a developer who needs to rotate it, without granting them access to TruffleHog Enterprise — set up trusted domains first so the link can be created.

> 📍 In TruffleHog Enterprise, go to Settings › Shared Secrets › Edit Trusted Domains and add your org's domain (e.g. `yourorg.com`). Then open a finding's detail page and click `Share` in the upper right corner, select an expiration window (e.g. `1 day`), and copy the generated URL to send to the recipient.

— Exact UI navigation and settings named step by step `product docs`

---

04 **Local notifiers for secret discovery alerts** NEW 90

Adds a `notifiers` field to the config YAML to configure alerts locally with options `name`, `sendunverified`, `sourcestonotify`, `customtext`, `includesecretlink`, `type`, `numworkers`, `trufflehogaddress`, `trufflehogscannergroup`, and `trufflehogscannertoken`. `sourcestonotify` accepts `all`, `sources in this config`, or a comma-separated list. Supports `stdout` (GA, open-source/on-premise), `email` / `jira` / `slack` / `webhook` (GA, cloud and on-premise), and `splunk` (alpha, on-premise) notifier types.

Alert a Slack channel whenever secrets are detected across all sources in your scan config, including a link to the secret.

```yaml
concurrency: '8'
loglevel: info
notifiers:
  - name: slack secrets notifications
    sendunverified: true
    sourcestonotify: sources in this config
    customtext: message to add to notifier
    includesecretlink: true
    type: slack
    numworkers: 16
    trufflehogaddress: https://gnarly-flying-
pancake.c1.prod.trufflehog.org
    trufflehogscannergroup: account-1.us-west-2
    trufflehogscannertoken: thog-agent-xxxxxxxxxxxxxxxxxxxxxxxxx
```

Print secret findings to stdout locally (open-source) without sending alerts to an external service.

```yaml
concurrency: '8'
loglevel: info
notifiers:
  - name: stdout secrets notifications
    sendunverified: true
    sourcestonotify: sources in this config
    type: stdout
    numworkers: 16
```

— Every config field named with two runnable YAML examples `product docs`

05 **Email notifications via AWS SES** NEW 85

TruffleHog Enterprise adds email notification support via AWS SES, configurable in `config.yaml` under `notifiers` with `connection '@type': type.googleapis.com/notifiers.email`, `ses.creds.key`, `ses.creds.secret`, `ses.region`, `recipients`, and `sender` fields. Adds a `ccleaker` option to automatically CC the person who leaked the credential when their email address is available, and supports configuring the notifier through the Integrations tab in the web UI as an alternative to local YAML config.

Route TruffleHog Enterprise secret-leak alerts to your security team via AWS SES, and CC the leaker automatically when their email is known.

```yaml
notifiers:
  - connection:
      '@type': type.googleapis.com/notifiers.email
    ses:
      creds:
        key: AKIAXXXXXXXXXXXXXXX
        secret: XXXXXXXXXXXXXXXXXXXXXXXXXX
      region: us-west-1
    recipients:
      - security@ourbusiness.org
    sender: <trufflehog> trufflehog@ourbusiness.org
    name: email secrets notifications
    type:
      notifier_type: email
    options:
      ccleaker: true
```

— Named config keys with a runnable YAML example product docs

06 **Pre-receive hook mode for self-hosted git servers** NEW

80

TruffleHog can now run as a server-side pre-receive hook on GitHub Enterprise Server, GitLab Self-Managed, and Bitbucket Data Center, blocking pushes without requiring anything installed locally by developers. Adds `--no-update`, `--only-verified`, and `--fail` flags for use in pre-receive hook scripts to control update checks, filter to verified secrets, and fail the push on findings.

Block pushes containing verified secrets on a self-hosted Git server without requiring developers to install anything locally.

```
$ #!/bin/bash
trufflehog git --bare --no-update --only-verified --fail file://
```

— Names exact flags and gives a runnable hook script product docs

07 **Per-detector filtering for AWS accounts and JDBC noise** NEW

80

Adds `detectors.config.aws.allowedAccountIds` and `detectors.config.aws.deniedAccountIds` list fields to the config YAML to restrict or block AWS credential verification by account ID, with denied IDs always taking precedence over allowed IDs. Adds `detectors.config.jdbc.ignorePattern` to suppress unwanted JDBC detector results via regular expressions.

Limit AWS credential verification to trusted production accounts while blocking canary token account IDs from generating alerts.

```yaml
detectors:
  config:
    aws:
      allowedAccountIds:
        - "111111111111"
        - "222222222222"
      deniedAccountIds:
        - "666666666666"
```

Suppress noisy JDBC false positives matching known test or placeholder connection strings.

```yaml
detectors:
  config:
    jdbc:
      ignorePattern:
        - "someP[aa]ttern[0-9]+"
```

— Exact config keys with runnable examples for both product docs

08 **Topology graph for credential blast radius** NEW 75

TruffleHog Enterprise gains Topology, an interactive graph that maps verified live secrets from the source service where they were found to the target services they grant access to. Every service node is categorized as Scanned, Not Monitored, or Non-Scannable; selecting an edge surfaces source/destination service, secret count, secret types, and a link to findings. The graph can be filtered by service category and triage state (refreshing on triage-state filter changes), and selecting an unmonitored scannable service links directly to its configuration page.

— Rich mechanism but UI-only, no exact commands product docs

09 **Analyze subcommand for live secret permissions** NEW 70

Adds a `trufflehog analyze` subcommand with self-discovery analyzers that perform non-state-changing API calls to determine a secret's resources and permissions, plus cloud analyze (enterprise-only) to examine cloud secrets against live environments for IAM posture, currently supporting GCP. The GitHub analyzer's output now also includes an `is_fine_grained` metadata field indicating whether a scanned

token is a fine-grained personal access token, and gist enumeration failures no longer abort the analyzer.

— Subcommand and field named, but no full runnable example `v3.97.6`

10 **Source archiving in Enterprise** NEW 60

TruffleHog Enterprise adds source archiving to stop scanning an integration while retaining all findings it has already reported; archived sources show a gray `archived` badge and are hidden from the integrations list by default, with a `Show archived` filter to surface them dimmed. Findings whose only sources are archived get a `verification frozen` badge, freezing the last known status without re-verifying and excluding them from MTTR calculations. Already-archived sources can be permanently deleted (destroying their findings), and archive actions are recorded in the activity log.

— Explains behavior fully but is UI-only, no exact path `product docs`

v3.97.6 NOTES ↗ CODE ↗ LATEST STABLE

TruffleHog v3.97.6 adds fine-grained token metadata to GitHub analyzer and makes gist enumeration non-fatal.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.97.6
https://github.com/trufflesecurity/trufflehog.git
# already have the repo? check out this version:
$ git checkout v3.97.6
```

Adds `is_fine_grained` metadata field to GitHub analyzer output, surfacing whether a scanned token is a fine-grained personal access token.

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ v0.165.15 NOTES ↗ RANK

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

> Infisical's v0.165.15 release adds credential masking to Privileged Access Management, a UI theme toggle, and removes a previously required stuck-timeout environment variable.

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

01 **Navbar theme toggle** NEW 40

A theme toggle was added to the navigation bar, letting users switch themes directly from the navbar.

— Clear UI location but no further detail v0.165.15

02 **Built-in credential masking in PAM** NEW 35

Infisical now masks privileged access credentials by default so they are not exposed in the UI when using Privileged Access Management (PAM).

— States purpose but no mechanism or exact UI location v0.165.15

03 **Removal of stuck-timeout environment variable** BREAKING 20

The environment variable used as a workaround for stuck timeouts has been removed, eliminating the need for that previously required workaround. The exact variable name is not specified.

— No variable name given, unclear migration impact v0.165.15

v0.165.15 NOTES ↗ LATEST STABLE

Infisical v0.165.15 adds built-in PAM masking, a navbar theme toggle, and a stuck-timeout env var removal.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.165.15
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.165.15
```

> Adds built-in PAM masking to protect privileged access credentials from being exposed in the UI.
>
> Adds a navbar theme toggle so users can switch themes directly from the navigation bar.
>
> Removes the stuck timeout environment variable, eliminating a previously required workaround.

WAS THIS USEFUL?  

◆ Web Application Security

Bright Security

Feedback ⓘ Release of September 22, 2026

NOTES ↗ RANK

Bright is a dynamic application & API security testing platform that offers automation, fast scans, minimal false positives, and comprehensive security testing.

> Bright adds a Fuzzing Mode to STAR Harness for memory-safety bug hunting, three new security tests (Email Spoofing, Source Map Exposure, MCP Tool Inventory Disclosure), automatic batching for scans over 2,000 entry points, and a new Optimizer scan-improvement dashboard, alongside a step-based Advanced Authentication Editor, expanded scan templates, and Scan Configuration API changes to entry point filtering.

└──▷ **WHAT SHIPPED · 12 FEATURES** most completely described first ? what's the number?

01 **entryPointFilter replaces entrypointsStatuses in Scan Configuration API**

DEPRECATED 90

Deprecates the `entrypointsStatuses` property in the Scan Configuration API in favor of the new `entryPointFilter` property, which accepts `securityStatus` (e.g. `'new'`) and `connectivityStatus` (e.g. `'ok'`) arrays for more flexible entry point selection.

Migrate a Scan Configuration API call from the deprecated `entrypointsStatuses` to the new `entryPointFilter` to select only new, reachable entry points.

```json
{
  "entryPointFilter": {
    "securityStatus": ["new"],
    "connectivityStatus": ["ok"]
  }
}
```

— Named API fields with a runnable migration example. Release of September 22, 2026

02 **Fuzzing Mode in STAR Harness** NEW 75

Adds a Fuzzing Mode to STAR Harness that hunts robustness and memory-safety bugs — crashes, hangs, type confusion, unhandled exceptions, unbounded allocation, and algorithmic complexity — in code past input validation, without requiring a scan, authentication, or Repeater.

— Detailed mechanism and bug classes, but no invocation steps given.

Release of September 22, 2026

03 **Enhanced Workflow Run details view** IMPROVED 70

Improved Workflow Run details now show execution progress as executed-out-of-total nodes, a node status breakdown, a warning indicator for problematic nodes, a View Scans action filtered by Workflow Run, and configurable table columns.

— Names UI elements and action, no exact path given. Release of September 22, 2026

04 **GitHub Marketplace and npm packages move to BrightSec org** 70

BREAKING

The GitHub Marketplace Actions and npm packages under the Neuralegion GitHub organization will move to `BrightSec` on October 20; workflows or integrations that reference Neuralegion directly will need to be updated.

— Exact date and org names given, clear migration action. Release of September 22, 2026

05 **Bright Optimizer scan-improvement dashboard** NEW 65

Introduces the Bright Optimizer, a new experience surfacing scan improvement opportunities such as manually triggered scans, low health scores, and long-running scans, configurable via Organization Settings -> General Details -> Optimizer Settings.

— UI path named as clear starting point. Release of September 22, 2026

06 **Task Overview dashboard widgets** NEW 65

Adds four new Task Overview widgets to Organization and Project dashboards — Task Results, Task Triggers, Task Actors, and Task Durations — with support for project, task type, and date-range filters.

— Names widgets and filters, UI-based starting point. Release of September 22, 2026

07 **Step-based Advanced Authentication Editor** IMPROVED 65

Replaces the raw JSON block in the Advanced Authentication Editor with a step-based editor supporting per-step view, edit, add, copy, delete, reorder, JSON validation, and sensitive value masking.

— Explains feature set, no exact navigation path given. Release of September 22, 2026

08 **Three new security tests: Email Spoofing, Source Map Exposure, MCP Tool Inventory Disclosure**

NEW 60

Adds an Email Spoofing test that checks a domain's public DNS records for missing or weak DMARC, SPF, DKIM, and BIMI configuration; a Source Map Exposure test that detects publicly accessible JavaScript source maps (`.js.map`) exposing original source code and internal application structure; and an MCP Tool Inventory Disclosure test that detects when an LLM-powered entry point can be made to reveal the list of MCP tools available to it.

— Named tests with mechanism but no direct config or command.

Release of September 22, 2026

---

09 **New and renamed built-in scan templates** IMPROVED 60

Adds new built-in Scan templates: Web Application , Web Application (without database), API Scan (without database), OWASP Top 10 for Web Apps (2025), and `LLM & GenAI Applications` . Also renames Light Scan to `CI/CD Light Scan` and Passive Scan to Passive Scan (production safe) to better describe their intended use.

— Names all templates and renames, clear UI selection point.

Release of September 22, 2026

---

10 **Extended coverage for ID Enumeration and Full Path Disclosure tests**

IMPROVED 60

Extends the ID Enumeration security test to cover list endpoints that return a JSON array at the root, detecting broken object-level authorization previously missed. Extends the Full Path Disclosure security test to recognize hosted CI/CD build roots including Azure DevOps, GitHub Actions, and Travis.

— Named test extensions, no direct action beyond scanning. Release of September 22, 2026

THINNER COVERAGE BELOW

11 **Automatic batching for scans over 2,000 entry points** IMPROVED 55

Workflows can now scan applications with more than 2,000 Entry Points by automatically dividing discovered entry points into sequential batches of up to 2,000 within a single Workflow Run.

— Describes mechanism, no config surface named. Release of September 22, 2026

---

12 **CVE filter dropdown in Issues tables** IMPROVED 50

Upgrades the CVE filter in Issues tables to a dropdown populated with CVEs detected in the current project, replacing manual text entry.

— UI change described, clear navigation point. Release of September 22, 2026

└──▷ ALSO FROM THESE RELEASES

└──▷ **BREAKING ON UPGRADE**

! The GitHub Marketplace Actions and npm packages under the `Neuralegion` GitHub organization will move to `BrightSec` on October 20; workflows or integrations that reference `Neuralegion` directly will need to be updated.

Release of September 22, 2026 | NOTES ↗ | LATEST | STABLE

Bright adds fuzzing mode to STAR Harness, three new security tests, 2,000+ entry point scanning, and a new Optimizer dashboard.

└──▷ **USE IT**

Migrate a Scan Configuration API call from the deprecated `entrypointsStatuses` to the new `entryPointFilter` to select only new, reachable entry points.

json

```
{
  "entryPointFilter": {
    "securityStatus": ["new"],
    "connectivityStatus": ["ok"]
  }
}
```

Deprecates the `entrypointsStatuses` property in the Scan Configuration API; use the `entryPointFilter` property instead, which accepts `securityStatus` and `connectivityStatus` arrays for more flexible entry point selection.

Adds `entryPointFilter` to the Scan Configuration API, supporting fields `securityStatus` (e.g. `'new'`) and `connectivityStatus` (e.g. `'ok'`) to replace `entrypointsStatuses`.

Adds a new Email Spoofing security test that checks a domain's public DNS records for missing or weak DMARC, SPF, DKIM, and BIMI configuration.

Adds a new Source Map Exposure security test that detects publicly accessible JavaScript source maps (`.js.map`) exposing original source code and internal application structure.

Adds a new MCP Tool Inventory Disclosure security test that detects when an LLM-powered entry point can be made to reveal the list of MCP tools available to it.

+11 more

└──▷ **BREAKING ON UPGRADE**

! The GitHub Marketplace Actions and npm packages under the `Neuralegion` GitHub organization will move to `BrightSec` on October 20; workflows or integrations that reference `Neuralegion` directly will need to be updated.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

ILSpy

Feedback ⓘ v11.1 NOTES ↗ RANK

.NET Decompiler with support for PDB generation, ReadyToRun, Metadata (&more) - cross-platform!

> ILSpy v11.1 introduces two new command-line testing tools for stress-testing the decompiler, extends `ilspycmd` to accept single-file bundles and partial navigation IDs, and adds C# 14 decompilation support alongside several smaller decompiler output improvements.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 **decompdiff mass-testing tool** NEW 78

Adds a `decompdiff` test tool that runs an assembly through two different versions of the decompiler via `ilspycmd` and produces a diff of the generated source code; supports the `--no-build` flag and can run against git refs.

— Names tool, flag and git-ref capability but no full command. v11.1

02 **nugetfuzz mass-testing tool** NEW 70

Adds a `nugetfuzz` test tool that downloads NuGet packages and runs them through a debug build of `ilspycmd`, recording crashes, assertion failures, decompiler warnings, and leaked compiler-generated identifiers in the resulting C# output.

— Mechanism and outputs described, but no exact invocation shown. v11.1

03 **ilspycmd bundle input and partial navigateto IDs** IMPROVED 70

`ilspycmd` now accepts single-file bundles as input (previously unsupported), and accepts member IDs without a full signature in `--navigateto`, e.g. `ilspycmd MyAssembly.dll --navigateto 'M:MyNamespace.MyClass.MyMethod'`.

Navigate directly to a member by ID without needing to supply a full signature — useful when you only have a partial ID from a crash report or reference.

```
$ ilspycmd MyAssembly.dll --navigateto
'M:MyNamespace.MyClass.MyMethod'
```

— Exact flag plus runnable example command shown. v11.1

THINNER COVERAGE BELOW

04 **C# 14 decompilation support** NEW 55

Decompiles C# 14 implicit span conversions (first-class span types) and C# 14 semi-auto properties using the `field` keyword.

— Names the exact C# 14 constructs but no user action needed. v11.1

---

05 **Responsive search and reference navigation** IMPROVED

50

Keeps the UI responsive while searching and drawing reference nodes, enabling search inside packages without stalling and navigation to results.

— Describes behaviour change but no concrete UI path or metric. v11.1

---

06 **UseLambdaSyntax extended to statement lambdas** IMPROVED

45

Extends the `UseLambdaSyntax` decompiler setting to emit statement lambdas.

— Names the setting but not where to toggle it. v11.1

---

07 **C# 7.1 default literal emission** IMPROVED 35

Emits C# 7.1 `default` literals where the surrounding context already supplies the type, instead of more verbose equivalents.

— Names the language feature but gives no before/after example. v11.1

---

08 **Decompiler syntax for ldtoken instruction** NEW 30

Invents a decompiler syntax for the `ldtoken` IL instruction.

— Names the IL instruction but no example of the new syntax. v11.1

---

09 **Nested conditional operator expansion** IMPROVED 30

Expands nested conditional operators back into if-else statements in decompiled output.

— States the transformation without a before/after example. v11.1

---

10 **Obfuscator-renamed local function detection** NEW 25

Detects local functions that have been renamed by obfuscators.

— Single-line description with no mechanism detail. v11.1

---

v11.1 NOTES ↗ LATEST STABLE

ILSpy v11.1 adds C# 14 decompilation, mass-testing tools, and `ilspycmd` bundle support

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v11.1 https://github.com/icsharpcode/ILSpy.git
# already have the repo? check out this version:
$ git checkout v11.1
```

└──▷ **TRY IT**

Navigate directly to a member by ID without needing to supply a full signature — useful when you only have a partial ID from a crash report or reference.

```
$ ilspycmd MyAssembly.dll --navigateto
'M:MyNamespace.MyClass.MyMethod'
```

Adds `nugetfuzz` test tool that downloads NuGet packages and runs them through a debug build of `ilspycmd`, recording crashes, assertion failures, decompiler warnings, and leaked compiler-generated identifiers in C# output.

Adds `decompdiff` test tool that runs an assembly through two different versions of the decompiler via `ilspycmd` and produces a diff of the generated source code; supports `--no-build` and can run against git refs.

Accepts single-file bundles as `ilspycmd` input (previously unsupported).

Accepts member IDs without a signature in `--navigateto` in `ilspycmd`.

Decompiles C# 14 implicit span conversions (first-class span types).

+7 more

WAS THIS USEFUL? 👍 👎

HyperDbg

Feedback ⓘ v0.24 NOTES ↗ RANK

State-of-the-art native debugging tools

> HyperDbg v0.24 expands memory inspection with new word, string and symbol-resolved read commands, adds user-mode I/O port access, extends the script engine with array support, and begins porting HyperDbg's kernel modules to Linux.

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

01 Memory-read commands for words, strings, and symbols NEW

77

Adds `dw` / `!dw` commands to read word (2-byte) values from virtual and physical memory, `da` / `!da` commands to read printable ASCII (null-terminated) strings from virtual and physical memory, and `dds`, `dqs`, and `dps` commands (plus physical memory counterparts) to read double-word, quad-word, and pointer-sized values resolved to symbol names in `module!symbol+offset` format.

Read an ASCII string from a virtual address to inspect null-terminated text in a target process without triggering OS debug APIs.

```
$ da 0xffff8000`deadbeef
```

Read pointer-sized values at a virtual address resolved to symbol names — useful for walking function pointer tables or vtables.

```
$ dps 0xfffff803`12340000
```

— Names every command and the symbol-resolution format, with runnable examples. v0.24

02 Linux port of HyperDbg kernel modules NEW

65

Ports all HyperDbg kernel modules to compile on Linux, including the `hyperhv` and `hyperlog` projects and core platform components (spinlocks, IRQL, process, time, and I/O), and adds GNU assembler (GAS) stub files to replace MASM assembly files for the Linux kernel module build.

— Details scope and build change but no runnable steps yet. v0.24

03 Script engine arrays and semantic test command NEW

60

Adds arrays as function parameters in the script engine, `char` and wide-char (`wchar_t`) array types, and a `test script-semantic` CLI command for running script engine semantic test cases.

— Names the new types and command but gives no usage example. v0.24

v0.24 NOTES ↗ LATEST STABLE

HyperDbg v0.24 adds new memory-read commands, user-mode I/O port access, script engine arrays, and Linux kernel module support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.24 https://github.com/HyperDbg/HyperDbg.git
# already have the repo? check out this version:
$ git checkout v0.24
```

└──▷ TRY IT

Read an ASCII string from a virtual address to inspect null-terminated text in a target process without triggering OS debug APIs.

```
$ da 0xffff8000`deadbeef
```

Read pointer-sized values at a virtual address resolved to symbol names — useful for walking function pointer tables or vtables.

```
$ dps 0xfffff803`12340000
```

Read a byte from an I/O port in user-debugger context to inspect hardware state (e.g. port 0x64, the keyboard controller status port).

```
$ uin 0x64 1
```

Adds `dw` and `!dw` commands to read word (2-byte) values from virtual and physical memory.

Adds `da` and `!da` commands to read printable ASCII (null-terminated) strings from virtual and physical memory.

Adds `dds`, `dqs`, and `dps` commands (plus physical memory counterparts) to read double-word, quad-word, and pointer-sized values resolved to symbol names in `module!symbol+offset` format.

> Adds `uin` and `uout` commands to read from and write to I/O ports (port-mapped I/O) from the user debugger.
>
> Adds `test script-semantic` CLI command for running script engine semantic test cases.
>
> +4 more

WAS THIS USEFUL? 👍 👎

Finite State Platform

Feedback ⓘ

September 22, 2026 - Scan queue API, Version counts, and Component license filters (2026

BLOG ↗ RANK

Finite State Platform is a software supply chain security tool that analyzes firmware and software artifacts for vulnerabilities and compliance risks.

> Finite State shipped a dedicated Scans page for live queue visibility along with new and extended public API endpoints for scan queue health, live snapshots, and richer scan metadata.

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

01 **Scan queue analytics endpoint** NEW 80

Adds `GET /public/v0/scans/analysis` endpoint returning Scan queue totals, success rate, and duration metrics for a selected period, usable for building security operations dashboards or SLA reports.

Pull weekly scan success rate and duration metrics to feed a security operations dashboard or SLA report.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/public/v0/scans/analysis'
```

— Named endpoint with runnable curl example and listed fields

September 22, 2026 - Scan queue API, Ve…

02 **Live scan queue snapshot endpoint** NEW 77

Adds `GET /public/v0/scans/ops-snapshot` endpoint returning the current live running and queued Scan snapshot, suitable for polling from CI pipelines or monitoring scripts to detect backlogs.

Poll the live scan queue from a CI pipeline or monitoring script to detect backlogs before they delay a release.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/public/v0/scans/ops-snapshot'
```

— Named endpoint with runnable curl example September 22, 2026 - Scan queue API, Ve…

03 **Extended scan listing endpoint** IMPROVED 70

Extends `GET /public/v0/scans` to include processing time, Finding count, and start time, and adds the ability to filter and sort by who initiated the Scan.

— Named endpoint and new fields but no usage example

September 22, 2026 - Scan queue API, Ve…

THINNER COVERAGE BELOW

04 **Scans page with live queue visibility** NEW 53

Adds a Scans page showing live Processing and Queued counts, scan history grouped by Project and Version, and volume and duration metrics.

— Describes UI but no exact navigation or config detail

September 22, 2026 - Scan queue API, Ve…

September 22, 2026 - Scan queue API, Version counts, and Component license filters (20

BLOG ↗ LATEST STABLE

Finite State adds a Scans page with live queue visibility and two new API endpoints for scan queue health and snapshots.

└──▷ **TRY IT**

Poll the live scan queue from a CI pipeline or monitoring script to detect backlogs before they delay a release.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/public/v0/scans/ops-snapshot'
```

Pull weekly scan success rate and duration metrics to feed a security operations dashboard or SLA report.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/public/v0/scans/analysis'
```

Adds `GET /public/v0/scans/analysis` endpoint returning Scan queue totals, success rate, and duration metrics for a selected period.

Adds `GET /public/v0/scans/ops-snapshot` endpoint returning the current live running and queued Scan snapshot.

Extends `GET /public/v0/scans` to include processing time, Finding count, start time, and filter/sort by who initiated the Scan.

Adds a Scans page showing live Processing and Queued counts, scan history grouped by Project and Version, and volume and duration metrics.

WAS THIS USEFUL?  

NETWORK SECURITY

◆ Network Analysis & Monitoring

NETRESEC

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ LATEST

STABLE

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

NETRESEC: FlowCarp assigns a single protocol or malware family label per flow using statistical classification (packet sizes, timing, message length distributions, request/response behavior, protocol structure), reducing the multi-alert overlap common with IDS signature matching.

└──▷ TRY IT

Train a custom detection model from sandbox PCAP captures of an unknown C2 protocol so FlowCarp can recognize it in future traffic.

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

Quickly identify C2 protocols in a PCAP during incident response without standing up a local FlowCarp instance.

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

Validate a newly built protocol model against captures from additional sandbox executions to confirm generalization.

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

▾ 1 more release in this issue · 2026-09-23

> changelog · Sep 23, 2026 CHANGELOG ↗ STABLE
>
> NETRESEC: Adds a free demo API endpoint at `https://demo.flowcarp.com` (via `curl --data-binary @<file>.pcap`) that accepts a raw PCAP and returns a JSON response containing detected flows, protocols, and alerts, backed by a commercial license covering 600+ protocols.
>
> └──▷ TRY IT
>
> Train a custom C2 protocol model from sandbox PCAPs of a newly discovered malware family, then classify an unknown capture against it.

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

Quickly triage a suspicious PCAP against 600+ known protocols via the free demo API and filter results with jq during incident response.

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

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Sophos Access Points

Feedback ⓘ

Fusion 2026.41NewReleased on September 22, 2026 NOTES ↗ RANK

Sophos Access Points provide secure network connectivity and VPN access management for remote workers and distributed organizations.

> Sophos Fusion 2026.41 adds a wired Network settings tab for AP6 access point profiles and advanced diagnostic logging for wireless troubleshooting.

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

01 **Network settings tab for AP6 profiles** NEW 80

Access point profiles gain a Network settings tab for configuring wired Ethernet port settings on AP6 access points, including Spanning Tree Protocol, IGMP snooping, LLDP-MED, and, on AP6 840E models, port trunking.

Enable port trunking and IGMP snooping for an AP6 840E access point profile to support multi-VLAN wired uplinks.

> 📍 In Sophos Fusion, open the access point profile for your AP6 840E, go to the Network settings tab, enable IGMP snooping and port trunking, then save the profile.

Configure port trunking and IGMP snooping on an AP6 840E via its access point profile.

> 📍 In Sophos Fusion, open the access point profile for your AP6 840E, navigate to the Network settings tab, and configure Spanning Tree Protocol, IGMP snooping, LLDP-MED, and port trunking as required.

— Names exact protocols and model but no defaults or limits

Fusion 2026.41NewReleased on September … changelog-20260923-b9e48cc9

02 **Advanced diagnostic logging in Fusion** NEW 60

Sophos Fusion adds advanced diagnostic logging, enabled via `Diagnostics > System logs`, to capture optional logs for wireless network troubleshooting.

Capture advanced diagnostic logs to investigate intermittent wireless connectivity issues on a managed AP6 deployment.

> 📍 In Sophos Fusion, navigate to Diagnostics › System logs and turn on advanced diagnostic logging.

Enable advanced diagnostic logging when troubleshooting a wireless network issue with Sophos support.

> 📍 In Sophos Fusion, go to Diagnostics › System logs and turn on advanced diagnostic logging.

— Exact menu path given but mechanism of logs unclear

`Fusion 2026.41NewReleased on September …` `changelog-20260923-b9e48cc9`

`Fusion 2026.41NewReleased on September 22, 2026` **NOTES ↗** **STABLE**

Sophos Fusion 2026.41 adds wired network settings (STP, IGMP, LLDP-MED, port trunking) for AP6 profiles and advanced diagnostic logging.

└──▷ HOW TO FIND IT

Enable port trunking and IGMP snooping for an AP6 840E access point profile to support multi-VLAN wired uplinks.

> 📍 In Sophos Fusion, open the access point profile for your AP6 840E, go to the Network settings tab, enable IGMP snooping and port trunking, then save the profile.

Capture advanced diagnostic logs to investigate intermittent wireless connectivity issues on a managed AP6 deployment.

> 📍 In Sophos Fusion, navigate to Diagnostics › System logs and turn on advanced diagnostic logging.

Adds a `Network settings` tab to access point profiles for configuring wired Ethernet port settings on AP6 access points, including Spanning Tree Protocol, IGMP snooping, LLDP-MED, and (on AP6 840E models) port trunking.

Adds advanced diagnostic logging via `Diagnostics > System logs` in Sophos Fusion to capture optional logs for wireless network troubleshooting.

WAS THIS USEFUL? 👍 👎

Sophos Email Security

Feedback ⓘ

2026.38Released September 21, 2026 NOTES ↗ RANK

Sophos Email Security is an enterprise email protection solution that detects and blocks threats like malware, phishing, and spam.

> Sophos Email Security 2026.38 adds custom protection policies for distribution lists and shared mailboxes, and automates Microsoft 365 journal onboarding.

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

01 **Custom policies for distribution lists and shared mailboxes**

NEW 50

Custom protection policies can now be applied directly to distribution lists and shared mailboxes, independently of the policies applied to individual members.

— Explains behaviour but no config path or limits given

2026.38Released September 21, 2026

02 **Automated Microsoft 365 journal onboarding** NEW 50

Sophos EMS now automates Microsoft 365 journal rule creation and related settings during onboarding, eliminating manual journal configuration once M365 permissions are authorized.

— Describes automation but no exact steps or UI path 2026.38Released September 21, 2026

2026.38Released September 21, 2026 NOTES ↗ LATEST STABLE

Sophos Email 2026.38 adds custom policies for distribution lists and shared mailboxes plus automated Microsoft 365 journal onboarding.

Enables custom protection policies to be applied directly to distribution lists and shared mailboxes, independently of individual member policies.

Automates Microsoft 365 journal rule creation and related settings during Sophos EMS onboarding, eliminating manual journal configuration after M365 permissions are authorized.

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ Adding Bots time series metrics API (2026-09-17)

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 shipped a new Bot Management API for configuring ContentGuard's handling of bot traffic, and separately exposed bot protection metrics through the existing Historical Stats API.

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

01 **Bot Management API for ContentGuard** NEW 95

A new Bot Management API, separate from the Next-Gen WAF API, gives programmatic control over how Fastly's ContentGuard responds to bot traffic. It supports setting an action for a bot category (such as AI crawlers, search engines, or headless browsers) so any bot later added to that category inherits the action automatically; overriding the action for an individual bot within a category, with the ability to set it back to `inherit` to restore the category default; and updating a workspace's protection mode between `off`, `log`, and `block`. It also adds a list-all-bots endpoint returning every bot across every category with each bot's `category_id` (to build a full bot-to-category map in one pass), and a `service_id` filter on the list workspaces endpoint to find the workspace serving a given service.

Find the ContentGuard workspace associated with a specific Fastly service before making policy changes.

```
$ curl -X GET 'https://api.fastly.com/bot-management/workspaces?
service_id=<service_id>' \
  -H 'Fastly-Key: <api_token>'
```

Set a workspace to `block` mode to actively block detected bot traffic rather than only logging it.

```
$ curl -X PATCH 'https://api.fastly.com/bot-
management/workspaces/<workspace_id>' \
  -H 'Fastly-Key: <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"protection_mode": "block"}'
```

Override the action for an individual bot to `log` while its category is set to `block`, then revert it to inherit the category default.

```
$ curl -X PATCH 'https://api.fastly.com/bot-
management/workspaces/<workspace_id>/bots/<bot_id>' \
  -H 'Fastly-Key: <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"action": "inherit"}'
```

— Names endpoints, fields, and modes with runnable curl examples.

Bot Management (2026-09-17)

THINNER COVERAGE BELOW

02 **Bot metrics in Historical Stats API** NEW 35

The Historical Stats API now exposes bot time series metrics served by ContentGuard, enabling per-service bot-protection data retrieval.

— Names the API but gives no endpoint, fields, or example.

Adding Bots time series metrics API (20…

Adding Bots time series metrics API (2026-09-17) BLOG ↗ STABLE

Fastly Historical Stats API now exposes bot time series metrics served by ContentGuard.

Bot Management (2026-09-17) BLOG ↗ LATEST STABLE

Fastly launches a Bot Management API to configure ContentGuard actions per bot category, individual bot, and workspace protection mode.

└──▷ TRY IT

Find the ContentGuard workspace associated with a specific Fastly service before making policy changes.

```
$ curl -X GET 'https://api.fastly.com/bot-management/workspaces?
service_id=<service_id>' \
  -H 'Fastly-Key: <api_token>'
```

Set a workspace to `block` mode to actively block detected bot traffic rather than only logging it.

```
$ curl -X PATCH 'https://api.fastly.com/bot-
management/workspaces/<workspace_id>' \
  -H 'Fastly-Key: <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"protection_mode": "block"}'
```

Override the action for an individual bot to `log` while its category is set to `block`, then revert it to inherit the category default.

```
$ curl -X PATCH 'https://api.fastly.com/bot-
management/workspaces/<workspace_id>/bots/<bot_id>' \
  -H 'Fastly-Key: <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"action": "inherit"}'
```

Introduces the Bot Management API to configure ContentGuard, separate from the Next-Gen WAF API, enabling programmatic control over how Fastly responds to bot traffic.

Supports setting an action for a bot category (such as AI crawlers, search engines, or headless browsers) so that any bot Fastly later adds to that category inherits the action automatically.

Supports overriding the action for an individual bot within a category, with the ability to set the action back to `inherit` to return it to the category default.

Supports updating a workspace to switch between `off`, `log`, and `block` protection modes.

Adds a list-all-bots endpoint that returns every bot across every category in a single request, with each bot reporting its `category_id` so a complete bot-to-category map can be built in one pass.

+1 more

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Chainguard

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗

RANK

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

> Chainguard launched a beta of Sovereign Artifacts for EU data residency and added dozens of new container images to its catalog, while also flagging an upcoming change to the SonarScanner CLI's working directory default.

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

01 SonarScanner working directory relocation

BREAKING

80

Starting October 12, the SonarScanner CLI image changes where it writes its working files, which affects pipelines that read `report-task.txt` from `<project>/.scannerwork`. To preserve the old behavior, pipelines can set `sonar.working.directory=<project>/.scannerwork` when invoking `sonar-scanner`, e.g. `docker run cgr.dev/chainguard/sonar-scanner-cli:latest sonar-scanner -Dsonar.working.directory=<project>/.scannerwork`.

Redirect the SonarScanner working directory to the old location so pipelines that read `report-task.txt` from `<project>/.scannerwork` continue to work after the October 12 change.

```
$ docker run cgr.dev/chainguard/sonar-scanner-cli:latest sonar-scanner -Dsonar.working.directory=<project>/.scannerwork
```

— Gives exact config key, path, and a runnable command to migrate

changelog-20260923-054520ad

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Chainguard launches Sovereign Artifacts beta for EU data residency, storing container and library artifacts in-region for ~30% faster pulls.

└──▷ TRY IT

Redirect the SonarScanner working directory to the old location so pipelines that read `report-task.txt` from `<project>/.scannerwork` continue to work after the October 12 change.

```
$ docker run cgr.dev/chainguard/sonar-scanner-cli:latest sonar-scanner -Dsonar.working.directory=<project>/.scannerwork
```

Adds `openjdk:26` and `ceph:19` container images.

Adds 26 new container images this week, including `clamav-prometheus-exporter`, `crossplane-azure-marketplaceordering`, `kserve-sklearnserver`, `commercial-spilo-fips`, `graphql-hive-gateway`, `jupyterhub-k8s-singleuser-sample`, `nuxeo-2025`, `chainguard-server-hypervisor-aws`, `mongodb-kubernetes`, `opentelemetry-python-instrumentation-fips`, `strimzi-kafka-fips`, `crossplane-aws-emrcontainers`, `kro`, `pyrra`, `starboard-exporter`, `krakend`, `mongodb-search`, and FIPS variants.

Adds `cursor-agent`, `chainguard-desktop-workstation-pc`, `commercial-avocado`, `rustfs`, `cloudzero-agent`, `trivy-k8s-wrapper`, `access-log-exporter`, `datahub-gms`, `kates-connect`, `kubeflow-trainer-controller-manager`, and `tracee` container images.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Plerion

Feedback ⓘ Tag an integration in Plerion: NOTES ↗ RANK

Plerion is a cloud security platform that discovers and manages vulnerabilities and misconfigurations across infrastructure.

> Plerion added role-scoped API access via a new OpenAPI endpoint, introduced integration-level tagging exposed through the API, shipped two new S3 misconfiguration checks, and decoupled SCIM provisioning from just-in-time user creation while escalating the severity of an existing cross-account S3 check.

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

01 New S3 checks for KMS encryption and Multi-Region Access Point

NEW 80

Adds check `PLERION-AWS-1182` : S3 bucket default encryption does not use a customer-managed KMS key — passes buckets encrypted with `aws:kms` or `aws:kms:dsse` under an owned key, fails buckets on AES256, no encryption, or an AWS-managed SSE-KMS key (including the console's AWS managed key (aws/s3)); Medium severity. Adds check `PLERION-AWS-1181` : S3 Multi-Region Access Point does not block public access — passes only when all four Block Public Access flags are enabled on the access point; Medium severity.

— Both check IDs and exact pass/fail logic specified Tag an integration in Plerion:

02 Standalone SCIM provisioning for user management NEW

80

Adds SCIM provisioning: point your identity provider at the SCIM base URL under Admin > Security > Single sign-on > SCIM provisioning and authenticate with a `readWrite` organization API key. The Create users at first sign-in just-in-time provisioning setting now operates independently, with no dependency on SCIM state.

Enable SCIM provisioning so your identity provider can automatically create, update, and deprovision Plerion users.

> 📍 In the Plerion console, go to Admin › Security › Single sign-on › SCIM provisioning, copy the SCIM base URL, then configure your identity provider to authenticate using a readWrite organization API key.

— Exact UI path and required API key role given Tag an integration in Plerion:

03 Per-integration tags with API exposure NEW

75

Adds a Tags section to each integration's settings page where you can define key-value tags that override AWS account tags with the same key — tag groups update within seconds. Every integration returned by the API now carries its tags along with the source of each tag.

Tag an integration via the UI when you cannot set tags in AWS Organizations — the tag immediately overrides the same key on the AWS account side and updates all groups using it.

> 📍 In the Plerion console, open the integration's settings page › scroll to the new Tags section › enter a key (autocompleted from keys already on other integrations) and a value › save.

— UI path and override behaviour clear, exact API field name absent

Tag an integration in Plerion:

---

04 **Severity increase for S3 cross-account access check** BREAKING

60

Raises PLERION-AWS-532 (S3 cross-account access) from Medium to High severity; all open findings on the check move to High on the next scan, and the check is now filed under data at risk, access control, and cross-account categories.

— Named check ID with clear before/after severity and category change

Tag an integration in Plerion:

Tag an integration in Plerion: NOTES ↗ STABLE

Plerion adds per-integration tags (API-accessible), two new S3 checks, and standalone SCIM provisioning

└──▷ **HOW TO FIND IT**

Tag an integration via the UI when you cannot set tags in AWS Organizations — the tag immediately overrides the same key on the AWS account side and updates all groups using it.

> 📍 In the Plerion console, open the integration's settings page › scroll to the new Tags section › enter a key (autocompleted from keys already on other integrations) and a value › save.

Enable SCIM provisioning so your identity provider can automatically create, update, and deprovision Plerion users.

> 📍 In the Plerion console, go to Admin › Security › Single sign-on › SCIM provisioning, copy the SCIM base URL, then configure your identity provider to authenticate using a readWrite organization API key.

Adds a `Tags` section to each integration's settings page where you can define key-value tags that override AWS account tags with the same key — tag groups update within seconds.

Exposes integration tags through the API: every integration returned by the API now carries its tags along with the source of each tag.

Adds new S3 check `PLERION-AWS-1182`: S3 bucket default encryption does not use a customer-managed KMS key — passes buckets encrypted with `aws:kms` or `aws:kms:dsse` under an owned key; fails buckets on AES256, no encryption, or an AWS-managed SSE-KMS key (including the console's `AWS managed key (aws/s3)`); Medium severity.

Adds new S3 check `PLERION-AWS-1181`: S3 Multi-Region Access Point does not block public access — passes only when all four Block Public Access flags are enabled on the access point; Medium severity.

Adds SCIM provisioning: point your identity provider at the SCIM base URL under Admin > Security > Single sign-on > SCIM provisioning and authenticate with a `readWrite` organization API key.

+2 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

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

> Huntress added new macOS detection coverage for quarantine-attribute stripping used by infostealer malware, introduced Screen Sharing session telemetry for macOS endpoints, and exposed a console toggle for per-host Defender scan settings.

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

01 **macOS quarantine-attribute removal detections** NEW

65

Three new macOS detections flag quarantine-attribute removal via xattr strip: a script/shell strip in common malware staging folders, a strip performed on a mounted disk image, and a combined correlation that flags a fake installer followed by a strip within five minutes. The detections cover AMOS, Poseidon, Odyssey, and MacSync infostealers and require Agent v0.14.196.

— Names mechanism, malware families and agent version, but no direct action.

changelog-20260923-a8927daf

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Huntress: Adds three macOS detections for quarantine-attribute removal (xattr strip via system call): script/shell strip in common malware staging folders, strip on a mounted disk image, and a combined fake-installer-plus-strip-within-five-minutes correlation — covering AMOS, Poseidon, Odyssey, and MacSync infostealers; requires Agent v0.14.196.

└──▷ HOW TO FIND IT

Enable Archive File Scanning or disable Network File Scanning for Defender-managed endpoints without touching Huntress defaults on every host.

> 📍 In the Huntress console, go to EDR → Managed Antivirus → Defender Configuration → Scans and toggle Network File Scanning or Archive File Scanning to your desired state, then save.

Adds Screen Sharing session telemetry on macOS — reports source address, session user, and authentication type for each session start, with per-host field suppression; requires Agent v0.14.196.

WAS THIS USEFUL?

Trend Vision One

Feedback

Mobile Security for Business app version 2.0.1651 (2026-09-23) BLOG ↗ RANK

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

> Trend Vision One's Mobile Security for Business app updated to version 2.0.1651, adding platform support and routing improvements.

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

01 **Android platform compatibility update** IMPROVED 25

Mobile Security for Business app version 2.0.1651 adds support for the latest Android platform requirements.

— Only states a bare compatibility claim, no specifics given.

Mobile Security for Business app versio…

> Mobile Security for Business app version 2.0.1651 (2026-09-23) BLOG ↗
>
> LATEST STABLE
>
> Mobile Security for Business 2.0.1651 adds MDM-provided phone number fallback and improved ZTNA/SWG routing.
>
> Adds support for the latest Android platform requirements.

WAS THIS USEFUL?

◆ SOAR

Tines

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

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

> Tines expanded its AI usage API with new queryable dimensions, added built-in functions for converting content to markdown, and fixed cases and dashboards to keep updating in background browser tabs.

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

01 **Background tab updates for cases and dashboards** IMPROVED 28

Cases and dashboards now stay up to date when viewed in background browser tabs.

– Explains behaviour change but no mechanism detail. changelog-20260923-c9c47f7f

02 **Markdown conversion built-in functions** NEW 26

New built-in functions are added for converting content to markdown format.

– Describes capability but not function names or usage. changelog-20260923-c9c47f7f

03 **New dimensions in AI usage API** IMPROVED 25

Tines adds new dimensions to the AI usage API, expanding what usage data can be queried and reported on.

– Names the API area but no specific dimensions or fields. changelog-20260923-c9c47f7f

└──▷ **ALSO FROM THESE RELEASES**

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Tines: Adds new dimensions to the AI usage API, expanding what usage data can be queried and reported on.

Adds new built-in functions for converting content to markdown format.

Cases and dashboards now stay up to date when viewed in background browser tabs.

WAS THIS USEFUL? 👍 👎

Shuffle

Feedback ⓘ v2.3.0-rc2 NOTES ↗ RANK

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

> Shuffle v2.3.0-rc2 adds graceful worker stop handling for AI Agent actions in hybrid deployments.

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

01 **Graceful stop for AI Agent hybrid workers** IMPROVED

35

Adds graceful stop handling for AI Agent actions during hybrid worker execution, preventing runaway processes when an agent-driven workflow terminates.

— Describes behaviour but no config, flag, or command given. v2.3.0-rc2

> v2.3.0-rc2 NOTES ↗ PRE-RELEASE
>
> Shuffle v2.3.0-rc2 adds graceful worker stop for AI Agent actions in hybrid deployments.
>
> └──▷ **GET THIS VERSION** ▾
>
> ```
> $ git clone --branch v2.3.0-rc2
> https://github.com/Shuffle/Shuffle.git
> # already have the repo? check out this version:
> $ git checkout v2.3.0-rc2
> ```

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Datadog Stratus Red Team

Feedback ⓘ v2.36.1 NOTES ↗

CODE ↗ RANK

Granular, Actionable Adversary Emulation for the Cloud

> Stratus Red Team v2.36.1 focuses on reliability for shared and automated environments, adding concurrent detonation support across all cloud providers and improving Azure API traceability while dropping the Azure CLI dependency.

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

01 Concurrent detonation support across all providers NEW

75

Enables concurrent detonation across AWS, Azure, GCP, Entra ID, Kubernetes and EKS by embedding a short correlation ID (`var.correlation.short`) in Terraform-managed resource names, preventing resource collisions when multiple users share an environment. Also adds an opt-out of the Terraform cache to prevent concurrency issues during parallel detonations.

– Names the exact Terraform variable and cache opt-out mechanism v2.36.1

02 Azure correlation ID propagation via clientRequestID IMPROVED

60

Propagates Azure correlation IDs via `clientRequestID` (migrated from the previous user-agent approach) so detonation activity is traceable across Azure API calls.

– Names the field and prior mechanism but no usage steps v2.36.1

THINNER COVERAGE BELOW

03 Azure CLI dependency removed via provider bump IMPROVED

55

Removes dependency on the `az` CLI by bumping the Azure (`azurerm`) Terraform provider to 4.81.0.

– Names exact provider and version but limited detail on impact v2.36.1

v2.36.1 NOTES ↗ CODE ↗ LATEST STABLE

Stratus Red Team v2.36.1 adds concurrent detonation support across all providers, KUBECONFIG multi-path handling, and Azure correlation ID propagation.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.36.1 https://github.com/DataDog/stratus-red-team.git
# already have the repo? check out this version:
$ git checkout v2.36.1
```

Propagates Azure correlation IDs via `clientRequestID` (migrated from user-agent) so detonation activity is traceable across Azure API calls.

Enables concurrent detonation across all providers (AWS, Azure, GCP, Entra ID, Kubernetes, EKS) by embedding a short correlation ID (`var.correlation.short`) in Terraform-managed resource names, preventing resource collisions when multiple users share an environment.

Removes dependency on the `az` CLI by bumping the Azure (`azurerm`) Terraform provider to 4.81.0.

Adds opt-out of the Terraform cache to prevent concurrency issues during parallel detonations.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Invite OpenSharing recipients by email (Beta) (2026-09-22) BLOG ↗ RANK

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

> Databricks expanded Unity Gateway with new model access, credential security and pricing controls, added a Shopify connector to Lakeflow Connect, opened a Beta backup schedule API for Lakebase, and let OpenSharing recipients be invited by email.

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

01 **Shopify connector in Lakeflow Connect** NEW 75

A new Beta Shopify connector in Lakeflow Connect ingests products, orders, customers, inventory, fulfillments, and discounts from Shopify into Databricks, authenticating via the Shopify OAuth client credentials grant. Workspace admins enable it from the Previews page.

— Names data entities, auth method and exact enablement path

Shopify connector (Beta) (2026-09-22)

02 **Lakebase backup schedule API** NEW 75

A new Beta API for Lakebase lets you read and set a root branch's automated snapshot cadences (daily, weekly, or monthly) and configure per-cadence retention periods.

— Names cadence options and retention mechanism but not endpoint path

Lakebase backup schedule API (Beta) (20…

THINNER COVERAGE BELOW

03 **Unity Catalog secret-backed model provider credentials**

IMPROVED 55

A model provider service credential in Unity Gateway can now be backed by a customer-owned Unity Catalog secret, so rotating the secret takes effect at query time with no service change.

— Explains zero-downtime rotation mechanism, no config path given

Back model provider service credentials…

04 **Invite OpenSharing recipients by email** NEW 50

OpenSharing now supports inviting recipients by email, sending a secure activation link instead of requiring manual exchange of a sharing identifier or credential file. This is a

Beta feature.

— Describes before/after but no concrete UI or API detail

Invite OpenSharing recipients by email …

05 **Provider discount multiplier for external model spend** NEW

50

Unity Gateway adds a price multiplier setting on model provider services so external model spend estimates reflect negotiated rates instead of published list prices.

— Names the setting but not its config location Configure provider discounts for extern…

06 **OpenAI GPT-6 Luna on Unity Gateway** NEW

45

OpenAI GPT-6 Luna is now accessible through Unity Gateway via Foundation Model APIs on a pay-per-token basis.

— Names model and pricing basis but little mechanism detail

OpenAI GPT-6 Luna is now available on U…

> Invite OpenSharing recipients by email (Beta) (2026-09-22) BLOG ↗ STABLE
>
> Databricks OpenSharing now supports inviting recipients by email with a secure activation link (Beta).
>
> Adds email-based invitation for OpenSharing recipients, sending a secure activation link instead of requiring manual exchange of a sharing identifier or credential file (Beta).

> Lakebase backup schedule API (Beta) (2026-09-22) BLOG ↗ STABLE
>
> Databricks adds a Beta API to programmatically manage Lakebase backup schedules with configurable cadence and retention.
>
> Adds a Beta snapshot schedule API for Lakebase that lets you read and set a root branch's automated snapshot cadences (daily, weekly, or monthly) and per-cadence retention periods.

> OpenAI GPT-6 Luna is now available on Unity Gateway (2026-09-22) BLOG ↗ STABLE
>
> OpenAI GPT-6 Luna is now accessible via Databricks Unity Gateway through Foundation Model APIs pay-per-token.

> Shopify connector (Beta) (2026-09-22) BLOG ↗ LATEST STABLE
>
> Databricks Lakeflow Connect adds a Shopify connector (Beta) for ingesting e-commerce data via OAuth.

Back model provider service credentials with a Unity Catalog secret (2026-09-21)

BLOG ↗ STABLE

Databricks lets you back model provider API keys with customer-owned Unity Catalog secrets, enabling zero-downtime key rotation.

Backs a model provider service credential with a customer-owned Unity Catalog secret so that rotating the secret takes effect at query time with no service change.

Configure provider discounts for external model spend (2026-09-19) BLOG ↗

STABLE

Unity Gateway lets you configure a price multiplier on model provider services to reflect negotiated rates for external model spend.

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ

Search Job API Architectural Improvements (Search) (2026-09-21) BLOG ↗ RANK

Sumo Logic is a cloud-based platform for monitoring, logging, and analyzing security and operational data from applications and infrastructure.

> Sumo Logic overhauled its Search Job Management API to be OpenAPI-compliant, stateless, and capable of cross-org automation, while removing session-cookie authentication in a breaking change.

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

01 Stateless, cookie-free Search Job Management API

BREAKING

67

The Search Job Management API no longer uses session cookies; any client that stored or sent cookies to authenticate requests must migrate to the `jobId -based` stateless request model. This lets long-running search jobs be polled without session cookies, making them compatible with AI agents and Model Context Protocol (MCP) frameworks.

— Names exact mechanism and migration requirement but no code sample

Search Job API Architectural Improvemen…

THINNER COVERAGE BELOW

02 OpenAPI-compliant Search Job Management API

IMPROVED

40

The Search Job Management API is now fully OpenAPI-compliant, enabling auto-generated SDKs in any language and native integration with Postman and CI/CD pipelines.

— States capability but no endpoint names or spec link

Search Job API Architectural Improvemen…

03 Parent-child and cross-org search automation

NEW

37

MSSPs and enterprise admins can now automate parent-child and cross-org searches via the Search Job Management API, a capability previously limited to the UI.

— Describes what changed but no config or call details

Search Job API Architectural Improvemen…

└──▷ **BREAKING ON UPGRADE**

! The Search Job Management API no longer uses session cookies — any client that stored or sent cookies to authenticate requests must migrate to the `jobId -based` stateless request model.

Search Job API Architectural Improvements (Search) (2026-09-21) BLOG ↗

LATEST STABLE

Search Job Management API goes OpenAPI-compliant, stateless, and adds parent-child org support for MSSP automation.

Search Job Management API is now fully OpenAPI-compliant, enabling auto-generated SDKs in any language and native integration with Postman and CI/CD pipelines.

Long-running search jobs can now be polled via the Search Job Management API without session cookies, making them compatible with AI agents and Model Context Protocol (MCP) frameworks.

MSSPs and enterprise admins can now automate parent-child and cross-org searches via the Search Job Management API, a capability previously limited to the UI.

└──▷ **BREAKING ON UPGRADE**

! The Search Job Management API no longer uses session cookies — any client that stored or sent cookies to authenticate requests must migrate to the `jobId-based` stateless request model.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

runZero is an asset discovery and inventory platform that identifies and maps all devices on networks without requiring agents.

> runZero's biggest update this window is OCSF 1.9.0 export support for SIEM and data lake ingestion, alongside finer-grained saved query permissions, automatic retirement of stale integration data, and broader cloud import coverage.

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

01 Custom role permissions for saved queries IMPROVED

60

Custom roles granting Saved queries read access can now view the Queries page and table without needing inventory access; roles with write or manage permissions can create, edit, copy, import, and delete saved queries.

– Clear permission behavior but no exact role UI path given changelog-20260923-d51cf482

THINNER COVERAGE BELOW

02 Automatic retirement of stale integration data IMPROVED

30

Software and Vulnerabilities collected from integrations are now retired when that integration's underlying data expires, keeping inventory data current with the source system.

– Describes behavior but no mechanism or config surface changelog-20260923-d51cf482

03 Legacy MAC prefix vendor lookup IMPROVED

20

runZero now populates MAC vendor information for legacy MAC prefixes, improving vendor identification for older devices.

– Bare statement with no scope or numbers changelog-20260923-d51cf482

04 Expanded AWS, Azure, and GCP asset types IMPROVED

20

AWS, Azure, and GCP imports now include additional asset types, though the specific new types are not named.

– No specific asset types or endpoints named changelog-20260923-d51cf482

> changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE
>
> runZero: Custom roles granting Saved queries read can now view the Queries page and table without inventory access; write or manage roles can create, edit, copy,

import, and delete saved queries.

└──▷ **HOW TO FIND IT**

Export the full asset inventory in OCSF 1.9.0 format for ingestion into a SIEM or data lake that understands the Open Cybersecurity Schema Framework.

> 📍 In the runZero console, go to Inventory › Export and select `OCSF 1.9.0`; assets appear as Device Inventory Info events and services as Live Evidence Info events. Alternatively, call the Export API endpoint to automate the pull.

Software and Vulnerabilities collected from integrations are now retired when that integration data expires.

runZero now populates MAC vendor information for legacy MAC prefixes.

AWS, Azure, and GCP imports now include additional asset types.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Aembit

Feedback ⓘ announced · Sep 23, 2026 PRESS ↗ RANK

Aembit is IAM for agentic AI and workloads. Enable secretless access and real-time policy enforcement across any environment to protect data and accelerate AI adoption.

> Aembit shipped a new Enterprise-Managed Authorization Credential Provider that brings Okta Cross App Access (XAA) support to AI agents, letting them reach MCP servers without per-app OAuth consent while falling back to standard OAuth where XAA isn't yet supported.

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

01 Enterprise-Managed Authorization Credential Provider for Okta XAA

NEW 72

Aembit adds an Enterprise-Managed Authorization Credential Provider implementing the Model Context Protocol Working Group's Enterprise-Managed Authorization extension, enabling Okta Cross App Access (XAA) for AI agents without per-service consent steps. It captures a user's SSO ID token and refresh token (requiring the `offline_access` scope) at login, refreshes them in the background, and performs a dual token exchange at request time so agents can reach downstream MCP servers without per-app OAuth consent screens. Aembit's Edge Controller brokers XAA between identity providers and upstream MCP servers that don't yet support it natively, generating and signing assertions using the tenant's own keys when the IdP can't perform the exchange directly, and automatically falls back to the standard OAuth Authorization Code flow for any upstream MCP server or IdP without XAA support, requiring no separate configuration. Policy-based access enforcement and auditability extend to these XAA-driven agent connections alongside existing OAuth capabilities, giving a common enforcement point across both, and letting organizations adopt XAA incrementally.

— Mechanism and named surfaces detailed, but no CLI/UI steps given

launch-20260923-77fa1fb6 How Aembit Makes Cross App Access Usabl…

> announced · Sep 23, 2026 PRESS ↗
>
> Aembit adds Okta Cross App Access support via a new Enterprise-Managed Authorization Credential Provider for AI agent identity.
>
> Adds an Enterprise-Managed Authorization Credential Provider that implements the Model Context Protocol Working Group's Enterprise-Managed Authorization extension, enabling Okta Cross App Access (XAA) for AI agents without per-service consent steps.

Extends policy-based access enforcement and auditability to Okta XAA-driven agent connections, letting security teams centrally control which agents act on whose behalf and under what policy.

Supports XAA alongside existing OAuth authorization capabilities, allowing organizations to adopt Cross App Access incrementally as ecosystem support expands rather than requiring an all-or-nothing migration.

Provides a common enforcement point for AI agent access to MCP servers, tools, and enterprise resources across both XAA and OAuth environments.

How Aembit Makes Cross App Access Usable Before the Rest of the Stack Catches Up (2026

BLOG ↗ LATEST STABLE

Aembit launches Enterprise-Managed Authorization Credential Provider to broker XAA/EMA for AI agents against SaaS MCP servers

Adds a new Enterprise-Managed Authorization (EMA) Credential Provider that captures a user's SSO ID token and refresh token (requiring the `offline_access` scope) at login, refreshes them in the background, and performs a dual token exchange at request time so AI agents can access downstream MCP servers without per-app OAuth consent screens.

Aembit's Edge Controller brokers XAA between identity providers and upstream MCP servers that don't yet support XAA natively — generating and signing assertions using the tenant's own keys when the IdP cannot perform the exchange directly.

Automatically falls back to the standard OAuth Authorization Code flow for any upstream MCP server or identity provider that doesn't support XAA, requiring no separate configuration.

WAS THIS USEFUL?  

Okta for AI Agents

Feedback ⓘ changelog · Sep 23, 2026

CHANGELOG ↗

Okta for AI Agents provides identity and access management capabilities enabling secure authentication and authorization for AI-powered applications and automated systems.

└──▷ **HOW TO FIND IT**

Block access attempts originating from known malware proxy networks at the global session policy level.

> 📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add the `MALWARE_PROXIES` IP service category to the zone, then reference that zone in Security › Global Session Policy to deny access for matching traffic.

Prevent analysts connecting via remote desktop from accessing sensitive apps, without affecting direct corporate network logins.

> 📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add the `REMOTE_DESKTOP` IP service category, then add a deny rule for that zone in the relevant app sign-in policy under Security › Authentication Policies.

└──▷ **BREAKING ON UPGRADE**

! Okta On-prem SCIM Server agent is replaced by Okta On-prem SCIM agent; existing deployments referencing the old agent name must migrate to the new agent.

WAS THIS USEFUL? 👍 👎

iDenfy

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗

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

└──▷ **TRY IT**

Retrieve only specific fields from a KYB company record to reduce payload size in automated compliance pipelines.

```
$ curl -X GET 'https://api.idenfy.com/kyb/company?
fields=name,registrationNumber,directors' \
  -H 'Authorization: Bearer <token>'
```

└──▷ **BREAKING ON UPGRADE**

! The iOS SDK minimum deployment target is raised to iOS 15.

WAS THIS USEFUL? 👍 👎

Didit

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

> Didit extended its camera-control options for identity verification, letting integrators lock or hide camera switching for liveness and document capture both through the SDK and via direct verification URLs.

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

01 **Camera control parameters for liveness and document capture**

IMPROVED 88

Didit added `liveness_camera` and `liveness_camera_switch` query parameters (and equivalent document parameters) for direct verification URLs, mirroring the SDK options `defaultLivenessCamera`, `showLivenessCameraSwitchButton`, `defaultDocumentCamera`, and `showDocumentCameraSwitchButton`. This lets a kiosk or operator device be locked to a specific camera (e.g. rear) and hide the switch button, whether integrating via the SDK or opening a verification URL like `https://verify.didit.me/session/<session_id>?liveness_camera=back&liveness_camera_switch=false` directly.

Lock a kiosk or operator device to the rear camera for liveness capture and hide the switcher so the subject cannot flip to the front lens.

javascript

```
didit.init({
  defaultLivenessCamera: 'back',
  showLivenessCameraSwitchButton: false,
  defaultDocumentCamera: 'back',
  showDocumentCameraSwitchButton: false
});
```

Open the verification URL directly (without the SDK) and force rear-camera liveness with the switcher hidden.

```
$ https://verify.didit.me/session/<session_id>?
liveness_camera=back&liveness_camera_switch=false
```

— Names exact SDK options, query params, and a runnable URL example.

changelog-20260923-32287544

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Didit: Adds `liveness_camera` and `liveness_camera_switch` (and equivalent document) query parameters for direct verification URLs, enabling the same camera control without going through the SDK.

└──▷ **USE IT**

Lock a kiosk or operator device to the rear camera for liveness capture and hide the switcher so the subject cannot flip to the front lens.

javascript

```
didit.init({
  defaultLivenessCamera: 'back',
  showLivenessCameraSwitchButton: false,
  defaultDocumentCamera: 'back',
  showDocumentCameraSwitchButton: false
});
```

Open the verification URL directly (without the SDK) and force rear-camera liveness with the switcher hidden.

```
$ https://verify.didit.me/session/<session_id>?
liveness_camera=back&liveness_camera_switch=false
```

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗

RANK

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

> ConductorOne shipped three breaking API changes around revoke matching, secret deletion, and spend-remedy status, alongside new console features for connector change history, saved C1AI reports, role mining, an Identities & NHI overview, and entitlement access graphs.

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

01 **Saved C1AI reports in Reporting** NEW 72

The C1AI Report assistant under Platform > Reporting can now save a natural-language reporting query as a named, described report for recurring use. A saved report's detail page shows its creation date, creator, last successful run, status, and a donut chart with a data table.

Save a C1AI report for recurring use by your access-review team.

> 📍 In the C1 console, go to Platform › Reporting, open the Report assistant, ask your reporting question, then select Save report and enter a name and description.

— Exact navigation and save steps given, with resulting report fields.

changelog-20260923-e7a9b8d7

02 **Connector configuration change history** NEW 60

A new Change history icon on a connector's detail page opens a per-field before-and-after timeline, letting admins audit who changed a connector's configuration and what changed.

Navigate to the Change history for a connector to audit who changed its configuration and what changed.

> 📍 In the C1 console, go to Connectors › [select a connector] › detail page, then click the Change history icon to view the per-field before-and-after timeline.

— Exact UI path given but no mechanism beyond the timeline view.

changelog-20260923-e7a9b8d7

THINNER COVERAGE BELOW

03 **Role mining shared-access filtering** NEW 35

Role mining gained a Filter section supporting multiple filters and a Select shared access section with a threshold slider, histogram, and summary, per the updated UI.

— Only a screenshot describes the UI; no navigation or mechanism given.

changelog-20260923-e7a9b8d7

04 **Identities & NHI overview page** NEW 35

A new overview page for Identities & NHI shows tabs and stat tiles for human users, service accounts, secrets, agents, and NHI counts.

— Described only from a screenshot; no path or interaction detail.

changelog-20260923-e7a9b8d7

05 **Access graph for requested entitlements** NEW 35

The access graph for a requested entitlement now displays its granted repositories, highlighting the searched-for repository as a match.

— Only a screenshot describes the feature; no steps to reach it.

changelog-20260923-e7a9b8d7

└──▷ **BREAKING ON UPGRADE**

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

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

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

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

└──▷ **HOW TO FIND IT**

Navigate to the Change history for a connector to audit who changed its configuration and what changed.

📍 In the C1 console, go to Connectors › [select a connector] › detail page, then click the Change history icon to view the per-field before-and-after timeline.

Save a C1AI report for recurring use by your access-review team.

📍 In the C1 console, go to Platform › Reporting, open the Report assistant, ask your reporting question, then select `Save report` and enter a name and description.

└──▷ **BREAKING ON UPGRADE**

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

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

The easiest way to add authentication and user management to your application. Purpose-built for React, Next.js, Remix, and “The Modern Web”.

> Clerk added SMS delivery event logging with five typed events for debugging message delivery, and introduced CIMD support letting MCP clients authenticate via an HTTPS client_id without pre-registration.

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

01 **CIMD support for MCP client authentication** NEW 80

Adds Client ID Metadata Document (CIMD) support so MCP clients can authenticate using an HTTPS URL as their `client_id` without pre-registration or a client secret. Enabled via Publish CIMD support under OAuth applications settings > Client onboarding in the Clerk Dashboard; an optional Client admission setting can restrict this to Pre-registered clients only, with trusted clients added from the Applications tab.

Enable CIMD support so MCP clients can authenticate using an HTTPS URL as their `client_id` without pre-registration or a client secret.

> 📍 1. In the Clerk Dashboard, go to OAuth applications settings. 2. Under Client onboarding, enable Publish CIMD support. 3. Optionally set Client admission to Pre-registered clients only and add trusted clients from the Applications tab.

— Exact dashboard settings and toggle names given as steps changelog-20260923-3af810e3

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Clerk: Adds SMS delivery event logging with five typed events — `sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, and `sms.unconfirmed` — available in the SMS logs view under Application Logs on all plans.

└──▷ **HOW TO FIND IT**

Enable CIMD support so MCP clients can authenticate using an HTTPS URL as their `client_id` without pre-registration or a client secret.

> 📍 1. In the Clerk Dashboard, go to OAuth applications settings. 2. Under Client onboarding, enable Publish CIMD support. 3. Optionally set Client admission to Pre-registered clients only and add trusted clients from the Applications tab.

Investigate a user-reported missing SMS verification code by filtering delivery events to that recipient's number.

📍 1. In the Clerk Dashboard, open the SMS logs view. 2. Filter by the destination phone number or by event type `sms.failed` or `sms.undeliverable`. 3. Select a message to see its full delivery trace and the normalized `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`).

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ Release Notes 2022 (2026-09-23) 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 JumpCloud Access Requests with JIRA webhook integration and Directory Insights tracking, GA'd SaaS License Management alongside a new SaaS Management shadow-IT discovery tool, added AI Agent identity management with configurable OAuth client authentication types and a private AI Assistant knowledge base, launched a Case Portal later extended to MTP-wide case management, rolled out five new Android device policies, and deprecated the reactivate-user GET endpoint plus the Dynamic Groups API `$le` / `$ge` operators.

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

01 **AI Agent identity management** NEW 95

Adds Identity Management > Agents for registering AI agent identities with OAuth client credentials, independent of human accounts: admins set the agent name, owner, and company email, generate OAuth client credentials, and choose a client authentication type of Client secret basic or Client secret post (the latter for platforms like Amazon Bedrock AgentCore that require credentials in the request body); agents are added to Agent Groups, assigned SSO applications, and authorized to MCP servers via the agent's Applications tab.

Register and configure a new AI agent identity with OAuth credentials so it can access AI Gateway MCP servers independently of any human account.

> 📍 In the Admin Portal: go to Identity Management › Agents and click + Add. Enter the agent name, owner, and company email on the Details tab. On the Authentication tab, generate OAuth client credentials and set the client authentication type to Client secret post if your external platform (e.g. Amazon Bedrock AgentCore) sends credentials in the request body. Under Identity Management › Agent Groups, create an Agent Group, add the agent, and assign SSO applications. On the agent's Applications tab, use Manage connections to authorize MCP servers on behalf of the agent.

Configure an agent to send OAuth credentials in the request body when the external platform (e.g. Amazon Bedrock AgentCore) requires `client_secret_post` rather than the default header-based method.

> 📍 In the JumpCloud Admin Portal, open the agent record, go to the Authentication tab, and under Client Authentication Type select Client secret post

instead of the default `Client secret basic`.

— Gives exact console path, OAuth auth-type options, and full setup flow.

Release Notes 2026 (2026-09-23) changelog-20260923-3d5122e5

02 **JumpCloud Access Requests** NEW 80

Introduces JumpCloud Access Requests, an approval flow for access requests submitted from the JumpCloud User Portal and managed in the admin console; supports webhook configuration to integrate with ticketing systems such as JIRA, automatically creating tickets and sending real-time notifications on new requests, and tracks Access Request events in Directory Insights. Approval flows are configured in the Admin Portal under Access Requests > Approval Flows, where a webhook can be pointed at a JIRA project URL to receive new-request events.

Create a privileged-access approval flow tied to a JIRA project so that every access request auto-creates a ticket and notifies approvers in real time.

📍 In the JumpCloud Admin Portal, go to Access Requests › Approval Flows, create a new flow, assign approver admins, then under Integrations configure a webhook pointing to your JIRA project URL to receive new-request events.

— Names JIRA webhook, Directory Insights tracking, and exact console path.

Release Notes 2025 (2026-09-23)

03 **Private Repository Windows EXE install validation** NEW 70

Windows EXE updates deployed via the Private Repository now require an install success exit code so JumpCloud can confirm installation succeeded; admins update a version through the application's Details tab using `Update Version` and the Upload File flow, entering the exit code before deploying to associated devices.

Push an updated application binary to all associated devices without rebuilding the app record or reassigning device groups.

📍 In the JumpCloud Admin Portal, go to Software Management › Private Repository, open the application's Details tab, click `Update Version`, upload the new file via the Upload File flow, and (for EXE) enter the install success exit code before deploying.

— Requires the exact new setting with the exact deploy workflow steps.

changelog-20260923-3d5122e5

04 **macOS Ventura upgrade delay/block policy** NEW 65

Adds a `Delay or Block MacOS Ventura` macOS policy, configured under Device Management > Policies in the Admin Portal, to set an enforcement window that delays or blocks the Ventura upgrade across managed macOS devices.

Block macOS Ventura upgrades fleet-wide while you validate compatibility — configure the delay/block policy before devices auto-upgrade.

> 📍 In the Admin Portal, go to Device Management › Policies, create a new macOS policy, and select `Delay or Block MacOS Ventura` to set the enforcement window across your managed macOS devices.

— Exact policy name and console path given, ready to configure.

Release Notes 2022 (2026-09-23)

---

05 **Unified MFA user group exclusions** NEW 65

JumpCloud Protect's Unified MFA supports User Group Exclusions, configured under Security > MFA > JumpCloud Protect > Unified MFA, allowing a user group to be excluded from Push MFA enforcement so members can still authenticate via TOTP-only.

Exclude a remote-region user group from Push MFA enforcement so those users can still authenticate via TOTP-only without disrupting their JumpCloud Protect experience.

> 📍 In the JumpCloud Admin Portal, go to Security › MFA › JumpCloud Protect › Unified MFA, open User Group Exclusions, and add the target group to the exclusion list.

— Exact console path and fallback authentication method named.

Release Notes 2025 (2026-09-23)

---

06 **Push Bombing Mitigation review workflow** NEW 60

Push Bombing Mitigation, configured under Security > MFA Configurations, blocks more than 1 push per resource per 60 seconds; admins can confirm it is enabled and then review blocked fraudulent push attempts by filtering blocked push events in Directory Insights.

When you suspect push-bombing attacks against admin accounts, verify that Push Bombing Mitigation is active and review blocked attempts in Directory Insights.

> 📍 In Admin Portal, go to Security › MFA Configurations, confirm Push Bombing Mitigation is enabled (blocks >1 push per resource per 60 seconds), then open Directory Insights and filter for blocked push events to review fraudulent attempts.

— Gives the exact threshold and console path to verify and review.

Release Notes 2023 (2026-09-23)

THINNER COVERAGE BELOW

07 **SentinelOne and Microsoft 365 integration extensions** NEW

40

Integrates SentinelOne agent deployment via JumpCloud Commands for macOS, Windows, and Linux endpoint protection, and extends the Microsoft 365 integration to include M365/Azure AD security groups.

— Names both integrations but no setup steps included. Release Notes 2022 (2026-09-23)

└──▷ **ALSO FROM THESE RELEASES**

└──▷ **BREAKING ON UPGRADE**

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; all scripts, automations, and integrations must be updated to use `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! The `$le` and `$ge` memberQuery filter operators in the Dynamic Groups API are deprecated and replaced by `$lte` and `$gte`; any API requests using the old operators must be updated.

! The .appimage Linux build is no longer supported as of April 18, 2024; users must migrate to the RPM or Debian build.

Release Notes 2022 (2026-09-23) BLOG ↗ STABLE

JumpCloud: Launches JumpCloud Case Portal for centralized view of support cases and ideas submitted by any admin in the organization, with filtered search and clickable case detail.

└──▷ **HOW TO FIND IT**

Block macOS Ventura upgrades fleet-wide while you validate compatibility — configure the delay/block policy before devices auto-upgrade.

📍 In the Admin Portal, go to Device Management › Policies, create a new macOS policy, and select Delay or Block MacOS Ventura to set the enforcement window across your managed macOS devices.

Enforce a Conditional Access Policy that denies resource access to any device without disk encryption enabled across macOS, Windows, and Linux.

📍 In the Admin Portal, go to Security › Conditional Access, create or edit a policy, add a Device Condition, and set the disk encryption requirement to Enabled (FileVault for

macOS, system drive for Windows, root disk for Linux).

Integrates SentinelOne agent deployment via JumpCloud Commands for macOS, Windows, and Linux endpoint protection.

Extends the Microsoft 365 integration to include M365/Azure AD security groups.

Release Notes 2024 (2026-09-23) BLOG ↗ STABLE

JumpCloud: Releases SaaS Management — discover, categorize, warn, and block shadow IT SaaS apps on user devices via the JumpCloud Go browser extension, with manual app addition and licensing tracking.

└──▷ **TRY IT**

Migrate an existing user reactivation call from the deprecated GET endpoint to the new POST endpoint to avoid 404 errors.

```
$ POST https://console.jumpcloud.com/systemusers/reactivate/{id}
Content-Type: application/json

{"activation_email": "user@example.com"}
```

Enforce MFA across all MSP admins in a multi-tenant portal from a single setting rather than configuring each admin account individually.

📍 In the Admin Portal, go to Account › Settings, then enable Global MFA Requirement under the Global MFA Requirement section.

└──▷ **BREAKING ON UPGRADE**

! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; all scripts, automations, and integrations must be updated to use `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.

! The `$le` and `$ge` memberQuery filter operators in the Dynamic Groups API are deprecated and replaced by `$lte` and `$gte`; any API requests using the old operators must be updated.

Release Notes 2025 (2026-09-23) BLOG ↗ STABLE

JumpCloud: Introduces JumpCloud Access Requests, enabling approval flows for access requests submitted from the JumpCloud User Portal and managed in the admin console.

└──▷ **HOW TO FIND IT**

Exclude a remote-region user group from Push MFA enforcement so those users can still authenticate via TOTP-only without disrupting their JumpCloud Protect experience.

> 📍 In the JumpCloud Admin Portal, go to Security › MFA › JumpCloud Protect › Unified MFA, open User Group Exclusions, and add the target group to the exclusion list.

Create a privileged-access approval flow tied to a JIRA project so that every access request auto-creates a ticket and notifies approvers in real time.

> 📍 In the JumpCloud Admin Portal, go to Access Requests › Approval Flows, create a new flow, assign approver admins, then under Integrations configure a webhook pointing to your JIRA project URL to receive new-request events.

Supports webhook configuration in Access Requests to integrate with ticketing systems such as JIRA, automatically creating tickets and sending real-time notifications on new requests.

Tracks Access Request events in Directory Insights.

Releases SaaS License Management (GA) to automate license tracking for supported SaaS applications and optimize license cost management.

Adds MTP admin support for creating and managing support cases across all tenant organizations from within the MTP account.

Release Notes 2026 (2026-09-23) BLOG ↗ STABLE

JumpCloud: Adds new Android policies for Display Settings (screen brightness and timeout), Autofill (system autofill control on Android 8.0+), Work Account (Google enterprise account requirement), Default App and Intent Handling (managed app as default handler for activities), and enhanced Bluetooth restrictions including outbound Bluetooth file sharing and personal-profile Bluetooth sharing controls via the Work & Personal Usage policy.

└──▷ **HOW TO FIND IT**

Register and configure a new AI agent identity with OAuth credentials so it can access AI Gateway MCP servers independently of any human account.

> 📍 In the Admin Portal: go to Identity Management › Agents and click + Add. Enter the agent name, owner, and company email on the Details tab. On the Authentication tab, generate OAuth client credentials and set the client authentication type to `Client secret post` if your external platform (e.g. Amazon Bedrock AgentCore) sends credentials in the request body. Under Identity Management › Agent Groups, create an Agent Group, add the agent, and assign SSO applications. On the agent's Applications tab, use Manage connections to authorize MCP servers on behalf of the agent.

Surface Windows patch compliance gaps across managed devices — without any manual data collection — and export the full result set for audit or ticketing.

📍 In the Admin Portal: go to Reports › Device Patch Report (Windows). Review per-device Installed, Pending, and Failed KB status in the portal preview (up to 500 rows), then click Download to export the complete organization result set as CSV.

Ground AI Assistant answers in org-specific runbooks and SOPs by uploading internal documentation into the private knowledge base.

📍 In the Admin Portal: go to AI Assistant › Knowledge Base and upload your internal documentation files (PDF, TXT, DOCX, or MD; max 10 MB per file, 25 files per org). AI Assistant will attribute the source document in any chat response that draws on your uploaded content.

changelog · Sep 23, 2026 CHANGELOG ↗ STABLE

JumpCloud: Adds Client secret basic and Client secret post client authentication type options on the agent's Authentication tab, letting agents match the exact OAuth credential-delivery method required by external platforms such as Amazon Bedrock AgentCore.

└──▷ HOW TO FIND IT

Push an updated application binary to all associated devices without rebuilding the app record or reassigning device groups.

📍 In the JumpCloud Admin Portal, go to Software Management › Private Repository, open the application's Details tab, click Update Version, upload the new file via the Upload File flow, and (for EXE) enter the install success exit code before deploying.

Configure an agent to send OAuth credentials in the request body when the external platform (e.g. Amazon Bedrock AgentCore) requires `client_secret_post` rather than the default header-based method.

📍 In the JumpCloud Admin Portal, open the agent record, go to the Authentication tab, and under Client Authentication Type select Client secret post instead of the default Client secret basic.

Requires an install success exit code when deploying Windows EXE updates via the Private Repository so JumpCloud can confirm installation succeeded.

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ

GA Announcement: Next Generation Certifications (2026-09-22) BLOG ↗ RANK

SailPoint is an identity governance and access management platform that automates user provisioning, access reviews, and compliance across enterprise systems.

> SailPoint's biggest move this window is taking Next Generation Certifications to GA — a rebuilt engine, automatic campaign start, a new email template set, and breaking changes to certification reporting endpoints — alongside a broad expansion of 21 new connectors including Infisical, StackHawk, FOSSA, and Redpanda Cloud Control Plane.

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

01 Certification reporting API endpoint changes BREAKING

90

The endpoint for extracting certification report data (CSV and PDF) via the `/campaign/` and `/certification/` reporting flow is changing; existing integrations must be updated per the Certification Reports Endpoint Moving migration guidance. Very high-volume offset-paginated calls to the `/campaign/` and `/certification/` endpoints may time out on tenants with approximately 4,000 or more records, and affected tenants are held from rollout until a fix ships.

— Names exact endpoints, a numeric threshold, and a migration resource.

GA Announcement: Next Generation Certif…

02 New certification email template set BREAKING

85

Introduces a new certification email template set alongside the existing templates (labeled Legacy in the UI) for side-by-side comparison during migration. A small number of template attributes are no longer supported in the new set (see Deprecated Variables in Certs Email Templates), and email templates larger than 133 KB will have images trimmed automatically to fit the new v2 templating engine size limit.

— Concrete size limit and migration resource named, but exact update steps not given.

GA Announcement: Next Generation Certif…

03 Next-generation certification engine at GA NEW

80

A new certification engine scales campaign generation to multi-million items, cutting large campaign generation time from days to minutes. Campaign templates, campaign notification templates, and scheduled campaigns migrate automatically during tenant cutover with no manual recreation. FedRAMP environments are excluded from this GA

rollout, and Role Composition campaigns continue running on the legacy engine until a later migration date.

— Explains scale, migration mechanics and exclusions but no admin action required.

GA Announcement: Next Generation Certif…

04 **21 new connectors added to identity governance** NEW

60

Adds connector support for Infisical (secrets management platform access), StackHawk (application security testing workflows), FOSSA (license compliance and dependency data), and Redpanda Cloud Control Plane (streaming data infrastructure), plus Anvil, Checkmk, Cloudiway, Condens, DataHub, Employment Hero, Freshcaller, Juniper Mist, Kashoo, LangChain, Nifty, noCRM, Qdrant, Scrive, Trigger, Vendasta, and YouSign.

— Names all 21 connectors but no config or setup detail given. product docs

05 **'Generate and Start' campaign option** NEW

60

Adds a Generate and Start option so small campaigns kick off immediately after generation without a manual start step, and extends this to campaigns launched from workflows, replacing the legacy manual-start pattern for targeted certifications.

— Named option and behaviour change described, but no exact UI path given.

GA Announcement: Next Generation Certif…

06 **Post-cutover certification UI and legacy access** IMPROVED

60

In-flight certifications automatically pick up the improved reviewer UI after tenant cutover with no reviewer action required, and administrators can find campaigns generated before cutover under Legacy Campaigns in the left navigation of the admin console.

After cutover, locate pre-migration campaigns that are no longer on the active list by navigating to the Legacy Campaigns view in the admin console.

📍 In the SailPoint admin console, go to Certifications › Legacy Campaigns in the left navigation to view and manage campaigns generated before your tenant cutover date.

— Gives exact navigation path but reviewer UI change itself is undescribed.

GA Announcement: Next Generation Certif…

THINNER COVERAGE BELOW

07 **Enriched Campaign Filters** NEW

45

Adds Enriched Campaign Filters (an in-progress feature) to replace the unsupported combination of search queries and campaign filters on a single campaign; campaigns

using both must migrate to filters only via this feature before cutover.

— Named feature and what it replaces, but marked in-progress with no mechanism detail.

GA Announcement: Next Generation Certif…

└──▷ **BREAKING ON UPGRADE**

! The endpoint for extracting certification report data (CSV and PDF) via the `/campaign/` and `/certification/` reporting flow is changing; existing API integrations that pull report data must be updated per the Certification Reports Endpoint Moving migration guidance.

! A small number of template attributes are no longer supported in the new certification email template set; customized notification templates using those attributes must be updated after migration (see Deprecated Variables in Certs Email Templates).

! Very high-volume offset-paginated calls to `/campaign/` and `/certification/` endpoints may time out on tenants with approximately 4,000 or more records — affected tenants are held from rollout until a fix ships.

! Campaigns configured with both a search query and campaign filters on the same campaign will not run on the new engine; these must be migrated to use campaign filters only (via Enriched Campaign Filters) before cutover.

! Email notification templates larger than 133 KB will have images trimmed automatically during migration to fit the new v2 templating engine size limit.

! Role Composition campaigns are excluded from the new engine at GA and continue to run on the legacy engine until a later migration date.

GA Announcement: Next Generation Certifications (2026-09-22) BLOG ↗ LATEST

STABLE

SailPoint Next Generation Certifications reaches GA, bringing multi-million-item campaign scale, a new reviewer UI, and a Generate and Start option.

└──▷ **HOW TO FIND IT**

After cutover, locate pre-migration campaigns that are no longer on the active list by navigating to the Legacy Campaigns view in the admin console.

📍 In the SailPoint admin console, go to Certifications › Legacy Campaigns in the left navigation to view and manage campaigns generated before your tenant cutover date.

Introduces a new certification engine that scales campaign generation to multi-million items, reducing large campaign generation time from days to minutes.

Adds a Generate and Start option so small campaigns kick off immediately after generation without requiring a manual start step.

Adds Generate and Start availability for campaigns launched from workflows, replacing the legacy manual-start pattern for targeted certifications.

Migrates campaign templates, campaign notification templates, and scheduled campaigns automatically as part of the engine cutover — no manual recreation

required.

Introduces a new certification email template set alongside the existing templates (labeled `Legacy` in the UI), enabling side-by-side comparison during migration.

+5 more

└─▷ **BREAKING ON UPGRADE**

- ! The endpoint for extracting certification report data (CSV and PDF) via the `/campaign/` and `/certification/` reporting flow is changing; existing API integrations that pull report data must be updated per the `Certification Reports Endpoint Moving` migration guidance.
- ! A small number of template attributes are no longer supported in the new certification email template set; customized notification templates using those attributes must be updated after migration (see `Deprecated Variables in Certs Email Templates`).
- ! Very high-volume offset-paginated calls to `/campaign/` and `/certification/` endpoints may time out on tenants with approximately 4,000 or more records — affected tenants are held from rollout until a fix ships.
- ! Campaigns configured with both a search query and campaign filters on the same campaign will not run on the new engine; these must be migrated to use campaign filters only (via Enriched Campaign Filters) before cutover.
- ! Email notification templates larger than 133 KB will have images trimmed automatically during migration to fit the new v2 templating engine size limit.
- ! Role Composition campaigns are excluded from the new engine at GA and continue to run on the legacy engine until a later migration date.

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ

Prevent device code phishing with a new in-browser nudge NOTES ↗ RANK

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

> Nudge Security shipped an in-browser warning that fires on OAuth device code authorization pages to block device code phishing attacks.

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

01 In-browser nudge for OAuth device code phishing NEW

75

A new in-browser nudge fires on every visit to an OAuth device code authorization page for Microsoft, Google, GitHub, GitLab, Salesforce, AWS, Docker, Zoom, Vercel, and OpenAI, warning employees never to enter a code provided by someone else. Every employee acknowledgment is recorded as a nudge response for audit and tracking, and the detection is configured and enabled in Browser Extension Settings.

Enable the device code phishing nudge so employees are warned in real time whenever they land on an OAuth device code authorization page.

> 📍 In the Nudge Security console, go to Browser Extension Settings, locate the device code phishing detection option, and turn on the nudge.

— Names exact config location and covered providers, but no API/flag.

Prevent device code phishing with a new…

Prevent device code phishing with a new in-browser nudge NOTES ↗ LATEST STABLE

Nudge Security adds an in-browser warning on OAuth device code pages to block device code phishing attacks.

└──▷ **HOW TO FIND IT**

Enable the device code phishing nudge so employees are warned in real time whenever they land on an OAuth device code authorization page.

> 📍 In the Nudge Security console, go to Browser Extension Settings, locate the device code phishing detection option, and turn on the nudge.

Adds a new in-browser nudge that fires on every visit to an OAuth device code authorization page for Microsoft, Google, GitHub, GitLab, Salesforce, AWS, Docker,

> Zoom, Vercel, and OpenAI, warning employees never to enter a code provided by someone else.
>
> Records every employee acknowledgment of the device code phishing nudge as a nudge response for audit and tracking.
>
> Configure detection and enable the device code phishing nudge in `Browser Extension Settings`.

WAS THIS USEFUL? 👍 👎

Passbolt

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

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

> Passbolt 5.16.0 introduces a beta Offline Mode in its browser extension, letting users access and autofill selected credentials when the server is unreachable.

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

01 **Offline Mode (Beta) in browser extension** NEW 45

Passbolt 5.16.0 adds Offline Mode (Beta) to the browser extension, giving users secure read-only access and autofill for selected credentials when the Passbolt server cannot be reached.



— Names the feature and behaviour but no config, scope limits or mechanism.

changelog-20260923-fa63499d

> changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE
>
> Passbolt 5.16.0 adds Offline Mode (Beta) for secure read-only access and autofill when the server is unreachable.

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password for Mac beta 8.12.40-23 (2026-09-23)

BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password for Mac's latest beta adds native macOS Autofill support for one-time passwords and streamlines developer workflows with GitHub Actions setup guidance and faster API credential creation.

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

01 **Save one-time passwords via macOS Autofill** NEW 81

1Password for Mac can now act as the system provider for one-time password codes through macOS Autofill. Enable it by going to System Settings > General > AutoFill & Passwords > Set Up Codes In and selecting 1Password.

Enable macOS to offer 1Password as the provider for one-time password codes during Autofill on your Mac.

> 📍 In System Settings, go to General › AutoFill & Passwords › Set Up Codes In and select 1Password.

— Exact config path and usage steps given, mechanism explained.

1Password for Mac beta 8.12.40-23 (2026…

THINNER COVERAGE BELOW

02 **GitHub Actions setup instructions in Developer Environments**

NEW 50

A new button in Developer > Environments > GitHub Actions destination lets users view setup instructions directly.

— UI path named but no detail on instruction content.

1Password for Mac beta 8.12.40-23 (2026…

03 **API Credential promoted on New Item screen** IMPROVED 37

API Credential is now positioned near the top of the New Item screen for faster access.

— Names the item type but gives no further mechanism.

1Password for Mac beta 8.12.40-23 (2026…

1Password for Mac beta 8.12.40-23 (2026-09-23) BLOG ↗ LATEST STABLE

1Password for Mac now supports saving one-time passwords via macOS Autofill, plus GitHub Actions setup instructions in Developer Environments.

└──▷ **HOW TO FIND IT**

Enable macOS to offer 1Password as the provider for one-time password codes during Autofill on your Mac.

> 📍 In System Settings, go to General › AutoFill & Passwords › Set Up Codes In and select `1Password`.

Adds support for saving one-time passwords in 1Password through macOS Autofill — configure via System Settings > General > AutoFill & Passwords > Set Up Codes In, selecting 1Password.

Adds a button in Developer > Environments > GitHub Actions destination to view setup instructions.

Promotes `API Credential` to near the top of the New Item screen for faster access.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter

Feedback ⓘ

09/22 September 22 – Web Application and API Update (2026-09-22) BLOG ↗ RANK

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

> DNSFilter's dashboard update adds visibility into why domains are blocked and who triggered CyberSight activity, tightens deletion safety for Roaming Clients, and scopes deployment access controls for read-only and policy-only users.

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

01 **Scoped deployment access for read-only and policy-only users**

IMPROVED 65

Scopes deployment access controls for read-only and policy-only users across Sites, Relays, and Roaming Clients — hiding Site Keys, disabling Relay deployment controls, and restricting Roaming Client installer access, each with a tooltip directing users to contact a dashboard admin.

– Names each restricted surface but gives no exact navigation path.

09/22 September 22 – Web Application an…

02 **'Blocked By' column in Unblock Requests grids** NEW

60

Adds a Blocked By column to the Active Requests and Request History grids showing the reason each domain was blocked, consistent with the Query Log's Method column, and also visible in the approval drawer.

Review why a domain was blocked without leaving the Unblock Requests queue — the new Blocked By column surfaces the blocking reason inline.

> 📍 In the dashboard, go to Unblock Requests › Active Requests or Request History and locate the Blocked By column in the grid, or open the approval drawer for any request to see the blocking reason.

– Names exact grids and column but not underlying blocking logic.

09/22 September 22 – Web Application an…

THINNER COVERAGE BELOW

03 **Last Logged In User in CyberSight Top Clients** NEW

55

Adds Last Logged In User as a default column and Friendly Name as an optional column in the CyberSight Top Clients View All grids on Activity Overview and AI Usage.

Identify the person behind a Roaming Client hit in CyberSight without drilling into individual device records.

> 📍 In the dashboard, go to CyberSight › Activity Overview (or AI Usage) › Top Clients › View All and review the Last Logged In User column; click the column chooser to optionally add Friendly Name.

— Exact column names and navigation given, no further mechanism.

09/22 September 22 – Web Application an…

04 **Type-to-confirm for Roaming Client deletion** IMPROVED

45

Introduces a type-to-confirm requirement for deleting Roaming Clients and for running cleanup deletions, preventing accidental mass deletions.

— Describes the safeguard but no UI steps or confirmation text given.

09/22 September 22 – Web Application an…

09/22 September 22 – Web Application and API Update (2026-09-22) BLOG ↗

LATEST STABLE

DNSFilter adds Blocked By column to Unblock Requests grids, type-to-confirm for Roaming Client deletion, and Last Logged In User in CyberSight Top Clients.

└──▷ **HOW TO FIND IT**

Review why a domain was blocked without leaving the Unblock Requests queue — the new Blocked By column surfaces the blocking reason inline.

> 📍 In the dashboard, go to Unblock Requests › Active Requests or Request History and locate the Blocked By column in the grid, or open the approval drawer for any request to see the blocking reason.

Identify the person behind a Roaming Client hit in CyberSight without drilling into individual device records.

> 📍 In the dashboard, go to CyberSight › Activity Overview (or AI Usage) › Top Clients › View All and review the Last Logged In User column; click the column chooser to optionally add Friendly Name.

Adds a Blocked By column to the Active Requests and Request History grids showing the reason each domain was blocked, consistent with the Query Log's Method column, and also visible in the approval drawer.

> Adds `Last Logged In User` as a default column and `Friendly Name` as an optional column in the CyberSight Top Clients `View All` grids on Activity Overview and AI Usage.
>
> Introduces type-to-confirm requirement for deleting Roaming Clients and running cleanup deletions, preventing accidental mass deletions.
>
> Scopes deployment access controls for read-only and policy-only users across Sites, Relays, and Roaming Clients — hiding Site Keys, disabling Relay deployment controls, and restricting Roaming Client installer access, each with a tooltip directing users to contact a dashboard admin.

WAS THIS USEFUL?  

tailscale

Feedback ⓘ View device posture status (2026-09-23)

BLOG ↗ RANK

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale's biggest window yet: it launched Aperture, a centralized AI gateway that routes LLM requests through tailnet identity with guardrails, connectors, dashboards and retention controls, alongside Tailscale PAM (beta), a full privileged access management platform with connectors, secrets management, Terraform IaC, and custom S3-backed session recording. It also documented Tailnet Lock's node-signing workflow, added device posture detail to the Machines page, and gave the GitHub Action a log-mode flag.

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

01 Aperture access for devices outside the tailnet NEW

95

Aperture can be reached from devices outside the tailnet via `ts-unplug`, a standalone proxy from the `ts-plug` repository built with `make ts-unplug` that accepts `-dir` and `-port` flags and forwards local traffic to an Aperture gateway, or via an Aperture CLI bridge activated from the Bridges settings menu that embeds a Tailscale node inside the CLI and auto-configures launched coding agents to use a local proxy endpoint. Both support `auth key` authentication for non-interactive approval in ephemeral environments such as CI/CD runners, enabling cross-tailnet and contractor/personal-device access to Aperture without installing Tailscale.

Run ts-unplug on a CI runner to expose Aperture on a fixed local port without installing Tailscale on the host.

```
$ ./build/ts-unplug -dir ./state -port <aperture-hostname>.
<tailnet-name>.ts.net
```

Verify the ts-unplug proxy is forwarding requests correctly before wiring it into a script or CI job.

```
$ curl -s http://localhost:8080/v1/messages \
  -H "Content-Type: application/json" \
  -d '{"model": "claude-haiku-4-5-20251001", "max_tokens": 25,
"messages": [{"role": "user", "content": "respond with:
hello"}]}'
```

Create and activate an Aperture CLI bridge for an interactive coding session on a device outside the tailnet.

> 📍 1. Run `aperture`, then press `s` to open Settings. 2. Select Bridges, press `a`, and enter a name (e.g. `ci-runner`). 3. Return to Settings, select Aperture Endpoints, press `a`, choose Bridge, select the bridge, and enter your Aperture URL (e.g. `http://<aperture-hostname>.<tailnet-name>.ts.net`). 4. Select the new endpoint to activate it, then open the printed authentication URL in a browser. 5. If required, approve the device in the Machines page of the admin console.

— Named flags and runnable commands provided. product docs

02 **External secret references in PAM service config** NEW

95

Tailscale PAM service configuration fields (upstream username, password, hostname, AWS access key ID, AWS secret access key) for SSH, database, and AWS PAM service types can reference external secrets instead of storing plaintext values, using `from:env:<VARIABLE>` to read environment variables at connector runtime, `from:file:<PATH>` to read a file on the connector host, `from:aws:ssm:<PATH>` and `from:aws:secretsmanager:<PATH>` to pull from AWS Systems Manager Parameter Store or Secrets Manager (requiring connector IAM read permissions), and `from:keeper:<RECORD>:<FIELD>` to pull from Keeper Secrets Manager via the `ksm` CLI installed on the connector host. A `jq_exp=<expression>` option on any `from:` reference extracts a single field from a JSON secret using dot notation, e.g. `from:aws:secretsmanager:production-database,jq_exp=.password`.

Configure a database PAM service to pull its username from an environment variable and its password from AWS SSM Parameter Store — no plaintext credentials in the service config.

```text
Username: from:env:RDS_USERNAME
Password: from:aws:ssm:/production/database/password
```

Extract individual fields from a JSON secret stored in AWS Secrets Manager so a single secret object can serve both username and password fields.

```text
Username: from:aws:secretsmanager:production-
database,jq_exp=.username
Password: from:aws:secretsmanager:production-
database,jq_exp=.password
```

Configure an SSH PAM service to retrieve its credentials from a local file and AWS Secrets Manager, keeping secrets out of Tailscale PAM configuration entirely.

```text
Username: from:file:/etc/ssh_username
Password: from:aws:secretsmanager:/staging/sshpassword
```

— Exact syntax and worked examples for every secret backend. product docs

03 **Custom S3 bucket for PAM session recordings** NEW 95

Tailscale PAM (Premium plan) can store session recordings in a customer-owned Amazon S3 bucket instead of Tailscale-managed storage, via a cross-account IAM role. Setup requires an IAM policy granting `s3:PutObject`, `s3:GetObject`, `s3:ListBucket`, and `s3:GetBucketLocation` on the bucket ARNs, a trust policy allowing Tailscale PAM's AWS account (`arn:aws:iam::235487987553:root`) to `sts:AssumeRole` (optionally scoped with an `sts:ExternalId` condition), and configuration in a new Recording storage section of the PAM page in the admin console (bucket name, role ARN, optional external ID, optional region). Switching to a custom bucket does not migrate existing recordings — previously stored recordings become unavailable for replay in the admin console.

Create the IAM policy granting Tailscale PAM read/write access to your S3 bucket, then attach it to a cross-account role Tailscale can assume — no long-lived AWS credentials required.

```json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:ListBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::my-pam-recordings",
        "arn:aws:s3:::my-pam-recordings/*"
      ],
      "Effect": "Allow",
      "Sid": "AllowAccessS3Bucket"
    }
  ]
}
```

Lock down the IAM role trust policy so only Tailscale PAM's account can assume it, using an external ID to prevent confused-deputy attacks.

```json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": ["arn:aws:iam::235487987553:root"]
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "my-unique-external-id"
        }
      }
    }
  ]
}
```

Point Tailscale PAM at your S3 bucket after creating the AWS resources, so new session recordings are written to your own account instead of Tailscale-managed storage.

📍 1. Open the PAM page of the Tailscale admin console. 2. In the Recording storage section, select Configure storage. 3. Enter the name of your S3 bucket (e.g. my-pam-recordings). 4. Enter the ARN of the IAM role (e.g. arn:aws:iam::123456789012:role/pam-recordings). 5. Enter the External ID you set in the IAM role trust policy. 6. Optionally enter the AWS region containing the bucket. 7. Select Save.

— Concrete IAM policy JSON and console steps given. product docs

04 **Tailnet Lock node-signing workflow** NEW 95

Tailnet Lock lets trusted nodes cryptographically sign new peers so no node joins the tailnet without approval, even if Tailscale's infrastructure is compromised, via a tailnet key authority (TKA) that tracks trusted signing keys so peers can verify signatures before connecting. It's managed with `tailscale lock init` (enables Tailnet Lock, generating ten disablement secrets and distributing trusted keys), `tailscale lock sign nodekey:<key> tlpub:<key>` (signs a new node from a trusted signing node), `tailscale lock status` (verifies provisioning state), `tailscale lock disable` (disables using a stored disablement secret), and `tailscale lock local-disable` (lets a node ignore Tailnet Lock without a disablement secret, as a recovery option). A `property:locked-out` filter on the Machines page finds nodes awaiting a signature,

and macOS, Windows, and iOS clients (with QR code support on iOS) support signing via app links as a CLI alternative.

Sign a newly joined node from a trusted signing node so it can participate in the tailnet — supply the node key and Tailnet Lock public key shown in the admin console.

```
$ tailscale lock sign nodekey:1abddef1 tlpub:abcdef12
```

— Every subcommand named with a runnable example. product docs

05 **Aperture AI gateway with identity-based provider routing** NEW

90

Aperture is a centralized AI gateway that routes LLM requests to upstream providers (OpenAI, Anthropic, Google) using Tailscale identity for authentication, so no API keys are distributed to developer devices or CI/CD systems. The gateway centrally manages provider API keys in server configuration and injects them per request; the `openai_chat` compatibility flag plus `baseurl` and `apikey` config fields let any provider exposing a `/v1/chat/completions` endpoint (Groq, Together AI, Fireworks, Mistral, DeepSeek, Perplexity, and others) work without code changes, and `auth_mode: passthrough` forwards a client-supplied Authorization header for subscription plans like Claude Max or ChatGPT Plus. Devices provisioned via approved single-use auth keys through the built-in tailnet connector get `custom:createdByAperture` and `custom:createdByAI` node attributes, and requests can also be routed to self-hosted LLMs inside a tailnet without exposing endpoints publicly.

Route any OpenAI-compatible provider (e.g. Groq) through Aperture by pointing to its base URL — clients then use model names as normal, with no provider credentials on their devices.

```yaml
baseurl: "https://api.groq.com/openai"
apikey: "<your-groq-key>"
compatibility: openai_chat
```

Use a Claude Max or ChatGPT Plus subscription plan with Aperture by enabling passthrough mode, which forwards the client-supplied Authorization header instead of injecting a centrally managed key.

```yaml
auth_mode: "passthrough"
```

— Config keys and provider list given; no endpoint documented. product docs

06 **Tailscale PAM (beta) platform launch** NEW 85

Tailscale PAM (beta) is a privileged access management platform, enabled from the `PAM` page of the admin console, which automatically creates the `tag:border0-managed` tag, an `autoApprovers` rule, an initial PAM application grant for `autogroup:admin`, and an OIDC trust credential viewable on the `Trust credentials` page. It proxies sessions to SSH, HTTP/HTTPS, RDP, VNC, Kubernetes, and arbitrary TCP-port services, plus databases including Amazon RDS, Elasticsearch, Google Cloud SQL, Microsoft SQL Server, MongoDB, MySQL, PostgreSQL, and Snowflake, all defined on the `Services` page (e.g. `Secure Shell - Standard Shell`). PAM connectors — deployable on Linux, AWS EC2, Docker, and Kubernetes, with a high-availability mode — act as identity-aware application proxies handling session termination, real-time policy enforcement, session recording, and secrets/credential injection via configurable secret stores; a WebAssembly-based browser client gives agentless access to SSH, MySQL/Postgres databases, and RDP directly from a browser at `my.tailscale.com` with no client install. Access is governed by grants in the tailnet policy file (editable on the `Access controls` page) supporting conditions on user identity, specific PAM services, device posture/MDM, time, and IP/country, and Tailscale syncs tailnet users and groups into PAM automatically, mapping Owner/Admin/IT admin to PAM Admin, Member/Billing admin to Client Access Only, and Network admin/Auditor to Read Only. All connector-to-device tunnels are encrypted with WireGuard, and PAM provides complete session auditing and recording.

Add a new PAM-managed SSH or database service to your tailnet so privileged access is brokered through Tailscale PAM.

> 📍 1. Open the Services page of the Tailscale admin console. 2. Select `Add service`, then `PAM service`. 3. Select a service type (e.g. SSH, PostgreSQL, RDP) and select `Continue`. 4. Follow the prompts and select a PAM connector for the service. 5. Select `Save` to activate the service.

Enable Tailscale PAM and create a connector so privileged SSH access is proxied through your tailnet without manually configuring tags or approvals.

> 📍 1. In the Tailscale admin console, go to PAM and select Enable. 2. Go to Connectors, select Add connector, choose your platform (Linux, AWS EC2, Docker, or Kubernetes), and follow the Launch new connector instructions. 3. Go to Services, select Add service › PAM service › Secure Shell - Standard Shell, choose your connector, and select Save.

Access a PAM-managed SSH or database service from a browser without installing the Tailscale client — useful for onboarding users or accessing services from unmanaged devices.

> 📍 1. In a browser, navigate to `my.tailscale.com`. 2. Log in with your Tailscale account. 3. Select the PAM service you want to access.

Grant a specific user access to a PAM service (e.g., a PostgreSQL database) by referencing them in a grant in the tailnet policy file — the default is no access until explicitly granted.

> 📍 In the Tailscale admin console, go to Access controls to open the tailnet policy file, then add a grant referencing the user and the desired PAM service.

— Extensive named surfaces across five docs pages. product docs

07 **Terraform provider for PAM resources** NEW 85

Tailscale PAM resources can be managed as Infrastructure as Code via the `borderzero/border0` Terraform provider, authenticated with the `BORDER0_TOKEN` environment variable. The `border0_connector` resource creates and manages PAM connectors, and the `border0_socket` resource creates and manages PAM services with `socket_type`, `connector_ids`, and `http_configuration.upstream_url` fields.

Provision a PAM connector and an HTTP service together in one Terraform config, authenticated via an env var to avoid storing secrets in HCL.

```
$ export BORDER0_TOKEN="<service-account-token>"
terraform init
terraform apply
```

Define a PAM connector and associate an HTTP service with it so an internal admin app is reachable through Tailscale PAM.

```hcl
terraform {
  required_providers {
    border0 = {
      source = "borderzero/border0"
    }
  }
}

provider "border0" {}

resource "border0_connector" "production" {
  name        = "production"
  description = "Tailscale PAM connector for production services"
}

resource "border0_socket" "internal_admin" {
  name          = "internal-admin"
  description   = "Internal admin application"
  socket_type   = "http"
  connector_ids = [border0_connector.production.id]

  http_configuration {
    upstream_url = "https://admin.internal.example.com"
  }
}
```

— Named resources with a runnable Terraform config. product docs

08 **Aperture usage dashboard, spending controls and exports** NEW

70

Aperture's admin dashboard adds a Usage Overview page with organization-wide adoption analytics — active users over time, usage histograms, and per-user/per-model/per-time-period cost tracking — plus configurable budgets, per-user spending limits, per-user quotas, and token purchasing. A Logs page groups related requests into sessions for admin-level conversation and coding-task tracing with full request/response data, and usage data can be exported to external authorization engines and log aggregators for observability and governance.

— Named UI pages but no API or config key given. product docs

09 **Aperture CLI, chat UI and coding-agent integrations** NEW

60

Aperture ships a CLI for discovering, configuring, and running coding agents through the Aperture proxy, and a built-in chat UI so users can talk to LLMs directly through their

browser without a separate client. Client integrations route requests through the Aperture proxy, authenticated via Tailscale identity, for Claude Code (including its GitHub Action, so CI runners authenticate with Tailscale identity instead of individual API keys), Claude Desktop, OpenAI Codex, OpenCode, and OpenAI-compatible tools including Gemini CLI, Roo Code, and Cline. An alpha Tailscale skill lets coding agents such as Claude Code, Cursor, or OpenAI Codex work with Tailscale configuration, the CLI, and the API.

— Lists many integrations, no commands or flags given. product docs

THINNER COVERAGE BELOW

10 **Aperture connectors for MCP servers and HTTP APIs** NEW

50

Aperture connectors proxy connections to external MCP servers and HTTP APIs, injecting authentication automatically so AI agents discover available services through the MCP tool list without managing individual credentials; MCP server proxying aggregates tools and resources from remote MCP servers with identity-based access control.

— Mechanism described, no config keys given. product docs

11 **Aperture data retention, privacy controls and SIEM export** NEW

50

A configurable retention policy controls how long Aperture keeps captured request/response bodies, including a zero data retention mode that prevents prompt or response content from ever being written to disk, alongside broader configurable data retention and privacy controls governing where request content is stored and exported. Captured logs and events can also be exported to external SIEM systems for durable storage and review.

— Zero-retention mode named, no config key given. product docs

12 **Custom DNS domains for PAM HTTP services** NEW

50

Tailscale PAM supports custom DNS domains for HTTP services, with wildcard CNAME routing and automatic TLS provisioning; when registering multiple custom domains, a Make this the default domain option designates which domain is primary for PAM HTTP services.

Expose a PAM HTTP service at a human-readable hostname under your own domain instead of the default PAM hostname.

📍 In the Tailscale admin console, go to PAM › Domains › Add domain. Enter your domain (e.g. `infra.example.com`), then add the two CNAME records shown: *. `infra.example.com` CNAME <your-pam-subdomain>.border0.io _acme-challenge. `infra.example.com` CNAME <domain>.<your-pam-subdomain>. `acme.border0.io` Once DNS propagates, select Save. Your PAM

HTTP service named `wiki` will be reachable at `https://wiki.infra.example.com`.

— Single toggle described; example shows CNAME setup steps. product docs

13 **Device posture detail on Machines page** IMPROVED 50

The Machines page in the admin console now shows, per machine, each posture assertion required by a device posture check, the machine's current value for that assertion, whether it is passing, and a count of policy file rules that affect the machine and require that posture.

Verify which posture assertions a specific machine is failing and how many policy rules depend on them, to diagnose access issues.

📍 In the admin console, go to Machines › select a machine › Device posture status. For each posture assertion, review the required value, the machine's current value, the pass/fail status, and the count of policy file rules that require that posture.

— Named UI location, no API or CLI equivalent given.

View device posture status (2026-09-23)

14 **Log-mode control in GitHub Action** NEW 45

The Tailscale GitHub Action (v4.2.0) adds a `log-mode` parameter to control whether log output is `grouped`, `ungrouped`, or `quiet`.

— Single flag with its three values named, no usage example.

Tailscale GitHub Action v4.2.0 (2026-09…

15 **Aperture guardrails for request inspection** NEW 40

Aperture guardrails inspect, modify, or block LLM requests at the gateway before they leave the network, enabling PII scrubbing, tool-declaration restrictions, and content policy enforcement.

— Describes function only, no config surface named. product docs

└──▷ **BREAKING ON UPGRADE**

! Switching recording storage location does not migrate existing recordings — recordings previously stored by Tailscale will no longer be available for replay in the admin console after switching to a custom S3 bucket.

View device posture status (2026-09-23) BLOG ↗ LATEST STABLE

Tailscale admin console now shows per-machine posture assertion details and policy rule counts on the Machines page.

└──▷ **HOW TO FIND IT**

Verify which posture assertions a specific machine is failing and how many policy rules depend on them, to diagnose access issues.

📍 In the admin console, go to Machines › select a machine › Device posture status. For each posture assertion, review the required value, the machine's current value, the pass/fail status, and the count of policy file rules that require that posture.

In the admin console, the Machines page now displays each posture assertion required by a device posture check, the machine's current value for that assertion, whether it is passing, and a count of policy file rules that affect the machine and require that posture.

Tailscale GitHub Action v4.2.0 (2026-09-22) **BLOG ↗** **STABLE**

Tailscale GitHub Action v4.2.0 adds `log-mode` parameter to control grouped, ungrouped, or quiet log output.

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ September 22, 2026

NOTES ↗ RANK

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across infrastructure.

> CertCentral now automatically revalidates domains every five days using the Persistent DNS TXT Record DCV method, removing the need to manually resubmit domains before their validation expires.

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

01 Automatic revalidation via Persistent DNS TXT Record DCV NEW

88

CertCentral adds automatic domain revalidation using the Persistent DNS TXT Record DCV method: as long as the Persistent DNS TXT record with the correct persistent URI value (Account or Unique) remains in place, DigiCert checks and revalidates the domain every five days, and validation is reusable for nine days before revalidation is required. This applies to domains for OV and EV TLS, Qualified Website Authentication Certificates (QWAC), QWAC PSD2, PKIo Private Services Server, and X9 PKI for TLS certificate types, with validated domains added to the Domains page where revalidation runs, eliminating the need to manually resubmit domains before expiry.

Enroll a domain in automatic 5-day revalidation so certificate issuance is never delayed by an expired domain validation — do this once via the CertCentral Domains page.

> 📍 In CertCentral, go to Domains, add or select your domain, choose Persistent DNS TXT Record as the DCV method, complete the initial validation, then leave the DNS TXT record with the correct persistent URI value (Account or Unique) in place. DigiCert will revalidate the domain automatically every five days.

— Names mechanism, cadence, cert types and UI path clearly September 22, 2026

September 22, 2026 NOTES ↗ LATEST STABLE

CertCentral now auto-revalidates domains via Persistent DNS TXT Record DCV every five days, preventing expired validations from blocking certificate issuance.

└──▷ **HOW TO FIND IT**

Enroll a domain in automatic 5-day revalidation so certificate issuance is never delayed by an expired domain validation — do this once via the CertCentral Domains page.

📍 In CertCentral, go to Domains, add or select your domain, choose `Persistent DNS TXT Record` as the DCV method, complete the initial validation, then leave the DNS TXT record with the correct persistent URI value (Account or Unique) in place. DigiCert will revalidate the domain automatically every five days.

Adds automatic domain revalidation using the Persistent DNS TXT Record DCV method, where DigiCert checks and revalidates domains every five days as long as the Persistent DNS TXT record with the correct persistent URI value (Account or Unique) remains in place.

Supports automatic revalidation for OV and EV TLS, Qualified Website Authentication Certificates (QWAC), QWAC PSD2, PKIo Private Services Server, and X9 PKI for TLS certificate types — with validated domains added to the Domains page where revalidation runs.

Persistent DNS TXT domain validation is reusable for nine days before revalidation is required, with the automatic process eliminating the need to manually resubmit domains before expiry.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

bytebase

Feedback ⓘ 3.22.1 NOTES ↗ RANK

Database governance built for humans and agents — controlling changes and access across every major database.

> Bytebase renamed its Worksheets service to a permission-scoped SavedQueryService with per-object sharing, shipped a new `bytebase recovery` break-glass CLI, added project-scoped instances, a SQL Editor Data Explorer, and login lockout, alongside GPT-5 model support and a string of security-hardening breaking changes to `/metrics`, `/v1/actuator/info`, and account sign-in enforcement.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 **bytebase recovery CLI for locked-out admins** NEW 90

Adds `bytebase recovery` CLI command as a break-glass path to re-enable password sign-in or reset a user's password directly against the metadata database, bypassing the UI.

Recover a locked-out admin by re-enabling password sign-in directly against the metadata database when the UI is inaccessible.

```
$ bytebase recovery
```

— Runnable command with clear mechanism and use case. 3.22.0

02 **SavedQueryService replaces WorksheetService** BREAKING 85

Renames `WorksheetService` to `SavedQueryService`, replacing REST paths `/v1/{parent=projects/*}/worksheets…` with `/v1/{parent=projects/*}/savedQueries…` and `…/savedQueryFolders…`, with no compatibility aliases. The `visibility` field (`PRIVATE` / `PROJECT_READ` / `PROJECT_WRITE`) is removed in favor of per-object IAM; on upgrade every existing saved query migrates to owner-private and must be re-shared explicitly. Custom roles using `bb.worksheets.*` permissions must be updated to the new per-verb `bb.savedQueries.{create, search, list, get, getIamPolicy, setIamPolicy, update, delete}` permission family.

— Full endpoint, field, and permission names given, but no single migration command.

3.22.0

03 **/metrics restricted to loopback with override flag** BREAKING

85

Restricts `/metrics` to loopback-only by default; adds `--metrics-remote-access` startup flag to restore remote scraper access when the endpoint is protected by a gateway, firewall, or NetworkPolicy. On upgrade, remote Prometheus scrapers pointed directly at the server stop collecting unless this flag is set.

Allow a remote Prometheus scraper to reach `/metrics` when Bytebase sits behind a gateway or NetworkPolicy.

```
$ bytebase --metrics-remote-access
```

— Named endpoint, flag, and exact upgrade impact given. 3.22.0

04 **Data Explorer in SQL Editor** NEW

70

Adds Data Explorer to SQL Editor for browsing table and collection data without writing a query, with native predicate filtering per engine — `WHERE` for SQL, filter document for MongoDB, query expression for Elasticsearch.

— Per-engine mechanism named, UI path implied but not exact. 3.22.0

05 **Terraform provider 3.22.1 field removals** BREAKING

70

Terraform provider 3.22.1 removes the `cluster` field in `data_sources` blocks and the computed `service_key` attribute on `bytebase_user` — delete any references from Terraform configs — and makes project webhook URLs write-only, so data sources now read back an empty `webhooks.url`.

— Every removed or changed field named with clear remediation. 3.22.0

06 **Per-object sharing for Saved Queries** NEW

60

Adds per-object sharing for Saved Queries: grant Viewer or Editor access to individual `user:` and `group:` principals via `SetIamPolicy`, and organize queries into folders with drag-and-drop and star favorites.

— Named API and principal types but no UI navigation or command. 3.22.0

07 **Project-scoped instances** NEW

60

Adds project-scoped instances: instances can now be assigned to a single project via an optional `parent` argument on the `bytebase_instance` Terraform resource (provider 3.22.1).

— Names field and resource but gives no full example. 3.22.0

THINNER COVERAGE BELOW

08 **Archive/restore blocked by active task runs** BREAKING

55

Archiving or restoring an instance now returns an error if it has task runs in `PENDING`, `AVAILABLE`, or `RUNNING` state; cancel or wait for those runs before archiving or restoring.

— Exact states named with clear precondition to resolve. 3.22.0

09 **Operational health metrics on /metrics** NEW

50

Adds operational health metrics to `/metrics`: schema sync health, runner cycles, license expiry, and instance capacity.

— Names four metric types but no thresholds or query examples. 3.22.0

10 **Password and identity domain enforcement for admins** BREAKING

50

Enforces `disallow_password_signin` and `enforce_identity_domain` uniformly for all accounts including workspace admins; an admin whose only sign-in path is blocked by these settings will be locked out on upgrade.

— Named settings but no config location or remediation steps. 3.22.0

11 **Legacy async export-data workflow removed** BREAKING

50

The legacy async export-data issue workflow is fully removed; integrations polling the old export-issue API must migrate to JIT access-grant exports or synchronous SQL Editor export.

— Migration alternatives named but no endpoint specifics. 3.22.0

12 **Authentication required on /v1/actuator/info** BREAKING

45

Requires authentication on `/v1/actuator/info`; monitoring scripts that poll this endpoint must now include a credential, or they will stop working after upgrade.

— Endpoint named but no credential format or migration steps. 3.22.0

13 **Login attempt lockout** NEW

45

Adds login attempt lockout: repeated failed sign-in attempts for password, emailed code, or MFA code lock the targeted credential for a cooldown period.

— Mechanism described but cooldown duration and config unnamed. 3.22.0

14 **GPT-5 model support** NEW 30

Adds support for GPT-5 models for AI features, alongside existing GPT-4o support.

— Names models but no config or scope detail given. 3.22.1

15 **Organization-policy endpoints removed for project-scoped instances** 20

BREAKING

Two Instance- and Database-level organization-policy API endpoints previously unsupported for project-scoped instances are removed.

— Endpoints not named; too thin to act on. 3.22.0

16 **Service account search in Grant Access** IMPROVED 10

Improves service account search within the Grant Access dialog.

— Only mentioned in summary, no mechanism or detail given. 3.22.1

└──▷ **BREAKING ON UPGRADE**

! `WorksheetService` is renamed to `SavedQueryService` with no compatibility aliases — all RPCs, messages, and REST paths (e.g. `/v1/{parent=projects/*}/worksheets…`) must be updated to their new equivalents.

! The `visibility` field (`PRIVATE` / `PROJECT_READ` / `PROJECT_WRITE`) is removed from saved queries; on upgrade every existing saved query migrates to owner-private and must be re-shared explicitly.

! Custom roles referencing `bb.worksheets.*` permissions must be updated to the new `bb.savedQueries.*` verb family.

! `/metrics` now only accepts requests from the server host (loopback); remote Prometheus scrapers pointed directly at the server stop collecting on upgrade unless `--metrics-remote-access` is set.

! `/v1/actuator/info` now requires authentication; unauthenticated monitoring scripts stop working on upgrade.

! `disallow_password_signin` and `enforce_identity_domain` now apply to workspace admins too — an admin whose only sign-in path is blocked by these settings will be locked out on upgrade.

! Terraform provider 3.22.1: the `cluster` field in `data_sources` blocks and the computed `service_key` attribute on `bytebase_user` are removed — delete any references from Terraform configs.

! Terraform provider 3.22.1: project webhook URLs are now write-only; data sources read back an empty `webhooks.url`.

! Two Instance- and Database-level organization-policy API endpoints previously unsupported for project-scoped instances are removed.

! Archiving or restoring an instance now returns an error if it has task runs in `PENDING`, `AVAILABLE`, or `RUNNING` state — cancel or wait for those runs before archiving/restoring.

! The legacy async export-data issue workflow is fully removed; integrations polling the old export-issue API must migrate to JIT access-grant exports or synchronous SQL Editor export.

3.22.1 NOTES ↗ STABLE

Bytebase 3.22.1 adds GPT-5 model support and improves service account search in Grant Access.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.22.1 https://github.com/bytebase/bytebase.git
# already have the repo? check out this version:
$ git checkout 3.22.1
```

Supports GPT-5 models for AI features alongside existing GPT-4o support.

3.22.0 NOTES ↗ STABLE

Bytebase 3.22.0 adds project-scoped instances, Data Explorer, login lockout, new `bytebase recovery` CLI, and per-object saved-query sharing.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 3.22.0 https://github.com/bytebase/bytebase.git
# already have the repo? check out this version:
$ git checkout 3.22.0
```

└──▷ TRY IT

Recover a locked-out admin by re-enabling password sign-in directly against the metadata database when the UI is inaccessible.

```
$ bytebase recovery
```

Allow a remote Prometheus scraper to reach `/metrics` when Bytebase sits behind a gateway or NetworkPolicy.

```
$ bytebase --metrics-remote-access
```

Adds `bytebase recovery` CLI command as a break-glass path to re-enable password sign-in or reset a user's password directly against the metadata database, bypassing the UI.

Restricts `/metrics` to loopback-only by default; adds `--metrics-remote-access` startup flag to restore remote scraper access when the endpoint is protected by a gateway, firewall, or NetworkPolicy.

Requires authentication on `/v1/actuator/info`; monitoring scripts that poll this endpoint must now include a credential.

Renames `WorksheetService` to `SavedQueryService` with new REST paths `/v1/{parent=projects/*}/savedQueries…` and `…/savedQueryFolders…`; per-object IAM replaces the `visibility` field.

Replaces `bb.worksheets.*` permissions with per-verb `bb.savedQueries.{create, search, list, get, getIamPolicy, setIamPolicy, update, delete}` permissions on custom roles.

+6 more

└─▷ **BREAKING ON UPGRADE**

- ! `WorksheetService` is renamed to `SavedQueryService` with no compatibility aliases — all RPCs, messages, and REST paths (e.g. `/v1/{parent=projects/*}/worksheets…`) must be updated to their new equivalents.
- ! The `visibility` field (`PRIVATE` / `PROJECT_READ` / `PROJECT_WRITE`) is removed from saved queries; on upgrade every existing saved query migrates to owner-private and must be re-shared explicitly.
- ! Custom roles referencing `bb.worksheets.*` permissions must be updated to the new `bb.savedQueries.*` verb family.
- ! `/metrics` now only accepts requests from the server host (loopback); remote Prometheus scrapers pointed directly at the server stop collecting on upgrade unless `--metrics-remote-access` is set.
- ! `/v1/actuator/info` now requires authentication; unauthenticated monitoring scripts stop working on upgrade.
- ! `disallow_password_signin` and `enforce_identity_domain` now apply to workspace admins too — an admin whose only sign-in path is blocked by these settings will be locked out on upgrade.
- ! Terraform provider 3.22.1: the `cluster` field in `data_sources` blocks and the computed `service_key` attribute on `bytebase_user` are removed — delete any references from Terraform configs.
- ! Terraform provider 3.22.1: project webhook URLs are now write-only; data sources read back an empty `webhooks.url`.

- ! Two Instance- and Database-level organization-policy API endpoints previously unsupported for project-scoped instances are removed.
- ! Archiving or restoring an instance now returns an error if it has task runs in `PENDING`, `AVAILABLE`, or `RUNNING` state — cancel or wait for those runs before archiving/restoring.
- ! The legacy async export-data issue workflow is fully removed; integrations polling the old export-issue API must migrate to JIT access-grant exports or synchronous SQL Editor export.

WAS THIS USEFUL? 👍 👎

Ubiq Security

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗

RANK

Learn how Ubiq discovers sensitive data, maps identity-to-data access, and enforces protection policies across applications, databases, APIs, and AI.

> Ubiq Security expanded structured encryption with four new dataset types (Integer, Date, Date & Time, Token) and extended date/datetime datasets to non-numeric character sets, shipped an Airflow integration for encrypting DAG data, launched a new Rust SDK, and added searchable integer encryption, IdP authentication, and usage-metadata tagging to its JavaScript SDK, plus a Go key-number lookup helper.

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

01 Integer, Date, Date & Time, and Token structured dataset types

NEW 94

Adds `Integer` structured dataset type supporting 32-bit (-99,999,999 to 99,999,999) and 64-bit (-9,999,999,999,999,999 to 9,999,999,999,999,999) whole numbers with encrypted output staying within integer ranges; `Date` for calendar dates (01/01/0001 to 11/28/2738) with no time component; `Date & Time` for timestamps to the second (1653-02-10 06:13:21 to 2286-11-20 17:46:39 UTC, encrypted values falling between 0018-01-16 11:01:21 and 3921-12-16 12:58:39 UTC); and `Token` producing fixed-length base62 tokens (64-character for strings under 40 chars, 128-character for strings under 80 chars). Typed library methods are exposed per language, including `encryptInt` / `encryptLong` / `encryptDate` in Java and `CipherInt64` / `CipherDate` / `CipherDateTime` in Go, supported from C/C++ 2.3.0.0, .NET 2.4.0, Go 2.3.0, Java 2.3.0, Node.js 2.4.0, PHP 2.1.0, and Python 2.4.0. A follow-up release extended `date` and `datetime` datasets to accept input character sets other than `0-9`.

Encrypt a 64-bit integer column in place without changing its column type — useful for protecting numeric IDs or financial figures in an existing schema.

java

```java
// Java: encrypt a long value using a 64-bit Integer dataset
long encrypted = ubiqEncrypt.encryptLong("my-integer-dataset",
1234567890L);
long decrypted = ubiqDecrypt.encryptLong("my-integer-dataset",
encrypted);
```

Encrypt a timestamp field in Go while keeping the value as a native date-time type, so no schema changes are needed.

```go
// Go: encrypt a time.Time value using a Date & Time dataset
encrypted, err := ubiqEncrypt.CipherDateTime("my-datetime-dataset", originalTimestamp)
decrypted, err := ubiqDecrypt.CipherDateTime("my-datetime-dataset", encrypted)
```

Create a Token dataset in the Dataset Creation Wizard when you need fixed-length, base62-safe tokenized values — for example, replacing PII in a database with 64-character tokens for values under 40 characters.

> 📍 In the Ubiq console, go to Datasets › Create Dataset, select `Structured` › `Token`, and set Token length to 64 or 128 characters.

— Exact ranges, method names per language, and worked examples given.

changelog-20260923-87e010a1

---

02 **Airflow Operators for structured encrypt/decrypt in DAGs** NEW

86

Adds `UbiqEncryptOperator` and `UbiqDecryptOperator`, importable from `plugins/operators/`, for performing structured encrypt and decrypt operations inside Airflow DAGs, accepting tabular data as a list of lists. Supports per-column dataset mapping via the `dataset_names` parameter to selectively encrypt specific columns while leaving others untouched, and integrates with Ubiq credentials via ubiq.configCredentials(config_file=..., profile=...) to scope each DAG to specific datasets and permissions.

Encrypt sensitive columns (SSN, Birthdate) from a BigQuery table pull inside an Airflow DAG, then decrypt in a follow-on task — leaving non-sensitive columns untouched.

```python
import ubiq_security as ubiq
from operators.ubiq_encrypt_operator import UbiqEncryptOperator
from operators.ubiq_decrypt_operator import UbiqDecryptOperator

creds = ubiq.configCredentials(config_file='ubiq-credentials',
profile='default')

enc_task = UbiqEncryptOperator(
    task_id='enc_task',
    credentials=creds,
    data_task=get_data.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate']
)

dec_task = UbiqDecryptOperator(
    task_id='dec_task',
    credentials=creds,
    data_task=enc_task.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate']
)

get_data >> enc_task >> dec_task
```

— Full DAG code example names operators, parameter, and credential call. product docs

03

Structured encryption builder and searchable integer encryption (JS SDK)

NEW

75

Adds buildStructuredAsync() on `CryptographyBuilder` to construct a structured encryption/decryption object supporting all structured dataset types, and `encryptNumberForSearchAsync` for searchable encryption over integer datasets (32-bit and 64-bit), letting an application encrypt a value and generate matching search tokens to query a database without revealing the plaintext.

Encrypt and search an integer field (e.g. a customer ID) stored in a database without revealing the plaintext — requires a pre-configured int32 dataset.

```javascript
const credentials =
ubiq.UbiqFactory.readCredentialsFromFile('./credentials',
'default');
const configuration = ubiq.UbiqFactory.defaultConfiguration();
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

const tweakFF1 = [];
const encrypted = await
ubiqEncryptDecrypt.EncryptNumberAsync('int32', 12345, tweakFF1);
const searchTokens = await
ubiqEncryptDecrypt.encryptNumberForSearchAsync('int32', 12345,
tweakFF1);
// Use searchTokens to query the database for all matching
encrypted values
```

— Method names, bit-widths, and a runnable code sample given. product docs

04 **IdP-based credential authentication (JS SDK)** NEW 65

Adds createCredentialsWithIdp('<username>', '<password>', '<host>') to authenticate against Ubiq via an identity provider instead of static API keys, for use with CryptographyBuilder.buildStructuredAsync().

Authenticate via an IdP instead of static API keys when running in environments with SSO or federated identity.

```javascript
const credentials =
ubiq.UbiqFactory.createCredentialsWithIdp('<username>',
'<password>', '<host>');
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)

.withConfigurationObject(ubiq.UbiqFactory.defaultConfiguration())
  .buildStructuredAsync();
```

— Function signature and working example provided. product docs

05 **Custom metadata tagging on usage reporting (JS SDK)** NEW

65

Adds addReportingUserDefinedMetadata() on the structured encryption object to tag encryption usage records with a JSON payload (e.g. tenant or job identifier) for audit and billing attribution.

Tag all encryption usage records with a tenant or job identifier for audit and billing attribution.

```javascript
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

ubiqEncryptDecrypt.addReportingUserDefinedMetadata('{"tenant":
"acme-corp", "job": "nightly-batch"}');
```

— Method name and example payload given. product docs

06 **Key number retrieval from ciphertext (Go SDK)** NEW

60

Adds ubiq.GetKeyNumber(ciphertext) in Go to retrieve the key number behind a ciphertext without performing a full decryption, useful for key-rotation auditing or ciphertext provenance checks.

Retrieve the key number behind a ciphertext in Go without performing a full decryption — useful for key-rotation auditing or ciphertext provenance checks.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf('Key number: %d\n', keyNum)
```

— Named function with runnable Go example, limited further detail.

changelog-20260923-87e010a1

THINNER COVERAGE BELOW

07 **New Rust SDK** NEW

45

Adds a Rust SDK (`ubiq-rust`) for encrypting structured or unstructured data, including date/time fields and search-optimised ciphertext.

Add the Ubiq Rust SDK to a new project to encrypt structured or unstructured data, including date/time fields and search-optimised ciphertext.

```toml
# Cargo.toml
[dependencies]
ubiq-rust = '*'
```

— Only the Cargo dependency is shown, no further usage detail.

changelog-20260923-87e010a1

08 **Multithreaded environment guidance for NodeJS SDK** IMPROVED

25

Documents that Ubiq objects must be created and destroyed per worker thread in multithreaded NodeJS environments, because NodeJS workers serialize/copy objects rather than sharing them.

— Documentation note only, no code example given. product docs

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Ubiq Security: `date` and `datetime` datasets now work with input character sets other than `0-9`.

└──▷ **USE IT**

Retrieve the key number behind a ciphertext in Go without performing a full decryption — useful for key-rotation auditing or ciphertext provenance checks.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf('Key number: %d\n', keyNum)
```

Add the Ubiq Rust SDK to a new project to encrypt structured or unstructured data, including date/time fields and search-optimised ciphertext.

```Cargo.toml
# Cargo.toml
[dependencies]
ubiq-rust = '*'
```

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.28.731 (2026-09-23) 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 biggest shift this window is a full self-hosted Kubernetes deployment path (Runlayer Operator on EKS/GKE) with workload identity, external dependency wiring, MCP Catalog/AI Watch/OTLP support, alongside Agent Accounts, new Claude Opus 5.5/GPT-6/Gemini model support, channel-only Slack agents, a Figma MCP connector, and client-provenance policy conditions.

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

01 **Microsoft Teams bot integration for agents** NEW 96

The ECS Terraform module adds Microsoft Teams support via five variables — `ms_teams_app_client_secret` (recommended via `TF_VAR_ms_teams_app_client_secret`), `MS_TEAMS_APP_CLIENT_ID`, `MS_TEAMS_RUNLAYER_TENANT_ID`, `MS_TEAMS_AZURE_SUBSCRIPTION_ID`, `MS_TEAMS_AZURE_RESOURCE_GROUP`, and `MS_TEAMS_APP_URL` — which together provision a per-agent Entra app registration and Azure Bot Service resource; the integration stays off unless all five values are set, and admin consent requires registering `https://<domain_name>/api/v1/integrations/ms-teams/admin-consent-callback` as a Redirect URI. `MS_TEAMS_APP_URL` is effectively immutable once agents are provisioned — each agent's Azure Bot resource keeps the messaging endpoint captured at creation time, so changing the domain orphans existing bots until they are deleted and recreated.

Enable Microsoft Teams bot integration on a self-managed ECS deployment by supplying all five required Teams variables to the Terraform module.

```hcl
module "runlayer" {
  source                       = "runlayer/ecs"
  version                      = ">= 6.64.0, < 7.0.0"
  ms_teams_app_client_secret   = var.ms_teams_app_client_secret
# via TF_VAR_ms_teams_app_client_secret
  # remaining four values supplied as environment variables:
  # MS_TEAMS_APP_CLIENT_ID              = "00000000-0000-0000-
0000-000000000000"
  # MS_TEAMS_RUNLAYER_TENANT_ID         = "00000000-0000-0000-
0000-000000000000"
  # MS_TEAMS_AZURE_SUBSCRIPTION_ID      = "00000000-0000-0000-
0000-000000000000"
  # MS_TEAMS_AZURE_RESOURCE_GROUP       = "runlayer-ms-teams-
bots"
  # MS_TEAMS_APP_URL                    = "https://<domain_name>"
  # ... other required variables
}
```

Enable Microsoft Teams integration by supplying all five required values; pass the client secret via environment variable to avoid storing it in state.

```hcl
# terraform.tfvars
MS_TEAMS_APP_CLIENT_ID          = "00000000-0000-0000-0000-
000000000000"
MS_TEAMS_RUNLAYER_TENANT_ID     = "00000000-0000-0000-0000-
000000000000"
MS_TEAMS_AZURE_SUBSCRIPTION_ID  = "00000000-0000-0000-0000-
000000000000"
MS_TEAMS_AZURE_RESOURCE_GROUP   = "runlayer-ms-teams-bots"
MS_TEAMS_APP_URL                = "https://runlayer.example.com"
# export TF_VAR_ms_teams_app_client_secret=<secret> before
running terraform apply
```

— Full variable list, redirect URI, tfvars example, and breaking caveat. product docs

02 **Workload identity mapping for Operator ServiceAccounts** NEW

90

Adds a `spec.serviceAccount` field in the `RunlayerInstance` CR to annotate per-component Kubernetes ServiceAccounts with IRSA role ARNs (`eks.amazonaws.com/role-arn`) on EKS or GSA emails (`iam.gke.io/gcp-service-account`) on GKE, with no long-lived cloud keys embedded, plus `spec.deploy.platformNamespace` to override the default `runlayer-`

`<customerId>` namespace. Dedicated ServiceAccounts are provisioned per component — `{instance}-backend`, `{instance}-worker`, `{instance}-audit-consumer`, `{instance}-siem-export`, `{instance}-session-materializer`, `{instance}-toolguard` — each independently annotatable, with a published permission matrix covering S3/GCS, Kinesis/Pub/Sub, DynamoDB checkpoints, Bedrock/Vertex LLM invocation, and `bedrock-agentcore:InvokeAgentRuntime`.

Wire IRSA roles to each Runlayer component on EKS so workloads assume cloud permissions without embedded credentials.

```yaml
spec:
  serviceAccount:
    backend:
      annotations:
        eks.amazonaws.com/role-arn:
arn:aws:iam::123456789012:role/runlayer-acme-backend
    worker:
      annotations:
        eks.amazonaws.com/role-arn:
arn:aws:iam::123456789012:role/runlayer-acme-worker
    auditConsumer:
      annotations:
        eks.amazonaws.com/role-arn:
arn:aws:iam::123456789012:role/runlayer-acme-audit-consumer
    siemExport:
      annotations:
        eks.amazonaws.com/role-arn:
arn:aws:iam::123456789012:role/runlayer-acme-siem-export
    sessionMaterializer:
      annotations:
        eks.amazonaws.com/role-arn:
arn:aws:iam::123456789012:role/runlayer-acme-session-materializer
```

Verify that the IRSA annotation is present on the ServiceAccount and that the pod actually assumes the expected role.

```
$ kubectl -n runlayer-acme get sa acme-prod-backend -o yaml
# expect eks.amazonaws.com/role-arn in annotations
kubectl -n runlayer-acme exec deploy/acme-prod-backend -c backend
-- aws sts get-caller-identity
# Arn should be the IRSA role
```

Deploy Runlayer into a custom namespace instead of the default `runlayer-<customerId>` when your cluster naming conventions require it.

```yaml
spec:
  deploy:
    platformNamespace: runlayer-acme-prod
```

— Exact CR fields, annotations, and verification commands given. product docs

03

Audit streaming, SIEM export, and session materializer pipelines

NEW

90

Enables optional audit event streaming to Kinesis or GCP Pub/Sub wired into `components.auditConsumer` with `streamBackend: kinesis` or `streamBackend: pubsub` (plus `spec.gcp.projectId` for GCP), and optional SIEM export as an independent consumer on the same stream wired into `components.siemExport` with `streamBackend: kinesis`, `kinesis.streamArn`, `kinesis.consumerArn`, `kinesis.checkpointTableName`, `kinesis.deadLetterBucket`, and `export.s3Bucket`. Also enables an optional session materializer component (`{instance}-session-materializer`) backed by a Kinesis hook-events stream, DynamoDB checkpoint table, S3 session-payloads bucket, and S3 DLQ on AWS, or an ordered Pub/Sub subscription and GCS buckets on GCP.

Wire audit streaming and SIEM export into a RunlayerInstance CR after provisioning the Kinesis stream, EFO consumers, DynamoDB checkpoint tables, and DLQ/destination buckets out-of-band.

```yaml
spec:
  components:
    auditConsumer:
      enabled: true
      streamBackend: kinesis
      kinesis:
        streamArn: arn:aws:kinesis:us-east-
1:123456789012:stream/myorg-audit-log
        consumerArn: arn:aws:kinesis:us-east-
1:123456789012:stream/myorg-audit-log/consumer/myorg-audit-log-
consumer
        checkpointTableName: myorg-audit-log-consumer-checkpoints
        deadLetterBucket: myaccount-myorg-audit-log-consumer-dlq
    siemExport:
      enabled: true
      streamBackend: kinesis
      kinesis:
        streamArn: arn:aws:kinesis:us-east-
1:123456789012:stream/myorg-audit-log
        consumerArn: arn:aws:kinesis:us-east-
1:123456789012:stream/myorg-audit-log/consumer/myorg-siem-export
        checkpointTableName: myorg-siem-export-checkpoints
        deadLetterBucket: myaccount-myorg-siem-export-dlq
      export:
        s3Bucket: myaccount-myorg-siem-export
```

— Full YAML example naming every stream and bucket field. product docs

04 **Kubernetes self-hosting via Runlayer Operator** NEW

83

Runlayer adds a self-hosted Kubernetes deployment path on EKS and GKE via the Runlayer Operator, requiring Kubernetes 1.29+ and Helm 3.8+, supporting one or many isolated tenants using the MCPServer CRD on the same cluster for multi-tenant platform topologies. AWS Bedrock is supported as an AI provider in the EKS path (requiring IAM permissions and Anthropic Bedrock use-case form approval); the operator and `runlayer-instance` Helm chart are published via a Customer Distribution OCI registry at `oci://088332244652.dkr.ecr.us-east-1.amazonaws.com/runlayer-helm-charts/runlayer-instance`, supporting cross-account ECR pulls for EKS and GAR/private mirror for GKE or air-gapped environments, with CRD manifests `crds/platform.runlayer.com_runlayerinstances.yaml` and `crds/deploy.runlayer.com_mcpservers.yaml` shipped for upgrades.

Apply updated CRDs before upgrading the operator to avoid schema drift between the controller and the cluster.

```
$ kubectl apply -f
crds/platform.runlayer.com_runlayerinstances.yaml \
             -f crds/deploy.runlayer.com_mcpservers.yaml
```

— Named CRDs, registry path, and version requirements given. product docs

05 **MCP Catalog API onboarding for Operator deployments** NEW

81

Adds the `MCP_CATALOG_API_KEY` secret, set on the platform namespace, to enable MCP Catalog onboarding; without it, `GET /api/v1/catalog/` fails and setup shows `Couldn't load setup status`. A `components.mcpCatalogApi` Helm/CR field can optionally enable an in-cluster MCP Catalog API, though docs note it should not be enabled for the standard customer path.

Inject the MCP Catalog API key into the platform namespace so onboarding can reach `GET /api/v1/catalog/` and setup status loads correctly.

```
$ kubectl create secret generic runlayer-catalog \
  --namespace runlayer-platform \
  --from-literal=MCP_CATALOG_API_KEY=<key-from-runlayer>
```

— Exact secret, endpoint, field name, and kubectl command given. product docs

06 **Client-provenance policy conditions** NEW

80

Adds `meta.client.kind` (values `dcr`, `cimd`, `pre_registered`, `manual`, `ema`, `agent`) and `meta.plugin.id` as policy condition keys so access policies can gate connectors by how an MCP client obtained its credentials and whether it reaches a connector directly or through a plugin.

Lock a connector so self-registered MCP clients can never reach it, regardless of whether they go direct or through a plugin.

📍 In the Runlayer console, go to Policies › New policy condition, set the field to `meta.client.kind`, operator to `equals`, and value to `dcr` (or any of: `cimd`, `pre_registered`, `manual`, `ema`, `agent`). Add a second condition on `meta.plugin.id` if you also want to block plugin-routed paths. Set the policy action to Deny and apply it to the target connector.

— Exact condition keys, enumerated values, console steps given. 1.28.687 (2026-09-22)

07 Private CA trust for MCP servers in ECS deployments NEW

78

The ECS Terraform module adds a `trusted_ca_pem` variable that injects a PEM certificate bundle (single or concatenated) so backend, shared worker, and MCP execution worker tasks trust private CA-signed MCP servers and HTTPS proxies.

Trust a corporate private CA for MCP servers and HTTPS proxies in your ECS Runlayer deployment by passing your CA bundle to the Terraform module.

hcl

```hcl
module "runlayer" {
  source          = "runlayer/ecs"
  version         = ">= 6.64.0, < 7.0.0"
  trusted_ca_pem  = file("${path.module}/company-ca.pem")
  # ... other required variables
}
```

Trust an internal private CA so ECS workers can reach MCP servers or HTTPS proxies signed by that CA.

hcl

```hcl
trusted_ca_pem = file("${path.module}/company-ca.pem")
```

— Named Terraform variable with runnable config snippets. product docs

08 New agent model support: Claude Opus 5.5, GPT-6 Sol/Luna, Gemini

NEW

74

Adds Claude Opus 5.5 as an agent model, available via the Anthropic API or AWS Bedrock with adaptive thinking effort settings; adds GPT-6 Sol and GPT-6 Luna as agent model options for the OpenAI provider; and supports Google Gemini through OpenAI-Compatible providers pointed at `generativelanguage.googleapis.com`, sending API keys in the format Google expects rather than as a Bearer token.

Route Gemini model calls through an OpenAI-Compatible provider so agents can use Google models without Bearer-token auth failures.

📍 In the provider settings, add an OpenAI-Compatible provider with base URL set to `generativelanguage.googleapis.com` and supply your Google API key. Runlayer will send it in the format Google expects.

— Named models and endpoint, with setup steps but no exact API call.

1.28.731 (2026-09-23)

09 **AI Watch binary distribution and caching** NEW 70

Adds AI Watch binary distribution support, requiring a Runlayer download token and a customer-owned object-storage cache with workload identity permissions; missing inputs surface as 'No AI Watch installer is currently resolved' in the UI. Optional caching is enabled via a private S3 bucket set with `BINARY_PACKAGES_S3_BUCKET` on `runlayer-app`, requiring `s3:ListBucket`, `s3:GetObject`, and `s3:PutObject` on the `{instance}-backend` and `-worker` IRSA roles.

— Named env var and IAM actions but no config example. product docs

---

10 **PostgreSQL and Redis wiring for self-hosted deployments** NEW 62

Maps PostgreSQL (Aurora/RDS or Cloud SQL) into the `RunlayerInstance` CR via `spec.database.*` and Secrets `runlayer-db` / `runlayer-redis`, and maps Redis (ElastiCache or Memorystore) via `spec.redis.*`, with `spec.redis.tls: true` when TLS/auth is enabled.

Enable Redis with TLS in the RunlayerInstance CR after provisioning an ElastiCache replication group with `transit_encryption_mode: required`.

yaml

```yaml
spec:
  redis:
    tls: true
```

— Named CR fields and secrets with a TLS example. product docs

---

11 **Configurable HPA scaling behavior for Operator components** NEW 60

Adds a configurable HPA `behavior` block with documented defaults: 1800s scale-down stabilization for fleet components and 300s for others, max 10% removal per minute, and 60s scale-up allowing up to 50% or 2 pods per minute; setting `behavior` replaces the entire default block.

— Concrete numeric defaults given, no override example. product docs

THINNER COVERAGE BELOW

12 **Channel-only Slack agents (block direct messages)** NEW 58

Adds a Direct messages to the agent / Inbound messaging toggle in an agent's Slack settings to block inbound DMs; when disabled, the agent replies that it does not accept direct messages and redirects users to mention it in a channel instead.

Restrict an agent to channel-only interaction so it never processes DMs or assistant-pane conversations.

📍 In the agent's Slack settings, go to Inbound messaging and turn off `Direct messages to the agent`.

Prevent employees from DMing a sensitive Slack agent directly — force all interactions into auditable channels.

📍 In the agent's settings, go to Slack notifications and disable `Allow direct messages`. Users who DM the agent will be prompted to mention it in a channel instead.

— Named toggle with exact UI steps, no API surface. `1.28.731 (2026-09-23)`

---

13 **Figma remote MCP connector** NEW 54

Adds a Figma integration connecting Figma's remote MCP server through Runlayer using a customer-registered OAuth app: register a Figma OAuth app, have Figma approve its client ID, paste credentials into the connector, and users connect with their own Figma account.

— OAuth setup steps named, no exact UI path or endpoint. `1.28.687 (2026-09-22)`

---

14 **Shadow AI access restrictions** NEW 41

Restricts saved Shadow AI skill and plugin files to resource managers and security or analytics reviewers; organization-wide Shadow AI REST APIs and MCP tools now require organization-wide audit log access.

— Named roles and access requirement, no endpoint or steps. `1.28.731 (2026-09-23)`

---

15 **Post-deploy connector security settings review** IMPROVED 40

Adds a recommended security settings dialog on the connector page so admins with the `Manage security policies` capability can review and apply per-tool scanner settings after deploying a connector, instead of only during the deploy wizard.

— Names required capability, no exact navigation path. `1.28.687 (2026-09-22)`

---

16 **OTLP observability support for Kubernetes** NEW 34

Adds OTLP observability support for Kubernetes deployments, requiring Kubernetes ≥ 1.29 and reachable OTLP endpoints; hostname endpoints may need extra NetworkPolicy CIDRs.

— Requirements named, no config example provided. `product docs`

---

17 **Queued agent messages** NEW 30

Agents now support queued messages: users can send follow-ups while an agent is working and they run in order; messages can be edited, reordered, or removed, and

processing can be resumed after stopping a response.

— Behavior described, no config or UI path given. 1.28.687 (2026-09-22)

18 **RDS autoscaling threshold variables for ECS** NEW 30

The ECS Terraform module adds `rds_writer_cpu_utilization_threshold`, `rds_reader_cpu_utilization_threshold`, and `rds_writer_acu_utilization_threshold` variables to control RDS autoscaling targets for ECS deployments.

— Named variables only, no defaults or example. product docs

19 **Agent Accounts with per-agent identity** NEW 26

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

— Describes capability but no config key or UI path. product docs

20 **Shareable filtered Policies view URLs** NEW 26

Adds shareable Policies view URLs so admins can bookmark and share filtered policy views, and connector notices now open Policies pre-filtered to the relevant connector.

— Feature described without endpoint or UI path. 1.28.687 (2026-09-22)

└──▷ **BREAKING ON UPGRADE**

! The `MS_TEAMS_APP_URL` (domain name) is effectively immutable once agents are provisioned: each agent's Azure Bot resource stores the messaging endpoint at creation time and is never rewritten, so changing the domain orphans existing bots until each is deleted and recreated.

! The `MS_TEAMS_APP_URL` (domain name) is effectively immutable after agents are provisioned — changing it orphans each agent's Azure Bot resource, which must be deleted and recreated to restore Teams connectivity.

! Changing `MS_TEAMS_APP_URL` after initial provisioning does not move already-provisioned bots: each agent's Azure Bot resource keeps the messaging endpoint captured at provisioning time and must be recreated.

1.28.731 (2026-09-23) BLOG ↗ LATEST STABLE

Runlayer 1.28.731 adds Claude Opus 5.5, GPT-6 Sol/Luna, channel-only Slack agents, and Shadow AI access controls.

└──▷ **HOW TO FIND IT**

Prevent employees from DMing a sensitive Slack agent directly — force all interactions into auditable channels.

📍 In the agent's settings, go to Slack notifications and disable Allow direct messages.

Users who DM the agent will be prompted to mention it in a channel instead.

Route Gemini model calls through an OpenAI-Compatible provider so agents can use Google models without Bearer-token auth failures.

📍 In the provider settings, add an OpenAI-Compatible provider with base URL set to `generativelanguage.googleapis.com` and supply your Google API key. Runlayer will send it in the format Google expects.

Restricts saved Shadow AI skill and plugin files to resource managers and security or analytics reviewers; organization-wide Shadow AI REST APIs and MCP tools now require organization-wide audit log access.

Adds support for channel-only Slack agents: admins can turn off direct messages to an agent in its Slack settings, redirecting DMs to channel mentions instead.

Adds Claude Opus 5.5 as an agent model, available via the Anthropic API or AWS Bedrock, with adaptive thinking effort settings.

Adds GPT-6 Sol and GPT-6 Luna as agent model options for the OpenAI provider.

Supports Google Gemini through OpenAI-Compatible providers pointed at `generativelanguage.googleapis.com`, sending API keys in the format Google expects rather than as a Bearer token.

`1.28.687 (2026-09-22)` BLOG ↗ STABLE

Runlayer 1.28.687 adds Figma MCP connector, client-provenance policy conditions, and queued agent messages.

└──▷ **HOW TO FIND IT**

Lock a connector so self-registered MCP clients can never reach it, regardless of whether they go direct or through a plugin.

📍 In the Runlayer console, go to Policies › New policy condition, set the field to `meta.client.kind`, operator to `equals`, and value to `dcr` (or any of: `cimd`, `pre_registered`, `manual`, `ema`, `agent`). Add a second condition on `meta.plugin.id` if you also want to block plugin-routed paths. Set the policy action to Deny and apply it to the target connector.

Adds `meta.client.kind` (values: `dcr`, `cimd`, `pre_registered`, `manual`, `ema`, `agent`) and `meta.plugin.id` as policy condition keys so access policies can gate connectors by how an MCP client obtained its credentials and whether it reaches a connector directly or through a plugin.

Adds a Figma remote MCP connector to the catalog — register a Figma OAuth app, have Figma approve its client ID, paste credentials into the connector, and users connect with their own Figma account.

> Adds a recommended security settings dialog on the connector page so admins with the `Manage security policies` capability can review and apply per-tool scanner settings after deploying a connector, instead of during the deploy wizard.
>
> Adds queued agent messages: send follow-ups while an agent is working and they run in order; messages can be edited, reordered, or removed, and resumed after stopping a response.
>
> Adds shareable Policies view URLs so admins can bookmark and share filtered policy views, and connector notices now open Policies pre-filtered to the relevant connector.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

clawscan

Feedback ⓘ v0.2.0 NOTES ↗ RANK

Composable security scanning harness for agent skills

> clawscan added user-defined scanner support (BYOS) with templating and secret redaction, a declarative JSON gate policy with an OpenClaw install adapter, and a `--platform` flag for smaller platform-specific npm tarballs, alongside two breaking changes: a Go 1.27.1 build requirement and removal of automatic `.clawscan.yml` discovery.

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

01 BYOS custom scanner support in profile config NEW

95

Supports user-defined scanners (BYOS) declared in profile config, with `{{target}}` templating, `secretEnv` redaction, and isolated working directories. Configured via `.clawscan.yml`, e.g. defining a `scanners` entry with `id`, `command: 'mytool scan {{target}}'`, and `secretEnv` keys like `MY_API_KEY`.

Define a custom scanner (BYOS) in your profile config so clawscan invokes your own tooling against `{{target}}` in an isolated directory.

yaml

```yaml
scanners:
  - id: myscanner
    command: 'mytool scan {{target}}'
    secretEnv:
      - MY_API_KEY
```

— Names config file, templating syntax, and runnable YAML example. v0.1.7

02 --platform flag for platform-specific npm tarballs NEW

85

Adds `--platform` flag to build smaller npm tarballs targeting one supported operating system and architecture, while retaining the default universal package. Example: `clawscan --platform linux/amd64`.

Build a lean, platform-specific npm tarball for deployment to a known target OS/arch, avoiding the overhead of the universal package.

```
$ clawscan --platform linux/amd64
```

— Names flag and includes a runnable example command. v0.2.0

THINNER COVERAGE BELOW

03 **Explicit config required; auto-discovery removed** BREAKING

55

Config file auto-discovery is removed; explicit config is now required. `.clawscan.yml` is no longer auto-discovered, so an explicit config path must be provided on every invocation.

— Names exact config file and the required behavior change. v0.1.7

04 **Declarative gate policy and OpenClaw install adapter** NEW

45

Adds declarative JSON gate policy for controlling artifact verdicts based on scan results, an OpenClaw install policy adapter for evaluating install-gate policies, and exit-code evaluation of install-gate policy into an artifact verdict.

— Names policy types but no config key or endpoint given. v0.1.7

05 **Docker container labeling for worker-owned cleanup** NEW

45

Labels worker-owned Docker containers with optional run and command IDs so supervisors can identify and clean up containers after cancellation.

— Describes mechanism but no flag or config key named. v0.2.0

06 **Scanner report preservation on name collisions** IMPROVED

40

Preserves every scanner report when sanitized target, profile, or custom scanner names collide with each other or generated numeric suffixes.

— Explains behavior but no actionable surface named. v0.2.0

07 **Go 1.27.1 required for source builds** BREAKING

40

Source builds now require Go 1.27.1, previously a lower version was sufficient.

— Names exact version but gives no further migration detail. v0.2.0

08 **ClawHub profile plugin support and aig scanner** NEW

35

Supports OpenClaw plugins across the ClawHub profile and adds `aig` to the production ClawHub profile.

— Names `aig` addition but lacks usage detail. v0.1.7

09 **Relyable scanner adapter** NEW 20

Adds a relyable scanner adapter providing functional re-derivation evidence.

— Bare name with minimal description of function. v0.1.7

└──▷ **BREAKING ON UPGRADE**

! Source builds now require Go 1.27.1 (previously a lower version was sufficient).

! `.clawscan.yml` is no longer auto-discovered; an explicit config path must now be provided on every invocation.

v0.2.0 NOTES ↗ LATEST STABLE

clawscan v0.2.0 adds `--platform` for single-OS npm tarballs and Docker container labeling for worker-owned cleanup.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.2.0 https://github.com/openclaw/clawscan.git
# already have the repo? check out this version:
$ git checkout v0.2.0
```

└──▷ **TRY IT**

Build a lean, platform-specific npm tarball for deployment to a known target OS/arch, avoiding the overhead of the universal package.

```
$ clawscan --platform linux/amd64
```

Adds `--platform` flag to build smaller npm tarballs targeting one supported operating system and architecture, while retaining the default universal package.

Labels worker-owned Docker containers with optional run and command IDs so supervisors can identify and clean up containers after cancellation.

Preserves every scanner report when sanitized target, profile, or custom scanner names collide with each other or generated numeric suffixes.

└──▷ **BREAKING ON UPGRADE**

! Source builds now require Go 1.27.1 (previously a lower version was sufficient).

v0.1.7 NOTES ↗ STABLE

clawscan v0.1.7 adds BYOS support, declarative JSON gate policy, OpenClaw install policy adapter, and requires explicit config.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.1.7 https://github.com/openclaw/clawscan.git
# already have the repo? check out this version:
$ git checkout v0.1.7
```

└──▷ **USE IT**

Define a custom scanner (BYOS) in your profile config so clawscan invokes your own tooling against `{{target}}` in an isolated directory.

.clawscan.yml

```
scanners:
  - id: myscanner
    command: 'mytool scan {{target}}'
    secretEnv:
      - MY_API_KEY
```

Supports user-defined scanners (BYOS) declared in profile config, with `{{target}}` templating, `secretEnv` redaction, and isolated working directories.

Adds declarative JSON gate policy for controlling artifact verdicts based on scan results.

Adds OpenClaw install policy adapter for evaluating install-gate policies.

Evaluates exit-code install-gate policy into artifact verdict.

Supports OpenClaw plugins across the ClawHub profile.

+3 more

└──▷ **BREAKING ON UPGRADE**

! `.clawscan.yml` is no longer auto-discovered; an explicit config path must now be provided on every invocation.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ changelog · Sep 23, 2026 CHANGELOG ↗ RANK

The free and flexible app for your private thoughts.

> Obsidian 1.14.2 lets you open files stored outside your vault and integrates with OS-level Open with actions, adds extra-large widgets for iOS 27, and stops syncing base font size through Obsidian Sync.

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

01 **Font size setting moved out of sync** BREAKING 65

Base font size is removed from `appearance.json` and moved to device local storage — any vault where font size was previously synced via Obsidian Sync will no longer sync this setting.

— Clear before/after but no migration steps provided changelog-20260923-e3fd1cda

THINNER COVERAGE BELOW

02 **OS 'Open with' integration** NEW 35

Obsidian can now be invoked through the operating system's Open with functionality, letting files be opened directly in Obsidian from the OS file manager.

— Only named in summary, no mechanism detail changelog-20260923-e3fd1cda

└──▷ BREAKING ON UPGRADE

! Base font size is removed from `appearance.json` and moved to device local storage — any vault where font size was previously synced via Obsidian Sync will no longer sync this setting.

changelog · Sep 23, 2026 CHANGELOG ↗ LATEST STABLE

Obsidian 1.14.2 adds open-from-outside-vault, OS Open with integration, and iOS 27 extra-large widgets

└──▷ HOW TO FIND IT

Open a Markdown incident report or runbook stored outside your vault without moving or copying it in.

📍 In Obsidian, open the Command Palette and run 'Open file from outside the vault...', then navigate to the file on your filesystem and select it.

└──▷ BREAKING ON UPGRADE

> ! Base font size is removed from `appearance.json` and moved to device local storage — any vault where font size was previously synced via Obsidian Sync will no longer sync this setting.

WAS THIS USEFUL? 👍 👎

Syslifters SysReptor

Feedback ⓘ 2026.82 NOTES ↗ RANK

A customizable and powerful penetration testing reporting platform for offensive security professionals. Simplify, customize, and automate your pentest reports with ease.

> SysReptor 2026.82 introduces AI agent instructions and skills, adds an Audit Trail for sensitive actions, and enables importing multiple files at once.

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

01 **AI agent instructions and skills** NEW 48

Projects can now include an `AGENTS.md` file to define per-project AI agent instructions alongside report content, and the AI agent gains skills support for extended, composable automation capabilities.

— Names config file but lacks mechanism detail on skills 2026.82

02 **Audit Trail for sensitive actions** NEW 23

Adds an Audit Trail to track privileged operations across the platform.

— Brief description with no scope or access details 2026.82

03 **Multi-file import support** IMPROVED 18

Supports importing multiple files at once.

— Single sentence with no target or format specifics 2026.82

2026.82 NOTES ↗ LATEST STABLE

SysReptor 2026.82 adds AI agent skills, AGENTS.md project instructions, and an Audit Trail for sensitive actions.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch 2026.82
https://github.com/Syslifters/sysreptor.git
# already have the repo? check out this version:
$ git checkout 2026.82
```

Adds `AGENTS.md` support in projects so per-project AI agent instructions can be defined alongside report content.

Adds skills support to the AI agent for extended, composable automation capabilities.

> Adds an Audit Trail for sensitive actions to track privileged operations across the platform.
>
> Supports importing multiple files at once.

WAS THIS USEFUL?  

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ WAF - WAF Release - 2026-09-22 (2026-09-22)

BLOG ↗ RANK

Connect, protect, and build everywhere.

> Cloudflare shipped per-PR Worker Previews with isolated Durable Objects and Containers, expanded Rules expression power with 32-argument concat(), added private-network MCP server support and Cursor Origin builds to Workers, and tightened WAF managed rules against SSRF and Jinja SSTI.

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

01 **New WAF managed rules for SSRF and SSTI** NEW 80

Adds four managed rules, each with action changed from Log to Block: `...5f21b651` blocks SSRF attempts using cloud and link-local non-standard IP notation, `...75cd912a` blocks SSRF attempts using local non-standard IP notation, `...0f0313d6` blocks SSRF attempts using jar HTTP loopback payloads, and `...a1ba83f6` blocks Server-Side Template Injection targeting Jinja Dangerous Globals chains.

— Names all four rule IDs and their exact behavior change

WAF - WAF Release - 2026-09-22 (2026-09…

02 **concat() supports up to 32 arguments** IMPROVED 75

Adds concat() support for up to 32 arguments in Rules expressions, enabling richer dynamic value construction such as combining hostname, ray ID, source IP, and request path into a single header value.

Build a dynamic hostname header in a Rules expression by combining multiple values — now up to 32 — with concat().

```text
concat(http.host, "-", cf.ray_id, "-", ip.src, "-",
http.request.uri.path)
```

— Includes a runnable example expression product docs

03 **Worker Previews for every pull request** NEW 70

Adds Worker Previews, giving every pull request its own URL, configuration, observability, and isolated Durable Objects and Containers.

— Describes scope and isolation but no exact command or path product docs

04 **Cursor Origin support in Workers Builds** NEW 60

Adds Cursor Origin repository support in Workers Builds for automatic deployments, branch previews, and build status reporting.

— Names the feature set but no setup steps given product docs

THINNER COVERAGE BELOW

05 **Private network MCP server support in portals** NEW 45

Adds private network MCP server support for MCP server portals, allowing portals to connect to MCP servers available only on a private network.

— Explains the capability but gives no config path or command product docs

WAF - WAF Release - 2026-09-22 (2026-09-22) BLOG ↗ LATEST STABLE

Cloudflare WAF adds four new managed rules blocking SSRF non-standard IP notations, jar loopback payloads, and Jinja SSTI.

Adds managed rule `...5f21b651` to block SSRF attempts using cloud and link-local non-standard IP notation (action changed from Log to Block).

Adds managed rule `...0f0313d6` to block SSRF attempts using jar HTTP loopback payloads (action changed from Log to Block).

Adds managed rule `...75cd912a` to block SSRF attempts using local non-standard IP notation (action changed from Log to Block).

Adds managed rule `...a1ba83f6` to block Server-Side Template Injection targeting Jinja Dangerous Globals chains (action changed from Log to Block).

WAS THIS USEFUL? 👍 👎

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

vigolium

Feedback ⓘ v0.5.0 NOTES ↗ RANK

Vigolium - High-fidelity vulnerability scanner fusing agentic AI with native speed, modularity, and precision

> Vigolium's biggest window addition is the `run probe` passive host-sweep command with TLS fingerprinting and a deterministic `surface_score`, alongside a wholesale standardization of machine output (unified `-j` envelope, NDJSON event streaming, standardized exit codes) that breaks several existing integrations. It also shipped a source-map ingestion module, expanded traffic/db querying and export flags, and gated 14 hygiene hardening advisories behind explicit opt-in.

└──▷ **WHAT SHIPPED · 37 FEATURES** most completely described first ? what's the number?

01 **Attack-surface scoring on HTTP records** NEW 100

Adds `http_records.surface_score` — a deterministic, absolute percentage of 17 attack-surface signals (auth boundary, non-standard port, dynamic origin, permissive CORS, leaked internals, API surface, input surface, etc.) written by the new `surface-scoring` passive module. Adds `db ls --min-surface`, `--sort surface_score`, and a `SURFACE` column to `db ls`; exposes `min_surface` / `min_surface_score` on the records APIs and in `-j` /JSONL, CSV export, and `inspect_record`. The same flush also populates `http_records.technology` with the full fingerprint stack for a host. `surface_score` is now a percentage of the 17-signal set rather than a 10-point-per-signal scheme, so scores from earlier scans are not comparable to new ones.

After a sweep, rank hosts by attack-surface score to prioritise which ones deserve a full scan.

```
$ vigolium db ls --sort surface_score --min-surface 40
```

— Mechanism, signal count, fields and query flags all named v0.4.6

02 **`run probe` passive host sweep command** NEW 95

Adds `vigolium run probe` (aliases: `probing`, `httpx`, `alive`, `sweep`) for passive-only host sweeps — one request per target, no fuzzing — enabling `-T hosts.txt -c 20` to reach thousands of hosts. Probe rows are labelled `probe` and filterable via `traffic --source probe`, separating sweeps from full scans within the same project.

Sweep thousands of hosts from a recon list with passive TLS fingerprinting and no fuzzing budget consumed.

```
$ vigolium run probe -T hosts.txt -c 20 --tls-probe -j
```

Group sweep traffic by status code to quickly triage redirect vs. live vs. dead targets in a project.

```
$ vigolium traffic --source probe --group-by status_code --group-limit 20
```

— Full command, aliases, flags and filtering path given v0.4.6

---

03 **NDJSON event stream via `--events ndjson`** NEW 95

Adds `--events ndjson` to `scan` / `run` / `scan-url` / `scan-request` for a live NDJSON event stream on stdout emitting `scan.started`, `phase.started|progress|finished`, `waf.block`, `waf.pacing`, `finding.new`, `error`, and `scan.finished` events, while the human console remains on stderr.

Stream structured scan events to a monitoring process in CI while keeping the human-readable log on stderr.

```
$ vigolium scan https://example.com --events ndjson 2>scan.log | jq 'select(.event == "finding.new")'
```

— Flag, event names and a runnable pipeline example given v0.4.5

---

04 **TLS handshake fingerprinting via `--tls-probe`** NEW 90

Adds `--tls-probe` to complete a TLS handshake per HTTPS endpoint and report negotiated version/cipher plus leaf certificate details (subject/CN, SANs, issuer, validity window, serial, MD5/SHA-1/SHA-256 fingerprints) under a `tls` key.

Sweep thousands of hosts from a recon list with passive TLS fingerprinting and no fuzzing budget consumed.

```
$ vigolium run probe -T hosts.txt -c 20 --tls-probe -j
```

— Exact flag and full field list named v0.4.6

---

05 **Hygiene module tagging and intensity gating** BREAKING 90

Adds `--module-tag hygiene` and `--module-id <id>` selectors to target the new `hygiene` module tag, plus a `dynamic-assessment.hygiene_modules` config key to force hardening advisories on at any intensity or off even at deep. The 14 `hygiene-tagged` hardening advisories (including `express-session-audit`) no longer run at the default intensity — they now require `--intensity deep`, `--module-tag hygiene`, or the config key.

Run hardening advisory checks (hygiene modules) on a target without needing full deep intensity — useful for a focused advisory-only sweep in CI.

```
$ vigolium scan -t example.com --module-tag hygiene
```

Force hygiene hardening advisories on at a non-deep intensity level, or suppress them entirely even at deep, via the config key.

```yaml
dynamic-assessment:
  hygiene_modules: true
```

— Flags, config key and exact advisory count named with examples `v0.5.0`

---

06 **Expanded `kit js-beautify` controls** IMPROVED 90

Adds `-o`, `--modules`, `--profile`, `--unpack-modules`, `--max-ast-nodes`, `--max-input-mb`, and `--deadline` flags to `kit js-beautify`; `-j` output gains `status`, `diagnostics`, and per-endpoint provenance. Adds `max_unpack_input_mb` (default 16 MiB) to bound the source-level unpack pass independently of the AST node ceiling.

Beautify and unpack a large bundled JS file with an increased AST node budget and a time deadline to prevent runaway scans.

```
$ vigolium kit js-beautify -o output.js --unpack-modules --max-ast-nodes 500000 --max-input-mb 50 --deadline 60s bundle.js
```

— Every new flag named with a runnable example `v0.4.9`

---

07 **Unified machine-output envelope** BREAKING 85

Unifies machine output under one `-j` envelope — `{schema_version, command, project_uuid, db_path, total, offset, limit, items, query}` — for every command, with each command's old row key kept as a deprecated alias; `db stats -j` is no longer an exception. Machine-output timestamps now carry exactly three fractional

digits plus an epoch-millisecond sibling, changing format for consumers relying on microsecond precision.

— Exact envelope schema named but migration path only implied `v0.4.5`

---

08 **Redirect chain tracking and mode control** IMPROVED 85

Adds `--redirect-mode off|same-host|same-apex|any` (default `any`) to control redirect-following scope, replacing the former `FollowRedirects` / `FollowHostRedirects` bool pair. Adds `--record-redirect-chain` to store each followed redirect hop as its own `http_records` row chained through `parent_uuid` (on by default under `run probe`), and adds `http_records.response_location` storing the 3xx `Location` header verbatim so `traffic --tree` can display redirect destinations without re-parsing stored response bytes. Redirects now halt at a login/SSO wall under every `--redirect-mode`, with each phase reporting the walls it hit; `--record-redirect-chain` no longer gates request/response re-pairing, only hop rows, changing behavior for setups relying on it to suppress re-pairing.

— Flags, columns and behavior change all named across releases `v0.4.9` `v0.4.6`

---

09 **Filter flags for `import -B`** BREAKING 85

Adds filter flags `--host`, `--path`, `--method`, `--status`, `--search`, `--from` / `--to`, `-n`, `--all-hosts`, and `--yes` to `import -B`, matching the filter set `traffic -B` already had. An unfiltered `import -B` is now refused unless `--all-hosts` is explicitly passed; existing scripts that called `import -B` without host filters will fail with a pre-flight count prompt.

— All new flags named with clear before/after behavior `v0.4.5`

---

10 **Scratch cleanup with `kit tmp-clean`** NEW 85

Adds `vigolium kit tmp-clean` subcommand to collect abandoned vigolium scratch from the system temp directory left by interrupted runs, with `--max-age` (default 6h) controlling what counts as abandoned.

Reclaim disk space from vigolium scratch directories left behind by killed or interrupted scans older than 2 hours.

```
$ vigolium kit tmp-clean --max-age 2h
```

— Subcommand, flag and default with runnable example `v0.5.0`

---

11 **Source-map ingestion module** NEW 80

Adds `sourcemap-ingest` active module that fetches exposed source maps, secret-scans the recovered originals, and ingests their routes into the scan, backed by a shared

`jstangle/sourcemap` package providing a canonical definition of a map reference, sibling candidate, and parsed map. Detection extends to CSS, `eval` strings, `SourceMap` / `X-SourceMap` response headers, and unreferenced `<bundle>.map` sibling files.

— Module and package named, no direct usage example v0.4.9

---

12 **Standardized exit codes** BREAKING 80

Standardises exit codes across all commands: `0` success, `1` error, `2` usage error, `3` `fuzz --fail-on-match`, `4` `--fail-on` gate tripped. Exit code `1` previously meant both scan crashed and scan found something; it now means only error — automation that tested `exit == 1` to detect findings must switch to `exit == 4` (for `--fail-on` gates) or `exit == 3` (for `fuzz --fail-on-match`).

— Exact codes and migration instruction given v0.4.5

---

13 **File output flag for JSON commands** NEW 80

Adds `-o` / `--output` flag to `traffic -j`, `finding -j`, and `db ls -j` to write the result document to a file and print a receipt to stdout.

Dump a filtered traffic listing to a JSON file in CI — the file gets the result document while the terminal gets a short receipt.

```
$ vigolium traffic -j --url https://api.example.com/v1/login --
url https://api.example.com/v1/token -o traffic_matches.json
```

Narrow a `db ls` query to two exact endpoints and save results to a file for further processing.

```
$ vigolium db ls -j --url https://api.example.com/auth --url
https://api.example.com/session -o db_export.json
```

— Flag named with two runnable examples v0.4.7

---

14 **Per-phase rate limiting** NEW 80

Adds a repeatable phase qualifier to the pace dials — e.g. `--rate-limit known-issue-scan=20` — so individual phases can be capped independently without affecting other modules.

Cap only the nuclei-backed known-issue phase to 20 req/s without throttling native modules when scanning a rate-sensitive target.

```
$ vigolium scan https://example.com --rate-limit known-issue-scan=20
```

— Exact flag syntax with runnable example v0.4.5

15 **`traffic --group-by` aggregation** NEW 75

Adds `traffic --group-by <field>` to count matched records by `host`, `method`, `status_code`, `response_content_type`, `source`, `scan_uuid`, `ip`, or `is_authenticated`, bounded by `--group-limit`.

Group sweep traffic by status code to quickly triage redirect vs. live vs. dead targets in a project.

```
$ vigolium traffic --source probe --group-by status_code --group-limit 20
```

— All grouping fields named with an example v0.4.6

16 **`version --json` schema drift detection** NEW 75

Adds `vigolium version --json`, which reports `schema_version` and `db_schema_version` so a consumer can detect contract drift at startup.

Assert that the vigolium schema version matches what your driver expects before starting a long scan pipeline.

```
$ vigolium version --json | jq -e '.schema_version == "expected-version" and .db_schema_version == "expected-db-version"'
```

— Command, fields and a runnable check given v0.4.5

17 **Database path validation error codes** BREAKING 70

Pointing a read command at a missing `--db` / `$VIGOLIUM_DB_PATH` path no longer creates the database and returns `total: 0` — it now exits with `source_missing`. A foreign SQLite file that previously gained vigolium's tables, or any valid SQLite file that is not a vigolium store, now exits with `source_incompatible` instead of the generic `failed` error.

— Exact error codes named, no example command v0.4.6 v0.4.7

18 **Legacy JSON envelope compatibility flag** BREAKING 70

Adds `--json-legacy-keys` flag (and `VIGOLIUM_JSON_LEGACY_KEYS=1` env var) to restore the deprecated `records` / `findings` / `scans` / `rows` / `stats` envelope keys for callers still migrating. The `-j` /JSONL envelope no longer emits these legacy keys by default.

— Flag and env var named, no worked example v0.4.6

---

19 **Scope-origin tracking across outputs** NEW 70

Adds `scans.scope_origin_mode` field, carrying the resolved host-scope mode through `--format jsonl`, `vigolium export`, and the REST scan object, backed by `config.ResolveCLIOriginMode` as the single fallback chain read by the matcher, banner, summary, snapshot, and scan row. A new Host scope reference section documents how `--scope-origin` also selects stored records.

— Field and function named but no runnable example v0.5.0

---

20 **Native tool shims for ffuf, nuclei, katana, gau, arjun** NEW 70

Adds `vigolium ffuf|nuclei|katana|gau|arjun` tool shims that accept each tool's own argv, print the translation, and run the native phase into the pinned database.

— Tools named, mechanism described, no example command v0.4.5

---

21 **Batch ingest via repeatable `-i` and directory walking** NEW 70

Adds batch ingest: `-i` is now repeatable and `--dir` / `--dir-glob` walk a directory, so N specs cost one process instead of N invocations.

— Flags named with benefit stated, no example v0.4.5

---

22 **Traffic body/header extraction subcommands** NEW 65

Adds `vigolium traffic body` and `vigolium traffic headers` subcommands to extract one side of a stored exchange to stdout or a file without re-sending the request.

— Subcommands named, no worked example given v0.4.7

---

23 **Exact URL filtering for `traffic` and `db ls`** NEW 65

Adds `--url` flag to `traffic` and `db ls` for exact URL selection, repeatable and OR-ed, applied before pagination.

Dump a filtered traffic listing to a JSON file in CI — the file gets the result document while the terminal gets a short receipt.

```
$ vigolium traffic -j --url https://api.example.com/v1/login --
url https://api.example.com/v1/token -o traffic_matches.json
```

Narrow a `db ls` query to two exact endpoints and save results to a file for further processing.

```
$ vigolium db ls -j --url https://api.example.com/auth --url
https://api.example.com/session -o db_export.json
```

— Flag behavior and examples given v0.4.7

---

24 **New `http_records` provenance fields** NEW 65

Adds `http_records.target`, `root_uuid`, and `chain_truncated` fields — exposing the submitted target, O(1) chain-root lookup, and whether a redirect chain stopped short.

— Fields named, no worked usage v0.4.9

---

25 **Stateless scan and startup performance** IMPROVED 65

A single-phase stateless scan (`run probe -t <host> -S`) is several times faster (~0.78 s to ~0.12 s) due to bounded in-memory dedup tiers and executor drain improvements. Startup is approximately 20 ms faster following the `go-runewidth` v0.0.27 to v0.0.30 upgrade.

— Concrete before/after numbers, nothing to configure v0.5.0

---

26 **Default concurrency unification** BREAKING 60

Default concurrency is now 25 everywhere; previously the CLI, pace config, REST API, and discovery source used inconsistent values of 50, 40, 50, and 50.

— Exact before/after numbers given, no action needed v0.4.7

---

27 **`--fuzz-wordlist` renamed to `--discovery-wordlist`** DEPRECATED 60

Renames `--fuzz-wordlist` to `--discovery-wordlist` (old spelling kept as a deprecated alias) to eliminate collision with `vigolium fuzz -w`.

— Exact old and new flag names given v0.4.5

THINNER COVERAGE BELOW

28 **Bundle module cap increase and prioritization** IMPROVED 55

Raises `maxBundleModules` default from 64 to 512, orders modules by endpoint likelihood before the cap, and gives each sub-scan its own AST node budget.

— Numbers and mechanism named, no direct example v0.4.9

29 **`--export-only` record-type filter** NEW 55

Adds `--export-only` on `scan` / `run` to limit `--format jsonl` output to chosen record types, reusing the `export --only` vocabulary.

— Flag named, no example or scope detail v0.4.6

30 **Composable Burp push flags on `finding`** IMPROVED 55

Enables `finding --push-to-burp` and `--to-repeater` to compose rather than being mutually exclusive, matching `replay` behaviour.

— Flags named, behavior change is thinly described v0.4.5

31 **`-S` renamed to `--stateless`** DEPRECATED 55

`-S` is now `--stateless` everywhere; on `server` / `ingest` the old meaning of `--scan-on-receive` is now a deprecated alias that emits a warning.

— Old and new flags named, minimal further detail v0.4.5

32 **`--templates-dir` flag for nuclei templates** NEW 55

Adds `--templates-dir` flag to pin the `nuclei-templates` directory per run.

— Flag named but no example or scope v0.4.5

33 **`loginsig` package relocated** BREAKING 50

`pkg/spitolas/loginsig` has moved to `pkg/authsig` — any code importing the old path will break.

— Exact old and new paths named v0.4.9

34 **`db stats` path and size fields** NEW 40

`db stats` now reports `path` and `size` fields for the source path and on-disk size.

— Fields named, very thin description v0.4.7

35 **`spa` technology tag** NEW 40

Adds `spa` to the known tech tags so modules can gate logic on client-side-routing detection.

— Tag named, thin mechanism v0.4.6

36 **`--no-response` alias for omitting bodies** NEW 40

Adds `--no-response` as an alias for `--omit-response` to skip storing response bodies on sweep runs.

— Flag named, minimal further detail v0.4.6

37 **Module registry growth** IMPROVED 25

Registry grows to 207 active and 117 passive modules.

— Just a count, no further detail v0.4.6

└──▷ **BREAKING ON UPGRADE**

! The 14 `hygiene -tagged` hardening advisories (including `express-session-audit`) no longer fire at the default intensity — scans relying on these findings at default intensity will no longer produce them without adding `--intensity deep`, `--module-tag hygiene`, or configuring `dynamic-assessment.hygiene_modules`.

! `pkg/spitolas/loginsig` has moved to `pkg/authsig` — any code importing the old path will break.

! `--record-redirect-chain` no longer gates request/response re-pairing, only hop rows — setups that relied on this flag to suppress re-pairing will behave differently.

! Pointing a read command at a missing `--db` / `$VIGOLIUM_DB_PATH` path no longer creates the database and returns `total: 0` — it now exits with `source_missing`. A foreign SQLite file that previously gained vigolium's tables now exits with `source_incompatible` instead.

! Default concurrency is now 25 everywhere; previously the CLI, pace config, REST API, and discovery source used inconsistent values of 50, 40, 50, and 50.

! `surface_score` is now a percentage of the 17-signal set rather than a 10-point-per-signal scheme; scores from earlier scans are not comparable to new ones.

! The `-j` /JSONL envelope no longer emits the legacy `records` / `findings` / `scans` / `rows` / `stats` keys by default; callers must pass `--json-legacy-keys` or set `VIGOLIUM_JSON_LEGACY_KEYS=1` to restore them.

! Exit code `1` previously meant both `scan crashed` and `scan found something`; it now means only `error`. Automation that tested `exit == 1` to detect findings must switch to `exit == 4` (for `--fail-on` gates) or `exit == 3` (for `fuzz --fail-on-match`).

! The `-j` envelope schema changed for every command: each command's payload is now under a unified `{schema_version, command, project_uuid, db_path, total, offset, limit, items, query}` structure; drivers that read the old per-command row key must migrate (the old key is retained only as a deprecated alias).

! Machine-output timestamps now carry exactly three fractional digits plus an epoch-millisecond sibling; consumers relying on microsecond-precision fractional timestamps will see a change in format.

! `--fuzz-wordlist` is renamed to `--discovery-wordlist`; the old spelling is kept as a deprecated alias, so scripts using `--fuzz-wordlist` will still run but should be updated.

! `-S` is now `--stateless` everywhere; on `server` / `ingest` the old meaning of `--scan-on-receive` is now a deprecated alias that emits a warning.

! An unfiltered `import -B` is now refused unless `--all-hosts` is explicitly passed; existing scripts that called `import -B` without host filters will fail with a pre-flight count prompt.

v0.5.0 NOTES ↗ LATEST STABLE

vigolium v0.5.0 adds scratch cleanup, hygiene module controls, and scope-origin tracking across JSONL, export, and REST

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.5.0 https://github.com/vigolium/vigolium.git
# already have the repo? check out this version:
$ git checkout v0.5.0
```

└──▷ TRY IT

Reclaim disk space from vigolium scratch directories left behind by killed or interrupted scans older than 2 hours.

```
$ vigolium kit tmp-clean --max-age 2h
```

Run hardening advisory checks (hygiene modules) on a target without needing full deep intensity — useful for a focused advisory-only sweep in CI.

```
$ vigolium scan -t example.com --module-tag hygiene
```

Force hygiene hardening advisories on at a non-deep intensity level, or suppress them entirely even at deep, via the config key.

```yaml
dynamic-assessment:
  hygiene_modules: true
```

Adds `vigolium kit tmp-clean` subcommand to collect abandoned vigolium scratch from the system temp directory left by interrupted runs, with `--max-age` (default 6h) controlling what counts as abandoned.

Adds `--module-tag hygiene` and `--module-id <id>` selectors to target the new `hygiene` module tag.

Adds `dynamic-assessment.hygiene_modules` config key to force hardening advisories on at any intensity or off even at deep.

Adds `scans.scope_origin_mode` field, carrying the resolved host-scope mode through `--format jsonl`, `vigolium export`, and the REST scan object.

Adds `config.ResolveCLIOriginMode` as the single fallback chain read by the matcher, banner, summary, snapshot, and scan row.

+4 more

└──▷ **BREAKING ON UPGRADE**

! The 14 `hygiene -tagged` hardening advisories (including `express-session-audit`) no longer fire at the default intensity — scans relying on these findings at default intensity will no longer produce them without adding `--intensity deep`, `--module-tag hygiene`, or configuring `dynamic-assessment.hygiene_modules`.

v0.4.9 NOTES ↗ STABLE

vigolium v0.4.9 adds source-map ingestion, SSO-wall redirect halting, new HTTP record fields, and major `kit js-beautify` flag expansion.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.4.9 https://github.com/vigolium/vigolium.git
# already have the repo? check out this version:
$ git checkout v0.4.9
```

└──▷ **TRY IT**

Beautify and unpack a large bundled JS file with an increased AST node budget and a time deadline to prevent runaway scans.

```
$ vigolium kit js-beautify -o output.js --unpack-modules --max-ast-nodes 500000 --max-input-mb 50 --deadline 60s bundle.js
```

Adds `sourcemap-ingest` active module that fetches exposed source maps, secret-scans the recovered originals, and ingests their routes into the scan.

Adds `http_records.target`, `root_uuid`, and `chain_truncated` fields — exposing the submitted target, O(1) chain-root lookup, and whether a redirect chain stopped short.

Adds `-o`, `--modules`, `--profile`, `--unpack-modules`, `--max-ast-nodes`, `--max-input-mb`, and `--deadline` flags to `kit js-beautify`; `-j` output gains `status`, `diagnostics`, and per-endpoint provenance.

Adds `max_unpack_input_mb` (default 16 MiB) to bound the source-level unpack pass independently of the AST node ceiling.

Adds shared `jstangle/sourcemap` package providing a single canonical definition of a map reference, sibling candidate, and parsed map.

+3 more

└──▷ **BREAKING ON UPGRADE**

! `pkg/spitolas/loginsig` has moved to `pkg/authsig` — any code importing the old path will break.

! `--record-redirect-chain` no longer gates request/response re-pairing, only hop rows — setups that relied on this flag to suppress re-pairing will behave differently.

v0.4.7 NOTES ↗ STABLE

vigolium v0.4.7 adds targeted traffic body/header extraction, per-command file output, and exact URL filtering.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.4.7 https://github.com/vigolium/vigolium.git
# already have the repo? check out this version:
$ git checkout v0.4.7
```

└──▷ **TRY IT**

Dump a filtered traffic listing to a JSON file in CI — the file gets the result document while the terminal gets a short receipt.

```
$ vigolium traffic -j --url https://api.example.com/v1/login --url
https://api.example.com/v1/token -o traffic_matches.json
```

Narrow a `db ls` query to two exact endpoints and save results to a file for further processing.

```
$ vigolium db ls -j --url https://api.example.com/auth --url
https://api.example.com/session -o db_export.json
```

Adds `vigolium traffic body` and `vigolium traffic headers` subcommands to extract one side of a stored exchange to stdout or a file without re-sending the request.

Adds `-o` / `--output` flag to `traffic -j`, `finding -j`, and `db ls -j` to write the result document to a file and print a receipt to stdout.

Adds `--url` flag to `traffic` and `db ls` for exact URL selection, repeatable and OR-ed, applied before pagination.

`db stats` now reports `path` and `size` fields for the source path and on-disk size.

└──▷ **BREAKING ON UPGRADE**

! Pointing a read command at a missing `--db` / `$VIGOLIUM_DB_PATH` path no longer creates the database and returns `total: 0` — it now exits with `source_missing`. A foreign SQLite file that previously gained vigolium's tables now exits with `source_incompatible` instead.

! Default concurrency is now 25 everywhere; previously the CLI, pace config, REST API, and discovery source used inconsistent values of 50, 40, 50, and 50.

v0.4.6 NOTES ↗ STABLE

vigolium v0.4.6 adds `run probe` host sweeps, TLS fingerprinting, redirect-chain tracking, and a deterministic `surface_score` on every record.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.4.6 https://github.com/vigolium/vigolium.git
# already have the repo? check out this version:
$ git checkout v0.4.6
```

└──▷ **TRY IT**

Sweep thousands of hosts from a recon list with passive TLS fingerprinting and no fuzzing budget consumed.

```
$ vigolium run probe -T hosts.txt -c 20 --tls-probe -j
```

After a sweep, rank hosts by attack-surface score to prioritise which ones deserve a full scan.

```
$ vigolium db ls --sort surface_score --min-surface 40
```

Group sweep traffic by status code to quickly triage redirect vs. live vs. dead targets in a project.

```
$ vigolium traffic --source probe --group-by status_code --group-limit 20
```

Adds `vigolium run probe` (aliases: `probing`, `httpx`, `alive`, `sweep`) for passive-only host sweeps — one request per target, no fuzzing, enabling `-T hosts.txt -c 20` to reach thousands of hosts.

Adds `--tls-probe` to complete a TLS handshake per HTTPS endpoint and report negotiated version/cipher plus leaf certificate details (subject/CN, SANs, issuer, validity window, serial, MD5/SHA-1/SHA-256 fingerprints) under a `tls` key.

Adds `--redirect-mode off|same-host|same-apex|any` (default `any`) to control redirect-following scope, replacing the former `FollowRedirects` / `FollowHostRedirects` bool pair.

Adds `--record-redirect-chain` to store each followed redirect hop as its own `http_records` row chained through `parent_uuid` (on by default under `run probe`).

Adds `http_records.response_location` column storing a 3xx `Location` header verbatim, enabling `traffic --tree` to display redirect destinations without re-parsing stored response bytes.

+11 more

└──▷ **BREAKING ON UPGRADE**

! `surface_score` is now a percentage of the 17-signal set rather than a 10-point-per-signal scheme; scores from earlier scans are not comparable to new ones.

! The `-j` /JSONL envelope no longer emits the legacy `records` / `findings` / `scans` / `rows` / `stats` keys by default; callers must pass `--json-legacy-keys` or set `VIGOLIUM_JSON_LEGACY_KEYS=1` to restore them.

v0.4.5 NOTES ↗ STABLE

vigolium v0.4.5 adds NDJSON event streams, tool shims, batch ingest, and a stable machine-output envelope across all commands

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.4.5 https://github.com/vigolium/vigolium.git
# already have the repo? check out this version:
$ git checkout v0.4.5
```

└──▷ **TRY IT**

Stream structured scan events to a monitoring process in CI while keeping the human-readable log on stderr.

```
$ vigolium scan https://example.com --events ndjson 2>scan.log | jq 'select(.event == "finding.new")'
```

Assert that the vigolium schema version matches what your driver expects before starting a long scan pipeline.

```
$ vigolium version --json | jq -e '.schema_version == "expected-version" and .db_schema_version == "expected-db-version"'
```

Cap only the nuclei-backed known-issue phase to 20 req/s without throttling native modules when scanning a rate-sensitive target.

```
$ vigolium scan https://example.com --rate-limit known-issue-scan=20
```

Adds `--events ndjson` to `scan` / `run` / `scan-url` / `scan-request` for a live NDJSON event stream on stdout emitting `scan.started`, `phase.started|progress|finished`, `waf.block`, `waf.pacing`, `finding.new`, `error`, and `scan.finished` events while the human console remains on stderr.

Adds a repeatable phase qualifier to the pace dials — e.g. `--rate-limit known-issue-scan=20` — so individual phases can be capped independently without affecting other modules.

Adds `vigolium ffuf|nuclei|katana|gau|arjun` tool shims that accept each tool's own argv, print the translation, and run the native phase into the pinned database.

Adds batch ingest: `-i` is now repeatable and `--dir` / `--dir-glob` walk a directory, so N specs cost one process instead of N invocations.

Adds `vigolium version --json`, which reports `schema_version` and `db_schema_version` so a consumer can detect contract drift at startup.

+6 more

└──▷ **BREAKING ON UPGRADE**

! Exit code `1` previously meant both scan crashed and scan found something; it now means only error. Automation that tested `exit == 1` to detect findings must switch to `exit == 4` (for `--fail-on` gates) or `exit == 3` (for `fuzz --fail-on-match`).

- ! The `-j` envelope schema changed for every command: each command's payload is now under a unified `{schema_version, command, project_uuid, db_path, total, offset, limit, items, query}` structure; drivers that read the old per-command row key must migrate (the old key is retained only as a deprecated alias).
- ! Machine-output timestamps now carry exactly three fractional digits plus an epoch-millisecond sibling; consumers relying on microsecond-precision fractional timestamps will see a change in format.
- ! `--fuzz-wordlist` is renamed to `--discovery-wordlist`; the old spelling is kept as a deprecated alias, so scripts using `--fuzz-wordlist` will still run but should be updated.
- ! `-S` is now `--stateless` everywhere; on `server` / `ingest` the old meaning of `--scan-on-receive` is now a deprecated alias that emits a warning.
- ! An unfiltered `import -B` is now refused unless `--all-hosts` is explicitly passed; existing scripts that called `import -B` without host filters will fail with a pre-flight count prompt.

WAS THIS USEFUL? 👍 👎

DockSec

Feedback ⓘ v2026.9.21 NOTES ↗ RANK

AI-powered Docker security scanner that explains vulnerabilities in plain English. An OWASP Lab Project.

> DockSec's latest run adds AI-driven exploit-chain correlation and EPSS-based prioritization to its findings, an `--fix` flag for inline remediation, Markdown CI reports with a PR comment mode, and supply-chain hardening via GHCR distribution, signed releases and self-scanning CI.

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

01 **Auto-remediation with --fix flag** NEW 75

Adds a `--fix` flag that emits remediation commands alongside each finding, letting a developer fix issues in the same run as the scan without switching context, e.g. `docksec scan --fix <target>` or `docksec scan myapp:latest --fix`.

Get actionable remediation commands in the same run as the scan, so you can fix issues without switching context.

```
$ docksec scan --fix <target>
```

Scan and automatically generate fix commands for every finding so a developer can remediate without leaving the terminal.

```
$ docksec scan --fix .
```

Run a scan and get actionable remediation commands printed alongside each finding — useful in a developer workflow where you want to fix issues immediately.

```
$ docksec scan --fix <target>
```

After scanning, apply the tool-suggested remediations directly instead of manually interpreting each finding.

```
$ docksec scan --fix docker-compose.yml
```

— Names exact flag and runnable command, mechanism described v2026.9.21 v2026.9.20 v2026.9.20rc3 v2026.9.20rc2

02 **Markdown report output for CI** NEW 70

Adds `--format markdown` report output sized for CI/CD pipeline logs and pull-request summaries, so results render natively in GitHub Actions logs or GitLab MR job summaries, e.g. `docksec scan --format markdown <target>` or redirected to a file with `docksec scan myapp:latest --format markdown > docksec-report.md`.

Emit a Markdown report in CI so scan results render natively in GitHub Actions logs or as a PR annotation.

```
$ docksec scan --format markdown <target>
```

Generate a markdown report in CI/CD so findings render natively in pull request comments or pipeline summaries.

```
$ docksec scan --format markdown <target>
```

Scan a Docker Compose project in CI and emit a markdown report that renders natively in pull-request comments.

```
$ docksec scan --format markdown docker-compose.yml
```

Emit a markdown report in CI so GitHub Actions or GitLab MR pipelines render findings natively in the job summary.

```
$ docksec scan myapp:latest --format markdown > docksec-report.md
```

— Exact flag and CI usage pattern given, but no rendering details v2026.9.21

v2026.9.20rc3 v2026.9.20rc2 v2026.9.20rc1

THINNER COVERAGE BELOW

03 **EPSS-based exploit-probability prioritization** NEW 50

Adds EPSS-based priority scoring to structured findings output, ranking findings by real-world exploit probability rather than raw severity volume alongside severity data.

Scan an image and immediately get a prioritized fix list ordered by EPSS score, with remediation commands printed inline.

```
$ docksec scan myapp:latest --fix
```

— Explains ranking basis but no scoring formula or thresholds v2026.9.20rc2 v2026.9.21

04 **Supply-chain distribution and verification** NEW 40

Publishes the container image to GHCR as a new distribution channel alongside existing packaging, adds signed releases for supply-chain verification of DockSec artifacts, and runs self-scanning CI with integration guides.

— Names three concrete surfaces but no verification steps given v2026.9.21 v2026.9.20

05 **AI exploit-chain correlation of findings** NEW 35

Adds AI-driven finding correlation and exploit-chain analysis that identifies exploit chains across findings, grouping related issues into attack paths rather than a flat list.

— Describes grouping behaviour but no command or output example v2026.9.21 v2026.9.20

06 **PR comment mode for scan results** NEW 30

Adds a PR comment mode that posts scan results directly as a pull-request comment.

— States capability but no flag or config to invoke it v2026.9.21

07 **Jinja2-templated report rendering** IMPROVED 25

Replaces string-substitution HTML templates with Jinja2 for report rendering.

— Names the templating engine but no user-facing effect described v2026.9.20

v2026.9.21 NOTES ↗ LATEST STABLE

DockSec v2026.9.21 adds `--fix`, AI exploit-chain correlation, EPSS prioritization, Markdown CI reports, and PR comment mode.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.9.21 https://github.com/OWASP/DockSec.git
# already have the repo? check out this version:
$ git checkout v2026.9.21
```

└──▷ TRY IT

Get actionable remediation commands in the same run as the scan, so you can fix issues without switching context.

```
$ docksec scan --fix <target>
```

Emit a Markdown report in CI so scan results render natively in GitHub Actions logs or as a PR annotation.

```
$ docksec scan --format markdown <target>
```

Adds `--fix` flag that emits remediation commands alongside each finding.

Adds `--format markdown` report output, sized for CI/CD pipeline logs and pull-request summaries.

Adds EPSS-based priority scoring to structured findings output, surfacing exploitability likelihood alongside severity.

Adds AI-driven finding correlation and exploit-chain analysis to group related issues into attack paths.

Adds PR comment mode that posts scan results directly as a pull-request comment.

+2 more

v2026.9.20 NOTES ↗ STABLE

DockSec: Adds AI-powered vulnerability correlation that identifies exploit chains across findings.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.9.20 https://github.com/OWASP/DockSec.git
# already have the repo? check out this version:
$ git checkout v2026.9.20
```

└──▷ TRY IT

Scan and automatically generate fix commands for every finding so a developer can remediate without leaving the terminal.

```
$ docksec scan --fix .
```

Replaces string-substitution HTML templates with Jinja2 for report rendering.

Adds signed releases and self-scanning CI with integration guides.

v2026.9.20rc2 NOTES ↗ PRE-RELEASE

DockSec: Introduces EPSS-based priority scoring so findings are ranked by real-world exploit probability rather than raw severity volume.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.9.20rc2
https://github.com/OWASP/DockSec.git
# already have the repo? check out this version:
$ git checkout v2026.9.20rc2
```

└──▷ TRY IT

Scan a Docker Compose project in CI and emit a markdown report that renders natively in pull-request comments.

```
$ docksec scan --format markdown docker-compose.yml
```

After scanning, apply the tool-suggested remediations directly instead of manually interpreting each finding.

```
$ docksec scan --fix docker-compose.yml
```

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 23, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue