THE CYBER TOOLCHAIN

NO. 085

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 25, 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 369 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 63 tools

▾ // THE BRIEF

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

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

Three unrelated things stand out today: Google SecOps put Gemini triage and investigation inside SOAR playbooks, so alert work can start

vim ? help

where cases already run; Docker Sandboxes moved agent work to managed cloud sandboxes without handing it the laptop; Dalfox made long XSS testing more automatable and resumable.

RESPOND

Start investigations inside the playbook instead of copying context into a separate AI workflow

The Triage and Investigation Agent is embedded directly into SOAR playbooks. That matters because enrichment, hypothesis-building, and next-step guidance can happen in the same case workflow analysts already use, instead of requiring a side-channel prompt-and-paste loop.

Google SecOps

PROTECT

Run coding agents in governed sandboxes instead of on a developer laptop

Docker Sandboxes added managed Cloud Sandboxes behind `--cloud`, lifecycle controls, network policy, a secrets store, and `sbx move`; GitGuardian SaaS added hooks that keep API credentials off the sandbox and block secrets before they reach models or tools. Together, they reduce the blast radius of agent execution and remove the habit of giving local credentials to untrusted automation.

Docker Sandboxes · GitGuardian SaaS

PROTECT

Make XSS scanning easier to resume, integrate, and trust in CI output

Dalfox added structured JSON/JSONL/TOML/SARIF output with confidence and CWE fields, managed out-of-band blind-XSS detection, richer parameter mining, and resumable state-file scans; StackHawk added mid-scan authentication re-verification. The practical gain is fewer abandoned or unauthenticated scans, and findings that can move into defect and vulnerability workflows without hand-normalizing scanner output.

Dalfox · StackHawk

GOVERN

Put approval and audit points in front of AI and MCP tool use

ConductorOne added a public REST API for AI governance, Terraform resources for MCP sources, and wider change-history tracking; bytebase added MCP access-policy controls with refusal auditing. Teams get a way

> to treat agent-access paths like managed access paths: declared, reviewed, and evidenced, rather than hidden in local configs or ad hoc tokens.
>
> ConductorOne

Does **GPT-5.5** read better? 👍 👎

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Dreadnode

Feedback ⓘ Approve or reject held tool calls in web chat

NOTES ↗ RANK

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

> Dreadnode's biggest window items are inline approve/reject controls for held tool calls in web chat and Workflows moving into open beta with CLI/SDK/UI pipelines and Slack integration, backed by a major agentic red-teaming probe expansion and broad startup-latency work that cuts agent cold-start time by 93%.

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

01 **Agentic red-teaming probe expansion** NEW 85

Adds `generate_attack` support for agentic-probes 2026 attacks in the `ai-red-teaming` agent, including transforms for argument injection, deserialization payloads, RAG poisoning, and covert channels, plus matching evidence-gated scorers. Eight new agentic red-teaming probe families cover tool-misuse RCE, advanced exfiltration, MINJA memory poisoning, MCP line-jumping, RAG optimization attacks, and indirect prompt injection via the new `IterInject` attack loop, alongside a new provisionable multiagent sandbox target `agentic-probes-sandbox` for evaluating them.

— Names exact agent, attack loop and sandbox target, no CLI invocation shown

Approve or reject held tool calls in we…

02 **Explicit autonomy step budgets** NEW 85

Adds `--max-steps N` with `--print` for explicit autonomy budgets; unattended workflows (`dn --auto`, `/auto`, `/background`, `--print`, guard sessions, and web chat Auto/Background) now default to 100 steps.

Cap an unattended red-team workflow at a specific step count to bound runaway autonomous execution.

```
$ dn --print --max-steps 50
```

— Exact flag, defaults, and a runnable command given

Approve or reject held tool calls in we…

03 **Startup and latency performance improvements** IMPROVED 75

Agent cold-start is 93% faster — median agent construction drops from ~5s to ~0.3s by deferring heavy imports (pandas, numpy, pyarrow, Optuna) until actually used. E2B sandbox pre-warming bakes server, LiteLLM, agent graph, and first-turn imports into the E2B runtime snapshot, eliminating the largest filesystem-load costs at startup. Trace storage uploads now run in a background pool so the first chat turn no longer waits on credential fetching or object uploads, and SDK scope validation now runs reads concurrently with exponential backoff and an 8-second shared deadline, reducing validation latency.

— Concrete before/after numbers, but nothing for readers to configure

Approve or reject held tool calls in we…

---

04 **Workflows in open beta** NEW 75

Introduces Workflows in open beta: define, run, and inspect multi-step agent pipelines from the CLI, SDK, and a new Workflows section in the UI, with project-scoped catalog, stable graphs, an inline CLI guide, and model selection at run time. SDK 2.0.47 adds model selection for local workflow runs.

— Names entry points and catalog features, no exact command shown

Approve or reject held tool calls in we… Workflows are now in open beta

---

05 **Inline tool-call approvals and policy status** NEW 70

Held tool calls now appear inline in the conversation transcript with full arguments, the policy reason, and Approve / Reject buttons, replacing the previous generic input card. Transcript headers also show tool policy state (active, disabled, failed, awaiting approval, or not recorded) for both live and past sessions.

— Describes UI flow and exact policy states, but no programmatic approval API

Approve or reject held tool calls in we…

---

06 **TUI runtime log switching** NEW 70

Press `F5` then `Tab` in the TUI console to switch between local process logs and the connected runtime's server logs, with retained history up to 2,000 entries, level filtering, copy, and save.

— Exact keystroke and concrete limits given Approve or reject held tool calls in we…

---

07 **Slack integration for agent sessions** NEW 70

Org owners connect a workspace via Settings → Integrations, then @mention an agent to stream a full session in-thread with stop support and a direct link back to the session in Agents → Sessions.

Connect your Dreadnode org to Slack so your team can trigger and watch agent sessions without leaving their chat workspace.

📍 In the Dreadnode UI, go to Settings → Integrations and connect your Slack workspace. Once connected, @mention the agent in any channel to stream a full

session in-thread. Use the link posted in-thread to open the session in Agents → Sessions.

— Exact UI path and interaction steps given `Workflows are now in open beta`

---

08 **Runtime capability health inspection** NEW 65

Adds a `runtime_info` built-in tool agents can call to inspect which capabilities are loaded, enabled, and healthy — including failed loads and component health — without exposing raw config or credentials. SDK 2.0.47 adds this runtime capability health inspection for agents.

— Names the tool and what it reports, no example call

`Approve or reject held tool calls in we…` `Workflows are now in open beta`

---

09 **websockets 15.0.1 minimum requirement** BREAKING 65

The SDK now requires `websockets` 15.0.1 or later; environments running `websockets` 14.x will fail at `fastmcp` initialization.

— Exact package versions given, no upgrade command shown `Workflows are now in open beta`

THINNER COVERAGE BELOW

---

10 **Tool catalog in Agent Policies** NEW 55

Adds a browsable Tool catalog in Agent Policies with searchable metadata, copyable IDs, schema fingerprints, and collapsible input schemas for every visible tool and its revision history.

— Names catalog fields, only a UI path to reach it

`Approve or reject held tool calls in we…`

---

11 **New /beta workbench information architecture** NEW 55

Adds a new workbench at `/beta` with a new information architecture across project shell, Hub catalogs (Capabilities, Environments, Datasets, Models), Agent Outputs, and Sessions.

— Lists new sections but not what changed inside each

`Approve or reject held tool calls in we…`

---

12 **Workflow run page and listing enhancements** IMPROVED 55

The workflow run page has a new header, node panel, transcript view, and graph direction controls. Associated workflow runs now also appear inline on the Workflows page with status, duration, and run ID.

— Names new UI panels, no config or API given `Approve or reject held tool calls in we…`

13 **Organization avatars** NEW 55

Owners can set an organization avatar in Settings by choosing one of five color presets, uploading a PNG, JPEG, or SVG, or falling back to the org name initial.

— Exact formats and options listed, UI path only Workflows are now in open beta

14 **Memory creation panel** NEW 40

Adds a memory creation panel and form on the Memories page so memories can be created directly from the UI.

— Names the page but no fields or limits Workflows are now in open beta

15 **Package-relative imports for capability tools** IMPROVED 35

Capability tools, hooks, and workers can now use standard package-relative imports instead of relying on flat module names.

— States the change with no migration example Approve or reject held tool calls in we…

16 **Unread release-notes badge** NEW 35

Adds an unread badge on the sidebar account button showing how many release notes have not been seen yet, clearing automatically when Release Notes is opened in Account Settings.

— Clear behaviour but minor, no config surface Approve or reject held tool calls in we…

17 **Workflow graph on capability detail page** NEW 25

The capability detail page now shows a workflow graph of how the capability's steps connect.

— Bare description with no interaction detail Approve or reject held tool calls in we…

└──▷ **BREAKING ON UPGRADE**

! The SDK now requires `websockets` 15.0.1 or later; environments running `websockets` 14.x will fail at `fastmcp` initialization.

Approve or reject held tool calls in web chat NOTES ↗ LATEST STABLE

Dreadnode adds inline tool-call approvals, eight new agentic probe families, a browsable Tool catalog, and TUI runtime log switching.

└──▷ **TRY IT**

Cap an unattended red-team workflow at a specific step count to bound runaway autonomous execution.

```
$ dn --print --max-steps 50
```

Adds `generate_attack` support for agentic-probes 2026 attacks in the `ai-red-teaming` agent, including transforms for argument injection, deserialization payloads, RAG poisoning, and covert channels, plus matching evidence-gated scorers.

Adds eight new agentic red-teaming probe families to the SDK: evidence-gated scorers and transforms covering tool-misuse RCE, advanced exfiltration, MINJA memory poisoning, MCP line-jumping, RAG optimization attacks, and indirect prompt injection via the new `IterInject` attack loop.

Adds a new provisionable multiagent sandbox target `agentic-probes-sandbox` for 2026 agentic-probe evaluation families.

Press `F5` then `Tab` in the TUI console to switch between local process logs and the connected runtime's server logs, with retained history up to 2,000 entries, level filtering, copy, and save.

Adds a browsable Tool catalog in Agent Policies with searchable metadata, copyable IDs, schema fingerprints, and collapsible input schemas for every visible tool and its revision history.

+14 more

Workflows are now in open beta NOTES ↗ STABLE

Dreadnode Workflows enter open beta with CLI/SDK/UI pipelines, Slack integration, and a built-in agent runtime health tool.

└──▷ **HOW TO FIND IT**

Connect your Dreadnode org to Slack so your team can trigger and watch agent sessions without leaving their chat workspace.

> 📍 In the Dreadnode UI, go to Settings → Integrations and connect your Slack workspace. Once connected, @mention the agent in any channel to stream a full session in-thread. Use the link posted in-thread to open the session in Agents → Sessions.

Adds a `runtime_info` built-in tool agents can call to inspect which capabilities are loaded, enabled, and healthy — including failed loads and component health — without exposing raw config or credentials.

Introduces Workflows in open beta: define, run, and inspect multi-step agent pipelines from the CLI, SDK, and a new Workflows section in the UI, with project-scoped catalog, stable graphs, inline CLI guide, and model selection at run time.

Adds Slack integration: org owners connect a workspace via Settings → Integrations, then @mention an agent to stream a full session in-thread with stop support and a direct link back to the session in Agents → Sessions.

Adds a memory creation panel and form on the Memories page so memories can be created directly from the UI.

Adds organization avatars: owners can set an avatar in Settings by choosing one of five color presets, uploading a PNG, JPEG, or SVG, or falling back to the org name initial.

└──▷ **BREAKING ON UPGRADE**

! The SDK now requires `websockets` 15.0.1 or later; environments running `websockets` 14.x will fail at `fastmcp` initialization.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian Detection Engine

Feedback ⓘ

changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

> GitGuardian's Detection Engine added ten new secret detectors and eight new analyzers, expanded validation for Splunk, Artifactory, Salesforce, and Scaleway credentials, and shipped two breaking changes to existing detectors.

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

01 **Salesforce OAuth2 Keys validation now uses instance-specific token endpoint**

IMPROVED 70

Expands the Salesforce OAuth2 Keys detector to capture the Salesforce instance host and validate credentials against that instance's OAuth token endpoint (production or sandbox) instead of the token introspection endpoint; returns `failed_to_check` when no host is found.

— Explains the before/after validation mechanism and named failure state.

changelog-20260925-60235c14

02 **Ten new secret detectors added** NEW 63

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

— Names all ten detectors but no config or usage detail given.

changelog-20260925-60235c14

03 **Scaleway Token detector catches additional env var names**

IMPROVED 60

Expands the Scaleway Token detector to catch `SCW_SECRET_KEY` and `SCALEWAY_SECRET_KEY` assignments, including when the word `scaleway` is absent from the document.

— Names exact env vars matched and the relaxed matching condition.

changelog-20260925-60235c14

04 **Figma analyzer reports folders:read scope replacing projects:read**

60

BREAKING

Updates the Figma personal access token analyzer to report the `folders:read` scope, which replaces the deprecated `projects:read` scope; workflows filtering on `projects:read` scope results will need to be updated.

— Names both scopes and the required workflow update. changelog-20260925-60235c14

THINNER COVERAGE BELOW

05 **Eight new secret analyzers added** NEW

58

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

— Lists all analyzers but no mechanism beyond naming. changelog-20260925-60235c14

06 **Custom host support for Splunk and Artifactory token checkers**

IMPROVED

50

Adds a checker for Splunk authentication tokens with custom host support, where the Splunk Authentication Token detector now also matches an optional Splunk host alongside the token, and a checker for Artifactory reference tokens with custom host support.

— Names two checkers and the host-matching behavior, no further detail.

changelog-20260925-60235c14

07 **Username Password detector supports more languages** IMPROVED

35

Expands the Username Password detector to match Spanish, Portuguese, French, and Turkish username/password assignment keywords.

— Thin one-line description with no other mechanism given. changelog-20260925-60235c14

└──▷ BREAKING ON UPGRADE

! The `generic_database_assignment_attached_port` detector has been merged into `generic_database_assignment`; any rules, suppressions, or integrations referencing `generic_database_assignment_attached_port` by name will no longer match.

! The Figma personal access token analyzer now reports the `folders:read` scope in place of the deprecated `projects:read` scope; workflows filtering on `projects:read` scope results will need to be updated.

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

GitGuardian ships 10 new detectors, 8 new analyzers, and expanded validation for Splunk, Salesforce, Artifactory, and more.

Adds a new detector for Brightcove OAuth credentials.

Adds a new detector for Infor ION API Gateway OAuth client credentials, capturing the authorization host and token endpoint when present.

Adds a new detector for generic Sanctum tokens.

Adds a new detector for Telnyx API keys with the `KEY` prefix.

Adds a new detector for Bifrost virtual keys, matched alone or together with the URL of the Bifrost deployment.

`+20 more`

└──▷ **BREAKING ON UPGRADE**

! The `generic_database_assignment_attached_port` detector has been merged into `generic_database_assignment`; any rules, suppressions, or integrations referencing `generic_database_assignment_attached_port` by name will no longer match.

! The Figma personal access token analyzer now reports the `folders:read` scope in place of the deprecated `projects:read` scope; workflows filtering on `projects:read` scope results will need to be updated.

WAS THIS USEFUL?  

GitGuardian SaaS

Feedback ⓘ announced · Sep 24, 2026

BLOG ↗ RANK

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

> GitGuardian shipped a Docker Sandbox Mixin Kit that auto-installs ggshield and wires AI agent hooks into Claude Code, Codex, Cursor, and GitHub Copilot sandboxes, keeping API credentials off the sandbox and blocking secrets before they reach models or tools.

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

01 **Docker Sandbox Mixin Kit for AI coding agents** NEW 88

GitGuardian publishes a Docker Sandbox Mixin Kit on Docker Hub and GitHub that automatically installs `ggshield` and configures agentic AI hooks when a sandbox starts, requiring no manual setup, and covers Claude Code, Codex, Cursor, and VS Code with GitHub Copilot as supported coding assistants. The kit adds `sbx secret set gitguardian` to supply a GitGuardian API key through Docker's host-side credential proxy, keeping the real token outside the sandbox microVM, and can be stacked with other kits by passing additional `--kits <targeted mixin>` arguments to `sbx`.

Store your GitGuardian API key securely outside the sandbox before launching, so the credential never lands inside the agent microVM.

```
$ sbx secret set gitguardian
```

— Names exact command, credential proxy mechanism, and stacking flag

launch-20260924-211340b3

02 **AI agent hooks for secrets detection** NEW 64

Adds a prompt-submission hook via `ggshield` that scans a developer's prompt for secrets and blocks it before it reaches the AI model, a pre-tool-use hook that scans file reads, shell commands, and MCP calls before execution and blocks the action when a secret is detected, and a post-tool-use hook that scans tool output and sends a desktop notification when a secret appears in the result.

— Describes each hook's trigger and action, but no direct command

launch-20260924-211340b3

announced · Sep 24, 2026 BLOG ↗

GitGuardian publishes a Docker Sandbox Mixin Kit that auto-installs ggshield AI hooks for Claude Code, Codex, Cursor, and GitHub Copilot agents.

└──▷ **TRY IT**

Store your GitGuardian API key securely outside the sandbox before launching, so the credential never lands inside the agent microVM.

```
$ sbx secret set gitguardian
```

Adds `sbx secret set gitguardian` to supply a GitGuardian API key through Docker's host-side credential proxy, keeping the real token outside the sandbox microVM.

Publishes a Docker Sandbox Mixin Kit on Docker Hub and GitHub that automatically installs `ggshield` and configures agentic AI hooks when a sandbox starts, requiring no manual setup.

Adds a prompt-submission hook via `ggshield` that scans a developer's prompt for secrets and blocks it before it reaches the AI model.

Adds a pre-tool-use hook via `ggshield` that scans file reads, shell commands, and MCP calls before execution and blocks the action when a secret is detected.

Adds a post-tool-use hook via `ggshield` that scans tool output and sends a desktop notification when a secret appears in the result.

+2 more

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

StackHawk

Feedback ⓘ changelog · Sep 25, 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.

> HawkScan 6.5.0 adds mid-scan authentication re-verification, hardens the one-line CLI installer with version pinning and checksum verification, and fixes a tech-flags bug that silently re-enabled unrelated scan flags.

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

01 Mid-scan authentication re-verification NEW

91

A new `app.authentication.sessionVerification` config block polls `authentication.testPath` during a scan to detect logout, with tunable `pollFrequency` (default 15), `activeScanPollFrequency` (default 5), and `maxUnauthenticatedRequestPercent` (default 5); if too much of the scan ran unauthenticated it now finishes with a coverage warning instead of silently returning a thin result.

Enable mid-scan session re-authentication so HawkScan detects and recovers from logout during an active scan, tightening the check cadence during attack phases.

yaml

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

— Names exact config keys and defaults with a runnable YAML example

changelog-20260925-a9739752

02 CLI installer version pinning and checksum verification NEW

81

The one-line installer (`curl -sSf https://download.stackhawk.com/hawk/install.sh | sh`) gains `HAWK_INSTALL_VERSION` and `HAWK_INSTALL_DIR` environment variables plus SHA-256 verification against the published manifest; `HAWK_INSTALL_DIR` defaults to `~/.hawk/bin`.

Pin the CLI installer to a specific version in CI so the pipeline is reproducible and does not silently pick up a newer release.

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

— Exact env vars, defaults, and a runnable pinned install command

changelog-20260925-a9739752

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

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

└──▷ **USE IT**

Enable mid-scan session re-authentication so HawkScan detects and recovers from logout during an active scan, tightening the check cadence during attack phases.

stackhawk.yml

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

Pin the CLI installer to a specific version in CI so the pipeline is reproducible and does not silently pick up a newer release.

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

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

Adds `HAWK_INSTALL_VERSION` and `HAWK_INSTALL_DIR` environment variables to the one-line `curl -sSf https://download.stackhawk.com/hawk/install.sh | sh` installer, with SHA-256 verification against the published manifest; `HAWK_INSTALL_DIR` defaults to `~/.hawk/bin`.

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 requested change, and writes the result.

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

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

Semgrep Guardian now reports scan status in AI coding agents, and Agentic Workflows issues are exportable as CSV via UI and public API.

Semgrep Guardian now reports status — setup confirmation, scan summaries, sign-in requirements, and server errors — directly inside AI coding agent turns.

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ

CodeQL 2.27.1 adds C and C++ query and Kotlin 2.4.20 support (2026-09-25)

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's biggest ships this window are CodeQL 2.27.1 (new queries, Kotlin 2.4.20 support, and expanded C/C++, Go and Rust data-flow models), a new enterprise-wide default policy for Copilot GA features, and proof-of-presence re-authentication for high-impact actions on Entra ID EMU enterprises.

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

01 Default policy for new Copilot features NEW

88

Copilot Business and Enterprise admins can set a global Default policy for new features under AI Controls > Copilot > Features & clients, with options Enabled, Disabled, and Let organizations decide, controlling whether current and future eligible Copilot features are available to users by default. The policy also extends to the Copilot Code Review policy on the Agents page and the MCP servers in Copilot policy. Starting October 22, any Copilot feature left Unconfigured at the enterprise level will automatically follow this global default, which may change feature availability for users without an explicit administrator decision.

Prevent all future GA Copilot features from being silently enabled for your enterprise users until an admin explicitly approves them.

> 📍 In GitHub, go to your enterprise settings › AI Controls › Copilot › Features & clients (`https://github.com/enterprises/YOUR-ENTERPRISE-SLUG/ai-controls/copilot/features`), then under Default policy for new features select Disabled .

— Exact admin path and setting values, includes effective date

Default Enablement of Copilot Features …

02 Data flow model expansions for C/C++, Go, and Rust IMPROVED

80

CodeQL 2.27.1 improves taint tracking with new flow models: taint flow for `boost::asio::ip::basic_resolver::resolve`, flow summaries for `BloombergLP::bdlbb::Blob`, and flow summaries for the `google::protobuf::MessageLite` C++ API. It adds Go 1.27 standard-library models including `bytes.CutLast`, `database/sql.ConvertAssign`,

`database/sql/driver.RowsColumnScanner.ScanColumn`, `net/url.URL.Clone`, `net/url.Values.Clone`, `strings.CutLast`, and the new `encoding/json/jsontext` package, and expands existing Go `strings` models covering `Clone`, `Cut`, `CutPrefix`, `CutSuffix`, `Fields`, `FieldsFunc`, `Join`, `Builder`, `Reader`, and `Replacer`. It also adds Rust models for `core::fmt::Write`, improving detection of tainted data written to formatted output buffers.

— Names every added API surface across three languages

CodeQL 2.27.1 adds C and C++ query and …

---

03 **Query accuracy refinements for Fastify, ASP.NET Core, and Actions**

IMPROVED

75

The `js/missing-rate-limiting` query now attributes routes correctly for Fastify servers configured through chainable methods such as fastify().withTypeProvider() and fastify().setValidatorCompiler(...). The `cs/web/missing-token-validation` query now recognizes ASP.NET Core's `AutoValidateAntiforgeryTokenAttribute` when registered as a global MVC filter through `AddControllersWithViews` and related methods. The `actions/unpinned-tag` query no longer flags actions pinned by a structurally valid `.github/workflows/actions.lock` entry for the enclosing workflow.

— Names exact query IDs, methods, and file paths changed

CodeQL 2.27.1 adds C and C++ query and …

---

04 **New CodeQL queries for C++ and C#** NEW

73

CodeQL 2.27.1 adds the `cpp/ambiguous-assignment-of-comparison` query, which flags expressions that assign a comparison result to a variable and then use that assignment as a truth value, and the `cs/linq/missed-firstordefault` query, which identifies `foreach` loops that can be rewritten using LINQ's `FirstOrDefault` method.

— Names exact query IDs and behavior, no run example shown

CodeQL 2.27.1 adds C and C++ query and …

05 **Proof-of-presence re-authentication for high-impact actions**

NEW 65

GitHub Enterprise Cloud adds proof-of-presence challenges before high-impact actions for Entra ID EMU enterprises, supporting two requirement levels — re-authentication (password-level IdP sign-in) and MFA (authenticator app or biometric as configured in Entra ID) — with custom IdP policies governing which actions trigger which level. After a successful challenge, the verified session stays elevated for two hours in the same browser session, matching the existing sudo mode session model.

— Describes levels and session duration, no admin path given

Require proof of presence for high-impa…

06 **Private NuGet registry replacement for dependency downloads**

IMPROVED 60

Private NuGet registries with the `Replaces base` option enabled in the organization-level private registry configuration now replace default NuGet feeds whenever CodeQL downloads dependencies.

— Names exact config option and effect, no example given

CodeQL 2.27.1 adds C and C++ query and …

THINNER COVERAGE BELOW

07 **Kotlin 2.4.20 support in CodeQL** IMPROVED 30

CodeQL 2.27.1 supports Kotlin 2.4.20 in its Java/Kotlin analysis.

— Bare version-support statement with no further detail

CodeQL 2.27.1 adds C and C++ query and …

└──▷ **BREAKING ON UPGRADE**

! Starting October 22, any Copilot feature left `Unconfigured` at the enterprise level will automatically follow the new global default policy (Enabled, Disabled, or Let organizations decide), which may change feature availability for users without an explicit administrator decision.

CodeQL 2.27.1 adds C and C++ query and Kotlin 2.4.20 support (2026-09-25)

BLOG ↗ LATEST STABLE

CodeQL 2.27.1 adds new C/C++ and C# queries, Kotlin 2.4.20 support, and expanded data flow models for Go, Rust, and JS.

Adds the `cpp/ambiguous-assignment-of-comparison` query to detect expressions that assign a comparison result to a variable and use the assignment as a truth value.

Adds the `cs/linq/missed-firstordefault` query to identify `foreach` loops that can be rewritten with LINQ's `FirstOrDefault` method.

Supports Kotlin 2.4.20 in CodeQL's Java/Kotlin analysis.

Adds taint flow models for `boost::asio::ip::basic_resolver::resolve`, flow summaries for `BloombergLP::bdlbb::Blob`, and flow summaries for the `google::protobuf::MessageLite` C++ API, improving C/C++ taint tracking.

Adds data flow models for Go 1.27 standard-library APIs including `bytes.CutLast`, `database/sql.ConvertAssign`, `database/sql/driver.RowsColumnScanner.ScanColumn`, `net/url.URL.Clone`, `net/url.Values.Clone`, `strings.CutLast`, and the new `encoding/json/jsontext` package.

`+6 more`

`Default Enablement of Copilot Features for Copilot Business and Enterprise (2026-09-25`

BLOG ↗ STABLE

GitHub Copilot admins can now set a global default policy for all new GA features under `AI Controls`.

└──▷ **HOW TO FIND IT**

Prevent all future GA Copilot features from being silently enabled for your enterprise users until an admin explicitly approves them.

📍 In GitHub, go to your enterprise settings › AI Controls › Copilot › Features & clients (`https://github.com/enterprises/YOUR-ENTERPRISE-SLUG/ai-controls/copilot/features`), then under `Default policy for new features` select `Disabled`.

Adds a `Default policy for new features` setting under the `AI Controls` > `Copilot` > `Features & clients` admin page, with options `Enabled`, `Disabled`, and `Let organizations decide`, controlling whether current and future eligible Copilot features are available to users by default.

Extends the global default policy to cover the Copilot Code Review policy on the `Agents` page and the MCP servers in Copilot policy, in addition to features managed on the `Features & clients` page.

└──▷ **BREAKING ON UPGRADE**

! Starting October 22, any Copilot feature left `Unconfigured` at the enterprise level will automatically follow the new global default policy (Enabled, Disabled, or Let organizations decide), which may change feature availability for users without an explicit administrator decision.

Require proof of presence for high-impact actions (2026-09-24) BLOG ↗

STABLE

GitHub Enterprise Cloud adds proof-of-presence re-authentication before high-impact actions for Entra ID EMU enterprises.

Supports two proof-of-presence requirement levels — re-authentication (password-level IdP sign-in) and MFA (authenticator app or biometric as configured in Entra ID) — allowing enterprises to set custom IdP policies governing which actions trigger which level.

After a successful proof-of-presence challenge, the verified session remains elevated for two hours in the same browser session, matching the existing sudo mode session model.

WAS THIS USEFUL?  

GitLab

Feedback ⓘ

New MCP tools help platform teams scale automation (2026-09-17) BLOG ↗ RANK

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

> GitLab 19.4 ships a broad new MCP server toolset spanning CI/CD, merge requests, repositories, projects, work items, and vulnerabilities, plus a new merge-request-created trigger for the Duo Agent Platform, while GitLab 19.1 deprecates the `stuck_or_timeout_failure` value for `retry:when` in favor of more specific reasons.

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

01 **Merge request MCP toolset** NEW 90

Adds `save_merge_request` to open and update merge requests, `get_merge_request` to return diffs, conflicts, or approvals via the `include` param (diffs default to stats-level detail unless `detail: full_patch` is passed), `list_merge_requests` scoped to group level, `save_merge_request_review` covering the full review flow (line-level or summary comments, replying in or resolving discussions, submitting a full review, triggering a Duo review, approving or unapproving), and `accept_merge_request` which merges immediately by default, waits for checks with `strategy: merge_when_checks_pass`, or adds to a merge train with `add_to_merge_train_when_checks_pass`.

— Every MCP tool and parameter named, no runnable command

New MCP tools help platform teams scale…

02 **CI/CD pipeline and job MCP tools** NEW 85

Adds `save_pipeline` MCP tool to run, retry, cancel, or rename a CI/CD pipeline, and `get_job` MCP tool that returns a job by default; passing `include: ['log']` streams the job trace paginated via `byte_offset` and `byte_limit` (512 KB per page) so an agent can fetch and diagnose a failed build log directly.

— Names tools and params but gives no runnable example

New MCP tools help platform teams scale…

03 **Created trigger for merge requests in Duo Agent Platform** NEW 80

Adds `Created` as a new trigger action for merge requests, configurable under AI > Triggers in a project, firing when someone opens an MR once GitLab generates the diff. Requires Premium or Ultimate with Duo Agent Platform enabled.

Enable a first-pass agent review the moment a merge request is opened by wiring the Created trigger in your project.

> 📍 In your project, go to AI › Triggers, select `Merge request` as the event, choose `Created` as the action, and save the flow.

— Includes exact UI path, tier requirement, and usage example

New MCP tools help platform teams scale…

---

04 **Vulnerability management MCP tools (Ultimate)** NEW

76

Adds `list_vulnerabilities` (Ultimate) to list project vulnerabilities filterable by severity and report type with cursor pagination, `get_vulnerability` (Ultimate) to return full details for a single vulnerability by ID, and `save_vulnerability` (Ultimate) consolidating dismiss, confirm, revert to detected, update severity, and create a linked issue into one tool.

— Names tools, filters, and consolidated actions clearly

New MCP tools help platform teams scale…

---

05 **Deprecation of stuck_or_timeout_failure retry reason**

DEPRECATED

70

Deprecates the `stuck_or_timeout_failure` value for `retry:when` in GitLab 19.1; existing jobs retain it for historical data only, and new jobs will not trigger retries on this reason unless the more specific replacement reasons are also configured.

— Names exact config key and value with migration behavior product docs

---

06 **Work item MCP tools superseding issue tools** NEW

65

Adds `get_work_item`, `list_work_items`, and `save_work_item` MCP tools covering single-item reads, group/project-scoped search, and create-or-update for any supported work item type, superseding the older `get_issue` and `create_issue` tools.

— Names tools and supersession but no param detail

New MCP tools help platform teams scale…

THINNER COVERAGE BELOW

---

07 **Repository exploration MCP tools** NEW

55

Adds `list_repository_tree`, `list_branches`, `list_tags`, `list_releases`, `list_commits`, `get_commit`, `add_commit`, and `fork_repository` MCP tools for repository exploration and direct interaction with history.

— Lists tool names only, no parameters or mechanism

New MCP tools help platform teams scale…

08 **Unified note MCP tool** NEW 52

Adds `save_note` MCP tool to comment on a work item or merge request and reply in an existing thread, superseding the separate `create_merge_request_note` and `create_workitem_note` tools.

— Describes consolidation, no parameters given New MCP tools help platform teams scale…

09 **Project and user MCP tools** NEW 47

Adds `get_project`, `list_projects`, `list_project_members`, and `get_user` MCP tools for reading project details, member roles, and user information for assignment and mentions.

— Names tools only, minimal mechanism described New MCP tools help platform teams scale…

└──▷ **BREAKING ON UPGRADE**

! The `stuck_or_timeout_failure` value for `retry:when` is deprecated in GitLab 19.1 and preserved for historical data only; new jobs will not trigger retries on this reason unless the more specific replacement reasons are also configured.

New MCP tools help platform teams scale automation (2026-09-17) BLOG ↗

LATEST STABLE

GitLab 19.4 adds CI/CD, merge request, vulnerability, and project MCP server tools plus a new merge-request-created trigger for agentic flows.

└──▷ **HOW TO FIND IT**

Enable a first-pass agent review the moment a merge request is opened by wiring the Created trigger in your project.

📍 In your project, go to AI › Triggers, select Merge request as the event, choose Created as the action, and save the flow.

Adds `save_pipeline` MCP tool to run, retry, cancel, or rename a CI/CD pipeline.

Adds `get_job` MCP tool that returns a job by default; passing `include: ['log']` streams the job trace paginated via `byte_offset` and `byte_limit` (512 KB per page) so an agent can fetch and diagnose a failed build log directly.

Adds `save_merge_request` MCP tool to open and update merge requests.

Adds `get_merge_request` MCP tool that returns diffs, conflicts, or approvals via the `include` param, with diffs defaulting to stats-level detail unless `detail: full_patch` is passed for full patch text.

Adds `list_merge_requests` MCP tool scoped to group level.

+10 more

WAS THIS USEFUL?  

LavaMoat

Feedback ⓘ types-v2.0.0 NOTES ↗ RANK

tools for sandboxing your dependency graph

> LavaMoat's v19/v2/v3 wave reworks the policy generation pipeline — splitting generatePolicy() from disk writes via a new `writePolicy` export, expanding the `LavaMoatPolicy.include` schema with location-based includes, and adopting `@endo/parser-pipeline` for module parsing — alongside smaller extensions to `harden` script configs and policy overrides.

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

01 generatePolicy/writePolicy split and parser-pipeline adoption

93

BREAKING

`@lavamoat/node`'s generatePolicy() is split into separate generate and write steps: it now resolves with `{ policy: LavaMoatPolicy, hasWarnings: boolean }` instead of a bare `LavaMoatPolicy`, and no longer accepts `write` or `writableFs` options. A new `writePolicy` export from `@lavamoat/node` writes the policy to disk separately. `@lavamoat/node` also switches to `@endo/parser-pipeline` for module parsing. Migration requires destructuring `policy` from the result and calling `writePolicy` explicitly.

Generate a policy and check for warnings before writing it to disk — required migration pattern after the generatePolicy() API change.

javascript

```
import { generatePolicy, writePolicy } from '@lavamoat/node';

const { policy, hasWarnings } = await generatePolicy({
projectRoot: '.' });
if (hasWarnings) console.warn('Policy generated with warnings');
await writePolicy({ policy, policyPath:
'./lavamoat/node/policy.json' });
```

— Exact API shape, flags removed, and migration code shown `react-native-lockdown-v3.0.0` `node-v2.0.0` `lavamoat-core-v19.0.0`

THINNER COVERAGE BELOW

02 Expanded LavaMoatPolicy.include with location-based includes

50

BREAKING

`LavaMoatPolicy.include` in the LavaMoat policy type definitions and in `lavamoat-core` and `@lavamoat/node` is expanded to support a wider set of values, including location-based includes for policy configuration. This is a breaking schema change: existing code typed against the old `LavaMoatPolicy.include` definition may no longer type-check, and existing policy files using the previous `include` format may need updating.

— Names the field but not the new value shapes `types-v2.0.0` `node-v2.0.0` `lavamoat-core-v19.0.0`

---

03 **Expanded script configs in harden package** IMPROVED

23

The `harden` package extends its script configs, expanding the configuration surface available for controlling how scripts are hardened.

— No specific config keys or mechanism given `harden-v0.8.0`

---

04 **Policy-override compact function in lavamoat-core** NEW

19

lavamoat-core v19 exposes a compact function for policy overrides.

— Mentioned only in summary, no name or usage given `lavamoat-core-v19.0.0`

└──▷ BREAKING ON UPGRADE

! The shape of `LavaMoatPolicy.include` has changed — existing code typed against the old `LavaMoatPolicy.include` definition may no longer type-check and will require updates.

! generatePolicy() in `@lavamoat/node` now resolves with `{ policy: LavaMoatPolicy, hasWarnings: boolean }` instead of a bare `LavaMoatPolicy`; callers must destructure `policy` from the result.

! generatePolicy() in `@lavamoat/node` no longer accepts `write` or `writableFs` options and no longer writes policy to disk; use the new `writePolicy` import from `@lavamoat/node` instead.

! generatePolicy() now resolves with `{ policy: LavaMoatPolicy, hasWarnings: boolean }` instead of a bare `LavaMoatPolicy` — callers must destructure `policy` from the resolved object.

! generatePolicy() no longer accepts `write` or `writableFs` options and no longer writes policy to disk; use `writePolicy` imported from `@lavamoat/node` instead.

! generatePolicy() in `@lavamoat/node` now resolves with `{ policy: LavaMoatPolicy, hasWarnings: boolean }` instead of a bare `LavaMoatPolicy` — callers must destructure `policy` from the result.

! generatePolicy() no longer accepts `write` or `writableFs` options and will not write policy to disk; migrate to `writePolicy` imported from `@lavamoat/node`.

! `LavaMoatPolicy.include` has an expanded schema — existing policy files using the previous `include` format may need to be updated.

types-v2.0.0 NOTES ↗ LATEST STABLE

LavaMoat types v2.0.0 expands `LavaMoatPolicy.include` with a breaking change to its shape.

└──▷ GET THIS VERSION

```
$ git clone --branch types-v2.0.0
https://github.com/LavaMoat/LavaMoat.git
# already have the repo? check out this version:
$ git checkout types-v2.0.0
```

Expands `LavaMoatPolicy.include` to support a wider set of values in the LavaMoat policy type definitions.

└──▷ BREAKING ON UPGRADE

! The shape of `LavaMoatPolicy.include` has changed — existing code typed against the old `LavaMoatPolicy.include` definition may no longer type-check and will require updates.

react-native-lockdown-v3.0.0 NOTES ↗ STABLE

LavaMoat react-native-lockdown v3.0.0 adopts @endo/parser-pipeline and splits generatePolicy() into separate generate and write steps

└──▷ GET THIS VERSION

```
$ git clone --branch react-native-lockdown-v3.0.0
https://github.com/LavaMoat/LavaMoat.git
# already have the repo? check out this version:
$ git checkout react-native-lockdown-v3.0.0
```

Adds `writePolicy` export from `@lavamoat/node` to write a policy to disk, separate from generatePolicy().

Switches `@lavamoat/node` to use `@endo/parser-pipeline` for module parsing.

└──▷ BREAKING ON UPGRADE

! generatePolicy() in `@lavamoat/node` now resolves with `{ policy: LavaMoatPolicy, hasWarnings: boolean }` instead of a bare `LavaMoatPolicy`; callers must destructure `policy` from the result.

! generatePolicy() in `@lavamoat/node` no longer accepts `write` or `writableFs` options and no longer writes policy to disk; use the new `writePolicy` import

from `@lavamoat/node` instead.

`node-v2.0.0` **NOTES ↗** **STABLE**

LavaMoat Node 2.0 adds location-based policy includes, a new `writePolicy` export, and switches to `@endo/parser-pipeline`.

└──▷ **GET THIS VERSION**

```
$ git clone --branch node-v2.0.0
https://github.com/LavaMoat/LavaMoat.git
# already have the repo? check out this version:
$ git checkout node-v2.0.0
```

└──▷ **BREAKING ON UPGRADE**

! generatePolicy() now resolves with `{ policy: LavaMoatPolicy, hasWarnings: boolean }` instead of a bare `LavaMoatPolicy` — callers must destructure `policy` from the resolved object.

! generatePolicy() no longer accepts `write` or `writableFs` options and no longer writes policy to disk; use `writePolicy` imported from `@lavamoat/node` instead.

`harden-v0.8.0` **NOTES ↗** **STABLE**

LavaMoat harden v0.8.0 extends script configs for tighter supply chain control.

└──▷ **GET THIS VERSION**

```
$ git clone --branch harden-v0.8.0
https://github.com/LavaMoat/LavaMoat.git
# already have the repo? check out this version:
$ git checkout harden-v0.8.0
```

Extends script configs in the `harden` package, expanding the configuration surface available for controlling how scripts are hardened.

`lavamoat-core-v19.0.0` **NOTES ↗** **STABLE**

lavamoat-core v19 expands `LavaMoatPolicy.include`, exposes policy-override compact function, and adopts @endo/parser-pipeline

└──▷ **GET THIS VERSION**

```
$ git clone --branch lavamoat-core-v19.0.0
https://github.com/LavaMoat/LavaMoat.git
# already have the repo? check out this version:
$ git checkout lavamoat-core-v19.0.0
```

└──▷ **USE IT**

Generate a policy and check for warnings before writing it to disk — required migration pattern after the generatePolicy() API change.

```javascript
import { generatePolicy, writePolicy } from '@lavamoat/node';

const { policy, hasWarnings } = await generatePolicy({ projectRoot:
'.' });
if (hasWarnings) console.warn('Policy generated with warnings');
await writePolicy({ policy, policyPath:
'./lavamoat/node/policy.json' });
```

└──▷ **BREAKING ON UPGRADE**

! generatePolicy() in `@lavamoat/node` now resolves with `{ policy: LavaMoatPolicy, hasWarnings: boolean }` instead of a bare `LavaMoatPolicy` — callers must destructure `policy` from the result.

! generatePolicy() no longer accepts `write` or `writableFs` options and will not write policy to disk; migrate to `writePolicy` imported from `@lavamoat/node`.

! `LavaMoatPolicy.include` has an expanded schema — existing policy files using the previous `include` format may need to be updated.

WAS THIS USEFUL? 👍 👎

sbomify

Feedback ⓘ announced · Sep 24, 2026 BLOG ↗ RANK

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

> sbomify's v26.9.0 wave touched both the platform and its CI action: a new `/api/v2/` API replacing the deprecated v1 surface, full SPDX 3 validation and automatic conversion, a Kubernetes Helm chart, native TeamCity/Jenkins/CircleCI/Travis CI support in the action, and a rebuilt findings/vulnerability backend with Trust Center certification badges.

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

01 Helm chart for Kubernetes deployment NEW

95

Ships a Helm chart at `charts/sbomify` deploying Caddy, gunicorn/uvicorn, Dramatiq workers, a cron scheduler, and a pre-traffic migration job; supports workload identity on AWS and GCP to eliminate static object-storage credentials.

Deploy sbomify to Kubernetes with a real backing database and object storage, without bundling ephemeral in-chart services.

```
$ helm install sbomify ./charts/sbomify \
  --namespace sbomify --create-namespace \
  --set app.baseUrl=https://sbom.example.com \
  --set database.host=pg.internal
```

— Runnable helm install example plus named components launch-20260924-d959c16b

02 Document publishing from sbomify-action pipelines NEW

90

Adds `DOCUMENT_FILE`, `DOCUMENT_TYPE`, `DOCUMENT_VERSION`, `DOCUMENT_NAME`, and `DOCUMENT_COMPLIANCE_SUBCATEGORY` env vars to publish arbitrary documents (pentest reports, compliance PDFs, NDAs) directly from a pipeline without manual upload.

Upload a signed pentest report from a GitHub Actions pipeline so it appears in the platform Trust Center without manual hand-off.

```yaml
- uses: sbomify/sbomify-action@v26.9.0
  env:
    TOKEN: ${{ secrets.SBOMIFY_TOKEN }}
    COMPONENT_ID: ${{ vars.DOCS_COMPONENT_ID }}
    DOCUMENT_FILE: reports/pentest-2026.pdf
    DOCUMENT_TYPE: pentest-report
    DOCUMENT_VERSION: "2026.1"
```

— Full env var list plus a runnable pipeline example launch-20260924-30bf0844

03 **Configurable SBOM upload size limit** NEW 90

Adds `DATA_UPLOAD_MAX_MEMORY_SIZE_MB` config key to set the SBOM upload size ceiling, defaulting to 100 MB and enforced at the point stated rather than silently rejected by the framework.

Raise the SBOM upload ceiling beyond the default 100 MB for environments ingesting very large Yocto image SBOMs.

```ini
DATA_UPLOAD_MAX_MEMORY_SIZE_MB=200
```

— Named config key, default value and a runnable setting launch-20260924-d959c16b

04 **Native CI platform support in sbomify-action** NEW 85

Adds native support for TeamCity (`teamcity` , priority 40), Jenkins (`jenkins` , priority 50), CircleCI (`circleci` , priority 60), and Travis CI (`travis` , priority 70) via a new `CIPlatform` abstraction for accurate VCS provenance (branch, repo URL, ref). On Jenkins it reads `CHANGE_BRANCH` in preference to `BRANCH_NAME` for PR builds; on Travis CI it reads `TRAVIS_PULL_REQUEST_BRANCH` in preference to `TRAVIS_BRANCH` ; CircleCI builds now read `CIRCLE_TAG` and `CIRCLE_PROJECT_REPONAME` so `VERSION_FROM_RELEASE_TAG` and `NORMALIZE_VERSION` work on tag-triggered jobs.

— Names every platform, priority and env var involved launch-20260924-30bf0844

05 **SPDX 3 schema validation and version handling** IMPROVED 80

sbomify-action adds SPDX 3.0 JSON-schema validation using the bundled 3.0.1 schema so SPDX 3 documents are no longer silently passed to upload without a check, and SPDX 3.1 documents receive an explicit `unsupported` error naming accepted versions instead of silence. On the platform, uploads are validated against the official SPDX 3.0.1 schema, naming validation failures rather than a generic parse error; SPDX 3.0/3.0.0 documents

are left intentionally lenient, and SPDX 3.1 receives an explicit refusal naming accepted versions (2.2, 2.3, 3.0.x).

— Names schema versions and exact accepted-version behaviour `launch-20260924-d959c16b` `launch-20260924-30bf0844`

06 **API v2 launch replacing v1** NEW 80

Launches `/api/v2/` with renamed resource paths (`artifacts` instead of `sboms`, `workspaces` instead of `teams`) to replace the deprecated `/api/v1/` surface, and exposes security advisory create/read/delete and workspace plugin enable/disable over the API, previously only reachable through the UI.

— Names the version, renamed resources and new endpoints' scope `launch-20260924-d959c16b`

07 **Container image WORKDIR default** IMPROVED 75

Adds `WORKDIR /workspace` to the container image so `docker run -v "$PWD:/workspace" … ghcr.io/sbomify/sbomify-action:v26.9.0` is a complete invocation with no `-w` flag required.

Run sbomify-action on any CI system (Jenkins, CircleCI, Travis, etc.) as a container — WORKDIR `/workspace` means no `-w` flag is needed.

```
$ docker run --rm \
  -v "$(pwd):/workspace" \
  -e SBOMIFY_TOKEN=$SBOMIFY_TOKEN \
  ghcr.io/sbomify/sbomify-action:v26.9.0
```

— Exact runnable docker command given `launch-20260924-30bf0844`

08 **Consistent API error codes** IMPROVED 65

All API error responses now carry an `error_code` field; previously around two hundred responses returned `null`, and thirty-four server-side errors returned HTTP 400 instead of 500.

— Gives concrete before/after counts but no direct action `launch-20260924-d959c16b`

09 **Vulnerability visibility toggle on Trust Center** IMPROVED 65

Adds a workspace toggle for vulnerability posture visibility on public Trust Center release pages, which previously exposed severity counts and all findings by ID, package, and version to anyone with the link with no opt-in.

— Describes the exposure fixed and the new opt-in toggle `launch-20260924-d959c16b`

10 **Findings storage rebuilt for filtering** IMPROVED 65

Stores findings in a dedicated database table enabling server-side search, severity, VEX state, and KEV filtering and pagination — replacing in-process Python aggregation over raw scan JSON blobs.

— Names the filter dimensions and the mechanism replaced launch-20260924-d959c16b

11 **ISO 27001 and SOC 2 certification badges** NEW 65

Adds ISO 27001 and SOC 2 Type I / Type II certification badge display to Trust Center workspaces, earned only when a file-backed document is attached to a published company-wide component.

— Names the specific certifications and the earning condition launch-20260924-d959c16b

12 **Deprecation signaling on v1 API** DEPRECATED 60

Every `/api/v1/` response now includes an RFC 9745 Deprecation header, and v1 operations are struck through in `/api/v1/docs`.

— Names the header standard and docs location launch-20260924-d959c16b

13 **Automatic SPDX 3 to scanner-format conversion** NEW 60

Automatically converts SPDX 3 SBOMs to a scanner-readable format for osv-scanner and Dependency Track rather than returning a skip instructing the user to run `syft convert` manually.

— Names target scanners but no direct command launch-20260924-d959c16b

14 **CSAF 2.0 discovery documents** NEW 60

Serves the three CSAF 2.0 discovery documents and populates the `security.txt` CSAF field derived from the workspace's own domain, making CSAF publications machine-discoverable for the first time.

— Names the standard and field but no path to try it launch-20260924-d959c16b

THINNER COVERAGE BELOW

15 **Vulnerability panel triage workflow** IMPROVED 55

Enables triage directly from the vulnerabilities panel with a user-chosen row count, replacing a multi-page, unfiltered five-at-a-time workflow.

— Names the UI area but not the exact controls launch-20260924-d959c16b

16 **Relative audit trail paths** IMPROVED 50

Audit trail paths under the working directory are now recorded relative to it (e.g. `src/requirements.txt`) rather than as absolute paths exposing the build machine's directory layout and username.

— Clear before/after but no config surface to act on `launch-20260924-30bf0844`

announced · Sep 24, 2026 BLOG ↗

sbomify-action v26.9.0 adds TeamCity/Jenkins/CircleCI/Travis CI support, document publishing via `DOCUMENT_FILE`, and SPDX 3 round-trip fidelity.

└──▷ **USE IT**

Upload a signed pentest report from a GitHub Actions pipeline so it appears in the platform Trust Center without manual hand-off.

.github/workflows/sbom.yml

```
- uses: sbomify/sbomify-action@v26.9.0
  env:
    TOKEN: ${{ secrets.SBOMIFY_TOKEN }}
    COMPONENT_ID: ${{ vars.DOCS_COMPONENT_ID }}
    DOCUMENT_FILE: reports/pentest-2026.pdf
    DOCUMENT_TYPE: pentest-report
    DOCUMENT_VERSION: "2026.1"
```

Run sbomify-action on any CI system (Jenkins, CircleCI, Travis, etc.) as a container — WORKDIR `/workspace` means no `-w` flag is needed.

```
$ docker run --rm \
  -v "$(pwd):/workspace" \
  -e SBOMIFY_TOKEN=$SBOMIFY_TOKEN \
  ghcr.io/sbomify/sbomify-action:v26.9.0
```

Adds `DOCUMENT_FILE`, `DOCUMENT_TYPE`, `DOCUMENT_VERSION`, `DOCUMENT_NAME`, and `DOCUMENT_COMPLIANCE_SUBCATEGORY` env vars to publish arbitrary documents (pentest reports, compliance PDFs, NDAs) directly from a pipeline without manual upload.

Adds native CI platform support for TeamCity (`teamcity`, priority 40), Jenkins (`jenkins`, priority 50), CircleCI (`circleci`, priority 60), and Travis CI (`travis`, priority 70) via a new `CIPlatform` abstraction, enabling accurate VCS provenance (branch, repo URL, ref) in SBOMs produced on those systems.

On Jenkins, reads `CHANGE_BRANCH` in preference to `BRANCH_NAME` for PR builds; on Travis CI reads `TRAVIS_PULL_REQUEST_BRANCH` in preference to `TRAVIS_BRANCH`, so

the SBOM ref is the source branch rather than the merge target or job name.

CircleCI builds now read `CIRCLE_TAG` and `CIRCLE_PROJECT_REPONAME` so `VERSION_FROM_RELEASE_TAG` and `NORMALIZE_VERSION` work correctly on tag-triggered jobs.

Adds SPDX 3.0 JSON-schema validation using the bundled 3.0.1 schema, so SPDX 3 documents are no longer silently passed to upload without a schema check.

+3 more

announced · Sep 24, 2026 BLOG ↗

sbomify: Adds `DATA_UPLOAD_MAX_MEMORY_SIZE_MB` config key to set the SBOM upload size ceiling, defaulting to 100 MB and enforced at the point stated rather than silently rejected by the framework.

└──▷ **TRY IT**

Deploy sbomify to Kubernetes with a real backing database and object storage, without bundling ephemeral in-chart services.

```
$ helm install sbomify ./charts/sbomify \
  --namespace sbomify --create-namespace \
  --set app.baseUrl=https://sbom.example.com \
  --set database.host=pg.internal
```

Raise the SBOM upload ceiling beyond the default 100 MB for environments ingesting very large Yocto image SBOMs.

.env

```
DATA_UPLOAD_MAX_MEMORY_SIZE_MB=200
```

Launches `/api/v2/` with renamed resource paths (`artifacts` instead of `sboms`, `workspaces` instead of `teams`, one path per resource) to replace the deprecated `/api/v1/` surface.

Exposes security advisory create/read/delete and workspace plugin enable/disable over the API, previously only reachable through the UI.

Ships a Helm chart at `charts/sbomify` in the repository, deploying Caddy, gunicorn/uvicorn, Dramatiq workers, a cron scheduler, and a pre-traffic migration job; supports workload identity on AWS and GCP to eliminate static object-storage credentials.

Validates SPDX 3.0.1 documents against the official SPDX 3.0.1 schema on upload, naming validation failures rather than surfacing a generic parse error; SPDX

> 3.0/3.0.0 documents are left intentionally lenient, and SPDX 3.1 receives an explicit refusal naming accepted versions (2.2, 2.3, 3.0.x).
>
> Automatically converts SPDX 3 SBOMs to a scanner-readable format for osv-scanner and Dependency Track rather than returning a skip instructing the user to run `syft convert` manually.
>
> +8 more

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Acunetix 360

Feedback ⓘ v26.9.1 – 24 Sep 2026 (2026-09-24)

BLOG ↗ RANK

Acunetix 360 scans web applications and APIs to identify security vulnerabilities through automated testing.

> Acunetix 360 added a PolyShell security check for an Adobe Commerce/Magento RCE flaw, along with secrets support in custom authentication scripts and extended scan data retention.

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

01 **PolyShell check for Adobe Commerce/Magento RCE** NEW 65

Adds a PolyShell security check (APSB25-94) that identifies Adobe Commerce and Magento installations where the unauthenticated guest-cart REST API allows attackers to upload PHP polyglot files, bypass image validation, and achieve remote code execution.

– Names exact vuln class and advisory ID but no scan config steps.

changelog-20260925-61c01d8d v26.9.1 – 24 Sep 2026 (2026-09-24)

THINNER COVERAGE BELOW

02 **Secrets in custom authentication scripts** NEW 30

Adds support for secrets in custom authentication scripts, enabling secure credential handling during authenticated scans.

– States capability but no mechanism or config detail. v26.9.1 – 24 Sep 2026 (2026-09-24)

03 **Extended scan data retention** IMPROVED 20

Extends scan data retention, allowing scan results to be stored for longer periods.

– No concrete duration or limits given. v26.9.1 – 24 Sep 2026 (2026-09-24)

v26.9.1 – 24 Sep 2026 (2026-09-24) BLOG ↗ STABLE

Acunetix 360 v26.9.1 adds secrets in custom auth scripts, extended scan data retention, and a PolyShell security check.

Adds support for secrets in custom authentication scripts, enabling secure credential handling during authenticated scans.

Extends scan data retention, allowing scan results to be stored for longer periods.

WAS THIS USEFUL? 👍 👎

DataDome

Feedback ⓘ Varnish module 3.6.0 (2026-09-25) BLOG ↗

RANK

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

> DataDome's Varnish module now supports Varnish 8.

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

01 **Varnish 8 support in Varnish module** NEW 35

The DataDome Varnish module 3.6.0 adds support for Varnish 8.

— States only the version added, no mechanism or setup detail

Varnish module 3.6.0 (2026-09-25)

> Varnish module 3.6.0 (2026-09-25) BLOG ↗ LATEST STABLE
>
> DataDome Varnish module 3.6.0 adds support for Varnish 8.

WAS THIS USEFUL? 👍 👎

Levo

Feedback ⓘ SaaS Platform Services -- 1.2.135.3 NOTES ↗ RANK

Continuous API Security Assurance â observe, detect, and protect your APIs.

> Levo shipped broader default test coverage in `levo test`, OTP-aware web app scanning, and a one-click WSO2 API Gateway push, alongside API inventory endpoint merging, AI-assisted vulnerability triage, and several reporting and UI improvements.

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

01 **Full test category coverage by default** IMPROVED 60

Default `levo test` runs now execute every test category instead of a subset, making a standard test run materially more complete without extra configuration.

– Names the command and before/after but no category list

SaaS Platform Services -- 1.2.135.3

THINNER COVERAGE BELOW

02 **One-click push to WSO2 API Gateway** NEW 50

Adds one-click push of endpoints to WSO2 API Gateway directly from the Levo platform.

– Names the integration but not the exact UI path SaaS Platform Services -- 1.2.135.2

03 **OTP support in web app scans** NEW 40

Web app scans can now prompt for a one-time code (OTP) mid-scan, enabling authenticated scanning of MFA-protected applications.

– No config flag or setup step named SaaS Platform Services -- 1.2.135.3

04 **Skipped status for tests that could not run** IMPROVED 40

Tests that could not run are now reported as skipped rather than failed, reducing noise in security testing results.

– Clear behavior change but no flag or config named SaaS Platform Services -- 1.2.135.2

05 **Endpoint merging in API inventory** NEW 35

Duplicate API endpoints can now be merged into a single canonical endpoint in the API inventory.

– UI action described but no exact navigation given SaaS Platform Services -- 1.2.135.3

06 **AI-assisted vulnerability triage** NEW 35

Adds AI-assisted triage directly from the vulnerability view.

— Location named but mechanism unexplained SaaS Platform Services -- 1.2.135.2

07 **API Surface report filters and columns** IMPROVED 33

The API Surface report can now exclude inactive APIs from results, and gains new columns and filters requested by teams.

— New columns and filters left unnamed SaaS Platform Services -- 1.2.135.3

08 **Full-screen Service Graph page** NEW 30

Adds a dedicated full-screen page for the Service Graph.

— Bare description with no further detail SaaS Platform Services -- 1.2.135.2

09 **Redesigned PDF scan report format** IMPROVED 30

Adds a redesigned PDF report format for scan results.

— No specifics on what changed in the redesign SaaS Platform Services -- 1.2.135.2

SaaS Platform Services -- 1.2.135.3 NOTES ↗ LATEST STABLE

Levo 1.2.135.3 adds full test-category coverage by default, OTP support in web scans, endpoint merging, and API Surface report filters.

Default `levo test` runs now cover every test category instead of a subset, making a standard test run materially more complete without extra configuration.

Web app scans can now prompt for a one-time code (OTP) mid-scan, enabling authenticated scanning of MFA-protected applications.

Duplicate API endpoints can now be merged into a single canonical endpoint in the API inventory.

Inactive APIs can now be excluded from API Surface reports.

The API Surface report gains new columns and filters requested by teams.

SaaS Platform Services -- 1.2.135.2 NOTES ↗ STABLE

Levo 1.2.135.2 adds WSO2 API Gateway integration, full-screen Service Graph, AI triage, and DAST second-factor support.

Adds one-click push of endpoints to WSO2 API Gateway from the Levo platform.

Adds a dedicated full-screen page for the Service Graph.

Adds AI-assisted triage directly from the vulnerability view.

Adds a redesigned PDF report format for scan results.

Reports tests that could not run as skipped rather than failed, reducing noise in

> security testing results.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

> ElcomSoft's iOS Forensic Toolkit 10.11 extends bootloader-level, low-level extraction to new Apple hardware, pushing past the A11-chip generation barrier for the first time.

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

01 Bootloader-level extraction for Apple Watch and Apple TV NEW

55

iOS Forensic Toolkit 10.11 adds bootloader-level extraction for Apple Watch Series 4 and Series 5 and for Apple TV 4K 2nd Generation, each producing a full file system image and decrypted keychain. This extends low-level extraction support beyond the A11 generation of Apple silicon for the first time.

— Names devices and output but no CLI/UI steps changelog-20260925-852e5e9d

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

iOS Forensic Toolkit 10.11 adds bootloader-level extraction for Apple Watch Series 4/5 and Apple TV 4K 2nd Gen, yielding full file system images and decrypted keychains.

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

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

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

WAS THIS USEFUL? 👍 👎

libewf

Feedback ⓘ 20140817 CODE ↗ RANK

Legacy version of libewf

> libewf 20140817 adds FUSE3 mounting, support for the L01 format, and updates to the multi-threaded ewftools-mt, alongside a build system change moving setuptools configuration into pyproject.toml.

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

01 **Setuptools config moved to pyproject.toml** IMPROVED 40

Moves setuptools configuration into `pyproject.toml`.

– Names the exact config file but no further build detail. 20140817

02 **FUSE3 mount support** NEW 25

libewf 20140817 adds fuse3 support, allowing EWF images to be mounted via FUSE3.

– Only named as added; no mechanism or usage detail given. 20140817

03 **L01 format support** NEW 25

libewf 20140817 adds support for the L01 format.

– Format named but no details on scope or usage. 20140817

04 **ewftools-mt changes** IMPROVED 21

libewf 20140817 includes changes to ewftools-mt, the multi-threaded ewftools variant.

– Mentioned only by name with no specifics on the changes. 20140817

20140817 CODE ↗ LATEST STABLE

libewf 20140817 adds fuse3 support, L01 format support, and ewftools-mt changes

└──▷ **GET THIS VERSION**

```
$ git clone --branch 20140817 https://github.com/libyal/libewf-legacy.git
# already have the repo? check out this version:
$ git checkout 20140817
```

Moves setuptools configuration into `pyproject.toml`.

WAS THIS USEFUL?  

◆ Reverse Engineering

CHIPSEC

Feedback ⓘ 2.0.9 NOTES ↗ CODE ↗ RANK

Platform Security Assessment Framework

> CHIPSEC 2.0.9 broadens platform coverage with an ACPI NVS fallback for `SBREG_BAR` discovery and delivers a cluster of fixes to UEFI firmware decoding accuracy and diagnostics.

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

01 **UEFI decode accuracy and diagnostics** IMPROVED 71

The `uefi decode` command gains missing UEFI section-type decoders per the PI and UEFI specifications, including correct handling of `EFI_SECTION_DISPOSABLE` (type `0x03`) and fixed DEPEX opcode decoding to PI 1.9. Output writing was changed from the set_log_file() redirect pattern to direct file I/O for `.lst` artifacts, eliminating log-file conflicts where NVRAM output clobbered the `.UEFI.lst` handler, and parse warning messages now include offset and GUID context so users can trace which firmware module triggered each warning.

— Names exact command, decoder types, opcode spec, and file handling fix 2.0.9

THINNER COVERAGE BELOW

02 **ACPI NVS fallback for SBREG_BAR discovery** IMPROVED 53

Adds an ACPI NVS fallback path for `SBREG_BAR` discovery, broadening platform coverage on systems where the primary P2SB unhide path is unavailable.

— Names mechanism and fallback trigger but no command example 2.0.9

2.0.9 NOTES ↗ CODE ↗ LATEST STABLE

CHIPSEC 2.0.9 improves UEFI firmware decode with new section decoders and adds an ACPI NVS fallback for `SBREG_BAR` discovery.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 2.0.9 https://github.com/chipsec/chipsec.git
# already have the repo? check out this version:
$ git checkout 2.0.9
```

Adds an ACPI NVS fallback for `SBREG_BAR` discovery, broadening platform coverage when the primary P2SB unhide path is unavailable.

> Adds missing UEFI section-type decoders per the PI and UEFI specifications, including correct handling of `EFI_SECTION_DISPOSABLE` (type `0x03`) and fixed DEPEX opcode decoding to PI 1.9.
>
> Improves `uefi decode` output by replacing the set_log_file() redirect pattern with direct file I/O for `.lst` artifacts, eliminating log-file conflicts where NVRAM output clobbered the `.UEFI.lst` handler.
>
> Adds offset and GUID context to UEFI parse warning messages so users can trace which firmware module triggered each warning.

WAS THIS USEFUL? 👍 👎

Finite State Platform

Feedback ⓘ

September 17, 2026 - Automatic resolution for vendor-authored distro patches (2026-09-17

BLOG ↗ RANK

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

> Finite State added automatic resolution of CVE findings when a Linux distro package already includes a vendor-authored patch, backed by advisory evidence in the audit trail.

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

01 **Automatic resolution for vendor-authored distro patches** NEW

60

Finite State now auto-resolves findings when a Linux distro package already contains a vendor-authored CVE patch, recording advisory evidence in the audit trail. It supports vendor-patch recognition for Debian, Ubuntu, Alpine, Wolfi, Chainguard, MinimOS, Red Hat Enterprise Linux, AlmaLinux, Rocky Linux, and openSUSE Leap packages.

— Names mechanism and ten supported distros but no UI/API path

September 17, 2026 - Automatic resoluti…

> September 17, 2026 - Automatic resolution for vendor-authored distro patches (2026-09-
>
> BLOG ↗ LATEST STABLE
>
> Finite State now auto-resolves findings when a Linux distro package already contains a vendor-authored CVE patch, with advisory evidence in the audit trail.
>
> Supports vendor-patch recognition for Debian, Ubuntu, Alpine, Wolfi, Chainguard, MinimOS, Red Hat Enterprise Linux, AlmaLinux, Rocky Linux, and openSUSE Leap packages.

WAS THIS USEFUL? 👍 👎

Frida

Feedback ⓘ 17.19.0 NOTES ↗ CODE ↗ RANK

Main repo for hosting release binaries

> Frida 17.19.0 adds PlayStation 5 platform support and a barebone WinNT ARM64 agent to its release assets.

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

01 **PS5 platform support** NEW 35

Frida now ships prebuilt binaries for the PlayStation 5 (PS5) platform.

— Names platform but no setup or usage detail 17.19.0

02 **Barebone WinNT ARM64 agent** NEW 35

A barebone WinNT ARM64 agent is now included among the CI-shipped release assets.

— Names asset but not its capabilities or use 17.19.0

17.19.0 NOTES ↗ CODE ↗ LATEST STABLE

Frida 17.19.0 ships PS5 support and adds a barebone WinNT ARM64 agent

└──▷ **GET THIS VERSION**

```
$ git clone --branch 17.19.0 https://github.com/frida/frida.git
# already have the repo? check out this version:
$ git checkout 17.19.0
```

Adds prebuilt Frida binaries for the PlayStation 5 (PS5) platform.

Adds a barebone WinNT ARM64 agent, now included in CI-shipped assets.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

NETRESEC

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

> NetreSec introduced FlowCarp, a statistical protocol classification tool for PCAP files, adding commands to build and apply custom C2 detection models plus a public demo API covering 600+ protocols.

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

01 **FlowCarp demo classification API** NEW 75

A public demo API at `demo.flowcarp.com` performs statistical C2 protocol classification on submitted PCAP files, supporting identification across 600+ protocols without running a local FlowCarp instance, e.g. `curl --data-binary @dump.pcap https://demo.flowcarp.com | jq`.

Submit a PCAP to the demo server for instant protocol identification across 600+ protocols without standing up a local FlowCarp instance.

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

— Names endpoint and protocol count but little mechanism detail

changelog-20260925-5c21cf84

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

FlowCarp adds statistical C2 protocol classification for PCAP files, with a demo API at `demo.flowcarp.com` supporting 600+ protocols.

└──▷ **TRY IT**

Build a custom model for a newly discovered C2 protocol using sandbox capture sessions, so FlowCarp can recognise it in future PCAPs.

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

Classify traffic in a new sandbox PCAP against the custom model to confirm which samples use the same C2 protocol.

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

Submit a PCAP to the demo server for instant protocol identification across 600+ protocols without standing up a local FlowCarp instance.

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

Adds `flowcarp --input <pcap> --output model:<Label>:<MALICIOUS|BENIGN>:<file>.fcp` to build a custom statistical protocol model from example traffic captures.

Adds `flowcarp --model <file>.fcp --input <pcap> --output -` to classify traffic in a new PCAP against a previously built protocol model.

WAS THIS USEFUL? 👍 👎

Netcap

Feedback ⓘ v0.9.13 NOTES ↗ CODE ↗ RANK

A framework for secure and scalable network traffic analysis - https://netcap.io

> Netcap v0.9.13 adds a DNP3 point-map flag for resolving field-device names during capture, rewrites the DNP3 decoder for ICS traffic, and moves several protocol decoders to per-message timestamping.

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

01 **DNP3 point map flag for capture** NEW 90

`net capture` gains a `--dnp3-point-map` flag (env `NC_DNP3_POINT_MAP`) accepting a CSV of outstation, group, index, and name to resolve raw DNP3 point indexes to human-readable labels, useful for correlating alerts to specific field devices in OT networks.

Resolve raw DNP3 point indexes to named tags during capture of an OT network PCAP — useful when correlating alerts to specific field devices.

```
$ net capture -read ot-traffic.pcap --dnp3-point-map points.csv
```

Do the same on a live interface with the point map supplied via environment variable, e.g. in a containerised sensor deployment.

```
$ NC_DNP3_POINT_MAP=/etc/netcap/points.csv sudo net capture -iface eth0
```

— Named flag and env var with two runnable command examples. v0.9.13

02 **Per-message timestamps for FTP, IMAP, IRC, HTTP, and IEC 62351**

IMPROVED 60

FTP, IMAP, and IRC messages are now timestamped from the individual packet that carried them rather than from the connection; HTTP requests and responses are likewise timestamped individually rather than once per connection; and IEC 62351 records now get their own per-record timestamp and traffic direction across all four decoding paths.

— Clear before/after per protocol but no configurable surface. v0.9.13

THINNER COVERAGE BELOW

03 **Full decoder listing in util -decoders** IMPROVED 50

`util -decoders` now reports every stream decoder, covering all 40+ registered decoders instead of a partial list.

— Names the command and count but no example run. v0.9.13

04 **Rewritten DNP3 decoder for ICS traffic** IMPROVED 45

The DNP3 decoder was rewritten to produce only the frames and objects actually present in captured traffic, improving accuracy for ICS/OT traffic analysis.

— States the behavior change but gives no numbers or example. v0.9.13

05 **Decoder selection based on evidence strength** IMPROVED 40

Decoder selection now picks the decoder that required the most evidence to match traffic, rather than the one bound to the lowest port number.

— Describes the mechanism change but no user-facing control. v0.9.13

06 **Full TACACS+ and SSH signature validation** IMPROVED 35

Netcap now validates full TACACS+ and SSH signatures including confusable edge cases.

— Brief statement with no mechanism or example given. v0.9.13

v0.9.13 NOTES ↗ CODE ↗ LATEST STABLE

Netcap v0.9.13 adds a DNP3 point-map flag, per-message timestamps, and a fully rewritten DNP3 decoder for ICS traffic.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.9.13 https://github.com/dreadl0ck/netcap.git
# already have the repo? check out this version:
$ git checkout v0.9.13
```

└──▷ TRY IT

Resolve raw DNP3 point indexes to named tags during capture of an OT network PCAP — useful when correlating alerts to specific field devices.

```
$ net capture -read ot-traffic.pcap --dnp3-point-map points.csv
```

Do the same on a live interface with the point map supplied via environment variable, e.g. in a containerised sensor deployment.

```
$ NC_DNP3_POINT_MAP=/etc/netcap/points.csv sudo net capture -iface eth0
```

Adds `--dnp3-point-map` flag (env `NC_DNP3_POINT_MAP`) to `net capture`, accepting a CSV of outstation, group, index, and name to resolve DNP3 point indexes to human-readable labels.

Adds `util -decoders` now reports every stream decoder, covering all 40+ registered decoders instead of a partial list.

Timestamps each FTP, IMAP, and IRC message from the individual packet that carried it, rather than from the connection.

Timestamps each HTTP request and response individually, rather than once per connection.

Gives IEC 62351 records their own per-record timestamp and traffic direction across all four decoding paths.

+3 more

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Sophos Switch

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

RANK

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

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

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

01 **Switch name as primary identifier and search key** IMPROVED

50

Switch names are now shown as the primary identifier on the switch, site, and stack management pages in Sophos Central, making it easier to identify switches at a glance. The search box on these same pages now supports searching by switch name in addition to existing search criteria.

— Names the exact pages affected but no config/API detail changelog-20260925-7c747c22

> changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE
>
> Sophos Switch now displays switch names as the primary identifier and supports name-based search on management pages.
>
> Switch names are now shown as the primary identifier on the switch, site, and stack management pages in Sophos Central, making it easier to identify switches at a glance.
>
> The search box on switch, site, and stack management pages now supports searching by switch name in addition to existing search criteria.

WAS THIS USEFUL? 👍 👎

Sublime Security

Feedback ⓘ

Platform v1.88.278 (Expanded Link Exclusions) (2026-09-24) BLOG ↗ RANK

Sublime Security is a cloud email security platform that detects and remediates advanced threats using AI-powered analysis.

> Sublime Security expanded Link Exclusions so they can match on more than just domain, giving analysts finer-grained control over which links skip inspection.

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

01 Expanded matching criteria for Link Exclusions IMPROVED

86

Link Exclusions can now match on path, query params, display text, scheme, or subdomain, not just domain. Create one via `Exclusions > Create` with `Scope` set to `Link`, using expressions like `strings.icontains(href_url.path, "/unsubscribe") and href.url.domain.root_domain == "acme.com"` to skip analysis of matching links (e.g., unsubscribe links from a specific domain).

Skip analysis of unsubscribe links from a specific domain so they never trigger link-inspection rules.

📍 In the Sublime Security console, go to Exclusions › Create, set Scope to `Link`, and enter the expression: strings.icontains(href_url.path, "/unsubscribe") and `href.url.domain.root` _domain == `acme.com`

— Names exact UI path and full expression syntax with example

Platform v1.88.278 (Expanded Link Exclu…

Platform v1.88.278 (Expanded Link Exclusions) (2026-09-24) BLOG ↗ LATEST

STABLE

Link Exclusions now match on path, query params, display text, scheme, or subdomain — not just domain.

└──▷ **HOW TO FIND IT**

Skip analysis of unsubscribe links from a specific domain so they never trigger link-inspection rules.

📍 In the Sublime Security console, go to Exclusions › Create, set Scope to `Link`, and enter the expression: strings.icontains(href_url.path, "/unsubscribe") and `href.url.domain.root` _domain == `acme.com`

> Adds Link Exclusion scope via `Exclusions > Create` with `Scope` set to `Link`, supporting expressions like `strings.icontains(href_url.path, "/unsubscribe") and href.url.domain.root_domain == "acme.com"`.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Chainguard

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

RANK

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

> Chainguard's biggest move this window is a new Skills Registry and full `chainctl skills` CLI suite for uploading, hardening, and distributing AI agent skills, alongside a Sovereign Artifacts beta storing Containers and Libraries in an EU region, several new container images, and a breaking default-user change in the `sonar-scanner-cli` image.

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

01 Skill hardening via chainctl skills harden and status NEW

98

`chainctl skills harden` uploads a local skill directory or an existing uploads-registry artifact, submits it for server-side hardening, and downloads the hardened result along with a `HARDENING.md` findings report (e.g. to `./hardened/<name>/`). It supports four input forms — positional path, `--folder`, an `uploads.cgr.dev` reference, or `--digest` with `--name` — plus `--wait` and `--timeout` to block until completion and `-o id` to return only the job ID for deferred checks. `chainctl skills status --id <job-id> --wait --timeout` tracks an in-progress hardening job using its 64-character job ID and can resume waiting for the result; hardened artifacts are published under a per-user namespace path `skills.cgr.dev/<org-uidp>/users/<user-namespace>/<skill-name>@sha256:<digest>`.

Harden a local skill directory and wait for the server-side result in one step — ideal for CI pipelines where you want to block until the hardened artifact and HARDENING.md report are ready.

```
$ chainctl skills harden ./harden-demo --group your-organization
--wait --timeout 30m
```

Submit a hardening job from CI without blocking, then check progress later using the captured job ID.

```
$ JOB_ID=$(chainctl skills harden ./harden-demo --group your-organization -o id)
chainctl skills status --group your-organization --id "$JOB_ID" --wait --timeout 30m
```

— Full mechanism, flags, paths and runnable commands given product docs

02 **Skill lifecycle management commands in chainctl skills** NEW

85

Additional `chainctl skills` subcommands round out the skill lifecycle: `chainctl skills push <name> --group <org> --tag <version>` uploads a skill to `uploads.cgr.dev`; `chainctl skills describe` inspects a full artifact reference (upload or hardened) by registry path, or a hardened ref to retrieve its install name; `chainctl skills pull` downloads a hardened artifact by exact reference into a local directory; `chainctl skills install` detects local agents and deploys a hardened skill, reporting the placement location; `chainctl skills validate` checks a skill directory's structure locally before submitting for hardening; `chainctl skills uninstall <install-name>` removes an installed skill from the local machine; `chainctl skills versions <hardened-repo>` lists available tags for a published hardened skill; and `chainctl skills delete <hardened-repo>:<tag>` removes a specific published version from the hardened registry.

Browse all hardened skills across user-namespace folders, then install and invoke a specific hardened result in Claude Code.

```
$ chainctl skills list --group your-organization --source skills --recursive
export HARDENED_REF='skills.cgr.dev/<org-uidp>/users/<user-namespace>/hello-world@sha256:<digest>'
chainctl skills install "$HARDENED_REF"
export INSTALLED_SKILL='<install-name-from-describe>'
```

Remove a specific published version from the hardened registry after listing its available tags.

```
$ HARDENED_REPO="${HARDENED_REF%@*}"
chainctl skills versions "$HARDENED_REPO"
export HARDENED_TAG='<tag-from-the-versions-output>'
chainctl skills delete "$HARDENED_REPO:$HARDENED_TAG"
```

— Every subcommand named with its purpose product docs

03 **chainctl skills list filtering and traversal** NEW 83

`chainctl skills list` gained a `--recursive` (`-r`) flag to traverse nested user-namespace folders and show all hardened skills with full paths in the `NAME` column, and a `--source` flag with values `skills` (hardened registry at `skills.cgr.dev`), `uploads` (original uploads at `uploads.cgr.dev`), or `all` (both registries labeled by source).

List all hardened skills across every user namespace in your organization, including version-tagged and digest-only artifacts, to find a completed hardening result.

```
$ chainctl skills list --group your-organization --source skills
--recursive
```

— Exact flags and values with a runnable example `product docs`

---

04 **sonar-scanner-cli default non-root user change** BREAKING 83

Starting October 12, the `sonar-scanner-cli` container image changes its default behavior to run as a non-root user, moving scanner output to `/tmp/.scannerwork`. Teams can temporarily restore the old root behavior with `docker run --user 0 cgr.dev/chainguard/sonar-scanner-cli`, or preserve the previous output path by setting `-Dsonar.working.directory=<project>/.scannerwork`.

Temporarily restore the previous `sonar-scanner-cli` non-root behavior after the October 12 default-user change, while you migrate pipelines to read `report-task.txt` from `/tmp/.scannerwork`.

```
$ docker run --user 0 cgr.dev/chainguard/sonar-scanner-cli
```

Migrate a Kubernetes deployment of `sonar-scanner-cli` to read scanner output from the new `/tmp/.scannerwork` path, or explicitly preserve the old path by setting the working directory flag.

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

— Exact before/after, date, and migration commands given `changelog-20260925-4ed83dfe`

---

05 **Chainguard Skills Registry service launch** NEW 65

Chainguard introduced the Skills Registry, a hosted service for publishing, hardening, and distributing agent skills scoped to an organization, using `uploads.cgr.dev` for

original uploads and `skills.cgr.dev` for hardened, published results, with both public and private registries served at `skills.cgr.dev/<your-org>/`.

Upload a local agent skill, submit it for hardening, and wait up to 30 minutes for a signed, hardened result — the typical first-time workflow for a new org skill.

```
$ chainctl skills entitlements create --parent your-organization
chainctl skills accept-terms --group your-organization
chainctl skills push hello-world --group your-organization --tag
v1.0.0
chainctl skills harden uploads.cgr.dev/your-organization/hello-
world:v1.0.0 \
  --group your-organization --wait --timeout 30m
```

Upload an agent skill and submit it for hardening so the secured version is published to your org-scoped registry path.

```
$ chainctl skills harden
```

— Names registry paths but overlaps with harden/status mechanism product docs

THINNER COVERAGE BELOW

06 **Agent Skills entitlement and terms commands** NEW 58

`chainctl skills entitlements create` and `chainctl skills accept-terms` enable the Agent Skills service for an organization, and `chainctl skills entitlements list` checks whether an organization already has a Skills entitlement.

— Named commands but no further mechanism described product docs

07 **Sovereign Artifacts EU data-residency beta** NEW 53

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

— Explains mechanism and benefit but no enrollment steps given

changelog-20260925-4ed83dfe

08 **New container images: openjdk:26, ceph:19, kuma:2.12** NEW

48

Chainguard added `openjdk:26` (released 2026-09-17, EOL 2027-03-17), `ceph:19` (released 2026-09-19, EOL 2027-03-19), and `kuma:2.12` (released 2026-09-09, EOL 2027-03-09) container images.

— Named images and dates but no usage detail changelog-20260925-4ed83dfe

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Chainguard adds EU data-residency beta for Containers and Libraries, plus 26+ new container images including FIPS variants.

└──▷ **TRY IT**

Temporarily restore the previous `sonar-scanner-cli` non-root behavior after the October 12 default-user change, while you migrate pipelines to read `report-task.txt` from `/tmp/.scannerwork`.

```
$ docker run --user 0 cgr.dev/chainguard/sonar-scanner-cli
```

Migrate a Kubernetes deployment of `sonar-scanner-cli` to read scanner output from the new `/tmp/.scannerwork` path, or explicitly preserve the old path by setting the working directory flag.

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

Launches Sovereign Artifacts beta for enrolled customers, storing Chainguard Containers and Libraries artifact bytes in an EU region with automatic backfill of existing artifacts and ~30% faster pulls from Europe.

Adds `openjdk:26` container image (released 2026-09-17, EOL 2027-03-17).

Adds `ceph:19` container image (released 2026-09-19, EOL 2027-03-19).

Adds `kuma:2.12` container image (released 2026-09-09, EOL 2027-03-09).

WAS THIS USEFUL? 👍 👎

Google distroless

Feedback ⓘ changes since 2026-09-10

CODE ↗

Language focused docker images, minus the operating system.

Distroless adds a Dart example using the base-nossl image and requires Dart 3.9 for the dart build CLI.

└─▷ **GET THIS VERSION**

```
$ git clone --branch commits-2026-09-10
https://github.com/GoogleContainerTools/distroless.git
# already have the repo? check out this version:
$ git checkout commits-2026-09-10
```

└─▷ **USE IT**

Bootstrap a minimal Dart distroless container using the base-nossl image, as shown in the new example.

examples/dart/bin/hello.dart

```
// examples/dart/bin/hello.dart
void main() {
  print('Hello, distroless!');
}
```

Requires Dart 3.9 for the `dart build` CLI when building distroless Dart containers.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Upwind

Feedback ⓘ

Introducing the new Configurations experience in Upwind (2026-09-24) BLOG ↗

RANK

Upwind is a runtime-powered CNAPP that protects cloud workloads by combining posture, detection and response, and identity security with real-time runtime context.

> Upwind rebuilt its Configurations experience around live, continuous compliance scoring with a new Score Journey view for prioritizing fixes by impact, alongside changes to how findings and frameworks are organized.

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

01 Live compliance scoring across frameworks and clouds NEW

75

Adds continuous, live compliance scoring evaluated against multiple frameworks simultaneously, including CIS, NIST, ISO, EU AI Act, and ISO/IEC 42001, across AWS, Azure, Google Cloud, Oracle Cloud, Kubernetes, and on-premises environments. Each cloud account is scored separately so compliance gaps are connected to a specific environment and responsible team rather than surfaced only as an organization-wide number.

·

When preparing for an ISO audit, navigate to the Configurations experience to view a live compliance score filtered to your ISO framework and share the view directly with auditors without exporting a spreadsheet.

> 📍 In Upwind, go to Configurations › select the ISO framework › filter by the relevant cloud account › share the view with stakeholders.

— Names frameworks and clouds; UI path given but no API

Introducing the new Configurations expe…

02 Score Journey remediation prioritization NEW

75

Score Journey shows, for each finding, how much remediating it would improve the overall compliance score, how many resources it affects, and a difficulty rating from Easy to Hard, so teams can prioritize fixes by impact.

·

Use Score Journey to identify the highest-impact, lowest-effort fixes before a compliance review — sort findings by potential score improvement and difficulty to build a defensible remediation plan.

📍 In Upwind, go to Configurations › Score Journey › sort findings by Score Impact descending, then filter by difficulty Easy to identify quick wins.

— Exact navigation and sort/filter steps provided

Introducing the new Configurations expe…

THINNER COVERAGE BELOW

03 **Findings separation and versioned framework library** IMPROVED

38

Separates configuration findings from security issues so teams can distinguish policy gaps from active risks, and maintains a versioned, continuously updated framework library — including recent additions such as ISO/IEC 42001 and the EU AI Act — displaying each framework's version and latest update date.

— No mechanism, location, or config detail given Introducing the new Configurations expe…

Introducing the new Configurations experience in Upwind (2026-09-24) BLOG ↗

LATEST STABLE

Upwind rebuilds its Configurations experience with live compliance scoring, per-account breakdowns, and a Score Journey to prioritize fixes by impact.

└──▷ **HOW TO FIND IT**

When preparing for an ISO audit, navigate to the Configurations experience to view a live compliance score filtered to your ISO framework and share the view directly with auditors without exporting a spreadsheet.

📍 In Upwind, go to Configurations › select the ISO framework › filter by the relevant cloud account › share the view with stakeholders.

Use Score Journey to identify the highest-impact, lowest-effort fixes before a compliance review — sort findings by potential score improvement and difficulty to build a defensible remediation plan.

📍 In Upwind, go to Configurations › Score Journey › sort findings by Score Impact descending, then filter by difficulty Easy to identify quick wins.

Adds continuous, live compliance scoring evaluated against multiple frameworks simultaneously, including CIS, NIST, ISO, EU AI Act, and ISO/IEC 42001, across AWS, Azure, Google Cloud, Oracle Cloud, Kubernetes, and on-premises environments.

Introduces Score Journey, which shows for each finding how much remediating it would improve the overall compliance score, how many resources it affects, and a

difficulty rating from Easy to Hard.

Scores each cloud account separately so every compliance gap is connected to a specific environment and responsible team rather than surfaced only as an organization-wide number.

Separates configuration findings from security issues so teams can distinguish policy gaps from active risks.

Maintains a versioned, continuously updated framework library — including recent additions such as ISO/IEC 42001 and the EU AI Act — displaying each framework's version and latest update date.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Detection Engineering

Tenzir

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

> Tenzir's v6.18.0/v1.40.0 release adds Azure managed identity support for Key Vault and introduces parallel processing across transforms and sources in the Tenzir Node.

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

01 **Parallel processing across transforms and sources** IMPROVED

30

Enables parallel processing across transforms and sources in the Tenzir Node.

— States the change but no mechanism or config details. changelog-20260925-aeadaa87

> changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE
>
> Tenzir v6.18.0/v1.40.0 adds Azure managed identity for Key Vault and parallel processing across transforms and sources.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Google SecOps

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

RANK

Meet your business challenges head on with AI and cloud computing services from Google, including security, data management, and hybrid & multi-cloud.

> Google SecOps shipped a major Gemini-powered agent framework — the Triage and Investigation Agent embedded directly into SOAR playbooks — alongside the launch of Google Unified Security with Chrome Enterprise Premium and multi-cloud CDIR, plus a wave of preview features spanning GoogleSQL search, case-level playbooks, and a new Detection Engineering Agent.

└──▷ **WHAT SHIPPED · 26 FEATURES** most completely described first ? what's the number?

01 Triage and Investigation Agent (TIN) for autonomous alert triage

NEW 98

TIN autonomously analyzes incoming security alerts, executes an investigation plan, and delivers a structured true-positive/false-positive verdict with findings and a summarized explanation; investigation volume and token usage are visible in the curated TIN dashboard. Agentic Automation embeds TIN as the first supported AI agent directly into SOAR playbooks as a discrete step alongside deterministic automation, pausing the playbook asynchronously while the agent works without occupying the execution queue. Agent step results expose `verdict` and `confidence_level` fields in the Expression Builder for conditional branching (remediation, auto-close, or manual review), and findings surface in the Gemini Investigations side drawer and AI Agent step results within Case Overview, tracked through three states: Not Run, Running, and Completed.

Investigate an alert hands-free: let the Triage and Investigation Agent determine whether an alert is a true or false positive, then review its findings and explanation in the TIN dashboard.

📍 In the Google SecOps console, open an alert and select Investigate with Triage and Investigation Agent. Once the agent completes its plan, navigate to the TIN dashboard to review investigation volume, token usage, and the structured verdict.

Add a TIN agent step to a phishing-alert playbook and branch automatically on verdict — routing confirmed threats to isolation and high-confidence false positives to auto-close, to reduce analyst noise.

📍 1. Go to Response › Playbooks and open your phishing alert playbook. 2. In the Step Selection panel, expand the AI Agents category and drag the TIN agent step onto the canvas. 3. In the configuration panel, set Action type to Automatic and enable Retry on failure. 4. Under Error Handling (If step fails), select Skip to the next step so the workflow continues even if the alert type is unsupported. 5. Drag a Flow block (Conditional) after the AI Agent step, open the Placeholder menu, select the AI Agent step results, and use the Expression Builder to evaluate the verdict field. 6. Wire the True Positive branch to an `Isolate Host` action and the False Positive (High Confidence) branch (confidence_level = high) to a `Close Alert` action.

Opt in the TIN agent for your tenant so automatic investigations trigger on supported SIEM detection engine alerts without manual intervention.

📍 1. Open any case in Case Overview. 2. Open the Gemini Investigations side drawer. 3. Navigate to the settings of the Gemini investigation side menu. 4. Locate the Triage and Investigation Agent (TIN) integration and toggle it on for your tenant.

— Full mechanism, named fields, and runnable playbook configuration steps `product docs`

02

Cloud Detection, Investigation, and Response (CDIR) across multi-cloud

NEW

90

CDIR integrates Security Command Center Enterprise with Google SecOps for unified case management, SIEM log searching, and SOAR workflows across Google Cloud, AWS, and Azure. It adds Cloud Asset Inventory collection from Google Cloud, AWS, and Azure into the Google SecOps Entity Graph, enriching alerts with asset context such as region, owner, and metadata tags; CDIR Curated Detections driven by Cloud Threat Intelligence (including detection types like `SCC Enhanced Privilege Escalation`) automatically map to Google SecOps cases; Security Command Center connects to AWS for configuration and resource data, with Azure connection support in preview; and SIEM log searching supports natural-language, YARA-L, or raw log search against ingested Google Cloud audit logs, AWS CloudTrail, GuardDuty events, and Azure AD activity logs.

Activate multi-cloud CDIR so SOC analysts can investigate privilege escalation and lateral movement across Google Cloud, AWS, and Azure within a single Google SecOps instance.

📍 1. Activate the Enterprise Tier of SCC-E using your existing Google SecOps instance. 2. Validate Security Command Center Detection Services are enabled. 3. Connect Security Command Center to AWS (and optionally Azure, currently in Preview). 4. Ingest Google Cloud logs, AWS logs, and Azure logs into Google SecOps. 5. Enable the CDIR Curated Detections.

— Many named surfaces across clouds with setup steps product docs

03 **GCP organization ID enrichment for ingested logs** NEW

80

Google SecOps auto-enriches direct-ingestion logs with the source GCP organization ID under `metadata.base_labels.ingestion_kv_labels`, letting analysts scope UDM Search results to a single org in multi-org environments via `ingestion_kv_labels["gcp_organization_id"] = "<YOUR_GCP_ORGANIZATION_ID>"`.

Filter UDM Search to a specific GCP organization in a multi-org environment, scoping investigation or detection results to only that org's logs.

📍 metadata.base_labels.ingestion_kv_labels["gcp_organization_id"] = `<YOUR_GCP_ORGANIZATION_ID>`

Filter UDM search results to events from a specific Google Cloud organization when multiple orgs feed into the same SecOps instance.

```
$ metadata.base_labels.ingestion_kv_labels.key =
'gcp_organization_id' AND
metadata.base_labels.ingestion_kv_labels.value = '123456789012'
```

— Exact field path and query syntax given, narrow single-purpose feature

changelog-20260925-750e87f5

04 **Chrome Enterprise Premium integration in Google SecOps** NEW

80

Chrome Enterprise Premium (CEP) integrates with Google SecOps for browser-sourced threat visibility, with curated detections enriched by Safe Browsing and Google Threat Intelligence (GTI). Dedicated dashboards surface Chrome security events and anomalies, and native playbook orchestration enables automated Chrome Enterprise response actions — blocking malicious URLs, blocking extensions, and quarantining devices — directly from Google SecOps.

Enable Chrome Enterprise Premium threat detections and browser telemetry ingestion into Google SecOps to catch malicious extensions, phishing, and credential theft originating from the browser.

📍 1. Configure the Chrome Enterprise Premium Reporting Connector to send data to Google SecOps. 2. Connect Chrome Enterprise Premium to Google SecOps. 3. Enable Chrome Enterprise Threats Curated Detections.

— Named enrichment sources and concrete response actions with setup steps product docs

05 **Customizable schedules for multi-event rules (GA)** IMPROVED

80

Customizable schedules for multi-event rules reach General Availability, with configurable first-run delay offsets from 1 minute to 48 hours, automated true-up runs at 4 hours and optionally 30 hours, and migration of legacy rules from the Rules Dashboard; custom roles need the `chronicle.rules.modifyRules` and `chronicle.ruleDeployments.update` IAM permissions.

— Exact timings and IAM permissions named, migration path noted product docs

06 **Gemini pane for in-product documentation Q&A** NEW

65

An in-product Gemini pane answers natural-language questions about Google SecOps by summarizing relevant documentation and linking to related pages. Responses are grounded via Google Search, surfacing a Google G icon, search query chips, and citations for verification, and individual or all chat sessions can be removed via Delete chat and Delete all chat sessions controls.

Ask the Gemini pane how to investigate user logins without leaving the SecOps console.

📍 1. In Google SecOps, open the Gemini pane. 2. Enter the prompt: How do I investigate user logins 3. Review the generated summary and follow any linked documentation pages for deeper context.

— Mechanism and controls named, example prompt given product docs

07 **Legacy Mandiant and open-source IOC feeds retired** DEPRECATED

65

After March 18, 2027, the `MANDIANT_ACTIVE_BREACH_IOC`, `MANDIANT_FUSION_IOC`, and `OPEN_SOURCE_INTEL_IOC` feeds will be removed; customers must migrate to the `GTI_IOC` feed before that date.

— Exact feed names, deadline, and replacement named product docs

08 **AI Protection and Model Armor in Security Command Center Enterprise**

NEW

63

AI Protection within Security Command Center Enterprise (SCC-E) provides discovery of AI inventory (models, datasets, applications), virtual red teaming for attack path analysis, and detection and response for AI-specific threats including LLM hijacking, data poisoning, and exfiltration. Model Armor adds model-agnostic prompt and response screening protecting against data loss and offensive content across Google and third-party AI models on multiple clouds.

— Named capabilities and threat types but no setup steps product docs

09 **User Preferences dialog** NEW

60

A new User Preferences dialog, reached via the user avatar menu, configures time zone, date/time format, notification settings (including email delivery), and accessibility options such as feedback message duration.

Tune notification delivery so alerts reach you by email and set your local time zone to avoid timestamp confusion in investigations.

📍 1. Click your user avatar (top-right). 2. Select User preferences. 3. Under Localization, set Time Zone and Date/Time Format to match your region. 4. Under Notifications, click the User Notifications toggle to ON, select the desired notification types, and choose email or both for delivery. 5. Under Accessibility, set the feedback message duration to your preference.

— Exact UI navigation path given but limited mechanism depth product docs

10 **chronicle.readonly OAuth scope loses write permissions**

60

BREAKING

Effective January 25, 2027, write permissions will be removed from the `chronicle.readonly` OAuth scope; any workflow using `chronicle.readonly` for write operations must be updated to use the `chronicle` OAuth scope.

— Exact scopes and deadline given, clear migration step product docs

THINNER COVERAGE BELOW

11 **Natural-language search building in Gemini pane** NEW

55

The Gemini pane and natural-language prompt support in the search field let practitioners build, edit, run, and iterate on searches — adjusting scope, time range, and filters — without writing queries manually.

Refine an investigation search interactively — expand the time window and add filters — using natural language instead of editing UDM query syntax.

> 📍 In the Google SecOps search field or Gemini pane, enter: 'Show authentication failures for user jsmith in the last 7 days, filtered to events originating outside the US.' Use the Gemini pane to iteratively adjust scope or add additional filters.

— Example prompt given but mechanism described only briefly product docs

12 **GoogleSQL query support in Search (preview)** NEW

55

Search now supports GoogleSQL queries (public preview) using both declarative SQL and Piped SQL syntax to query UDM events, entity graphs, detection rules, and case management telemetry tables.

— Names syntaxes and query targets, no runnable example product docs

13 **Self-service Bindplane Enterprise license download (preview)**

NEW

55

Enterprise Plus and Google Unified Security customers can download a Bindplane Enterprise license (preview) directly from `SIEM Settings > Collection Agents`.

— Exact UI path named, minimal mechanism beyond that product docs

14 **Investigation Management UI: side-by-side view and resizable panels**

IMPROVED

53

The revamped Investigation Management experience adds a Side-by-side view layout (public preview) to the Alerts & Detections tab, showing alert metadata and Gemini insights in an adjacent pane, plus resizable Case preview and Alert and detection preview side panels adjustable via mouse or keyboard shortcuts.

— Names exact tab and panels, UI path is a clear starting point product docs

15 **Case-level playbooks and reaction triggers (preview)** NEW

52

Case playbooks (preview) let playbooks and manual actions run across an entire case container rather than individual alerts, while reaction triggers (preview) automatically

fire playbooks in response to real-time case or alert updates such as changes to assignee, tags, alert priority, or newly added entities.

— Mechanism described, no exact UI or config path product docs

16 **Relative time filter operators: Past, Previous, Current** NEW

50

Three distinct relative time filter operators — Past, Previous, and Current (public preview) — eliminate ambiguity between rolling windows and calendar-aligned periods across dashboards and Search.

— Named operators and purpose, no example of use product docs

17 **Detection Engineering Agent (preview)** NEW

48

An AI assistant accessible via Model Context Protocol (MCP) tools that extracts threat intelligence and drafts YARA-L detection rules, in public preview.

— Names protocol and function, no walkthrough given product docs

18 **Grok filter `match_all` option** IMPROVED

45

Parser syntax's Grok filter gains a `match_all` option that extracts all non-overlapping pattern occurrences in a field instead of only the first match.

— Names exact option and behaviour change, no example product docs

19 **Event simulation for detection coverage testing (preview)**

NEW

45

Event simulation (public preview) programmatically delivers threat sequences into the live ingestion pipeline to evaluate full-funnel detection coverage, from UDM normalization through multi-event correlation.

— Mechanism described, no concrete command or config product docs

20 **Unroll processor in Data Processing Pipelines** NEW

43

The new Unroll processor splits log entries containing arrays or slices into multiple discrete log events before parsing and ingestion.

— Named processor and behaviour, no configuration example product docs

21 **Mandiant Frontline Threats rule packs in Content Hub** NEW

42

Adds Curated Detections rule packs: Mandiant Frontline Threats for Linux, `Mandiant Frontline Threats for MacOS`, and Mandiant Frontline Threats for Google Cloud.

— Names exact rule packs, no further mechanism product docs

22 **Natural-language YARA-L rule generation** NEW 35

Practitioners can describe detection logic in plain text via a Gemini prompt and receive a ready-to-use YARA-L rule.

— Described only in one sentence, no example or steps product docs

23 **Gemini-assisted playbook creation and editing** NEW 33

Analysts can create and modify SOAR playbooks using natural-language prompts via Gemini assistance.

— One-line capability, no steps or example given product docs

24 **Operations feed in Emerging Threats Center** NEW 32

Adds Operations in the Emerging Threats Center feed, providing granular threat intelligence from frontline investigations scoped to individual organizations, alongside the existing global Campaigns.

— Brief description, no usage detail product docs

25 **Gemini threat intelligence assistant** NEW 30

Gemini-powered threat intelligence assistance answers general security domain questions and provides summaries about threat actors, IOCs, and other threat intelligence topics.

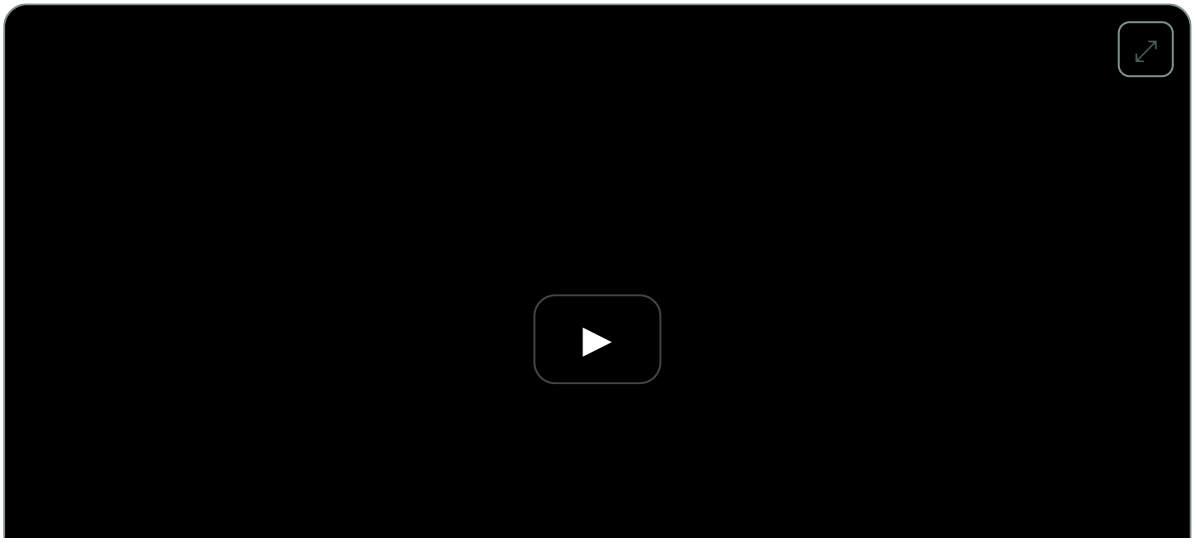
— Prose-only description, no named surface or example product docs

26 **AI-generated case summary widget** NEW 25

A Gemini case summary widget presents an AI-generated summary of a case along with suggested remediation steps.

— Thin one-line description, no mechanism or example product docs

└──▷ ALSO FROM THESE RELEASES





└──▷ **BREAKING ON UPGRADE**

! Effective January 25, 2027, write permissions will be removed from the `chronicle.readonly` OAuth scope; any workflow using `chronicle.readonly` for write operations must be updated to use the `chronicle` OAuth scope.

! After March 18, 2027, the `MANDIANT_ACTIVE_BREACH_IOC`, `MANDIANT_FUSION_IOC`, and `OPEN_SOURCE_INTEL_IOC` feeds will be removed; migrate to the `GTI_IOC` feed before that date.

> changelog · Sep 25, 2026 | CHANGELOG ↗ | STABLE
>
> Google SecOps auto-enriches direct-ingestion logs with the source GCP organization ID under `metadata.base_labels.ingestion_kv_labels`.
>
> └──▷ **HOW TO FIND IT**
>
> Filter UDM Search to a specific GCP organization in a multi-org environment, scoping investigation or detection results to only that org's logs.
>
> 📍 metadata.base_labels.ingestion_kv_labels["gcp_organization_id"] = `<YOUR_GCP_ORGANIZATION_ID>`

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

RANK

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

> Databricks shipped automatic sensitive-data classification in Unity Catalog, GA key-value tagging across dashboards/notebooks/apps/Genie Agents, GA metric-view sharing via OpenSharing, and Beta ABAC enforcement for time-travel queries, alongside a round of AI/BI dashboard, Genie Agents, and Genie One usability enhancements.

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

01 **Per-widget dashboard filter override** NEW 70

Dashboard widgets can be configured to ignore selected dashboard-level filters while those filters continue to apply to the rest of the page, via the widget's settings panel Ignore filters option.

Prevent a KPI widget from being affected by page-level date or region filters so it always shows the full baseline metric.

> 📍 In the dashboard editor, select the widget you want to isolate, open its settings panel, and enable Ignore filters so the widget displays unfiltered data while all other widgets on the page continue to respond to dashboard filters.

— Concrete behaviour with exact UI steps to enable it

Dashboard enhancements and bug fixes (2…

02 **Genie Agents monitoring column picker and page-level citations**

NEW 70

Genie Agents' Monitor tab gains a column picker on the agent monitoring table to choose which columns are displayed, and (Beta) page-level citations so that when an agent answers using content from an attached volume file, the citation points to the specific page of the source file.

Narrow the monitoring table to only the columns relevant to your team's triage workflow.

> 📍 In the Genie Agent, go to Monitor tab › agent monitoring table › column picker, then select or deselect the columns you want to display.

— Two named additions with navigation steps for one

Genie Agents enhancements (2026-09-24)

03 **Metric view sharing via OpenSharing (GA)** NEW 70

Metric views can now be included in a share via OpenSharing and queried with MEASURE(...) on serverless compute; this capability is now generally available.

— Names the query construct and compute requirement

Sharing metric views with OpenSharing i…

04 **ABAC enforcement on time travel queries** NEW 65

Starting with Databricks Runtime 19 (Beta), ABAC row filter and column mask policies are now enforced on Unity Catalog managed tables when using time travel queries.

— Names the runtime version and policy scope enforced

Time travel queries now enforce ABAC po…

THINNER COVERAGE BELOW

05 **Dataset results table and metric view preview grid enhancements**

NEW 55

The dataset results table and metric view preview grid gain column pinning and formatting, a search box, a column explorer, natural-language filtering, and a selection details panel.

— Lists five concrete additions but no usage steps

Dashboard enhancements and bug fixes (2…

06 **Tags for dashboards, notebooks, apps, and Genie Agents (GA)**

NEW 50

Governed and ungoverned key-value tags can now be applied to dashboards, notebooks, Databricks apps, and Genie Agents to organize and categorize workspace assets; this capability is now generally available.

— Names asset types and tag model but no exact UI steps

Tags for dashboards, notebooks, Databri…

07 **Genie One chat pinning and MCP write actions** NEW 50

Genie One lets you pin a chat to keep it at the top of the sidebar, and (Beta) supports MCP write actions in the Genie One mobile app, enabling write operations from connected external tools.

— Names two additions but no usage steps given Genie One enhancements (2026-09-24)

08 **Level of detail (LOD) expressions in metric view editor** NEW 45

The metric view editor adds include-style level of detail (LOD) expressions for building metric calculations.

— Thin one-line addition with no mechanism detail

Dashboard enhancements and bug fixes (2…

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

35

Databricks can automatically classify and tag sensitive data within Unity Catalog using the new Data Classification capability.

— Names the capability but no mechanism or navigation given changelog-20260925-22113423

└──▷ **ALSO FROM THESE RELEASES**

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Databricks adds automatic sensitive data classification and tagging within Unity Catalog.

Automatically classify and tag sensitive data in Unity Catalog with the new Data Classification capability.

Tags for dashboards, notebooks, Databricks apps, and Genie Agents are now generally av

BLOG ↗ STABLE

Databricks GA: apply key-value tags to dashboards, notebooks, apps, and Genie Agents for workspace asset organization.

Dashboard enhancements and bug fixes (2026-09-24) BLOG ↗ STABLE

Databricks AI/BI dashboards gain column pinning, natural-language filtering, LOD expressions, and per-widget filter overrides.

└──▷ **HOW TO FIND IT**

Prevent a KPI widget from being affected by page-level date or region filters so it always shows the full baseline metric.

📍 In the dashboard editor, select the widget you want to isolate, open its settings panel, and enable Ignore filters so the widget displays unfiltered data while all other widgets on the page continue to respond to dashboard filters.

Adds column pinning and formatting, search, a column explorer, natural-language filtering, and a selection details panel to the dataset results table and metric view preview grid.

Adds include-style level of detail (LOD) expressions to the metric view editor.

Enables configuring a dashboard widget to ignore selected dashboard filters while those filters continue to apply to the rest of the page.

`Genie Agents enhancements (2026-09-24)` BLOG ↗ STABLE

Genie Agents gains column picker for monitoring table and page-level citations for volume file answers.

└──▷ **HOW TO FIND IT**

Narrow the monitoring table to only the columns relevant to your team's triage workflow.

> 📍 In the Genie Agent, go to Monitor tab › agent monitoring table › column picker, then select or deselect the columns you want to display.

Adds a column picker to the agent monitoring table on the Monitor tab, letting you choose which columns are displayed.

Adds page-level citations (Beta) so that when a Genie Agent answers using content from an attached volume file, citations point to the specific page of the source file.

`Genie One enhancements (2026-09-24)` BLOG ↗ STABLE

Genie One adds chat pinning and MCP write actions in the mobile app (Beta).

Adds the ability to pin a chat to keep it at the top of the Genie One sidebar.

Adds MCP write actions support in the Genie One mobile app (Beta), enabling write operations from connected external tools.

`Sharing metric views with OpenSharing is now generally available (2026-09-24)` BLOG ↗ STABLE

Databricks GA: share metric views via OpenSharing and query them with MEASURE(...) on serverless compute.

`Time travel queries now enforce ABAC policies (Beta) (2026-09-23)` BLOG ↗ STABLE

Databricks Runtime 19 adds ABAC row filter and column mask enforcement for Unity Catalog time travel queries (Beta).

WAS THIS USEFUL? 👍 👎

Sumo Logic

Feedback ⓘ

Mobot Now Answers Ingest Volume and Billing Questions (Manage) (2026-09-25)

BLOG ↗ RANK

Sumo Logic is a cloud-based platform for monitoring, logging, and analyzing applications and infrastructure data in real-time.

> Sumo Logic added a Data Volume skill to Mobot, its Platform Optimization Agent, letting users ask natural-language questions about ingest volume and billing instead of writing manual queries.

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

01 **Data Volume skill in Mobot agent** NEW 65

Mobot's Platform Optimization Agent gains a Data Volume skill that answers natural-language questions about ingest volume, tier usage, source-category breakdowns, and spend attribution. It is available to all users with no opt-in required.

Ask Mobot to surface your noisiest sources so you can prioritize ingestion cost reduction.

> 📍 In Mobot, type: What are my top 10 source categories by daily ingest volume?

Reconcile tier usage for finance reporting without writing a manual query.

> 📍 In Mobot, type: How much data went to the Infrequent tier last month?

— Names the skill and query types but not underlying mechanism.

Mobot Now Answers Ingest Volume and Bil…

Mobot Now Answers Ingest Volume and Billing Questions (Manage) (2026-09-25)

BLOG ↗ LATEST STABLE

Mobot gains a Data Volume skill to answer natural-language ingest and billing questions without manual queries.

└──▷ **HOW TO FIND IT**

Ask Mobot to surface your noisiest sources so you can prioritize ingestion cost reduction.

> 📍 In Mobot, type: What are my top 10 source categories by daily ingest volume?

Reconcile tier usage for finance reporting without writing a manual query.

📍 In Mobot, type: How much data went to the Infrequent tier last month?

Adds the Data Volume skill to Mobot's Platform Optimization Agent, enabling natural-language questions about ingest volume, tier usage, source-category breakdowns, and spend attribution — available to all users with no opt-in required.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ changelog · Sep 25, 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 added case management capabilities directly in the storyboard, letting teams track cases with change control, create cases visually, and use formulas for case details.

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

01 **Case tracking and creation in storyboard** NEW 50

Tines now enables tracking cases with change control directly in the storyboard, adds formula support for case details, and enables visual case creation in the storyboard.

.
.
.
.

— Names case tracking, formulas, and visual creation but no exact config or API surface.

changelog-20260925-c9c47f7f

changelog · Sep 25, 2026 CHANGELOG ↗ STABLE

Tines: Enables tracking cases with change control directly in the storyboard.

Adds formula support for case details.

Enables visual case creation in the storyboard.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

> Huntress shipped granular control over policy rollouts and antivirus configuration, plus new detections for macOS quarantine-bypass techniques and a legacy Screen Sharing exploit, alongside tools to cut Unsecured Credentials alert noise and two new Exchange Online controls.

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

01 Network and archive file scan toggles for Defender NEW

82

Adds Network File Scanning and Archive File Scanning toggles for Microsoft Defender in Managed Antivirus, found under `EDR > Managed Antivirus > Defender Configuration > Scans`; these settings were previously fixed to Huntress defaults and are now configurable per policy.

Enable Archive File Scanning and disable Network File Scanning in Microsoft Defender for a specific policy to ensure malware inside ZIPs and RARs is caught without the performance overhead of scanning mapped drives.

📍 In the Huntress console, go to EDR › Managed Antivirus › Defender Configuration › Scans. Set Network File Scanning to off and confirm Archive File Scanning is on, then save the configuration.

— Exact config path and toggle names with runnable steps `changelog-20260925-292ff4c5`

02 Modified Deployments for policy library rollout NEW

75

Adds Modified Deployments, an org-level setting in Managed ISPM that lets admins hand-pick which security controls (low, medium, or high impact) from the Huntress policy library roll out and on which days; new controls added to the library trigger a notification but are held until explicitly approved, supporting compliance and change control requirements.

Switch an org from automatic Managed Deployments to Modified Deployments so no new controls roll out without explicit approval — suitable before a compliance audit or major change freeze.

📍 In the Huntress console, go to your organization settings and select Modified Deployments. Choose the impact levels (low, medium, or high) you want available,

then set the deployment schedule days. New library controls will notify you but will not deploy until you approve each one.

— Mechanism and approval gate explained with setup steps

Managed ISPM: Customize control rollout… | changelog-20260925-292ff4c5

03 **Credential Files dashboard and allowlisting** NEW 73

Adds a new Credential Files page under the Process Insights dashboard in Managed EDR to view and manage every file triggering an Unsecured Credentials report in one place; adds a preview mode showing which files would generate an Incident Report before reporting is enabled; and adds allowlisting so reviewed and accepted files are excluded from future Unsecured Credentials Incident Reports.

Audit noisy Unsecured Credentials alerts before re-enabling reporting — review flagged files and allowlist accepted ones to keep Incident Reports focused.

📍 In the Huntress console, go to Managed EDR › Process Insights › Credential Files. Review the listed files, select those you have accepted, and choose Allowlist to exclude them. Then enable Unsecured Credential Incident Reports.

Re-enable Unsecured Credentials reporting for an org that turned it off due to noise, using the new preview and allowlist to scope reports to genuinely risky files first.

📍 In the Huntress console, navigate to Process Insights › Credential Files. Use the preview mode to review which files would generate Incident Reports without turning reporting on. Allowlist any files you have reviewed and accepted, then enable Incident Report generation once the list is scoped to real risks.

— Names the dashboard, preview, and allowlist mechanism with steps

Unsecured Credentials: New Credential F… | changelog-20260925-292ff4c5

04 **Exchange Online controls against Shadow Workflows** NEW 60

Adds two new Microsoft Exchange Online controls against Shadow Workflows: Ensure Direct Send is disabled and 'Tenant Allow/Block list should contain zero entries in the allow list'; both are available for manual scheduling but excluded from automatic Managed Deployments rollout.

— Both controls named exactly with rollout caveat | changelog-20260925-292ff4c5

THINNER COVERAGE BELOW

05 **Last Logged-In User field for macOS endpoints** IMPROVED 55

Extends the Last Logged-In User field on the agent overview page to macOS endpoints (previously Windows-only), reflecting the most recent interactive console/GUI login; the field populates on Agent v0.14.196 or later.

— Field named with exact behaviour and version requirement changelog-20260925-292ff4c5

06 **Detection for Screen Sharing pre-auth exploit** NEW

45

Adds a detection for Screen Sharing sessions that attach to a root user session via a legacy authentication handshake — the signature of a pre-auth flaw in Apple's Screen Sharing service disclosed July 2026; requires Agent v0.14.196.

— Specific CVE-style detail but purely automatic detection changelog-20260925-292ff4c5

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Huntress: Adds Network File Scanning and Archive File Scanning toggles for Microsoft Defender in Managed Antivirus, found under `EDR > Managed Antivirus > Defender Configuration > Scans`; previously fixed to Huntress defaults.

└──▷ **HOW TO FIND IT**

Switch an org from automatic Managed Deployments to Modified Deployments so no new controls roll out without explicit approval — suitable before a compliance audit or major change freeze.

📍 In the Huntress console, go to your organization settings and select Modified Deployments. Choose the impact levels (low, medium, or high) you want available, then set the deployment schedule days. New library controls will notify you but will not deploy until you approve each one.

Re-enable Unsecured Credentials reporting for an org that turned it off due to noise, using the new preview and allowlist to scope reports to genuinely risky files first.

📍 In the Huntress console, navigate to Process Insights › Credential Files. Use the preview mode to review which files would generate Incident Reports without turning reporting on. Allowlist any files you have reviewed and accepted, then enable Incident Report generation once the list is scoped to real risks.

Enable Archive File Scanning and disable Network File Scanning in Microsoft Defender for a specific policy to ensure malware inside ZIPs and RARs is caught without the performance overhead of scanning mapped drives.

📍 In the Huntress console, go to EDR › Managed Antivirus › Defender Configuration › Scans. Set Network File Scanning to off and confirm Archive File Scanning is on, then save the configuration.

Adds a detection for Screen Sharing sessions that attach to a root user session via a legacy authentication handshake — the signature of a pre-auth flaw in Apple's Screen Sharing service disclosed July 2026; requires Agent v0.14.196.

Adds three macOS quarantine-attribute-strip detections (requires Agent v0.14.196): quarantine removed from a file in a common malware staging folder by a script or shell; quarantine removed from a file on a mounted disk image; a fake installer mounted and a quarantine strip on the same machine within five minutes.

Extends the `Last Logged-In User` field on the agent overview page to macOS endpoints (previously Windows-only), reflecting the most recent interactive console/GUI login; populates on Agent v0.14.196 or later.

Adds two new Microsoft Exchange Online controls against Shadow Workflows — `Ensure Direct Send is disabled` and 'Tenant Allow/Block list should contain zero entries in the allow list' — available for manual scheduling but excluded from automatic Managed Deployments rollout.

`Managed ISPM: Customize control rollouts with Modified Deployments` **NOTES ↗** **STABLE**

Huntress Managed ISPM adds Modified Deployments, letting orgs hand-pick security controls and approval gates before anything rolls out.

Adds Modified Deployments, an org-level setting in Managed ISPM that lets you hand-pick which security controls (low, medium, or high impact) from the Huntress policy library roll out and on which days, with no changes deployed until explicitly approved.

New controls added by Huntress to the policy library trigger a notification under Modified Deployments but are held until you approve them, supporting strict compliance and change control requirements.

`Unsecured Credentials: New Credential Files Dashboard and Allowlisting`

NOTES ↗ **STABLE**

Huntress Unsecured Credentials gets a new Credential Files dashboard, pre-enable preview, and allowlisting to cut alert noise.

└──▷ **HOW TO FIND IT**

Audit noisy Unsecured Credentials alerts before re-enabling reporting — review flagged files and allowlist accepted ones to keep Incident Reports focused.

📍 In the Huntress console, go to Managed EDR › Process Insights › Credential Files. Review the listed files, select those you have accepted, and choose `Allowlist` to exclude them. Then enable Unsecured Credential Incident Reports.

> Adds a new Credential Files page under the Process Insights dashboard in Managed EDR to view and manage every file triggering an Unsecured Credentials report in one place.
>
> Enables preview of which Credential Files would generate an Incident Report before turning Unsecured Credential reporting on.
>
> Adds allowlisting for Credential Files so reviewed and accepted files are excluded from future Unsecured Credentials Incident Reports.

WAS THIS USEFUL? 👍 👎

Microsoft Defender

Feedback ⓘ changelog · Sep 24, 2026

BLOG ↗ RANK

Microsoft Defender is built-in antivirus and threat protection software for Windows that detects and removes malware and security threats.

> Microsoft Defender adds AI-generated email detonation summaries powered by Microsoft Security Copilot to speed up SOC investigations.

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

01 AI email detonation summaries via Security Copilot NEW

45

Defender now delivers an email detonation summary powered by Microsoft Security Copilot that provides AI-generated explanations of URL and file sandboxing results, intended to accelerate email investigation and hunting.

·
·

— Describes purpose and benefit but no exact UI path or config. 2026-09-24-e947d3e1

> changelog · Sep 24, 2026 BLOG ↗
>
> Microsoft Defender adds AI-generated email detonation summaries via Microsoft Security Copilot for faster SOC investigations.

WAS THIS USEFUL? 👍 👎

Iru

Feedback ⓘ What's New in Iru | September 18th NOTES ↗ RANK

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

> Iru extended device vulnerability management to Windows, added OS update status visibility for Windows devices, and gave admins control over enrollment invite email addresses for Iru Identity.

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

01 Vulnerability Management and Response on Windows NEW

40

Iru now supports deploying Vulnerability Management and Response on Windows devices.



— Names platform but no mechanism or scope detail. What's New in Iru | September 18th

02 OS update status view for Windows NEW

35

Adds a clear status view of OS updates for every Windows device.

— Describes feature but lacks specifics on where to find it.

What's New in Iru | September 18th

03 Enrollment invite email control for Iru Identity IMPROVED

35

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

— Clear capability but no config path or field named. What's New in Iru | September 18th

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

Iru adds Windows vulnerability management, OS update status visibility, and enrollment invite email control.

Supports deploying Vulnerability Management and Response on Windows devices.

Adds a clear status view of OS updates for every Windows device.

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

WAS THIS USEFUL?  

◆ Security Data Pipeline

Cribl Stream

Feedback ⓘ changelog · Sep 24, 2026 BLOG ↗

RANK

Cribl Stream is a data processing platform that routes, transforms, and enriches data from multiple sources to destinations in real-time.

> Cribl Stream ships a new Databricks Zerobus Destination that streams telemetry directly into Delta Tables over gRPC, skipping the usual object-storage staging step.

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

01 Databricks Zerobus Destination for Delta Tables NEW

85

A new Databricks Zerobus Destination routes events from Cribl Stream pipelines directly into Databricks Delta Tables via Databricks Zerobus Ingest over gRPC using JSON records, eliminating the need for an intermediate S3/Blob staging bucket and a separate ingestion job. It supports `catalog.schema.table` targeting so different telemetry streams can be routed to specific Unity Catalog Delta Tables from within existing pipeline configuration, and uses Databricks service-principal authentication scoped to a Unity Catalog table to align with existing governance policies. The destination exposes connection testing and metrics (throughput, errors, insertion latency) in the Stream operating experience, plus batching and flush controls to tune delivery throughput and latency for high-volume workloads.

— Names auth model, targeting scheme and controls but no UI path or exact flag names.

2026-09-24-9fc12d57

changelog · Sep 24, 2026 BLOG ↗ LATEST STABLE

Cribl Stream adds a Databricks Zerobus Destination to route telemetry directly into Delta Tables via gRPC without intermediate object storage.

Adds the Databricks Zerobus Destination, routing events from Cribl Stream pipelines directly into Databricks Delta Tables via Databricks Zerobus Ingest over gRPC with JSON records, eliminating the need for an intermediate S3/Blob staging bucket and separate ingestion job.

Supports `catalog.schema.table` targeting so operators can route different telemetry streams to specific Unity Catalog Delta Tables from within existing Stream pipeline configuration.

Supports Databricks service-principal authentication scoped to a Unity Catalog table,

aligning data movement with existing Databricks governance policies.

Exposes connection testing and destination metrics (throughput, errors, insertion latency) for the new Databricks destination inside the Cribl Stream operating experience.

Provides batching and flush controls on the Databricks Zerobus Destination to let operators tune delivery throughput and latency for high-volume enterprise telemetry workloads.

WAS THIS USEFUL?  

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

runZero is an asset discovery and inventory platform that identifies and maps all devices on networks without requiring agents.

> runZero's latest release adds OCSF 1.9.0 and CycloneDX CBOM exports to the asset inventory, API-based vulnerability suppression, OAuth support for its MCP service, and a broad set of role/permission and custom-scripting improvements alongside numerous smaller UI refinements.

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

01 **OCSF 1.9.0 export for asset inventory** NEW 85

Adds `OCSF 1.9.0` export to the asset inventory, exporting assets as Device Inventory Info events and observed services as Live Evidence Info events, accessible via the inventory export menu and the Export API; unmapped fields appear under the event's `unmapped` object.

— Names exact export standard, event types, and API surface changelog-20260925-87ef25bf

02 **OAuth support for MCP service** NEW 65

Adds OAuth support to the MCP service, including using Export Tokens and Organization Keys as OAuth Client Secrets in addition to API Clients.

— Names credential types but no setup steps given changelog-20260925-87ef25bf

03 **Custom roles and permission scoping for API keys and MCP** NEW 65

Enables assigning custom roles to API keys, default user roles, group permissions, SSO defaults, and MCP connections. Users can create API clients with custom permissions capped by their own access rights, and account-level API clients now support granular authentication scopes in addition to roles. The Annotator role has also been added to the list of default roles selectable in SSO settings.

— Names every affected surface but no exact config keys changelog-20260925-87ef25bf

04 **Custom integration scripting enhancements** IMPROVED 60

Scripts can now run in console instances in addition to Explorers and the CLI. Users can author and test custom integration scripts in-product using an AI model of choice, and scripts now support settings shared across protocols such as HTTP and TLS. Users can also submit scripts for inclusion as native runZero sources.

— Names execution surfaces and protocols but no exact commands

changelog-20260925-87ef25bf

THINNER COVERAGE BELOW

05 **API vulnerability suppression** NEW 50

Adds API support for suppressing and unsuppressing vulnerabilities, vulnerability groups, and findings.

— Names entities covered but not the endpoint itself changelog-20260925-87ef25bf

06 **Paginated analysis task reports page** IMPROVED 50

The reports page now displays a paginated list of all analysis task-based reports with search and delete capabilities.

— Names page and capabilities, exact UI path given changelog-20260925-87ef25bf

07 **CycloneDX CBOM export** NEW 40

Adds a Cryptographic Bill of Materials (CBOM) export in CycloneDX JSON format along with a detailed report.

— Names format but no access path or mechanism detail changelog-20260925-87ef25bf

08 **Dashboards usage column in query library** NEW 40

Adds the Dashboards column to the query library, showing how many personal dashboards use a saved query and whether it is also used on other dashboards.

— Clear UI location but limited mechanism detail changelog-20260925-87ef25bf

09 **Task reprocessing and large-change summaries** IMPROVED 40

Allows users to reprocess tasks directly from the user interface, and task summaries now indicate when a task involved more than 1,000 asset changes.

— Names UI action and the 1,000-change threshold changelog-20260925-87ef25bf

10 **AI-generated content labeling for EU compliance** NEW 40

runZero now labels AI Threads and AI Artifacts as AI-generated following EU guidance, embedding the label in exported reports, downloaded files, and generated custom-integration scripts.

— Names affected outputs but no configuration control described

changelog-20260925-87ef25bf

11 **Promiscuous ARP probe toggle for scans** NEW 35

Adds a configuration option on the scan page to toggle the promiscuous ARP probe.

— Clear UI location but no further mechanism detail changelog-20260925-87ef25bf

12 **WS-Discovery hostname contribution** IMPROVED 35

WS-Discovery XAddr endpoints now contribute to asset hostnames.

— Names the protocol mechanism but no user action changelog-20260925-87ef25bf

13 **Mermaid diagrams in AI Threads** NEW 30

Adds Mermaid-based charting and diagram support to the AI Threads feature.

— Bare description of the added capability changelog-20260925-87ef25bf

14 **Close Issue from Issue Details page** IMPROVED 30

Allows users to close an Issue from the Issue Details page.

— Names exact UI location, nothing further described changelog-20260925-87ef25bf

15 **Sharing standard reports via External Reporting** IMPROVED 30

Users can now share standard reports through the External Reporting feature.

— Names the feature but no mechanism or limits changelog-20260925-87ef25bf

16 **Sort Findings by name** IMPROVED 25

Users can now sort Findings by name.

— Trivial UI addition with clear location changelog-20260925-87ef25bf

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

runZero adds OCSF 1.9.0 and CycloneDX CBOM exports, API vulnerability suppression, OAuth for MCP, and custom-role API keys.

Adds `OCSF 1.9.0` export to the asset inventory, exporting assets as Device Inventory Info events and observed services as Live Evidence Info events via the inventory export menu and the Export API; unmapped fields appear under the event's `unmapped` object.

Adds a Cryptographic Bill of Materials (`CBOM`) export in CycloneDX JSON format along with a detailed report.

Adds API support for suppressing and unsuppressing vulnerabilities, vulnerability groups, and findings.

Adds OAuth support to the MCP service, including using Export Tokens and Organization Keys as OAuth Client Secrets in addition to API Clients.

Enables assigning custom roles to API keys, default user roles, group permissions, SSO defaults, and MCP connections.

> +17 more

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Okta for AI Agents

Feedback ⓘ changelog · Sep 25, 2026

CHANGELOG ↗ RANK

> Okta 2026.09.0 introduces unified identity management for AI agents alongside MCP server registration via Dynamic Client Registration and the OIN catalog, plus malware-proxy network detection, an NFC authenticator for shared workstations, Linux support for Okta Verify, and a low-code sign-in customizer.

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

01 New IP service categories in Enhanced Dynamic Network Zones

NEW 90

Adds a `REMOTE_DESKTOP` IP service category, letting admins include or exclude remote-desktop traffic in global session and app sign-in policies, and adds a `MALWARE_PROXIES` IP service category, powered by Okta CyberDefense, covering proxy services tied to malware and botnet activity (911 S5, NSOCKS, iProxy, BHProxies, and others).

Block access for sessions originating from known malware proxy networks by excluding `MALWARE_PROXIES` traffic in a global session policy.

> 📍 In the Admin Console, go to Security › Networks › Enhanced Dynamic Network Zones, create or edit a zone, add the `MALWARE_PROXIES` IP service category to the Exclude list, then reference that zone in your Global Session Policy rule to deny access.

Block or step-up remote-desktop sessions by tagging `REMOTE_DESKTOP` traffic in an Enhanced Dynamic Network Zone and referencing it in an app sign-in policy.

> 📍 In the Admin Console, go to Security › Networks › Enhanced Dynamic Network Zones, create or edit a zone, add the `REMOTE_DESKTOP` IP service category to the Exclude list, then apply that zone to an app sign-in policy rule to deny or require step-up for remote desktop traffic.

— Named categories with concrete policy steps to apply them changelog-20260925-2fd4fb13

02 MCP server registration via DCR, OIN catalog, and tool discovery

IMPROVED 70

Supports Dynamic Client Registration (DCR) when manually registering an MCP server — Okta automatically registers a client with the provider and populates credentials; enables adding MCP servers directly from an Okta Integration Network (OIN) catalog entry without entering connection details manually; adds tool discovery letting admins test credentials and browse available tools at registration time; and improves the MCP server registration UI with updated scope handling and tool visibility.

— Names several registration mechanisms but no exact UI path `changelog-20260925-2fd4fb13`

---

03 **Failed WebAuthn enrollment logging in System Log** NEW

70

Logs failed WebAuthn (FIDO2) enrollment attempts in the System Log using the `user.mfa.factor.activate` event, with debug data including AAGUID, `isBackupEligible`, and matched authenticator groups.

— Names exact event and debug fields, no query steps `changelog-20260925-2fd4fb13`

---

04 **Entitlement import safeguards per app** NEW

65

Adds entitlement import safeguards per app, configurable with percentage-based or absolute count thresholds to prevent accidental role or license removal on user unassignment, with an optional block on imports that modify or delete entitlement schemas.

— Explains mechanism and thresholds, no exact config path `changelog-20260925-2fd4fb13`

---

05 **Okta Integration Wizard for custom app integrations** NEW

60

Introduces the Okta Integration Wizard (OIW) for creating and deploying custom app integrations with SSO, SCIM, Entitlement Management, Universal Logout, API service integration, and custom API Integration Actions — usable as a template for multiple app instances.

— Lists supported capabilities, no setup steps `changelog-20260925-2fd4fb13`

---

06 **Higher Access Request limits** IMPROVED

60

Increases Access Request limits: users per task or question to 25 (from 10), entitlement bundles in an access level condition to 1,000 (from 100), groups in an access level condition to 1,000 (from 500), request type configuration lists per org to 250 (from 100), and request types per org to 750 (from 500).

— Concrete before/after numbers, no config path `changelog-20260925-2fd4fb13`

THINNER COVERAGE BELOW

---

07 **NFC authenticator for shared Windows workstations** NEW

55

Introduces the NFC authenticator for frontline workers on Windows shared workstations — users tap an NFC badge and enter a PIN to satisfy MFA without a phone, password, or shared account.

— Clear mechanism, no enrollment steps given changelog-20260925-2fd4fb13

08 **Okta Verify support for Linux desktops** NEW 55

Extends Okta Verify support to Linux desktop devices (Ubuntu 24.04 and 26.04) with phishing-resistant hardware-backed Okta FastPass authentication.

— Names OS versions, no install/config steps changelog-20260925-2fd4fb13

09 **Low-code sign-in customization** NEW 55

Adds Low-Code Sign In Customization, letting admins modify text, colors, and images in the Sign-In Widget and Interstitial authentication using a JSON templating language with syntax highlighting and suggestions, without editing code.

— Describes tool, no exact config key or command changelog-20260925-2fd4fb13

10 **High availability and incremental import for Generic Databases connector**

IMPROVED 55

Adds On-prem Connector for Generic Databases support for high availability by assigning multiple Okta On-Premises SCIM Agents to a single app instance, and adds incremental import support that retrieves only users and entitlement assignments changed since the last successful import.

— Explains mechanism, no config steps changelog-20260925-2fd4fb13

11 **Filtering, search, and CSV export on Applications page** NEW 55

Adds filtering by app type and status, search by name or client ID, last-modified-date view, and CSV export to the Applications page.

— Names UI capabilities and export format changelog-20260925-2fd4fb13

12 **Provisioning support for WordPress, Ivanti, Chef, Vercel, Sophos**

NEW 50

Adds provisioning support for `WordPress`, Ivanti, Progress Chef, Vercel, and Sophos Cloud integrations.

— Names five integrations, no setup detail changelog-20260925-2fd4fb13

13 **Unified identity management for AI agents** NEW 45

Adds unified identity management for AI agents: admins can set an external identifier for uniqueness, configure creation and matching criteria for imported agents, and preview agents before registration.

— Describes capabilities but no config keys or UI path changelog-20260925-2fd4fb13

14 **Email alerts for AD/LDAP agent disruption** NEW 45

Adds email notifications for Active Directory and LDAP agent disruption and recovery, configurable in the Admin Console.

— Names config location but not exact path changelog-20260925-2fd4fb13

15 **Desktop MFA factor discovery for Windows** NEW 30

Adds Desktop MFA Factor Discovery for Windows, showing users only their enrolled MFA factors instead of all available factors.

— Thin one-line description changelog-20260925-2fd4fb13

16 **Higher realm assignment limit per profile source** IMPROVED 30

Increases the maximum realm assignments per profile source from 30 to 100.

— Single number change, no context changelog-20260925-2fd4fb13

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Okta 2026.09.0 adds AI agent identity management, MCP server DCR and OIN catalog, NFC authenticator, Linux Okta Verify, malware proxy detection, and a low-code sign-in customizer.

└──▷ **HOW TO FIND IT**

Block access for sessions originating from known malware proxy networks by excluding `MALWARE_PROXIES` traffic in a global session policy.

📍 In the Admin Console, go to Security › Networks › Enhanced Dynamic Network Zones, create or edit a zone, add the `MALWARE_PROXIES` IP service category to the Exclude list, then reference that zone in your Global Session Policy rule to deny access.

Block or step-up remote-desktop sessions by tagging `REMOTE_DESKTOP` traffic in an Enhanced Dynamic Network Zone and referencing it in an app sign-in policy.

📍 In the Admin Console, go to Security › Networks › Enhanced Dynamic Network Zones, create or edit a zone, add the `REMOTE_DESKTOP` IP service category to the Exclude list, then apply that zone to an app sign-in policy rule to deny or require step-up for remote desktop traffic.

Adds a `REMOTE_DESKTOP` IP service category to Enhanced Dynamic Network Zones, letting admins include or exclude remote-desktop traffic in global session and app sign-in policies.

Adds a `MALWARE_PROXIES` IP service category to Enhanced Dynamic Network Zones, powered by Okta CyberDefense, covering proxy services tied to malware and botnet activity (911 S5, NSOCKS, iProxy, BHProxies, and others).

Supports Dynamic Client Registration (DCR) when manually registering an MCP server — Okta automatically registers a client with the provider and populates credentials.

Enables adding MCP servers directly from the Okta Integration Network (OIN) catalog entry, without entering connection details manually.

Logs failed WebAuthn (FIDO2) enrollment attempts in the System Log using the `user.mfa.factor.activate` event, with debug data including AAGUID, `isBackupEligible`, and matched authenticator groups.

+16 more

WAS THIS USEFUL? 👍 👎

iDenfy

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

> iDenfy's Sprint 135 release adds Bank Card Verification, a Blocklist API endpoint, and field-level filtering on the KYB company endpoint, alongside CSV, branding, and capture-flow refinements and an iOS SDK update that raises the minimum deployment target to iOS 15.

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

01 **Field selection for KYB company endpoint** NEW 85

A new `fields` parameter on the KYB company endpoint lets callers request only the data fields they need (e.g. `name`, `registrationNumber`, `directors`), reducing response payload.

Fetch only specific KYB company fields to avoid processing unnecessary data in your integration.

```
$ curl -X GET 'https://api.idenfy.com/kyb/company?
fields=name,registrationNumber,directors' -H 'Authorization:
Bearer <token>'
```

— Named endpoint, parameter, and runnable curl example changelog-20260925-928497ee

02 **iOS SDK 9.2.2 update** BREAKING 60

iOS SDK 9.2.2 adds a recommended document badge and iOS 27 support; the minimum deployment target has been raised to iOS 15 as of this version.

— Names version, OS support, and breaking minimum target changelog-20260925-928497ee

THINNER COVERAGE BELOW

03 **Original-language columns in KYC CSV reports** NEW 50

KYC CSV reports gain an Original language data section that exports document field values in their native script as separate columns.

— Names the exact report section but no column list changelog-20260925-928497ee

04 **Branding logo uploads and auto-prefill from website** IMPROVED 45

Branding configuration now supports PNG logo uploads, and the branding theme can be prefilled automatically from your website, pulling in colors and logo without manual palette configuration.

– Names both additions but no config keys or UI path changelog-20260925-928497ee

05 **Injection detection extended to external images** IMPROVED 35

Injection detection was extended experimentally to images captured outside the iDenfy capture flow.

– Notes experimental scope but no accuracy or coverage detail changelog-20260925-928497ee

06 **KYB shared identity verification across companies** NEW 30

KYB flows gain shared identity verification across companies, enabling cross-company identity reuse within KYB verification.

– Explains purpose but no config or scope detail changelog-20260925-928497ee

07 **Real-time framing feedback in document capture** IMPROVED 30

Document capture adds real-time framing feedback during the capture flow.

– States the change but no mechanism detail changelog-20260925-928497ee

08 **Bank Card Verification capability** NEW 25

A new Bank Card Verification product was added to the verification product suite.

– Named but no mechanism or workflow detail changelog-20260925-928497ee

09 **Blocklist API endpoint** NEW 20

A new Blocklist API endpoint was added to the platform, per the sprint changelog summary; no request/response details were provided.

– Mentioned only in summary, no mechanism or path given changelog-20260925-928497ee

└──▷ **BREAKING ON UPGRADE**

! iOS minimum deployment target raised to iOS 15 as of iOS SDK 9.2.2.

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

iDenfy Sprint 135 adds Bank Card Verification, a Blocklist API endpoint, KYB `fields` filtering, CSV native-script columns, and PNG logo uploads.

└──▷ **TRY IT**

Fetch only specific KYB company fields to avoid processing unnecessary data in your integration.

```
$ curl -X GET 'https://api.idenfy.com/kyb/company?
fields=name,registrationNumber,directors' -H 'Authorization: Bearer
<token>'
```

New `fields` parameter on the KYB company endpoint lets callers request only the data fields they need, reducing response payload.

KYC CSV reports gain an `Original language data` section that exports document field values in their native script as separate columns.

New Bank Card Verification capability added to the verification product suite.

New KYB shared identity verification across companies, enabling cross-company identity reuse within KYB flows.

PNG logo uploads now supported for branding configuration.

+4 more

└─▷ **BREAKING ON UPGRADE**

! iOS minimum deployment target raised to iOS 15 as of iOS SDK 9.2.2.

WAS THIS USEFUL? 👍 👎

Didit

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

> Didit's Web SDK gained native camera-control options, bringing it to parity with the native SDKs.

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

01 **Camera-control options in Web SDK** NEW 78

Web SDK 0.4.0 adds `defaultDocumentCamera`, `defaultLivenessCamera`, `showDocumentCameraSwitchButton`, and `showLivenessCameraSwitchButton` options, bringing camera-control parity with the native SDKs. The same behaviour can be applied via matching URL query parameters (e.g. `liveness_camera=back` and `liveness_camera_switch=false`) when opening the verification URL directly without the SDK.

Apply the same rear-camera, no-switcher behaviour when opening the verification URL directly without the SDK.

```
$ https://verify.didit.me/session/<session_id>?
liveness_camera=back&liveness_camera_switch=false
```

— Names exact flags and query parameters with a working URL example

changelog-20260925-32287544

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Web SDK 0.4.0 adds native camera-control options — defaultDocumentCamera, defaultLivenessCamera, and switcher-hide flags — plus matching URL query parameters.

└──▷ **TRY IT**

Apply the same rear-camera, no-switcher behaviour when opening the verification URL directly without the SDK.

```
$ https://verify.didit.me/session/<session_id>?
liveness_camera=back&liveness_camera_switch=false
```

Adds `defaultDocumentCamera`, `defaultLivenessCamera`, `showDocumentCameraSwitchButton`, and `showLivenessCameraSwitchButton`

> options to the Web SDK, bringing camera-control parity with the native SDKs.

WAS THIS USEFUL?  

ConductorOne

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

RANK

Govern workforce access, deploy AI agents, and get work done with C1.

> ConductorOne's window is dominated by breaking changes — split secrets deletion from credential revocation, removed a spend-remedy status endpoint, and tightened MCP toolset access and revoke-criteria matching — alongside new EU data residency, a public REST API for AI governance, Terraform resources for MCP sources, and wider change-history tracking.

└──▷ **WHAT SHIPPED · 18 FEATURES** most completely described first ? what's the number?

01 Secrets deletion split from credential revocation

BREAKING

90

`DELETE /api/v1/secrets/{vault_id}` now only deletes 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, so callers that relied on one call must now issue both.

Revoke a credential after deleting a shared secret — two calls are now required where one used to suffice.

```
$ curl -X DELETE
'https://<tenant>.conductor.one/api/v1/secrets/{vault_id}' \
  -H 'Authorization: Bearer <token>'

curl -X POST 'https://<tenant>.conductor.one/api/v1/app-secrets-
admin/revoke' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"vault_id": "<vault_id>"}'
```

— Named endpoints, exact before/after, and a runnable curl example

changelog-20260925-e7a9b8d7

02 Task API changes: spend-remedy status endpoint removed

75

BREAKING

`GET /api/v1/tasks/{task_id}/spend-remedy-status` has been removed — migrate to the task APIs. JS/TS clients must now 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` value.

— Names removed endpoint and new value but no migration example

changelog-20260925-e7a9b8d7

03 **EU data residency instance** NEW 65

Introduces an EU data residency instance with region-scoped login, redirect URIs, API and webhook endpoints, MCP server, and CLI configuration, selected at domain registration.

— Lists exact region-scoped surfaces and where region is chosen

changelog-20260925-e7a9b8d7

04 **Change history and justification tracking** NEW 60

Adds a Change history icon to connectors, policies, entitlement routing rules, SSO and branding settings, and AI Gateway spend controls, showing a before-and-after timeline per field. Connectors can now also require a justification whenever their configuration changes, captured alongside that change history.

— Names every area covered but no exact UI steps given changelog-20260925-e7a9b8d7

THINNER COVERAGE BELOW

05 **Terraform resources for MCP sources** NEW 55

C1's Terraform provider adds resources for hosted and external MCP sources, covering their details, authentication, and tool enablement policy.

— Names provider and resource scope, no exact resource syntax given

changelog-20260925-e7a9b8d7

06 **Revoke criteria no longer matches by resource-type name alone** 50

BREAKING

A revoke criteria scoped only by a legacy resource-type name now fails the automation step closed instead of silently matching every app sharing that name; teams must re-scope any automation still relying on name-only matching.

— Clear before/after and remediation note, no code surface named

changelog-20260925-e7a9b8d7

07 **Public REST API for AI governance** NEW 50

The public REST API now exposes classifiers, tool gates, classifier bindings, templates, and the agent classifier policy for AI governance management.

— Names the resources but no endpoint paths or examples changelog-20260925-e7a9b8d7

08 **Set assignee step in Transformation rules** NEW 50

Transformation rules gain a Set assignee step so findings can be auto-assigned to a specific person; manual assignment still takes precedence.

— Names the step and precedence rule but not full rule config changelog-20260925-e7a9b8d7

09 **MCP access requires explicit toolset grant** BREAKING 45

MCP access now requires an explicit toolset grant on the server — app membership or admin role alone no longer implies access.

— States the new requirement but no config path to grant it changelog-20260925-e7a9b8d7

10 **C1AI diagnoses stuck tasks** NEW 45

Super Admins can ask C1AI to diagnose a stuck, failed, or waiting task, tracing approvals, provisioning, access-profile origin, and connector handoff.

— Explains what it traces but not how to invoke it changelog-20260925-e7a9b8d7

11 **Apps and Entitlements tabs in My resources** NEW 45

The My resources page gains Apps and Entitlements tabs alongside its existing groups, vaults, credential stores, and secrets tabs.

— Exact page and tabs named, straightforward UI path changelog-20260925-e7a9b8d7

12 **Reviewer justification in access-review revoke tickets** IMPROVED 40

Access-review revoke tasks now automatically include the reviewer's justification in ticketing system output, with new template variables available for custom templates.

— Mentions new template variables but does not name them changelog-20260925-e7a9b8d7

13 **Conflict monitors support uneven group sizes** IMPROVED 40

Conflict monitors now support groups of very different sizes (e.g., 90 and 11) as long as the total stays under 1024.

— Gives concrete numeric example and limit, no usage steps changelog-20260925-e7a9b8d7

14 **Entitlement request forms in Slack** IMPROVED 40

Request forms configured on entitlements now render and submit directly in the Slack modal for the C1 app for Slack.

— Names the surface but no steps to configure it changelog-20260925-e7a9b8d7

15 **Unified guided setup for MCP integrations** IMPROVED

35

Registered MCP integrations — both catalog and external — now share a single guided setup flow covering connection, access, tool approval, and authentication.

— Describes flow stages but no exact UI steps changelog-20260925-e7a9b8d7

16 **Customer agents can converse with c1aw agents** NEW

35

Customer agents connected to the C1 MCP endpoint can now discover and converse with c1aw agents, with admin session management.

— Names the mechanism but no setup steps given changelog-20260925-e7a9b8d7

17 **Campaigns include pending account invitations** IMPROVED

30

Campaigns can now include users with pending account invitations instead of silently dropping them.

— States the fix but no config or setting to enable it changelog-20260925-e7a9b8d7

18 **Resource description column in access reviews** NEW

30

Adds a Resource description column to access-review tables, showing custom descriptions or connector-synced metadata as a fallback.

— Simple UI addition, thinly described changelog-20260925-e7a9b8d7

└──▷ ALSO FROM THESE RELEASES

.
.
.

└──▷ BREAKING ON UPGRADE

! A revoke criteria scoped only by a legacy resource-type name now fails the automation step closed 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 25, 2026 CHANGELOG ↗ LATEST STABLE

ConductorOne adds EU data residency, AI task diagnostics, public REST API for AI governance, and new MCP Terraform resources alongside breaking API changes.

└──▷ **TRY IT**

Revoke a credential after deleting a shared secret — two calls are now required where one used to suffice.

```
$ curl -X DELETE
'https://<tenant>.conductor.one/api/v1/secrets/{vault_id}' \
  -H 'Authorization: Bearer <token>'

curl -X POST 'https://<tenant>.conductor.one/api/v1/app-secrets-
admin/revoke' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"vault_id": "<vault_id>"}'
```

Save a C1AI reporting conversation so your team can revisit and refresh the same query without re-prompting.

> 📍 In the ConductorOne console, open C1AI and ask your reporting question, then select `Save report`, give it a name and description, and find it later under Platform › Reporting where you can refresh it or edit it with the Report assistant.

Adds `POST /api/v1/app-secrets-admin/revoke` to queue a credential revoke separately from secret deletion, after `DELETE /api/v1/secrets/{vault_id}` is scoped to delete-only.

Exposes classifiers, tool gates, classifier bindings, templates, and the agent classifier policy through the public REST API for AI governance management.

Adds Terraform resources for hosted and external MCP sources, covering details, authentication, and tool enablement policy via C1's Terraform provider.

Introduces an EU data residency instance with region-scoped login, redirect URIs, API and webhook endpoints, MCP server, and CLI configuration — selected at domain registration.

Adds a `Set assignee` step to Transformation rules so findings can be auto-assigned to a specific person (manual assignment takes precedence).

+14 more

└──▷ **BREAKING ON UPGRADE**

! A revoke criteria scoped only by a legacy resource-type name now fails the automation step closed 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? 👍 👎

Auth0

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

> Auth0 launched the Identity Conversion Suite — combining Anonymous Sessions for pre-auth user tracking and Experiment Center for auth-flow A/B testing — alongside a new organization member access-level model and the My Organization API for embeddable member management and bulk invitations.

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

01 **Custom user ID assignment via pre-user-registration action**

NEW 90

Using the `onExecutePreUserRegistration` action with api.user.setUserId(crypto.randomUUID()), developers can assign a UUID-based custom user ID at signup to keep identity consistent with an existing system, without building a self-hosted signup page.

Assign a UUID-based custom user ID at signup to keep identity consistent with an existing system, without building a self-hosted signup page.

javascript

```
exports.onExecutePreUserRegistration = async (event, api) => {
  const crypto = require('crypto');
  api.user.setUserId(crypto.randomUUID());
};
```

— Runnable code snippet with exact API calls. changelog-20260925-66040154

02 **Member access levels for organizations** NEW 65

Adds a member access level model (`none` , `readonly` , `limited` , `full`) that computes each member's permissions automatically from their organization connection access levels.

— Names four levels and mechanism, but automatic so little direct action.

changelog-20260925-66040154

THINNER COVERAGE BELOW

03 **Experiment Center for auth flow A/B testing** NEW 50

Also part of the Identity Conversion Suite, Experiment Center enables A/B tests and gradual rollouts on authentication flow changes — such as new sign-up fields, passkeys, or MFA rules — targeting a slice of traffic before full rollout.

— Describes behavior but no concrete surface to act on. product docs

04 **Anonymous Sessions for pre-auth user tracking** NEW

45

Part of the new Identity Conversion Suite bundle, Anonymous Sessions tracks and manages users from their first pre-authentication interaction and carries that activity forward to their authenticated profile upon sign-up or login.

— Explains mechanism but no config or API surface named. product docs

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Auth0: Adds a member access level model (none, readonly, limited, full) that computes each member's permissions automatically from their organization connection access levels.

.
.
.
.

└──▷ **USE IT**

Assign a UUID-based custom user ID at signup to keep identity consistent with an existing system, without building a self-hosted signup page.

javascript

```
exports.onExecutePreUserRegistration = async (event, api) => {
  const crypto = require('crypto');
  api.user.setUserId(crypto.randomUUID());
};
```

Activate the My Organization API to start using embeddable member management UI components and the bulk invitation endpoint.

📍 In the Auth0 Dashboard, go to Applications › APIs, locate My Organization API, and toggle it to enabled.

Activate the My Organization API by navigating to `Applications > APIs` in the Auth0 Dashboard.

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> JumpCloud's September 2026 release centers on a new unified Patch Management dashboard with a Windows Device Patch Report and bulk remediation, plus AI Assistant knowledge base uploads, configurable OAuth client authentication for agent integrations, and early moves into Android policy expansion and AI Gateway analytics.

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

01 **OAuth client authentication type for agent integrations** NEW 95

Adds a configurable OAuth client authentication type for agent credentials in the Authentication tab — `Client secret basic` (credentials sent in the authorization header, the default) or `Client secret post` (credentials sent in the request body) — to match external platforms such as Amazon Bedrock AgentCore.

Switch an agent integration to send OAuth credentials in the request body to satisfy Amazon Bedrock AgentCore's token endpoint requirements.

> 📍 In the Admin Portal, open the agent › Authentication tab › Client Authentication Type and select `Client secret post` (sends client ID and secret in the request body) instead of the default `Client secret basic`.

— Exact field names, defaults, and step-by-step navigation given

changelog-20260925-f4d0af19

02 **Device Patch Report for Windows** NEW 70

Adds a Device Patch Report (Windows) in the Admin Portal showing per-device Installed, Pending, and Failed status for each KB, with a preview of up to 500 rows and full CSV export.

— Concrete fields and limits, no exact nav path changelog-20260925-f4d0af19

THINNER COVERAGE BELOW

03 **Android policy expansion** NEW 10

The changelog summary notes an expansion of Android policy capabilities, but no bullet, example, or specific policy detail was provided.

— Only named in summary, no supporting detail changelog-20260925-f4d0af19

04 **AI Gateway analytics** NEW 10

The changelog summary mentions new AI Gateway analytics, but no further description, metric, or UI location was given.

— Only named in summary, no supporting detail `changelog-20260925-f4d0af19`

└──▷ **ALSO FROM THESE RELEASES**

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

JumpCloud adds unified Patch Management dashboards, Android policy expansion, AI knowledge base uploads, and AI Gateway analytics in September 2026.

└──▷ **HOW TO FIND IT**

Upload an internal runbook so AI Assistant can answer org-specific questions about your incident response process without exposing data to a shared model.

📍 In the Admin Portal, navigate to AI Assistant › Knowledge Base and select `Upload File`. Choose a PDF, TXT, DOCX, or MD file up to 10 MB. Repeat for up to 25 files per org. Once uploaded, AI Assistant will cite the source document when its responses draw on your content.

Switch an agent integration to send OAuth credentials in the request body to satisfy Amazon Bedrock AgentCore's token endpoint requirements.

📍 In the Admin Portal, open the agent › Authentication tab › Client Authentication Type and select `Client secret post` (sends client ID and secret in the request body) instead of the default `Client secret basic`.

Replace a deployed custom Windows EXE in the Private Repository without reassigning device groups, and confirm successful installation via exit code.

📍 In the Admin Portal, go to the application's Details tab › Update Version, upload the new EXE via `Upload File`, enter the install success exit code, then save. JumpCloud deploys the update to all devices already associated with the application.

Adds bulk remediation in the Windows Patch dashboard to install missing critical patches or non-compliant device updates immediately or on the policy schedule.

Adds a Device Patch Report (Windows) in the Admin Portal showing per-device Installed, Pending, and Failed status for each KB, with preview of up to 500 rows and full CSV export.

Adds a unified Patch Management dashboard in the Admin Portal with an Overview tab for cross-platform Windows and Apple patch compliance, plus dedicated Apple

and Windows tabs switchable from one page.

Adds configurable OAuth client authentication type for agent credentials in the Authentication tab — `Client secret basic` (credentials in authorization header, default) or `Client secret post` (credentials in request body) — to match external platforms such as Amazon Bedrock AgentCore.

WAS THIS USEFUL?  

Ping Identity

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

RANK

Read docs, explore use cases, learn best practices

> Ping Identity's PingOne and PingOne Advanced Identity Cloud window centers on early-access identity verification upgrades (NFC document reading, Recognize auto-enrollment), new outbound group sync provisioning, custom client IDs for imported OIDC apps, AI-agent bot detection enhancements, and a round of admin console UI refreshes, alongside a REST User-Agent override and the ThreatMetrix-to-LexisNexis DDP node rename.

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

01 Outbound group sync provisioning for external apps NEW

78

Adds an Outbound Name field (defaulting to the CN part of the group name) and a Group Settings / Sync Status tab under the Provisioning tab, enabling outbound group sync provisioning of external groups to third-party applications such as AWS.

— Names exact new UI field and tabs; no API endpoint given. product docs

02 Custom outbound User-Agent for REST service calls NEW

75

The `userAgent` property can be set in an external REST service call configuration to override the default outbound User-Agent header, useful when a third-party API enforces a particular User-Agent string.

Override the default outbound user agent for a specific external REST service call by setting the `userAgent` property, useful when a third-party API enforces a particular User-Agent string.

```json
{
  "userAgent": "MyCompanyIntegration/1.0"
}
```

— Named config property with a runnable example. changelog-20260925-fb4a2526

03 AI agent (CUA) detection in Protect Bot Detection IMPROVED

73

Enhances the PingOne Protect Bot Detection predictor with new rules to identify specific computer-using agent (CUA) types, giving granular visibility into AI-driven activity; requires Protect (Signals) SDK for Web v5.6.10 and, for some workforce flows, PingID Device Trust agent v1.2 or later. Also releases PingID Device Trust agent v1.2 with improved data collection and Protect (Signals) SDK for Web v5.6.11.

— Names exact SDK/agent versions and dependency requirements. product docs

04 **Custom client ID and secret for imported OIDC apps** NEW

60

Enables setting a custom client ID and client secret via the PingOne API when importing OIDC-based applications and custom resources from solutions such as PingFederate or third-party tools — applies only to newly imported resources, not existing ones.

— Names API capability and scope limit but no endpoint. product docs

THINNER COVERAGE BELOW

05 **NFC-based identity document verification in PingOne Verify** NEW

55

PingOne Verify adds NFC verification (early access), letting users read secure NFC chips on government-issued IDs and passports via their mobile device for contactless identity document verification.

— Explains mechanism but no setup or API detail. product docs

06 **UI refresh for environment creation, overview, and audit** IMPROVED

55

Updates the environment creation workflow with a modern, simplified, license-driven UI that restricts service and product selection to what's included in the customer's agreement and distinguishes paid from trial environments; updates the environment Overview page UI with a new look providing quick access to environment settings and resources; and updates the Audit UI with a more modern, intuitive experience and a simplified time zone selector.

— Describes three UI changes but no navigation paths. product docs

07 **ThreatMetrix nodes renamed to LexisNexis DDP** BREAKING

50

The ThreatMetrix nodes are renamed to LexisNexis Dynamic Decision Platform (DDP) nodes; any journeys or configurations referencing ThreatMetrix node names will need to be updated.

— Names the rename but gives no migration steps. changelog-20260925-fb4a2526

08 PingOne Recognize auto-enrollment from Verify NEW

30

Adds PingOne Recognize auto-enrollment from PingOne Verify (early access), enabling automatic device recognition enrollment during a Verify flow.

— Brief early-access mention, no setup detail. product docs

└──▷ BREAKING ON UPGRADE

! The ThreatMetrix nodes are renamed to LexisNexis Dynamic Decision Platform (DDP) nodes — any journeys or configurations referencing ThreatMetrix node names will need to be updated.

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

└──▷ USE IT

Override the default outbound user agent for a specific external REST service call by setting the `userAgent` property, useful when a third-party API enforces a particular User-Agent string.

json

```
{
  "userAgent": "MyCompanyIntegration/1.0"
}
```

└──▷ BREAKING ON UPGRADE

! The ThreatMetrix nodes are renamed to LexisNexis Dynamic Decision Platform (DDP) nodes — any journeys or configurations referencing ThreatMetrix node names will need to be updated.

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ

New Capability: Purview, Cyera & BigID Data Classification Labels Integration in Data Ac

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 migrating IdentityIQ to a Jakarta EE 10 foundation for self-hosted deployments, alongside extending Identity Graph to all Business+ tenants at no extra cost; it also added native DSPM classification imports, self-service SaaS connector log downloads, and richer workflow version history.

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

01 IdentityIQ 9.0 migration to Jakarta EE 10

BREAKING

80

IdentityIQ 9.0 moves to Jakarta EE 10 as its application-server foundation, adding support for Apache Tomcat 10.1.40, Tomcat 11.0.6, JBoss EAP, and WebSphere Liberty LTS runtimes. It introduces an Upgrade Pre-Check Utility that inventories Java EE customizations and generates proposed Jakarta EE updates for review and testing, and consolidates future capabilities (Metadata Campaigns, time-based access controls, modernized application definitions) onto this single forward release line. Breaking: Apache Tomcat 9.0.x features APR/native connectors and Tomcat Native 1.3.x will not be included in the upcoming Tomcat 9.1.x stream, so installations relying on these must audit their configuration before moving to 9.1.x.

— Names runtimes and a pre-check tool but no exact commands

IdentityIQ 9.0 and Jakarta EE: a modern…

02 Data classification import in Data Access Security

NEW

80

Adds a Classification Label Import wizard in Data Access Security to onboard Microsoft Purview, Cyera, or BigID as data classification sources, or upload a CSV, eliminating duplicate classification runs. Supports scheduled imports so labels stay current, configurable per application, and each application can run either the native classification engine or import results from a third-party DSPM vendor. Configuration lives under Compliance → Data Classification → Classification Import in the SailPoint ISC console.

Bring Purview classification labels into ISC entitlements so your access reviews reflect actual data sensitivity — without re-running classification.

📍 In the SailPoint ISC console: 1. Go to Data Access Security and onboard your data source application. 2. Select `Import data classification` during application setup. 3. Run discovery and permission collection. 4. Navigate to Compliance › Data Classification › Classification Import and configure Microsoft Purview as the vendor with a scheduled import cadence.

— Names vendors and exact config path, but only a UI navigation

`New Capability: Purview, Cyera & BigID …`

---

03 **SaaS connector log download in ISC UI** NEW 70

Adds a log download section at `Admin → Sources → Source Details → Advanced → Logging`, letting administrators download connector-scoped log snapshots covering ~24 hours of history without CLI access or a support case. Debug log level changes are reflected in downloaded snapshots, so admins can enable debug mode and re-download to capture more detail during an incident.

Troubleshoot a SaaS connector connectivity failure without opening a support case or using the CLI.

📍 1. In the ISC admin console, go to Admin → Sources → <your SaaS source> → Source Details → Advanced → Logging. 2. Enable `Debug` log level. 3. Reproduce or wait for the connectivity failure. 4. Click `Download` to retrieve the connector-scoped log snapshot (up to ~24 hours of history). 5. Review the log for error details; if the snapshot is still insufficient, escalate to SailPoint support.

— Exact UI path and time window named, but no API/CLI

`New Capability: SaaS Connector Log Down…`

---

04 **Identity Graph availability for Business+ tenants** NEW 65

Identity Graph (previously available as Observability & Insights) is extended to all Business+ customers at no additional charge, with no licensing, configuration, or implementation required. It provides unified interactive visibility across humans, non-humans, and agents, visualizes agent-human relationships, and flags overpermissioned identities by comparing access side-by-side; it's reached from the Identity Graph icon on the MySailPoint homepage.

Navigate to Identity Graph from the MySailPoint homepage to immediately visualize a user's full access landscape and assess blast radius after a suspected compromise.

📍 1. Log in to your MySailPoint homepage. 2. Click the Identity Graph icon (red circle on the left sidebar). 3. Search for the suspected compromised identity. 4. Review the blast radius view to see all access paths, entitlements, and connected

agents. 5. Use Direct action controls within the graph to remediate access without leaving the view.

— Clear UI entry point but no config or API detail

New Availability: The SailPoint Identit…

05 **Workflow version history comparison** IMPROVED 60

The workflow History tab lets admins open any history row to view event details such as Workflow Update , Version Created , or Version Forked , and follow a version-oriented event to compare two workflow versions side-by-side, including across forks.

Audit who changed a specific workflow step and what version that change produced, without digging through Admin Activity logs.

📍 1. Open the NERM Classic Admin console and navigate to the workflow you want to audit. 2. Select the History tab. 3. In the search box, enter the step label or step ID you are investigating. 4. Click a Version Created or Workflow Update row to open event details and confirm the actor, timestamp, and affected step. 5. From a version-oriented row, click through to open a side-by-side comparison of the before and after versions.

— Named event types but limited mechanism detail Enhancement: Workflow Versioning — Hist…

└──▷ BREAKING ON UPGRADE

! Apache Tomcat 9.0.x features APR/native connectors and Tomcat Native 1.3.x will not be included in the upcoming Tomcat 9.1.x stream; installations relying on these must audit their configuration before moving to 9.1.x.

New Capability: Purview, Cyera & BigID Data Classification Labels Integration in Data

BLOG ↗ LATEST STABLE

SailPoint Data Access Security gains native integration with Microsoft Purview, Cyera, and BigID to import data classification labels via a new wizard.

└──▷ HOW TO FIND IT

Bring Purview classification labels into ISC entitlements so your access reviews reflect actual data sensitivity — without re-running classification.

📍 In the SailPoint ISC console: 1. Go to Data Access Security and onboard your data source application. 2. Select Import data classification during application setup. 3. Run discovery and permission collection. 4. Navigate to Compliance › Data Classification › Classification Import and configure Microsoft Purview as the vendor with a scheduled import cadence.

Adds a Classification Label Import wizard in Data Access Security to onboard Microsoft Purview, Cyera, or BigID as data classification sources, or upload a CSV, eliminating duplicate classification runs.

Supports scheduled imports so classification labels stay current without manual re-runs, configurable per application in Data Access Security.

Enables each Data Access Security application to run either the native classification engine or import results from a third-party DSPM vendor (Purview, Cyera, BigID).

Exposes classification import configuration under `Compliance → Data Classification → Classification Import` in the SailPoint ISC console.

`New Capability: SaaS Connector Log Download (2026-09-24)` BLOG ↗ STABLE

SailPoint ISC admins can now download SaaS connector log snapshots directly from the UI without CLI access or a support case.

└──▷ **HOW TO FIND IT**

Troubleshoot a SaaS connector connectivity failure without opening a support case or using the CLI.

📍 1. In the ISC admin console, go to Admin → Sources → <your SaaS source> → Source Details → Advanced → Logging. 2. Enable `Debug` log level. 3. Reproduce or wait for the connectivity failure. 4. Click `Download` to retrieve the connector-scoped log snapshot (up to ~24 hours of history). 5. Review the log for error details; if the snapshot is still insufficient, escalate to SailPoint support.

Adds a log download section in the SaaS connector UI at `Admin → Sources → Source Details → Advanced → Logging`, letting administrators download connector-scoped log snapshots covering ~24 hours of history.

Debug log level changes are reflected in downloaded snapshots, so admins can enable debug mode and re-download to capture more detail during an incident.

`Enhancement: Workflow Versioning — History tab (2026-09-24)` BLOG ↗ STABLE

SailPoint: Enables opening any history row to view event details such as `Workflow Update`, `Version Created`, or `Version Forked`.

└──▷ **HOW TO FIND IT**

Audit who changed a specific workflow step and what version that change produced, without digging through Admin Activity logs.

📍 1. Open the NERM Classic Admin console and navigate to the workflow you want to audit. 2. Select the `History` tab. 3. In the search box, enter the step label or step ID you are investigating. 4. Click a `Version Created` or `Workflow Update` row to open event details and confirm the actor, timestamp, and affected step. 5. From a version-

oriented row, click through to open a side-by-side comparison of the before and after versions.

Enables following a version-oriented event to compare two workflow versions side-by-side, including across forks.

`IdentityIQ 9.0 and Jakarta EE: a modern application-server foundation for self-hosted`

BLOG ↗ STABLE

IdentityIQ 9.0 migrates to Jakarta EE 10, enabling support for Tomcat 10.1/11.0, JBoss EAP, and WebSphere Liberty LTS runtimes.

Supports Jakarta EE 10 as the new application-server foundation, enabling IdentityIQ to run on current LTS runtimes: Apache Tomcat 10.1.40, Tomcat 11.0.6, JBoss EAP, and WebSphere Liberty.

Introduces the IdentityIQ 9.0 Upgrade Pre-Check Utility, which inventories Java EE customizations and generates proposed Jakarta EE updates for team review and testing.

Delivers a single forward IdentityIQ release line, consolidating all future capabilities — including Metadata Campaigns, time-based access controls, and modernized application definitions — onto the Jakarta EE 10 stack.

└──▷ **BREAKING ON UPGRADE**

! Apache Tomcat 9.0.x features APR/native connectors and Tomcat Native 1.3.x will not be included in the upcoming Tomcat 9.1.x stream; installations relying on these must audit their configuration before moving to 9.1.x.

`New Availability: The SailPoint Identity Graph is Coming to Business+ Tenants (2026-09`

BLOG ↗ STABLE

SailPoint: Enables unified visibility across all identity types — humans, non-humans, and agents — in a single interactive graph view.

└──▷ **HOW TO FIND IT**

Navigate to Identity Graph from the MySailPoint homepage to immediately visualize a user's full access landscape and assess blast radius after a suspected compromise.

📍 1. Log in to your MySailPoint homepage. 2. Click the Identity Graph icon (red circle on the left sidebar). 3. Search for the suspected compromised identity. 4. Review the blast radius view to see all access paths, entitlements, and connected agents. 5. Use Direct action controls within the graph to remediate access without leaving the view.

> Visualizes agent-human relationships and flags overpermissioned identities by comparing access side-by-side.
>
> Extends Identity Graph (previously available as Observability & Insights) to all Business+ customers at no additional charge, with no licensing, configuration, or implementation required.

WAS THIS USEFUL? 👍 👎

Microsoft Entra

Feedback ⓘ changelog · Sep 24, 2026 BLOG ↗

RANK

Microsoft Entra is Azure's identity and access management platform enabling secure authentication, authorization, and conditional access for cloud and hybrid environments.

> Microsoft Entra's Global Secure Access now integrates with Microsoft Purview to enforce data-security policies at the network layer, blocking sensitive data from reaching shadow AI tools.

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

01 Purview integration with Global Secure Access NEW

53

Generally available integration between Microsoft Purview and Microsoft Entra Global Secure Access enforces context-aware classification and data-loss policies at the network layer, blocking sensitive files and text from being uploaded to unsanctioned AI tools in real time. Coverage extends to both human traffic and on-behalf-of (OBO) agentic traffic.

— Explains mechanism and scope but no config/UI path given 2026-09-24-f3e05cbd

> changelog · Sep 24, 2026 BLOG ↗
>
> Microsoft Purview and Entra Global Secure Access now GA: enforce data-security policies at the network layer to block sensitive data from reaching shadow AI.

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ changelog · Sep 25, 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 shipped multi-IdP enterprise SSO per Organization, made Client ID Metadata Documents generally available, and added SMS delivery event logging with a dedicated logs view.

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

01 Client ID Metadata Documents GA NEW

85

Enables the Publish CIMD support toggle in the OAuth applications settings page under Client onboarding, with a Client admission control (Pre-registered clients only) and per-client scope management via the Applications tab — making Client ID Metadata Documents generally available for all applications.

Allow MCP or other public OAuth clients to authenticate using a URL as their client_id without pre-registering a secret, then restrict flows to only reviewed clients.

> 📍 In the Clerk Dashboard, go to OAuth applications settings › Client onboarding, enable Publish CIMD support, then set Client admission to Pre-registered clients only and add trusted client ID URLs from the Applications tab.

— Exact toggle, control and tab names plus dashboard steps changelog-20260925-ef501d89

02 Multiple enterprise connections per Organization IMPROVED

80

Removes the one-enterprise-connection-per-Organization limit across the Clerk Dashboard and Backend API, allowing multiple SAML or OIDC connections — each with its own domain, protocol, and Directory Sync — on a single Organization. Supports multi-directory membership resolution: a user's active status and Organization membership are recalculated across all enabled directories, with deactivation or membership removal triggered only when every managing directory has deprovisioned them.

Federate a post-acquisition org whose employees sign in through two separate IdPs by adding a second SAML or OIDC connection to the same Organization.

> 📍 In the Clerk Dashboard, go to the SSO connections page, select the target Organization, and add a second enterprise connection with its own domain and Directory Sync configuration.

— Names mechanism, protocols, and dashboard path but no API example

changelog-20260925-ef501d89 | Multiple enterprise connections per Org…

03 **SMS delivery event logs** NEW 80

Adds SMS delivery event logging with five normalized event types — `sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, and `sms.unconfirmed` — each carrying a `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`) and full trace linking across a message's lifecycle. A dedicated SMS logs view, filterable by event type or destination phone number, is available on all plans as part of Application Logs.

— Named event types and fields but only a UI path, no example changelog-20260925-ef501d89

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Clerk adds multi-IdP enterprise SSO per Organization, GA Client ID Metadata Documents, and SMS delivery event logs

└──▷ **HOW TO FIND IT**

Allow MCP or other public OAuth clients to authenticate using a URL as their client_id without pre-registering a secret, then restrict flows to only reviewed clients.

> 📍 In the Clerk Dashboard, go to OAuth applications settings › Client onboarding, enable Publish CIMD support, then set Client admission to Pre-registered clients only and add trusted client ID URLs from the Applications tab.

Federate a post-acquisition org whose employees sign in through two separate IdPs by adding a second SAML or OIDC connection to the same Organization.

> 📍 In the Clerk Dashboard, go to the SSO connections page, select the target Organization, and add a second enterprise connection with its own domain and Directory Sync configuration.

Enables Publish CIMD support toggle in the OAuth applications settings page under Client onboarding, with a Client admission control (Pre-registered clients only) and per-client scope management via the Applications tab — making Client ID Metadata Documents generally available for all applications.

Adds SMS delivery event logging with five normalized event types — `sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, and `sms.unconfirmed` — each carrying a `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`) and full trace linking across a message's lifecycle.

Adds a dedicated SMS logs view filterable by event type (e.g. `sms.failed`) or destination phone number, available on all plans as part of Application Logs.

Removes the one-enterprise-connection-per-Organization limit across the Clerk Dashboard and Backend API, allowing multiple SAML or OIDC connections — each with its own domain, protocol, and Directory Sync — on a single Organization.

Supports multi-directory membership resolution: a user's active status and Organization membership are recalculated across all enabled directories, with deactivation or membership removal triggered only when every managing directory has deprovisioned them.

Multiple enterprise connections per Organization (2026-09-24) BLOG ↗

STABLE

Clerk Organizations can now link multiple SAML or OIDC enterprise connections, enabling multi-directory user provisioning.

Supports linking multiple SAML or OIDC enterprise connections to a single Organization, allowing more than one directory to provision the same user.

WAS THIS USEFUL? 👍 👎

Aembit

Feedback

changelog · Sep 25, 2026

CHANGELOG ↗

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 new credential automation across Kubernetes, GitHub Actions, and SAML-based OAuth flows, including automatic mTLS certificate issuance for workloads via its Secrets Operator.

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

01 mTLS certificate issuance via Secrets Operator

BREAKING

95

The Secrets Operator can issue mTLS certificates to a Kubernetes workload without the private key ever leaving the cluster, using `credentialType: X509Svid`; generation, signing, and renewal are handled automatically, writing a `kubernetes.io/tls` Secret instead of the generic Secret type used by other credential types. Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret, a breaking change for existing schedules.

Issue mTLS certificates to a Kubernetes workload without the private key ever leaving the cluster — Secrets Operator handles generation, signing, and renewal automatically.

yaml

```yaml
credentialType: X509Svid
```

— Names config field, Secret type, and exact breaking-change mechanism.

changelog-20260925-58a29786

02 SAML attribute mapping into OAuth 2.0 credentials

NEW

60

A SAML user's identity can now be bound into an OAuth 2.0 credential by reading their NameID or a group attribute directly from the SAML assertion, configured as a dynamic claim in the Tenant UI's OAuth 2.0 Authorization Code Credential Provider form using values like `${saml.response.subject.nameId}` or `${saml.response.attributes['groups']}`.

Bind a SAML user's identity into an OAuth 2.0 credential by reading their NameID and a group attribute directly from the SAML assertion.

📍 In the Aembit Tenant UI, open the OAuth 2.0 Authorization Code Credential Provider form, add a dynamic claim, and set its value to

${ `saml.response.subject.nameId` } or
${ `saml.response.attributes` ['groups']}.

— Named claim syntax given, but only UI steps, no API. changelog-20260925-58a29786

└──▷ **BREAKING ON UPGRADE**

! Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret (X509Svid writes a `kubernetes.io/tls` Secret; all other credential types write a generic Secret).

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

└──▷ **USE IT**

Issue mTLS certificates to a Kubernetes workload without the private key ever leaving the cluster — Secrets Operator handles generation, signing, and renewal automatically.

yaml

```yaml
credentialType: X509Svid
```

Let a GitHub-hosted Actions runner pick a specific IAM role when an Access Policy holds multiple AWS STS Federation Credential Providers.

.github/workflows/deploy.yml

```yaml
- uses: aembit/aembit-github-action@v1.3.0
  with:
    aws-access-key-id: AKIA_PROD_DEPLOY
```

Bind a SAML user's identity into an OAuth 2.0 credential by reading their NameID and a group attribute directly from the SAML assertion.

📍 In the Aembit Tenant UI, open the OAuth 2.0 Authorization Code Credential Provider form, add a dynamic claim, and set its value to
${ `saml.response.subject.nameId` } or
${ `saml.response.attributes` ['groups']}.

└──▷ **BREAKING ON UPGRADE**

! Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret (X509Svid writes a `kubernetes.io/tls` Secret; all other credential types write a generic Secret).

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

All the latest releases from the 1Password team

> 1Password's CLI beta cuts down on repeated authorization prompts for scripted/agent use on macOS and speeds up commands for Environments accounts.

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

01 **Shared desktop app auth session for CLI on macOS** IMPROVED

40

On macOS, successive CLI invocations by coding agents or scripts now share a single desktop app authorization session instead of prompting for each command, reducing repeated auth prompts.

— Explains behavior change but no flag or config named changelog-20260925-16e950f8

02 **Faster CLI commands for Environments accounts** IMPROVED

25

CLI commands run faster for accounts using 1Password Environments.

— Bare performance claim with no mechanism or numbers changelog-20260925-16e950f8

└──▷ **ALSO FROM THESE RELEASES**

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

1Password CLI 2.40.0-beta.02 reduces repeated auth prompts on macOS and speeds up Environments account commands.

On macOS, successive CLI invocations by coding agents or scripts now share a single desktop app authorization session instead of prompting for each command.

CLI commands run faster for accounts using 1Password Environments.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

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

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

01 **DNS-over-HTTPS fallback for remote workers** NEW 80

Setting `ENABLE_DOH_FALLBACK=true` in `dns_agent.conf` enables DNS-over-HTTPS fallback so filtering stays enforced for remote workers whose VPN tunnel blocks DNS-over-TLS (port 853).

Enable DNS-over-HTTPS fallback so filtering stays enforced for remote workers whose VPN tunnel blocks DNS-over-TLS (port 853).

```env
ENABLE_DOH_FALLBACK=true
```

— Exact config key and env var given, mechanism explained changelog-20260925-4d627d69

02 **New DNSFilter Root CA G1 certificate check** BREAKING 75

v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices without the updated `Combined_Certificates.mobileconfig` profile will display a Fix Installation Issues prompt in the menu bar after upgrading, and the new profile must be pushed via MDM to clear it.

— Names exact profile file and MDM remediation step changelog-20260925-4d627d69

└──▷ BREAKING ON UPGRADE

! The macOS Roaming Client now requires macOS 13 Ventura or later; devices running macOS 12 or earlier will not run v2.4.6.

! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices without the updated `Combined_Certificates.mobileconfig` profile will display a Fix Installation Issues prompt in the menu bar after upgrading and must have the new profile pushed via MDM to clear it.

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

└──▷ USE IT

Enable DNS-over-HTTPS fallback so filtering stays enforced for remote workers whose VPN tunnel blocks DNS-over-TLS (port 853).

dns_agent.conf

```
ENABLE_DOH_FALLBACK=true
```

└──▷ **BREAKING ON UPGRADE**

! The macOS Roaming Client now requires macOS 13 Ventura or later; devices running macOS 12 or earlier will not run v2.4.6.

! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices without the updated `Combined_Certificates.mobileconfig` profile will display a Fix Installation Issues prompt in the menu bar after upgrading and must have the new profile pushed via MDM to clear it.

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

> Firezone this window added recognition of a new device-attestation certificate claim, exposed gateway lifecycle management via its own binary, and improved Android connectivity recovery after sleep.

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

01 **Resume-from-sleep connectivity restoration on Android**

IMPROVED 23

Firezone Android improves restoration of connectivity when the device resumes from sleep.

— Only a brief summary, no mechanism or scope given Android 1.5.15 (2026-09-24)

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Firezone: Recognizes `iru-id` SAN claims in device attestation certificates as an alias for `kandji-id`.

└──▷ **TRY IT**

Authenticate a gateway and manage its service lifecycle without needing the separate `firezone` CLI binary.

```
$ firezone-gateway authenticate
firezone-gateway enable-service
firezone-gateway disable-service
```

Android 1.5.15 (2026-09-24) BLOG ↗ STABLE

Firezone Android 1.5.15 adds `iru-id` SAN claim support in device attestation and improves resume-from-sleep connectivity restoration.

Recognizes `iru-id` SAN claims in device attestation certificates as an alias for `kandji-id`.

WAS THIS USEFUL? 👍 👎

Twingate

Feedback ⓘ AI Usage Overview NOTES ↗ RANK

Identity-based access for users, services, and AI agents that deploys in minutes, scales to every resource, and finally lets you retire your VPN.

> Twingate added an AI Usage Overview to DNS filtering, letting admins track AI service activity across devices and export per-device usage data for policy and license reviews.

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

01 **AI Usage Overview in DNS filtering** NEW 72

A new AI Usage Overview section in DNS Filtering surfaces which AI services appear in device DNS activity, letting admins expand any AI service to view its matched domains. It supports a per-device CSV export (via Export CSV) to support policy discussions and license reviews.

Identify which AI services employees are accessing and export a per-device breakdown to CSV for license review or policy discussions.

> 📍 In the Twingate console, go to DNS Filtering › AI Usage Overview, expand any AI service to view its matched domains, then select Export CSV to download the per-device activity report.

— Names exact UI path and export action, lacks technical mechanism

changelog-20260925-b7769956 AI Usage Overview

AI Usage Overview NOTES ↗ STABLE

Twingate DNS filtering gains an AI Usage Overview section to track AI service activity across devices, with per-device CSV export.

└──▷ **HOW TO FIND IT**

Identify which AI services employees are accessing and export a per-device breakdown to CSV for license review or policy discussions.

> 📍 In the Twingate console, go to DNS Filtering › AI Usage Overview, expand any AI service to view its matched domains, then select Export CSV to download the per-device activity report.

Adds an AI Usage Overview section to DNS filtering that surfaces which AI services appear in device DNS activity, with per-device CSV export for policy and license review follow-up.

WAS THIS USEFUL? 👍 👎

tailscale

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale's tailnet now supports subscribing to network events for routing notifications, along with smarter binary handling and finer control over log verbosity.

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

01 **Ungrouped log output mode** NEW 65

A `log-mode=ungrouped` config option suppresses verbose grouped log output, useful for container deployments where logs need to stay parseable by a downstream log aggregator.

Suppress verbose grouped log output in a container deployment to keep logs parseable by a downstream log aggregator.

env

```
log-mode=ungrouped
```

— Names exact config value but lacks broader documentation of other log-mode options.

changelog-20260925-234c8f4a

THINNER COVERAGE BELOW

02 **Skip redundant binary downloads on Linux** IMPROVED 40

Tailscale skips re-downloading binaries on Linux when a file with a matching SHA already exists, avoiding unnecessary network transfer during updates.

— Explains behavior but gives no command or config to trigger it.

changelog-20260925-234c8f4a

03 **Tailnet event subscriptions for notifications** NEW 30

Tailscale now supports subscribing to tailnet events and routing notifications through an integration or app, though the exact mechanism and supported integrations are not detailed.

— No integration names, endpoints, or setup steps given. changelog-20260925-234c8f4a

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

tailscale: Supports subscribing to tailnet events and routing notifications through an integration or app.

└──▷ **USE IT**

Suppress verbose grouped log output in a container deployment to keep logs parseable by a downstream log aggregator.

env

```
log-mode=ungrouped
```

Skips re-downloading binaries on Linux when a file with a matching SHA already exists.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Ubiq Security

Feedback ⓘ changelog · Sep 25, 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 shipped Apache Airflow encrypt/decrypt operators, a new Rust SDK (Beta), a Kong API gateway tokenization plugin, and Grafana Loki as a new SIEM destination, alongside broad dataset type expansion (integers, dates, tokens, generic strings) across its C, JavaScript/Node.js, Python, PHP, and .NET libraries, plus consolidation of IDP/SCIM integrations and Access Graph and Security History improvements.

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

01 JavaScript/Node.js structured encryption API expansion NEW

95

Expands dataset type support in the Node.js 2.4.0 library to include `int32`, `int64`, `date`, `datetime`, `token` (64/128-byte), and generic strings with Base32, Base64, or no encoding, and adds the corresponding JavaScript async API: `EncryptNumberAsync` / `DecryptNumberAsync` and `encryptNumberForSearchAsync` for int32 (range -99,999,999 to 99,999,999) and int64 (range -9,999,999,999,999,999 to 9,999,999,999,999,999); `EncryptDateTimeAsync` / `DecryptDateTimeAsync` / `encryptDateTimeForSearchAsync` (plaintext range 1653-02-10 to 2286-11-20 UTC); `EncryptDateAsync` / `DecryptDateAsync` / `encryptDateForSearchAsync`; Token dataset `EncryptAsync` / `DecryptAsync` / `encryptForSearchAsync` for 64 or 128 character base62 tokens; and generic string `EncryptAsync` / `DecryptAsync` / `encryptForSearchAsync` with base32/base64 encoding and configurable padding. Also adds `buildStructuredAsync` on `CryptographyBuilder` via `withCredentialsObject` and `withConfigurationObject`, UbiqFactory.createCredentialsWithIdp(<username>, <password>, <host>) for IDP-based credentials, UbiqFactory.readConfigurationFromFile(configurationFile), and `addReportingUserDefinedMetadata` to attach a JSON string (max 1024 characters) to usage records.

Encrypt and search over an integer field in a database without revealing plaintext values — useful for encrypted PII lookups on numeric identifiers.

```text
const credentials =
ubiq.UbiqFactory.readCredentialsFromFile('./credentials',
'default');
const configuration = ubiq.UbiqFactory.defaultConfiguration();
const ubiqEncryptDecrypt = await (new
ubiq.CryptographyBuilder()).withCredentialsObject(credentials).wi
thConfigurationObject(configuration).buildStructuredAsync();
const intVal = 12345678;
const encrypted = await
ubiqEncryptDecrypt.EncryptNumberAsync('int32', intVal, []);
const searchValues = await
ubiqEncryptDecrypt.encryptNumberForSearchAsync('int32', intVal,
[]);
```

Authenticate using an Identity Provider instead of static API keys, for orgs managing credentials through SSO/IDP.

```text
const credentials =
ubiq.UbiqFactory.createCredentialsWithIdp(<username>, <password>,
<host>);
const ubiqEncryptDecrypt = await (new
ubiq.CryptographyBuilder()).withCredentialsObject(credentials).bu
ildStructuredAsync();
```

Tag usage records with environment or team metadata for audit and billing attribution.

```text
const ubiqEncryptDecrypt = await (new
ubiq.CryptographyBuilder()).withCredentialsObject(credentials).wi
thConfigurationObject(configuration).buildStructuredAsync();
ubiqEncryptDecrypt.addReportingUserDefinedMetadata('{"team":
"security", "env": "prod"}');
```

— Runnable code for every named function, with ranges and formats.

changelog-20260925-917e6c4d

02 **Encrypt/decrypt operators for Apache Airflow** NEW 80

Adds `UbiqEncryptOperator` and `UbiqDecryptOperator` for Apache Airflow DAGs, letting pipeline tasks encrypt and decrypt named dataset columns (e.g. SSN, Birthdate) in place without touching non-sensitive fields, using credentials loaded via ubiq_security.configCredentials(config_file=..., profile=...).

Encrypt sensitive columns (SSN, Birthdate) pulled from BigQuery in a DAG, then decrypt them in a downstream task — without touching non-sensitive columns.

python

```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')

get_data = BigQueryGetDataOperator(
    task_id='get_data',
    dataset_id=DATASET_NAME,
    table_id=TABLE_NAME,
    max_results=10,
    selected_fields='name,email,ssn,birthday',
)

enc_task = UbiqEncryptOperator(
    task_id='enc_task',
    credentials=creds,
    data_task=get_data.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate'],
)

dec_task = UbiqDecryptOperator(
    task_id='ubiq_dec_task',
    credentials=creds,
    data_task=enc_task.task_id,
    dataset_names=['', '', 'SSN', 'Birthdate'],
)

get_data >> enc_task >> dec_task
```

— Full runnable DAG example with operator and credential calls.

changelog-20260925-917e6c4d changelog-20260925-2f75a2ec

03 **Go key-number retrieval functions** NEW 80

Adds `CipherAndKeyNumber`, `GetCurrentKeyNumber`, and `GetKeyNumber` functions to the Go library to retrieve the key number used to produce a ciphertext without decrypting it, useful for auditing key rotation coverage across a dataset.

Inspect which encryption key version was used to produce a ciphertext — without decrypting — to audit key rotation coverage across a dataset.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Key number: %d\n", keyNum)
```

Inspect the key number used to produce a ciphertext in Go without performing a full decryption — useful for key rotation audits.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Encrypted with key number: %d\n", keyNum)
```

— Runnable code sample with all three named functions. changelog-20260925-917e6c4d changelog-20260925-2f75a2ec

04 **Dataset type expansion in C, PHP, Python and .NET libraries**

NEW 75

Expands dataset type support to `int32`, `int64`, `date`, `datetime`, `token` (64/128-byte), and generic strings with Base32, Base64, or no encoding across the C 2.3.0.0, PHP 2.1.0, Python 2.4.0, and .NET 2.4.2 libraries. C 2.3.0.0 also adds support for the dataset creation wizard and lets `date` / `datetime` datasets accept input character sets other than `0-9`; .NET 2.4.2 raises a clear error when a string method is called on an integer, date, or datetime dataset. Also publishes a new structured dataset type reference covering Date & Time, Date, Integer, Token, Formatted String, and Generic String types with their options and supported ranges.

— Names types, encodings and library versions but no direct command.

changelog-20260925-917e6c4d changelog-20260925-2f75a2ec

05 **Grafana Loki as a SIEM destination** NEW 70

Adds Grafana Loki as a SIEM destination for Ubiq activity and usage events, for enterprise customers (previously Splunk only), configured from Settings > Integrations by adding a Loki connection with endpoint and credentials after Ubiq support enables the feature flag.

Enable Grafana Loki as a SIEM destination so Ubiq activity events flow into your existing Loki stack alongside Splunk.

📍 In the Ubiq console, go to Settings › Integrations, select `Grafana Loki`, enter your Loki endpoint and credentials, then save. Contact Ubiq support to enable the feature flag for your enterprise account.

Enable Grafana Loki as a SIEM destination for Ubiq activity events so SOC dashboards receive tokenization audit trails alongside other log streams.

📍 1. Contact Ubiq support to enable Grafana Loki for your enterprise account. 2. In the Ubiq dashboard, go to Settings › Integrations. 3. Add a new Grafana Loki connection and supply your Loki endpoint and credentials. 4. Save and verify events appear in your Loki data source.

— Exact UI path and setup steps given. `changelog-20260925-917e6c4d` `changelog-20260925-2f75a2ec`

06 **Rust SDK (Beta)** NEW 65

Adds a new Ubiq Rust SDK (Beta), published on crates.io as `ubiq-rust`, supporting structured and unstructured encryption, integer and date/time datasets, `encrypt_for_search`, and key zeroization.

Add the `ubiq-rust` crate to a Rust project to start encrypting data with the new beta SDK, including search-compatible encryption.

```toml
ubiq-rust = "1.0.0"
```

Bootstrap a new Go project using the Rust SDK to encrypt a sensitive field with `encrypt_for_search`, enabling ciphertext lookups across a dataset.

```toml
# Cargo.toml
[dependencies]
ubiq-rust = "1.0.0"
```

— Crate name and capability list, plus install snippet. `changelog-20260925-917e6c4d` `changelog-20260925-2f75a2ec`

THINNER COVERAGE BELOW

07 **IDP and SCIM integration consolidation** NEW 55

Consolidates IDP and SIEM connections onto a single Integrations page, adding support for self-signed IDPs, SCIM provisioning, and direct IDP-group-to-Access-Group mapping, all configurable from that one page.

– Names capabilities but no exact config or endpoint. changelog-20260925-917e6c4d

changelog-20260925-2f75a2ec

08 **PHP bulk key/config prefetch** IMPROVED 50

Adds bulk prefetch of dataset configuration and encryption keys in a single call in the PHP 2.1.0 library, reducing first-use latency for many datasets.

– Names mechanism and library version, no code sample. changelog-20260925-917e6c4d

changelog-20260925-2f75a2ec

09 **.NET structured encryption uses BITLEN per updated NIST FF1** IMPROVED 50

Updates structured encryption in the .NET library to use the `BITLEN` function in place of `LOG`, per NIST FF1 updated recommendations.

– States the mechanism change but no reader action. changelog-20260925-917e6c4d

10 **Kong API gateway tokenization plugin** NEW 45

Adds a new Ubiq plugin for the Kong API gateway providing field-level tokenization and detokenization of JSON request and response bodies with no application code changes.

– Describes capability but no config or endpoint shown. changelog-20260925-917e6c4d

11 **Access Graph extended to IDP groups with bidirectional filtering** IMPROVED 45

Extends the Access Graph to include IDP groups and enables bidirectional filtering so access can be traced from either an identity or a dataset.

– Describes new filtering behavior, no exact navigation given.

changelog-20260925-917e6c4d

12 **Security History event attribution and SCIM failure logging** IMPROVED 40

Adds attribution of Security History events to the IDP identity that performed them, and records SCIM sync and attribute-mapping failures as events.

– Describes what changed but no way to act on it directly. changelog-20260925-917e6c4d

13 **Dataset Management dropdown search and length display** IMPROVED 40

Adds search and multi-select to the Dataset Association dropdown in Dataset Management; generic string datasets now display their maximum input length.

– Names the UI control but gives no exact navigation path. changelog-20260925-917e6c4d

14 New Airflow and Self-Signed IDP library guides NEW

25

Adds new library guides for Apache Airflow and Self-Signed IDP authentication, alongside updated guides for Go, Java, Node.js, PHP, C/C++, BigQuery, and Snowflake.

— Only a list of updated documentation, no technical detail. changelog-20260925-917e6c4d

changelog-20260925-2f75a2ec

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Ubiq Security: Adds new Ubiq encrypt and decrypt operators for Apache Airflow DAGs to protect data in pipeline tasks.

└──▷ **USE IT**

Inspect which encryption key version was used to produce a ciphertext — without decrypting — to audit key rotation coverage across a dataset.

go

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Key number: %d\n", keyNum)
```

Add the `ubiq-rust` crate to a Rust project to start encrypting data with the new beta SDK, including search-compatible encryption.

Cargo.toml

```toml
ubiq-rust = "1.0.0"
```

Enable Grafana Loki as a SIEM destination so Ubiq activity events flow into your existing Loki stack alongside Splunk.

📍 In the Ubiq console, go to Settings › Integrations, select Grafana Loki, enter your Loki endpoint and credentials, then save. Contact Ubiq support to enable the feature flag for your enterprise account.

Consolidates IDP and SIEM connections onto a single Integrations page, adding support for self-signed IDPs, direct IDP-group-to-Access-Group mapping, and new Apache Airflow and Self-Signed IDP authentication library guides.

Adds bulk prefetch of dataset configuration and encryption keys in a single call in the PHP 2.1.0 library, reducing first-use latency.

Publishes a new structured dataset type reference covering Date & Time, Date, Integer, Token, Formatted String, and Generic String types with options and supported ranges.

changelog · Sep 25, 2026 CHANGELOG ↗ STABLE

Ubiq adds Rust SDK, Kong gateway plugin, Airflow operators, Grafana Loki SIEM, and new dataset types across C, Node.js, Python, PHP, .NET, and Go libraries.

└──▷ **USE IT**

Inspect the key number used to produce a ciphertext in Go without performing a full decryption — useful for key rotation audits.

go

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Encrypted with key number: %d\n", keyNum)
```

Bootstrap a new Go project using the Rust SDK to encrypt a sensitive field with `encrypt_for_search`, enabling ciphertext lookups across a dataset.

Cargo.toml

```toml
# Cargo.toml
[dependencies]
ubiq-rust = "1.0.0"
```

Enable Grafana Loki as a SIEM destination for Ubiq activity events so SOC dashboards receive tokenization audit trails alongside other log streams.

📍 1. Contact Ubiq support to enable Grafana Loki for your enterprise account. 2. In the Ubiq dashboard, go to Settings › Integrations. 3. Add a new Grafana Loki connection and supply your Loki endpoint and credentials. 4. Save and verify events appear in your Loki data source.

Adds `CipherAndKeyNumber`, `GetCurrentKeyNumber`, and `GetKeyNumber` functions to the Go library to retrieve the key number behind a ciphertext without decrypting it.

Adds new Ubiq Rust SDK (Beta), available on crates.io as `ubiq-rust`, supporting structured and unstructured encryption, integer and date/time datasets, `encrypt_for_search`, and key zeroization.

> Adds a new Ubiq plugin for the Kong API gateway providing field-level tokenization and detokenization of JSON request and response bodies with no application code changes.
>
> Adds Grafana Loki as a SIEM destination for Ubiq activity and usage events (enterprise customers; previously Splunk only).
>
> Expands dataset type support in the C 2.3.0.0 library to include `int32`, `int64`, `date`, `datetime`, `token` (64/128-byte), and generic strings with Base32, Base64, or no encoding, including support for the dataset creation wizard.
>
> +11 more

WAS THIS USEFUL? 👍 👎

SandboxAQ AQtive Guard

Feedback ⓘ

changelog · Sep 25, 2026 CHANGELOG ↗ RANK

SandboxAQ AQtive Guard is a quantum-safe cryptography platform protecting applications against current and future quantum computing threats.

> SandboxAQ AQtive Guard added RHEL 10 as a supported deployment target for its on-premises packages.

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

01 **RHEL 10 support for on-premises packages** NEW 35

AQtive Guard on-premises packages can now be deployed on RHEL 10.

— Names OS target but no install steps or version specifics changelog-20260925-256719a4

> changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE
>
> AQtive Guard on-premises packages now available for RHEL 10.

WAS THIS USEFUL? 👍 👎

Varonis

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

The leading data and AI security platform that proactively reduces risk, enforces policies, and stops active threats with our AI-native automation.

> Varonis expanded infrastructure-as-code deployment options for its Private Cloud Data Collector, added private networking and connectivity validation for Snowflake, and broadened SaaS coverage with GitHub Enterprise Cloud monitoring, deeper Confluence and Jira access analysis, and more flexible on-demand scan scheduling.

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

01 **Private networking for Snowflake integration** NEW 70

Snowflake integrations can be configured to use AWS PrivateLink or Azure Private Link directly from the Varonis integration wizard, with self-service endpoint URL generation and connectivity validation.

– Describes wizard flow and validation but not exact fields changelog-20260925-213595ed

02 **Private Cloud Data Collector deployment and secrets management**

NEW 65

The Varonis Private Cloud Data Collector can now be deployed to AWS ECS using Terraform, enabling infrastructure-as-code management of collector environments. Collectors can also be configured to use AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault for customer-managed secret storage.

– Names IaC tooling and secret stores but no exact commands changelog-20260925-213595ed

THINNER COVERAGE BELOW

03 **Prioritized on-demand scan scheduling** IMPROVED 55

When scheduling on-demand scans, users can now choose between a Prioritized Scan for faster results or a standard queued rescan.

– Names the two scan options but not the exact UI path changelog-20260925-213595ed

04 **GitHub Enterprise Cloud monitoring** NEW 50

Varonis can monitor GitHub Enterprise Cloud via a GitHub App with least-privilege permissions and enterprise-wide coverage.

– States mechanism (GitHub App) but no setup steps changelog-20260925-213595ed

05 **Confluence historical version sensitive-content analysis** NEW

50

Varonis analyzes historical versions of Confluence pages for sensitive content, automatically tagging pages that contain sensitive data in prior versions.

— Explains tagging mechanism but no config or path given changelog-20260925-213595ed

06 **Jira access tracking for overexposed resources** NEW

45

Varonis tracks access to Jira issues, comments, attachments, and worklogs to identify overexposed or broadly accessible resources.

— Describes capability but no navigation or detection detail changelog-20260925-213595ed

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Varonis: Deploy the Varonis Private Cloud Data Collector to AWS ECS using Terraform, enabling infrastructure-as-code management of collector environments.

Configure Private Cloud Data Collectors to use AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault for customer-managed secret storage.

Configure Snowflake integrations that use AWS PrivateLink or Azure Private Link directly from the Varonis integration wizard, with self-service endpoint URL generation and connectivity validation.

Monitor GitHub Enterprise Cloud via a GitHub App with least-privilege permissions and enterprise-wide coverage.

Analyze historical versions of Confluence pages for sensitive content, with automatic tagging of pages that contain sensitive data in prior versions.

Track access to Jira issues, comments, attachments, and worklogs to identify overexposed or broadly accessible resources.

+1 more

WAS THIS USEFUL? 👍 👎

bytebase

Feedback ⓘ 3.23.0 NOTES ↗ RANK

Database governance built for humans and agents — controlling changes and access across every major database.

> Bytebase 3.23.0 adds MCP access-policy controls with refusal auditing, OIDC-based workload identity support for the Terraform provider, inline plan-review comments, and a multi-engine query-plan visualizer in SQL Editor, alongside a large set of breaking API and permission changes.

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

01 OIDC workload identity auth for Terraform provider

BREAKING

95

The Terraform provider can now authenticate as a workload identity instead of a long-lived service account key, via new provider fields `workload_identity_email`, `workload_identity_token`, and `workload_identity_token_file`. The `bytebase_workload_identity` resource gains `provider_type = "OIDC"` and an optional `jwks_url` so any OIDC issuer (not just GitHub/GitLab) can be used, but now requires `issuer_url`, `allowed_audiences`, and `subject_pattern` on every resource. Existing GitHub/GitLab workload identities whose `subject_pattern` matches every repo/project (e.g. `*`, `repo:*`, `project_path:*`) stop authenticating until narrowed, and pipelines requesting an audience not in `allowed_audiences` stop authenticating.

Authenticate the Terraform provider as a workload identity (e.g. from GitHub Actions) instead of a long-lived service account key.

hcl

```
provider "bytebase" {
  url                          = "https://bytebase.example.com"
  workload_identity_email      = "ci-
deployer@bytebase.example.com"
  workload_identity_token_file =
"/var/run/secrets/bytebase/token"
}
```

Register any OIDC provider (e.g. HashiCorp Vault, Azure AD) as a Bytebase workload identity when the issuer does not publish a standard discovery document.

```hcl
resource "bytebase_workload_identity" "vault" {
  name              = "vault-ci"
  issuer_url        =
"https://vault.example.com/v1/identity/oidc"
  jwks_url          =
"https://vault.example.com/v1/identity/oidc/.well-known/keys"
  allowed_audiences = ["bytebase"]
  subject_pattern   = "my-org/my-repo"
  provider_type     = "OIDC"
}
```

— Names exact fields, resource, and breaking migration requirements. 3.23.0

02 **MCP settings endpoint and client token refresh** BREAKING

85

`GET /v1/mcp/info` (`WorkspaceService.GetMCPInfo`) is removed; MCP settings now come from `GET /v1/actuator/info` (`ActuatorService.GetActuatorInfo`), which requires authentication and drops the `modes`, `methods`, `engines`, and `data_masking_available` fields with no replacement. MCP clients connected before 3.22.0 cannot refresh their tokens and fail with `invalid_grant` once their current access token expires; the server must be removed and re-added in each client.

Query the new MCP settings endpoint (requires authentication) after removing the now-deleted `GET /v1/mcp/info` call.

```
$ curl -H 'Authorization: Bearer <token>' \
  https://bytebase.example.com/v1/actuator/info \
  | jq '.mcpSetting'
```

— Exact endpoints, removed fields, and remediation step named. 3.23.0

03 **MCP access policy control and refusal auditing** NEW

83

Adds MCP access policy control under Integration > MCP > Access policy with three levels — Disabled, Read-only, and Read-write — applied immediately to already-connected clients. New workspaces now default to Read-only (workspaces that never set a policy stay Read-write). Requests refused by permission checks or blocked by the MCP access policy are now recorded in the audit log as `WARNING`, and enabling Enable audit logging to stdout (self-hosted) streams every refusal to stdout.

— Exact UI path and policy levels given, mechanism partly described. 3.23.0

04 **Audit log `user` field renamed to `actor`** BREAKING

78

The audit-log `user` field is renamed `actor` in the API and exports, with values now typed as `users/{email}`, `serviceAccounts/{email}`, or `workloadIdentities/{email}`; filters like `user == 'users/bb@bytebase.com'` must be updated to `actor == 'users/bb@bytebase.com'`.

— Exact field name, value formats, and filter migration given. 3.23.0

05 **Query-plan visualizer in SQL Editor** NEW

78

SQL Editor's result pane gains a query-plan visualizer with text, diagram, grid, and summary views for PostgreSQL, SQL Server, and Spanner, including `EXPLAIN FORMAT JSON`, `XML`, `YAML`, and `ANALYZE` on PostgreSQL. The query API's `mssql_explain_format` field is renamed `explain_format` to support the visualizer across engines, breaking integrations that used the old field name.

— Names engines, view formats, and renamed field, but no walkthrough. 3.23.0

06 **DML permission required for data-modifying PostgreSQL queries**

70

BREAKING

PostgreSQL queries in SQL Editor that modify data — including `SELECT` with a data-modifying `WITH` clause and `EXPLAIN ANALYZE` of a write — now require DML permission; `bb.sql.select` alone is no longer sufficient.

— Names permission and exact query patterns affected. 3.23.0

07 **Access control for Oracle database link queries** BREAKING

70

Oracle database link queries (`schema.table@link`) in SQL Editor, export, or MCP now require the same grant as querying the linked database directly; links to unmanaged or external instances are refused before the query runs.

— Names query syntax, surfaces affected, and refusal behavior. 3.23.0

08 **API pagination, atomicity, and validation changes** BREAKING

70

`ListTaskRuns` is now paginated, so callers that assumed a single full response must page through results. The four `BatchGet` API methods are now all-or-nothing, returning an error instead of partial results when part of a request fails. Update API methods now reject requests whose body disagrees with `update_mask`. The login-page restriction endpoint is renamed, and listing identity providers now requires authentication.

— Names several endpoints and their new required behavior. 3.23.0

09 **Pod annotations via Helm value** NEW 65

Adds `bytebase.podAnnotations` Helm value to set pod annotations on the Bytebase pod (self-hosted), for example for Prometheus scraping.

— Exact Helm key and example use case given. 3.23.0

THINNER COVERAGE BELOW

10 **Inline SQL line comments on plans** NEW 50

The Plan detail editor now supports inline comments on specific SQL lines, with thread replies, Resolve/Reopen actions, an open-thread counter, and thread navigation.

— Describes what and where, no config or API named. 3.23.0

11 **Block self-approval of last-edited plans** NEW 50

New project setting Allow approval by the person who last changed the plan lets teams prevent the person who last edited a plan from approving their own change.

— Names exact setting; no further mechanism given. 3.23.0

12 **Chat history for SQL Editor AI assistant** NEW 50

SQL Editor's AI assistant now keeps chat history — the last 20 conversations per workspace and user — with inline-editable conversation titles.

— Gives a concrete limit but no config or command. 3.23.0

13 **Skip metadata lock check for gh-ost migrations** NEW 40

Adds a Skip metadata lock check opt-in option to MySQL gh-ost migrations.

— Names the option but not its effect or mechanism. 3.23.0

14 **Approval rules and risk levels for OceanBase** NEW 35

Enables `statement.sql_type` approval rules and risk levels for OceanBase changes.

— Names the field but gives no further detail. 3.23.0

15 **Service accounts blocked from web console sign-in** BREAKING 25

Service accounts can no longer sign in to the Bytebase web console.

— States the change with no mechanism or migration path. 3.23.0

└──▷ BREAKING ON UPGRADE

- ! `GET /v1/mcp/info` (`WorkspaceService.GetMCPInfo`) is removed; read `mcpSetting` from `GET /v1/actuator/info` (`ActuatorService.GetActuatorInfo`) instead. The `modes`, `methods`, `engines`, and `data_masking_available` fields have no replacement.
- ! The audit-log `user` field is renamed `actor` in the API and exports; update filters from `user == 'users/bb@bytebase.com'` to `actor == 'users/bb@bytebase.com'`.
- ! `mssql_explain_format` is renamed `explain_format` in the query API; integrations using the old field name will break.
- ! `bytebase_workload_identity` in the Terraform provider now requires `issuer_url`, `allowed_audiences`, and `subject_pattern`; existing resources missing these fields will fail to apply.
- ! Existing GitHub and GitLab workload identities whose subject pattern matches every repository or project (e.g. `*`, `repo:*`, `project_path:*`) stop authenticating until narrowed to at least the owner or group.
- ! Workload identity pipelines requesting an audience not on the identity's allowed list stop authenticating; the audience must be added to the list or the pipeline must request an allowed audience.
- ! New workspaces default to MCP access policy `Read-only`; workspaces that have never set a policy remain `Read-write` until an admin explicitly sets one.
- ! MCP clients connected before 3.22.0 cannot refresh their tokens and will fail with `invalid_grant` once the current access token expires; the server must be removed and re-added in each client.
- ! PostgreSQL queries in SQL Editor that use a data-modifying `WITH` clause or `EXPLAIN ANALYZE` on a write now require DML permission; `bb.sql.select` alone is no longer sufficient.
- ! Oracle database link queries (`schema.table@link`) now require the same grant as querying the linked database directly; links to unmanaged or external instances are refused.
- ! Service accounts can no longer sign in to the Bytebase web console.
- ! `ListTaskRuns` is now paginated; callers that assumed a single full response must be updated to page through results.
- ! The four `BatchGet` API methods are now all-or-nothing; a request that partially fails now returns an error instead of partial results.
- ! Update API methods now reject requests whose body disagrees with `update_mask`.
- ! The login-page restriction endpoint is renamed; callers using the old name will break.
- ! Listing identity providers now requires authentication; unauthenticated calls will be rejected.

3.23.0 NOTES ↗ LATEST STABLE

Bytebase 3.23.0 adds MCP access policies, inline plan-review comments, any-OIDC workload identities, and a query-plan visualizer in SQL Editor.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 3.23.0 https://github.com/bytebase/bytebase.git
# already have the repo? check out this version:
$ git checkout 3.23.0
```

└──▷ **USE IT**

Authenticate the Terraform provider as a workload identity (e.g. from GitHub Actions) instead of a long-lived service account key.

```hcl
provider "bytebase" {
  url                          = "https://bytebase.example.com"
  workload_identity_email      = "ci-deployer@bytebase.example.com"
  workload_identity_token_file = "/var/run/secrets/bytebase/token"
}
```

Register any OIDC provider (e.g. HashiCorp Vault, Azure AD) as a Bytebase workload identity when the issuer does not publish a standard discovery document.

```hcl
resource "bytebase_workload_identity" "vault" {
  name              = "vault-ci"
  issuer_url        = "https://vault.example.com/v1/identity/oidc"
  jwks_url          =
"https://vault.example.com/v1/identity/oidc/.well-known/keys"
  allowed_audiences = ["bytebase"]
  subject_pattern   = "my-org/my-repo"
  provider_type     = "OIDC"
}
```

Query the new MCP settings endpoint (requires authentication) after removing the now-deleted `GET /v1/mcp/info` call.

```
$ curl -H 'Authorization: Bearer <token>' \
  https://bytebase.example.com/v1/actuator/info \
  | jq '.mcpSetting'
```

Adds `workload_identity_email`, `workload_identity_token`, and `workload_identity_token_file` provider fields so the Terraform provider can authenticate as a workload identity instead of a service account key.

Adds `provider_type = "OIDC"` and optional `jwks_url` to `bytebase_workload_identity` in the Terraform provider, enabling any OIDC issuer

as a workload identity.

Adds `GET /v1/actuator/info` (`ActuatorService.GetActuatorInfo`) as the new source for MCP settings, replacing the removed `GET /v1/mcp/info`.

Renames the audit-log `user` field to `actor` in the API and exports, with values now typed as `users/{email}`, `serviceAccounts/{email}`, or `workloadIdentities/{email}`.

Replaces `mssql_explain_format` with `explain_format` in the query API for the inline query-plan visualizer.

+12 more

└─▷ **BREAKING ON UPGRADE**

- ! `GET /v1/mcp/info` (`WorkspaceService.GetMCPInfo`) is removed; read `mcpSetting` from `GET /v1/actuator/info` (`ActuatorService.GetActuatorInfo`) instead. The `modes`, `methods`, `engines`, and `data_masking_available` fields have no replacement.
- ! The audit-log `user` field is renamed `actor` in the API and exports; update filters from `user == 'users/bb@bytebase.com'` to `actor == 'users/bb@bytebase.com'`.
- ! `mssql_explain_format` is renamed `explain_format` in the query API; integrations using the old field name will break.
- ! `bytebase_workload_identity` in the Terraform provider now requires `issuer_url`, `allowed_audiences`, and `subject_pattern`; existing resources missing these fields will fail to apply.
- ! Existing GitHub and GitLab workload identities whose subject pattern matches every repository or project (e.g. `*`, `repo:*`, `project_path:*`) stop authenticating until narrowed to at least the owner or group.
- ! Workload identity pipelines requesting an audience not on the identity's allowed list stop authenticating; the audience must be added to the list or the pipeline must request an allowed audience.
- ! New workspaces default to MCP access policy `Read-only`; workspaces that have never set a policy remain `Read-write` until an admin explicitly sets one.
- ! MCP clients connected before 3.22.0 cannot refresh their tokens and will fail with `invalid_grant` once the current access token expires; the server must be removed and re-added in each client.
- ! PostgreSQL queries in SQL Editor that use a data-modifying `WITH` clause or `EXPLAIN ANALYZE` on a write now require DML permission; `bb.sql.select` alone is no longer sufficient.
- ! Oracle database link queries (`schema.table@link`) now require the same grant as querying the linked database directly; links to unmanaged or external instances are refused.

- ! Service accounts can no longer sign in to the Bytebase web console.
- ! `ListTaskRuns` is now paginated; callers that assumed a single full response must be updated to page through results.
- ! The four `BatchGet` API methods are now all-or-nothing; a request that partially fails now returns an error instead of partial results.
- ! Update API methods now reject requests whose body disagrees with `update_mask`.
- ! The login-page restriction endpoint is renamed; callers using the old name will break.
- ! Listing identity providers now requires authentication; unauthenticated calls will be rejected.

WAS THIS USEFUL? 👍 👎

Microsoft Purview

Feedback ⓘ changelog · Sep 24, 2026

BLOG ↗ RANK

Microsoft Purview is a unified data governance platform that discovers, classifies, and manages data across organizations.

> Microsoft Purview's September 2026 update centers on GA network-layer DLP to block sensitive data from reaching unsanctioned AI tools, alongside eDiscovery support for SharePoint embedded containers, granular SharePoint archiving and permanent Priority Cleanup in Data Lifecycle Management, and scale/audit improvements to auto-labeling policies.

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

01 **Network-layer DLP for shadow AI via Global Secure Access** NEW

65

Network-layer data loss prevention via Microsoft Purview and Microsoft Entra Global Secure Access is now generally available, enforcing context-aware classification policies at the network layer to block sensitive files and text from reaching unsanctioned AI tools. Coverage extends to both human traffic and on-behalf-of (OBO) agentic traffic.

— Mechanism and scope named; no setup steps given 2026-09-24-b5f8a874

02 **eDiscovery support for SharePoint embedded containers** NEW

65

Microsoft Purview eDiscovery now supports search, hold, review, and export of content in user-owned SharePoint embedded containers, covering apps like Microsoft Loop, Copilot Pages, Copilot Notebooks, and Outlook newsletters mapped to a user, without requiring the container URL from a SharePoint administrator.

— Named apps and workflow but no exact steps 2026-09-24-b5f8a874

03 **Selective SharePoint content archiving in Data Lifecycle Management**

NEW

65

Administrators in Microsoft Purview Data Lifecycle Management can now archive inactive SharePoint content at the individual-content level without archiving the entire site. Archived content stays subject to retention, legal hold, and eDiscovery, while being removed from Microsoft 365 Copilot indexing until reactivated.

Archive inactive SharePoint content site-selectively to drop it from Copilot indexing while keeping it under retention and legal hold.

📍 In the Microsoft Purview compliance portal, go to Data Lifecycle Management › SharePoint archiving, select the target content (not the full site), apply the archive action, and confirm the content remains discoverable in eDiscovery but excluded from Microsoft 365 Copilot indexing until reactivated.

— Behavior and scope clear; UI path given as example 2026-09-24-b5f8a874

04 **Priority Cleanup for permanent content deletion** NEW

65

Priority Cleanup in Microsoft Purview Data Lifecycle Management permanently deletes approved content — including stale Teams recordings and transcripts — removing it from eDiscovery, SharePoint search, and Microsoft 365 Copilot indexing.

Remove stale Teams recordings from Copilot indexing and eDiscovery to reduce AI data sprawl in regulated environments.

📍 In the Microsoft Purview compliance portal, go to Data Lifecycle Management › Priority Cleanup, create a new cleanup policy targeting Teams recordings and transcripts, obtain approval, and execute to permanently delete the content from SharePoint search, eDiscovery, and Microsoft 365 Copilot.

— Clear scope and effect; UI steps provided as example 2026-09-24-b5f8a874

05 **Auto-labeling policy scale and audit improvements** IMPROVED

65

Purview auto-labeling policy simulations now scale to support up to 20 million items and up to 50,000 sites through adaptive scopes. Administrators can also edit an auto-labeling policy without re-running simulation, and new audit insights and reporting show policy coverage and processing activity.

— Numeric limits named but no command or navigation given 2026-09-24-b5f8a874

THINNER COVERAGE BELOW

06 **HTML conversion option in eDiscovery export** NEW

30

eDiscovery export gains an optional HTML conversion to produce a more human-readable version of content for downstream legal tools.

— One-line addition with no further detail 2026-09-24-b5f8a874

changelog · Sep 24, 2026 BLOG ↗

Purview gains network-layer DLP for shadow AI, eDiscovery for Copilot content, selective SharePoint archiving, and Priority Cleanup.

└──▷ HOW TO FIND IT

Remove stale Teams recordings from Copilot indexing and eDiscovery to reduce AI data sprawl in regulated environments.

> 📍 In the Microsoft Purview compliance portal, go to Data Lifecycle Management › Priority Cleanup, create a new cleanup policy targeting Teams recordings and transcripts, obtain approval, and execute to permanently delete the content from SharePoint search, eDiscovery, and Microsoft 365 Copilot.

Archive inactive SharePoint content site-selectively to drop it from Copilot indexing while keeping it under retention and legal hold.

> 📍 In the Microsoft Purview compliance portal, go to Data Lifecycle Management › SharePoint archiving, select the target content (not the full site), apply the archive action, and confirm the content remains discoverable in eDiscovery but excluded from Microsoft 365 Copilot indexing until reactivated.

Adds network-layer data loss prevention via Microsoft Purview and Microsoft Entra Global Secure Access (now generally available), enforcing context-aware classification policies at the network layer to block sensitive files and text from reaching unsanctioned AI tools, covering both human and on-behalf-of (OBO) agentic traffic.

Extends Microsoft Purview eDiscovery to support search, hold, review, and export of content in user-owned SharePoint embedded containers, enabling investigators to find content from Microsoft Loop, Copilot Pages, Copilot Notebooks, Outlook newsletters, and other apps mapped to a user without needing the container URL from a SharePoint administrator.

Adds an optional HTML conversion in eDiscovery export to produce a more human-readable version of content for downstream legal tools.

Enables administrators in Microsoft Purview Data Lifecycle Management to archive inactive SharePoint content at the individual-content level without archiving the entire site, keeping archived content subject to retention, legal hold, and eDiscovery while removing it from Microsoft 365 Copilot indexing until reactivated.

Adds Priority Cleanup in Microsoft Purview Data Lifecycle Management to permanently delete approved content — including stale Teams recordings and transcripts — removing it from eDiscovery, SharePoint search, and Microsoft 365 Copilot indexing.

+3 more

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ changelog · Sep 25, 2026

CHANGELOG ↗ RANK

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating certificate lifecycle, and ensuring compliance across organizations.

> DigiCert CertCentral shipped automated domain revalidation via a persistent DNS TXT record, moved TLS certificate issuance to new G5 root hierarchies by default, and tightened domain validation to require Multi-Perspective Issuance Corroboration (MPIC) from six remote network perspectives — alongside UI updates for certificate validity period selection and common name/SAN replacement.

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

01 **Certificate validity period selection with custom options**

NEW 30

A validity period selection screen offers five options when ordering a certificate: 1 year, 2 years, 3 years, a custom expiration date, or a custom length.

— Only a screenshot caption describes the options, no steps changelog-20260925-5b329386

02 **Replace common name and SANs on a certificate** IMPROVED

15

CertCentral includes a UI flow for replacing the common name and SANs on a certificate, though the specific mechanism is not described beyond the interface screenshot.

— Only a screenshot filename names this capability changelog-20260925-5b329386

└──▷ **ALSO FROM THESE RELEASES**

└──▷ **BREAKING ON UPGRADE**

! TLS certificate issuance now defaults to the G5 root hierarchies (`DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5`); customers managing trust stores, hard-coding root or intermediate CA certificate trust, or pinning root or intermediate CA certificates must update their configurations or face service disruptions.

! MPIC now requires six remote network perspectives and four corroborations per domain query; customers restricting DNS queries or website traffic by IP address, region, query volume, or rate must verify their configurations support MPIC traffic from six remote perspectives.

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

└──▷ **HOW TO FIND IT**

Set up the Persistent DNS TXT record so CertCentral can automatically revalidate your domain every five days without manual intervention.

📍 In CertCentral, go to Domains, select your domain, choose `Persistent DNS TXT Record` as the DCV method, add the displayed TXT record to your DNS zone, and submit for initial validation. Leave the record in place permanently — CertCentral will revalidate the domain on its five-day cycle automatically.

└──▷ **BREAKING ON UPGRADE**

! TLS certificate issuance now defaults to the G5 root hierarchies (`DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5`); customers managing trust stores, hard-coding root or intermediate CA certificate trust, or pinning root or intermediate CA certificates must update their configurations or face service disruptions.

! MPIC now requires six remote network perspectives and four corroborations per domain query; customers restricting DNS queries or website traffic by IP address, region, query volume, or rate must verify their configurations support MPIC traffic from six remote perspectives.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ 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 changes this window tighten agent governance and infrastructure: AI Watch Enforce now blocks unauthorized MCP launchers, Shadow AI data access is restricted to reviewers, and the Terraform ECS module gains per-role Aurora alarms, GPU autoscaling, an authenticated sandbox egress proxy, and trusted-CA injection. Alongside these, agents gained a Figma MCP connector, a Google AI Studio provider, inline tool-call approvals, and finer-grained Slack messaging and monitoring controls.

└──▷ **WHAT SHIPPED · 28 FEATURES** most completely described first ? what's the number?

01 Aurora CPU/ACU alarm thresholds in Terraform ECS module NEW

95

Adds `rds_writer_cpu_utilization_threshold` to fire a writer-cpu CloudWatch alarm when Aurora WRITER `CPUUtilization` exceeds the threshold for 3 consecutive 5-minute periods, evaluated on the `Role = WRITER` dimension rather than the cluster average; `rds_reader_cpu_utilization_threshold` to fire a per-instance reader-cpu alarm gated on `AuroraReplicaLag` so it tracks role swaps on failover and avoids masking a saturated reader behind an idle one; and `rds_writer_acu_utilization_threshold` to fire a writer-acu alarm when Aurora Serverless v2 `ACUUtilization` hits or exceeds the threshold for 6 consecutive 5-minute periods, only created when `database_config.min_capacity < max_capacity`.

Set per-role Aurora CPU thresholds so writer and reader alarms fire independently — useful when a pinned writer saturates while an idle replica would mask a cluster-average alarm.

```hcl
rds_writer_cpu_utilization_threshold = 75
rds_reader_cpu_utilization_threshold = 85
rds_writer_acu_utilization_threshold = 90
```

— Names every alarm input with exact thresholds and gating conditions. product docs

02 GPU utilization autoscaling for ToolGuard fleet BREAKING

90

Adds a second target-tracking autoscaling policy on the `RunlayerToolGuard/GPUUtilization` CloudWatch metric (target 60%, `Environment` dimension, `Average` statistic), complementing the existing CPU policy so the fleet scales on whichever signal is higher; metric publishing frequency was accelerated from every 5 minutes to every 60 seconds for finer-grained decisions, alongside a 900-second scale-in cooldown. Predictive scaling is automatically disabled in unsupported regions with a plan warning, and `terraform plan` now warns if `container_insights = "enabled"` is not paired with the GPU policy. Upgrading changes the ToolGuard launch template on every GPU profile and triggers an ASG rolling instance refresh, replacing hosts and restarting tasks (each restart reloads the model), so expect a capacity dip per replacement batch.

— Names exact metric, dimensions, statistic, and rollout impact. product docs

---

03 **`aiwatch config show` CLI command** NEW 90

Adds `aiwatch config show` command, displaying managed org API keys as a prefix (e.g. `rl_org_Ab1C2d3...`) and enrollment keys as prefix plus last four characters, to confirm which key a device is using without exposing the full secret.

Verify which org API key or enrollment key AI Watch is using on a managed device without exposing the full secret.

```
$ aiwatch config show
```

Confirm which org API key or enrollment key AI Watch is using on a managed device before troubleshooting enrollment issues.

```
$ aiwatch config show
```

— Exact runnable command with output format described. changelog-20260925-4be31508 1.28.780 (2026-09-24)

---

04 **AI Watch Enforce MCP launcher validation** BREAKING 85

AI Watch Enforce now treats a local `runlayer run` MCP launcher as Runlayer-managed only when its `--host` points at your Runlayer deployment and names an active connector; launchers without `--host`, aimed at another host, or naming an unknown/disabled connector (or carrying unsupported options) are blocked, and can no longer be allowlisted from the Shadow AI activity view.

— Names exact command and flag with clear before/after behavior.

changelog-20260925-4be31508 1.28.780 (2026-09-24)

---

05 **ToolGuard health-check grace period input** NEW 85

Adds `runlayer_tool_guard_health_check_grace_period_seconds` input (accepting 0–2147483647) to override the ECS health-check grace period for the ToolGuard service, defaulting to 900 seconds to cover GRID driver install and model load.

Pin the ToolGuard ECS health-check grace period to int32 max as a stopgap when ECS loops on task replacement due to `/health` returning 503 during worker-recycle drain.

```hcl
runlayer_tool_guard_health_check_grace_period_seconds =
2147483647
```

— Exact input name, range, and default explained. product docs

06 **Authenticated sandbox egress proxy** NEW 85

Adds `sandbox_egress_mode` Terraform variable (values `proxy` and `proxy_only`) and `sandbox_egress_proxy_egress_ports` / `sandbox_egress_proxy_image_uri` variables to route ECS sandbox internet traffic through an authenticated egress proxy enforcing per-agent network access policies.

Roll out the authenticated egress proxy in stages to enforce per-agent network access policies on sandbox internet traffic.

```hcl
sandbox_egress_mode                = "proxy_only"
sandbox_egress_proxy_egress_ports  = [443, 8443]
sandbox_egress_proxy_image_uri     = ".../runlayer/sandbox-
egress-proxy:1.28.500"
```

— Names exact variables and values, with a staged-rollout example. product docs

07 **Trusted CA bundle injection via `trusted_ca_pem`** NEW 80

Adds `trusted_ca_pem` input accepting a PEM bundle of additional trusted CA certificates, merged with the image public roots for backend, shared worker, and MCP execution worker HTTPS clients; private keys are rejected.

Inject a private CA bundle so backend and worker HTTPS clients trust an internal PKI without rebuilding the container image.

```hcl
trusted_ca_pem = file("${path.module}/certs/internal-ca-
bundle.pem")
```

Trust a corporate private CA for MCP server TLS connections in an ECS self-hosted deployment.

```hcl
trusted_ca_pem = file("${path.module}/company-ca.pem")
```

— Exact config key with usage examples for private CA trust. `product docs`

08 **OpenAI models on Amazon Bedrock support** NEW 75

Adds agents-on-Amazon-Bedrock support for OpenAI models, including `gpt-6` Astra and `global.` / `us.` / `eu.` `-prefixed` OpenAI model IDs; a bare `openai.gpt-6…` ID is automatically served under its `global.` inference profile.

— Names exact model ID prefixes and fallback behavior. `changelog-20260925-4be31508`

09 **Per-agent Slack sending controls** NEW 75

Adds Direct messages to people and Channel posts switches (both default off) under Sending in an agent's Slack settings to control whether the agent can initiate DMs or post to channels.

Lock down an agent so it only responds to channel mentions and thread follow-ups, preventing users from bypassing audit trails via direct messages.

📍 In the agent's Slack settings, go to Inbound messaging controls and turn off Direct messages to the agent and Replies. Under Sending, confirm Direct messages to people and Channel posts are set to off.

— Exact toggle names, defaults, and location given. `product docs`

10 **Artifacts sharing settings and inline display** IMPROVED 70

Artifacts an agent creates with the Runlayer MCP `create_artifact` tool now follow that agent's Artifact sharing setting. Artifacts (interactive apps, images, videos, documents) are surfaced inline in agent chat directly under the agent's answer and can be opened full-page from the caption, and interactive HTML Artifact previews are now supported inside Slack using existing Runlayer permissions (requires embeds enabled for the workspace's Runlayer app).

— Names the tool and Slack embed requirement but no config key.

`changelog-20260925-4be31508` `1.28.780 (2026-09-24)`

11 **JumpCloud MDM deployment for AI Watch** NEW 70

Deploys AI Watch to macOS (`.pkg`) and Windows (`.msi`) through JumpCloud Private Repo, launched from the MDM configuration wizard with pre-filled configuration profiles

and Private Repo Command Line Options for the signed MSI; each package bundles Detect, endpoint modes, and Sessions in one install per OS.

— Names package formats and deployment path but no exact command. 1.28.780 (2026-09-24)

12 **Sessions tab run visibility improvements** IMPROVED 70

Adds a Slack badge on an agent's Sessions tab linking directly to the Slack message or thread that started a given run; run details for Slack channel monitor-triggered runs now show the channel, the message reacted to, and why it matched the monitor's condition; and the agent run panel on the Sessions tab now leads with live status, elapsed time, cost, and failure reason alongside the Retry action.

— Names UI elements added but no exact navigation path beyond Sessions tab.

1.28.780 (2026-09-24)

13 **`enable_channels=[]` parameter for MCP agent creation** NEW 70

Supports `enable_channels=[]` parameter when creating an agent through MCP to skip Slack setup.

Skip Slack setup when programmatically creating an agent via MCP so it is immediately usable without requiring Slack connector access.

```
$ enable_channels=[]
```

— Exact parameter name and effect, directly usable. 1.28.780 (2026-09-24)

14 **Minimum AWS provider version requirement** BREAKING 65

The Terraform ECS module now requires AWS provider `>= 6.64.0, < 7.0.0` for alarm warm-up used by ToolGuard registration monitoring; the dependency lock file must be updated before running `terraform plan`.

— Exact version bounds and required upgrade step given. product docs

15 **Google AI Studio provider for Gemini models** NEW 65

Adds a dedicated Google AI Studio provider so agents can run on Google Gemini models — configure via Settings → AI Providers with your Gemini API key and model IDs.

Add the Google AI Studio provider so agents can be backed by Gemini models.

> 📍 In Runlayer, go to Settings → AI Providers, select Google AI Studio, paste your Gemini API key, and enter the model IDs you want available to agents.

— Exact UI navigation path given, though no API details. 1.28.780 (2026-09-24)

16 **MS Teams `MS_TEAMS_APP_URL` immutability** BREAKING 60

`MS_TEAMS_APP_URL` (the domain name) is effectively immutable after agents are provisioned — changing it orphans existing Azure Bot resources until each is deleted and recreated.

Enable Microsoft Teams integration by supplying all five required environment values to the ECS module.

hcl

```
ms_teams_app_client_secret     = var.ms_teams_app_client_secret
# remaining values passed via TF_VAR_* or 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://<domain_name>"
```

— Names the exact variable and describes the consequence of changing it. product docs

17 **Inline tool-call approval decisions in agent chats** NEW 60

Enables inline approve/deny of agent tool-call approval requests directly in Agent Playground, Runlayer Assistant, and Home Build, with the agent continuing in the same conversation after the decision.

— Names surfaces but no exact UI control or API. 1.28.780 (2026-09-24)

18 **Slack channel monitor reply-only-when-useful behavior** 60

BREAKING

Slack channel monitors now reply only when the agent has something useful to add, with silent runs still appearing in activity; the previous reply-behavior setting was removed with no option to restore the prior post-every-run behavior.

— Clear before/after behavior but no opt-out config. 1.28.780 (2026-09-24)

THINNER COVERAGE BELOW

19 **AI Watch client config backups** NEW 55

Adds AI Watch client config backups for all supported clients (Cursor, VS Code, Codex, Gemini CLI, and others), retaining only the 5 most recent backups per file.

— Names retention limit and clients but no restore command. changelog-20260925-4be31508

20 **Recommended scanner settings for 39 more connectors** NEW 55

Adds recommended per-tool security scanner settings for 39 more connectors, reviewable and applicable via a dialog on the connector page after adding it from the catalog.

— Gives exact connector count and UI location for the dialog. 1.28.780 (2026-09-24)

21 **Dedicated security group for tunnel VPC endpoints** NEW 50

Adds a dedicated security group input for ECR API, ECR DKR, CloudWatch Logs, and Secrets Manager interface endpoint ENIs used by tunnel tasks when connecting through an existing VPC.

— Endpoints named but exact input variable name not given. product docs

22 **Figma remote MCP server integration** NEW 45

Adds Figma's remote MCP server to the catalog, now generally available, with OAuth app registration and per-user Figma account connection.

— Names the integration but no exact setup steps given. changelog-20260925-4be31508

23 **Top spenders drilldown for model usage** NEW 40

Adds a Top spenders drilldown linking to a searchable, paginated view of every user's model spend or token usage.

— Describes the view but no exact navigation path given. changelog-20260925-4be31508

24 **Delete skills still in use by agents** NEW 40

Allows deleting a skill even when agents use it, with a confirmation dialog listing the affected agents.

— Describes the dialog but no exact UI path or API. 1.28.780 (2026-09-24)

25 **Persistent named Runlayer Assistant sessions** NEW 40

Adds persistent, named Runlayer Assistant sessions resumable from History as you navigate Runlayer.

— Names the History feature but no exact path or limits. 1.28.780 (2026-09-24)

26 **On-demand connector tool discovery for Pi agents** IMPROVED 35

Pi agents discover connector tools on demand, keeping unused tool definitions out of the model's context.

— Explains the mechanism briefly but no configuration surface. 1.28.780 (2026-09-24)

27 **Accountable owner on agent accounts** NEW 35

Every agent account now has an accountable owner shown in the account summary, with the creator set as owner by default.

— States the default but no way to change owner is described. 1.28.780 (2026-09-24)

28 **Browser extension for AI web chat monitoring** NEW 25

Adds a browser extension to monitor and protect AI web chats in Chrome, Edge, and Firefox.

— Only a one-line summary with no mechanism or setup detail. product docs

└──▷ BREAKING ON UPGRADE

! AI Watch Enforce now blocks MCP launchers whose `--host` does not point at your Runlayer deployment, names an unknown or disabled connector, or carries unsupported options — and these launchers can no longer be allowlisted from the Shadow AI activity view.

! Saved Shadow AI skill and plugin files are now restricted to resource managers and security or analytics reviewers; organization-wide Shadow AI REST APIs and MCP tools now require organization-wide audit log access.

! The module now requires AWS provider `>= 6.64.0, < 7.0.0`; any older AWS provider version pinned in your dependency lock file must be upgraded before running `terraform plan`.

! Upgrading changes the ToolGuard launch template on every GPU profile and triggers an ASG rolling instance refresh, replacing hosts and restarting tasks (each restart reloads the model); expect a capacity dip per replacement batch.

! `MS_TEAMS_APP_URL` (the domain name) is effectively immutable after agents are provisioned — changing it orphans existing Azure Bot resources until each is deleted and recreated.

! The reply-behavior setting for Slack channel monitors is removed; monitors now only reply when the agent has something useful to add, with no option to restore the previous post-every-run behavior.

! AI Watch Enforce no longer allows `runlayer run` MCP launchers without `--host`, aimed at another host, naming an unknown or disabled connector, or carrying unsupported options — and they can no longer be allowlisted from the Shadow AI activity view.

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Runlayer: Adds Figma's remote MCP server to the catalog (now generally available), with OAuth app registration and per-user Figma account connection.

└──▷ **TRY IT**

Verify which org API key or enrollment key AI Watch is using on a managed device without exposing the full secret.

```
$ aiwatch config show
```

Adds `create_artifact` tool behavior alignment: artifacts an agent creates with the Runlayer MCP `create_artifact` tool now follow that agent's Artifact sharing setting.

Adds agents-on-Amazon-Bedrock support for OpenAI models (`gpt-6` Astra and `global.` / `us.` / `eu.` `-prefixed` OpenAI model IDs); a bare `openai.gpt-6…` ID is automatically served under its `global.` inference profile.

Adds Top spenders drilldown linking to a searchable, paginated view of every user's model spend or token usage.

Adds AI Watch client config backups for all supported clients (Cursor, VS Code, Codex, Gemini CLI, and others), retaining only the 5 most recent backups per file.

└──▷ **BREAKING ON UPGRADE**

! AI Watch Enforce now blocks MCP launchers whose `--host` does not point at your Runlayer deployment, names an unknown or disabled connector, or carries unsupported options — and these launchers can no longer be allowlisted from the Shadow AI activity view.

! Saved Shadow AI skill and plugin files are now restricted to resource managers and security or analytics reviewers; organization-wide Shadow AI REST APIs and MCP tools now require organization-wide audit log access.

`1.28.780 (2026-09-24)` **BLOG ↗** **STABLE**

Runlayer 1.28.780 adds Google AI Studio, JumpCloud MDM deployment, on-demand agent tools, and inline approval decisions in agent chats.

└──▷ **TRY IT**

Confirm which org API key or enrollment key AI Watch is using on a managed device before troubleshooting enrollment issues.

```
$ aiwatch config show
```

Skip Slack setup when programmatically creating an agent via MCP so it is immediately usable without requiring Slack connector access.

```
$ enable_channels=[]
```

Add the Google AI Studio provider so agents can be backed by Gemini models.

📍 In Runlayer, go to Settings → AI Providers, select `Google AI Studio`, paste your Gemini API key, and enter the model IDs you want available to agents.

Adds `aiwatch config show` command, displaying managed org API keys as a prefix (e.g. `rl_org_Ab1C2d3...`) and enrollment keys as prefix plus last four characters, to confirm which key a device is using.

Deploys AI Watch to macOS and Windows through JumpCloud from the MDM configuration wizard, with pre-filled configuration profiles and Private Repo Command Line Options for the signed MSI.

AI Watch Enforce now treats a local `runlayer run` MCP launcher as Runlayer-managed only when its `--host` points at your Runlayer deployment and names an active connector; launchers without `--host`, aimed at another host, or naming an unknown/disabled connector are blocked.

Adds a dedicated Google AI Studio provider so agents can run on Google Gemini models — configure via Settings → AI Providers with your Gemini API key and model IDs.

Enables inline approve/deny of agent tool-call approval requests directly in Agent Playground, Runlayer Assistant, and Home Build, with the agent continuing in the same conversation after the decision.

+12 more

└──▷ **BREAKING ON UPGRADE**

! The reply-behavior setting for Slack channel monitors is removed; monitors now only reply when the agent has something useful to add, with no option to restore the previous post-every-run behavior.

! AI Watch Enforce no longer allows `runlayer run` MCP launchers without `--host`, aimed at another host, naming an unknown or disabled connector, or carrying unsupported options — and they can no longer be allowlisted from the Shadow AI activity view.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Docker Sandboxes

Feedback ⓘ announced · Sep 24, 2026

BLOG ↗ RANK

Docker Sandboxes isolates applications in lightweight containers, providing secure, reproducible environments for development and deployment.

> Docker Sandboxes shipped Cloud Sandboxes — managed cloud compute for agents behind a new `--cloud` flag, with its own lifecycle commands, network policy, secrets store, environment files, and templates — plus a `sbx move` command to shift work between laptop and cloud, a centralized MCP gateway, and the open Sandbox Kit Specification v3 for declaring agent authority as reviewable OCI config.

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

01 **Cloud Sandboxes lifecycle via `--cloud` flag** NEW 98

Adds a global `--cloud` flag (`sbx` v0.45.0+) routing commands to Docker-managed cloud infrastructure instead of local execution, plus `sbx login` to authenticate against a Docker Agentic Platform subscription. Lifecycle is managed with `sbx --cloud create`, `sbx --cloud run` (flags `--name`, `--new`, `--detached`, `--cpus`, `--memory`, `--platform`, `--ttl`, `--on-timeout` with values `stop` / `restart` / `delete`, `--env`, `--allow-network`), `sbx --cloud attach <sandbox-name>`, `sbx --cloud exec`, `sbx --cloud stop`, `sbx --cloud rm --force`, `sbx --cloud ls`, and `sbx --cloud diagnose` (checks CLI health, sign-in, and cloud API connectivity without a local daemon). Access includes SSH via `sbx --cloud setup ssh` / `ssh <sandbox-id>@sbx_cloud` with `--detach-keys` (default `Ctrl+\`), file transfer via `sbx --cloud cp`, and port exposure via `sbx --cloud ports <sandbox> --publish` / `--unpublish` returning a public HTTPS URL. Sandboxes default to 2 CPUs/4 GiB with named T-shirt sizes (micro through xl) on `linux/amd64` or `linux/arm64`, expire after one hour by default, and can be extended with `sbx --cloud ttl +<duration>` up to a 24-hour cap.

Run a long-running agent task unattended in a cloud sandbox with 4 CPUs and 8 GiB RAM, then clean up automatically — useful in CI pipelines.

```
$ sbx --cloud create --name ci-task --cpus 4 --memory 8g --ttl 2h
--on-timeout delete claude
sbx --cloud exec ci-task make test
sbx --cloud rm --force ci-task
```

Expose a web service running inside a cloud sandbox to get a public HTTPS URL for sharing or webhook testing.

```
$ sbx --cloud ports cloud-project --publish 8080/tcp
```

Spin up a cloud sandbox with selective network access and run a Claude agent against a cloned repo — without tying up local compute.

```
$ sbx --cloud create --name cloud-project --allow-network
github.com:443 claude
sbx --cloud exec cloud-project git clone
https://github.com/docker/welcome-to-docker.git
/home/agent/workspace/project
sbx --cloud attach cloud-project
```

Retrieve output produced by a cloud agent and clean up the sandbox when done.

```
$ sbx --cloud cp cloud-project:/home/agent/workspace/review.md
./review.md
sbx --cloud rm cloud-project
```

— Extensive named commands, flags, and defaults given `launch-20260924-6c1aef17`

02 **Cloud environment files (`sbx --cloud env`)** NEW 95

Adds `sbx --cloud env <file>` to manage cloud sandboxes from a `.sbxenv.yaml` file supporting agents, kits, environment variables, CPU/memory limits, and provider credentials (a subset of local config fields). `sbx --cloud env plan <file>` previews a configuration including inherited cloud credentials before provisioning; `sbx --cloud env run --detached <file>` creates a sandbox without attaching; `sbx --cloud env rm <file>` tears down a sandbox and removes only the secrets it provisioned. A `snapshot: true` credential entry resolves a host command or vault reference once and stores it as a literal secret, preventing dynamic refresh, and `DOCKER_ACCESS_TOKEN` scopes cloud environment state to a specific Docker identity.

Preview a cloud sandbox plan and then launch it detached — useful in CI where you want approval before provisioning and no interactive terminal.

```
$ sbx --cloud env plan ./cloud.sbxenv.yaml
sbx --cloud env run --detached ./cloud.sbxenv.yaml
```

Define a cloud sandbox environment file with resource limits and a project-scoped environment variable, then run a command inside it.

```yaml
# cloud.sbxenv.yaml
project: cloud-project
memory: 4g
env:
  PROJECT_NAME: example
```

Clean up a cloud sandbox and its provisioned secrets when an unattended run is complete.

```
$ sbx --cloud env rm ./cloud.sbxenv.yaml
```

— Names config file, subcommands and env var fully product docs

03 **Cloud network policy (`sbx --cloud policy`)** NEW 95

Adds `sbx --cloud policy init` to set an account-scoped outbound default (`allow-all`, `balanced`, `deny-all`), with `--sandbox <sandbox>` to scope deny-all to one sandbox. `sbx --cloud policy allow network <host:port>` (optionally `--sandbox <sandbox>`) creates allow rules; `sbx --cloud policy ls` / `ls <sandbox>` inspects policy documents; `sbx --cloud policy log <sandbox>` reviews enforcement decisions; `sbx --cloud policy rm network --resource <host:port> --sandbox <sandbox>` removes rules; `sbx --cloud policy reset --force` restores the platform default. The cloud-only policy store applies deny rules over allow across account and sandbox scopes, but rejects local governance profiles and the `--protocol` option, and does not support L7 (HTTP method/path) filtering like local sandboxes do.

Lock down all outbound traffic for your cloud sandboxes by default, then carve out only the endpoints your AI agent actually needs.

```
$ sbx --cloud policy init deny-all
sbx --cloud policy allow network api.github.com:443
sbx --cloud policy allow network api.anthropic.com:443
```

Create a new cloud sandbox with network rules baked in at creation time, so no post-creation policy steps are needed in CI.

```
$ sbx --cloud create --name cloud-project \
  --allow-network api.github.com:443 \
  --deny-network example.com claude
```

Audit what network connections a sandbox is actually making and whether policy allowed or denied them.

```
$ sbx --cloud exec cloud-project curl -I https://api.github.com
sbx --cloud policy log cloud-project
```

— Full command set, scopes, and precedence rules named product docs

04 **Cloud secrets store (`sbx --cloud secret`)** NEW 95

Adds `sbx --cloud secret set <service>` to store API keys/tokens for built-in services (anthropic, cursor, droid, github, google, groq, mistral, nebius, openai, openrouter, xai), including an `--oauth` flow for `openai`/Codex. `--sandbox <name>` scopes a credential to one sandbox, taking precedence over the account-scoped default. `sbx --cloud secret set-custom --name --host --env --header --format` injects arbitrary credentials for non-built-in services via cloud proxy header injection; `sbx --cloud secret ls` lists metadata without exposing values; `sbx --cloud secret rm <service> --sandbox <name> --host <host> --force` removes credentials non-interactively. This is a separate cloud-only secret store from the local `sbx secret` store, and the underlying Cloud Sandboxes secrets proxy injects keys per-request so agents never hold the actual value.

Store an Anthropic API key in the cloud secret store, then launch a Claude Code agent in a named cloud sandbox — keeping credentials out of the sandbox filesystem.

```
$ sbx --cloud secret set anthropic
sbx --cloud run claude --name cloud-project
```

Inject a custom API credential for a non-built-in service so the cloud proxy sets the Authorization header on all requests to that host.

```
$ sbx --cloud secret set-custom --name project-api \
  --host api.example.com --env PROJECT_API_TOKEN \
  --header Authorization --format 'Bearer %s'
```

Scope an OpenAI API key to a single sandbox so it takes precedence over any account-scoped credential, then remove it without an interactive prompt in CI.

```
$ sbx --cloud secret set openai --sandbox cloud-project
# later, in a cleanup script:
sbx --cloud secret rm openai --sandbox cloud-project --force
```

Store a cloud-scoped Anthropic API key before launching an agent — local credentials are not shared with cloud sandboxes.

```
$ sbx --cloud secret set anthropic
```

— Every flag and built-in service name is carried through `launch-20260924-6c1aef17`

---

05 **`sbx move` for local/cloud snapshot transfer** NEW 95

`sbx move <sandbox> --to cloud` transfers a local sandbox filesystem snapshot to create a new cloud sandbox; `sbx move <sandbox> --to local` does the reverse. `--name` sets the destination name prefix (cloud appends a unique suffix); `--ttl` sets destination expiration (default typically one hour); `--on-timeout stop` / `delete` controls expiration behavior; `--force` skips confirmation when local workspace mounts or network restrictions are present; `--platform linux/amd64` / `linux/arm64` and a pre-transfer compatibility check verify destination support before transfer. `sbx cp` moves files separately when workspace mounts or cloud volumes aren't carried by the move, and `sbx move <sandbox-name> --cloud --force` migrates a local sandbox to the cloud backend.

Move a local sandbox to the cloud with a 2-hour TTL and stop-on-expiry so work is preserved rather than deleted.

```
$ sbx move local-project --to cloud --name cloud-project --ttl 2h
--on-timeout stop
```

Pull a cloud sandbox back to local for offline work, then verify both environments and check port bindings on the new local copy.

```
$ sbx move cloud-project --to local --name local-copy && sbx --
cloud ls && sbx ports local-copy
```

Migrate a local sandbox to the cloud backend for a workflow that needs cloud compute; confirm or force through the prompt when running in CI with no terminal.

```
$ sbx move <sandbox-name> --cloud --force
```

Copy files out of a cloud sandbox that has no host mount, before removing the source after a move.

```
$ sbx cp <sandbox-name>:/workspace/output ./local-output
```

Hand off a long-running refactor you started locally to cloud compute so you can close your laptop without interrupting the agent.

```
$ sbx move my-project --to cloud
```

— Direction, flags and compatibility checks all named `launch-20260924-6c1aef17`

---

06 **Sandbox templates: save, export, import, and run from images**

NEW

95

`sbx template save <sandbox-name> <name:tag>` snapshots a sandbox's filesystem as a reusable template image in the local image store; `--output <file>.tar` exports it as a tar; `sbx template load <file>.tar` imports a tar directly, bypassing any registry; `sbx template ls` lists and `sbx template rm <name:tag>` deletes saved templates. `-t` / `--template` on `sbx run` launches a sandbox from a saved or registry-hosted template (e.g. `sbx run -t my-template:v1 claude`), and `sbx secret set --registry <registry> --password-stdin` stores pull credentials for private registries (GHCR, ECR, ACR, self-hosted Nexus) beforehand; cached layers persist across sandbox creation/deletion until `sbx reset`. Cloud sandboxes have a counterpart, `sbx --cloud template save` with `--template` / `--image-ref`, to capture a running cloud sandbox as a reusable cloud template.

Snapshot a configured sandbox as a reusable template, then spin up a fresh sandbox from it — ideal for sharing a pre-tooled environment across a team.

```
$ sbx template save my-sandbox my-template:v1
sbx run -t my-template:v1 claude
```

Export a template to a tar and load it on another machine that has no registry access, then launch a sandbox from it.

```
$ # On the source machine
sbx template save my-sandbox my-template:v1 --output my-
template.tar

# On the target machine
sbx template load my-template.tar
sbx run -t my-template:v1 claude
```

Store credentials for a private GitHub Container Registry image before pulling it as a template, so the pull is authenticated rather than anonymous.

```
$ gh auth token | sbx secret set --registry ghcr.io --password-
stdin
sbx run --template ghcr.io/my-org/my-template:v1 claude
```

— Full subcommand set and registry flow named `product docs`

---

07 **Sandbox Kit Specification v3** NEW 95

Introduces the open source (Apache 2.0, `docker/sandbox-kit-spec`) OCI-native Kit format declaring exactly what authority a sandbox agent may exercise, stored in a single OCI manifest annotation (`vnd.docker.sandbox.kit.descriptor`) so Kits build with `docker buildx build`, pull with `docker pull`, and are pinnable by digest. Two kinds compose at launch — a `workload` Kit (root filesystem) and any number of `mixin` Kits (CLI, network rule, or credential overlays) — using versioned capability types like `com.docker.sandbox/network-policy@2` (explicit `allow` / `deny` rules scoped to `hosts`, `methods`, `paths`, deny taking precedence, e.g. blocking `DELETE /repos/**`) and `com.docker.sandbox/credential@1` (`proxyManaged: true` injects real tokens into named-domain requests while exposing only a sentinel in the sandbox). A `kind: set` descriptor runs build-time coherence checks across a Kit collection, and authority-diff gating auto-applies updates within a normalized grant set but requires explicit approval for any widening, including removing a `deny` rule. Two conformance suites verify artifact and runtime compliance.

Run a sandboxed agent with the GitHub CLI mixin to give it scoped GitHub access defined by the Kit's network-policy and credential declarations.

```
$ cd examples
sbx run ./hello --kit ./gh
```

Author a mixin Kit that allows most GitHub API methods but explicitly denies DELETE on `/repos`/**, keeping a token that can open PRs from deleting repositories.

```yaml
capabilities:
  - type: com.docker.sandbox/network-policy@2
    config:
      runtime:
        allow:
          - github.com
          - hosts: [api.github.com]
            methods: [GET, HEAD, POST, PATCH, PUT, DELETE]
        deny:
          - hosts: [api.github.com]
            methods: [DELETE]
            paths: [/repos/**]
  - type: com.docker.sandbox/credential@1
    optional: true
    config:
      service: github
      phase: runtime
      apiKey:
        name: GH_TOKEN
      proxyManaged: true
      inject:
        - {domain: api.github.com, header: Authorization, format:
'Bearer %s'}
```

— Spec mechanism, capability types, and conformance detailed `launch-20260924-0c0143b0`

08 Centralized MCP gateway and server management NEW

71

Adds `sbx --cloud mcp load <server-name> --sandbox <sandbox>` to connect a Docker Agentic Platform MCP server into a running cloud sandbox, and `sbx --cloud mcp ls` (account-wide) / `sbx --cloud mcp ls <sandbox>` to list servers reported by cloud sandbox gateways. A centralized MCP gateway lets agents running in cloud, locally, or in other clients (e.g. ChatGPT desktop) share one set of connected MCP servers (Jira, Linear, Grafana, incident.io, or any streamable HTTP endpoint).

— Commands named but no runnable example given `launch-20260924-6c1aef17`

09 Experimental persistent volumes for cloud sandboxes NEW

69

Adds `sbx --cloud volume create` and a `--volume <name>:<path>` flag at sandbox creation for experimental persistent volumes that snapshot on sandbox exit, useful for caching dependencies across runs.

Create a persistent volume for dependency caching and mount it into a new cloud sandbox so repeated runs skip reinstalling packages.

```
$ sbx --cloud volume create dependency-cache
sbx --cloud create --name cloud-project --volume dependency-
cache:/workspace/cache claude
sbx --cloud exec cloud-project sudo chown agent:agent
/workspace/cache
```

— Flags named but marked experimental with little mechanism `product docs`

---

10 **OpenCode: GitHub Copilot models and model switching** NEW

67

Supports GitHub Copilot models in OpenCode by configuring a GitHub credential, with no separate device login required inside the sandbox; `sbx run --model` selects a local model or inference endpoint, and the model's supported thinking levels are available as OpenCode variants cycled with Ctrl+T.

— Names flag and shortcut, no worked example `product docs`

---

11 **Kits: pre-built agent images for Cloud Sandboxes** NEW

60

Introduces Kits — pre-configured, pre-built sandbox images for leading coding agents (Claude Code, Codex, Copilot, Antigravity, Open Code, Hermes) — launchable with a single `sbx --cloud run <kit>` command; requires `sbx` version 0.45.1 or later.

Launch a Claude Code agent in the cloud before leaving for the day so it keeps running overnight without your laptop.

```
$ sbx --cloud run claude
```

— Names Kits and command but little mechanism `launch-20260924-6c1aef17`

THINNER COVERAGE BELOW

---

12 **Local and cloud sandboxes fully separated** BREAKING

57

Local and cloud sandboxes maintain fully separate resources — secrets, templates, volumes, and network policy — with independent stores (`sbx secret` vs `sbx --cloud secret`, `sbx ls` vs `sbx --cloud ls` showing independent lists), and local-only flags are rejected rather than silently ignored when used with `--cloud`. Existing local workflows must be reconfigured before moving to cloud.

— States the migration impact but no exact steps launch-20260924-6c1aef17

13 **Web console for Cloud Sandboxes** NEW 37

Adds a web console for spinning up and managing Cloud Sandboxes without the CLI, including Kit selection and a Run button.

— UI path only, no mechanism or config detail launch-20260924-6c1aef17

└──▷ **BREAKING ON UPGRADE**

! Local and cloud sandboxes maintain separate secrets, templates, and network policies — existing local workflows must be reconfigured before moving to cloud.

announced · Sep 24, 2026 BLOG ↗

Docker Sandbox Kit Specification v3 introduces OCI-native Kit format for declaring agent authority as versioned, reviewable config

└──▷ **TRY IT**

Run a sandboxed agent with the GitHub CLI mixin to give it scoped GitHub access defined by the Kit's network-policy and credential declarations.

```
$ cd examples
sbx run ./hello --kit ./gh
```

Author a mixin Kit that allows most GitHub API methods but explicitly denies DELETE on `/repos/**`, keeping a token that can open PRs from deleting repositories.

```yaml
capabilities:
  - type: com.docker.sandbox/network-policy@2
    config:
      runtime:
        allow:
          - github.com
          - hosts: [api.github.com]
            methods: [GET, HEAD, POST, PATCH, PUT, DELETE]
        deny:
          - hosts: [api.github.com]
            methods: [DELETE]
            paths: [/repos/**]
  - type: com.docker.sandbox/credential@1
    optional: true
    config:
      service: github
      phase: runtime
      apiKey:
        name: GH_TOKEN
      proxyManaged: true
      inject:
        - {domain: api.github.com, header: Authorization, format:
'Bearer %s'}
```

Introduces the Docker Sandbox Kit Specification v3, open source under Apache 2.0 at `docker/sandbox-kit-spec`, defining a portable OCI image format for declaring exactly what authority a sandbox agent may exercise.

Stores all Kit declarations in a single OCI manifest annotation (`vnd.docker.sandbox.kit.descriptor`), so Kits build with `docker buildx build`, pull with `docker pull`, and are pinnable by digest.

Supports two Kit kinds: a `workload` Kit (supplies the root filesystem) and a `mixin` Kit (an overlay adding a CLI, network rule, or credential binding), composed at launch with one workload and any number of mixins.

Defines versioned capability types — including `com.docker.sandbox/network-policy@2` and `com.docker.sandbox/credential@1` — each with a normative page describing what a conforming runtime must enforce.

Network-policy capabilities support explicit `allow` and `deny` rules scoped to `hosts`, `methods`, and `paths`, with `deny` taking precedence, so a mixin can permit most of an API while blocking specific destructive operations (e.g. `DELETE /repos/**`).

+4 more

announced · Sep 24, 2026 BLOG ↗

Docker Sandboxes adds Cloud Sandboxes: managed microVM compute in the cloud with `sbx --cloud run` and `sbx move --to cloud`.

└──▷ **TRY IT**

Launch a Claude Code agent in the cloud before leaving for the day so it keeps running overnight without your laptop.

```
$ sbx --cloud run claude
```

Hand off a long-running refactor you started locally to cloud compute so you can close your laptop without interrupting the agent.

```
$ sbx move my-project --to cloud
```

Adds `sbx --cloud run <kit>` to launch a sandbox on Docker-managed cloud compute, with size options from Micro (1 vCPU / 2 GiB, $0.07/hr) to XL (16 vCPU / 32 GiB, $1.12/hr) and sessions up to 24 hours.

Introduces Kits: pre-configured, pre-built sandbox images for leading coding agents (Claude Code, Codex, Copilot, Antigravity, Open Code, Hermes), launchable with a single `sbx --cloud run <kit>` command.

Adds a centralized MCP gateway so agents running in cloud, locally, or in other clients (e.g. ChatGPT desktop) can share a single set of connected MCP servers (Jira, Linear, Grafana, incident.io, or any streamable HTTP endpoint).

Adds a Cloud Sandboxes secrets proxy that injects stored keys or tokens per-request so agents never hold the actual secret value.

Adds network policies that define which endpoints agents can reach, configurable once and applied across all cloud sandboxes.

+2 more

└──▷ **BREAKING ON UPGRADE**

! Local and cloud sandboxes maintain separate secrets, templates, and network policies — existing local workflows must be reconfigured before moving to cloud.

WAS THIS USEFUL? 👍 👎

ai-safe2-framework

Feedback ⓘ 2026-09-24_CLI_0.8.0

NOTES ↗ RANK

Ai-Safe2-Framework provides governance, risk, and compliance controls for securing agentic AI systems and non-human identities.

> AI SAFE² CLI 0.8.0 introduces evidence-bound remediation planning for AISM assessments, adding a source-template command and a planning command with machine-readable contracts, human decision gates, and safety guarantees.

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

01 **Remediation planning command with decision gates and safety checks**

NEW 98

`safe2 aism plan` validates a proposed implementation path and produces canonical JSON in the `aism-remediation-plan-v1` output contract plus a human-readable decision card. The `--previous` flag preserves action history across reassessments and detects completed actions that regress or disappear; `--strict` returns a failing exit code unless the plan reaches the `ready_for_human_decision` gate. Every plan carries a three-state human decision gate (`hold` , `review` , `ready_for_human_decision`) with `remediation_authorized` and `conformance_claim` fixed to `false` , alongside cycle detection and required dependency ordering for remediation action graphs, calendar-date validation for residual-risk reviews, and safe paired-output handling that removes a partial plan if decision-card creation fails.

— Command, flags, output contract, and safety mechanisms all named. 2026-09-24_CLI_0.8.0

02 **Remediation source template command and contract** NEW

81

`safe2 aism remediation-init` creates a remediation source template bound to the exact assessment, system identity, and assessment scope. It conforms to the new `aism-remediation-source-v1` versioned input contract, which defines actions, evidence, assumptions, alternatives, ownership, and accepted risk.

— Names command and contract; scope of what it binds to is described.

2026-09-24_CLI_0.8.0

2026-09-24_CLI_0.8.0 NOTES ↗ LATEST STABLE

> AI SAFE² CLI 0.8.0 adds evidence-bound AISM remediation planning with human decision gates and machine-readable output contracts.

└──▷ **GET THIS VERSION**

```
$ git clone --branch 2026-09-24_CLI_0.8.0
https://github.com/CyberStrategyInstitute/ai-safe2-framework.git
# already have the repo? check out this version:
$ git checkout 2026-09-24_CLI_0.8.0
```

Adds `safe2 aism remediation-init` command to create a remediation source template bound to the exact assessment, system identity, and assessment scope.

Adds `safe2 aism plan` command to validate a proposed implementation path and produce canonical JSON (`aism-remediation-plan-v1`) plus a human-readable decision card.

Adds `--previous` flag to `safe2 aism plan` to preserve action history across reassessments and detect completed actions that regress or disappear.

Adds `--strict` flag to `safe2 aism plan` to return a failing exit code unless the plan reaches the `ready_for_human_decision` gate.

Introduces `aism-remediation-source-v1` as the versioned input contract defining actions, evidence, assumptions, alternatives, ownership, and accepted risk.

+5 more

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗ RANK

The free and flexible app for your private thoughts.

> Obsidian's latest update adds a way to view files outside your vault, deeper OS integration, Kanban folder-grouping improvements, and removes the fuzzy-search file-count ceiling, while a font-size storage change breaks sync for that setting.

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

01 **Open file from outside the vault** NEW 70

A new `Open file from outside the vault...` command lets you pick and view any file on your computer via the system file picker without adding it to your vault. This is Desktop only.

Open a Markdown incident-report file that lives outside your vault — useful when triaging files dropped into a shared folder that is not part of any vault.

> 📍 In Obsidian, open the Command Palette and run 'Open file from outside the vault...', then navigate to the target file in the system file picker.

— Names exact command and platform scope, includes usage steps.

changelog-20260925-e3fd1cda

02 **Kanban folder-grouping improvements** IMPROVED 65

When a Kanban view is grouped by `file.folder`, cards can now be moved between folder columns, and creating a note in a column now places it in the corresponding folder.

— Names exact grouping key and two concrete behavior changes. changelog-20260925-e3fd1cda

03 **Fuzzy matching in link suggestions and quick switcher**

IMPROVED 65

Fuzzy matching in editor link suggestions (`[[`) and the quick switcher is now always used, removing the previous 10,000-file vault limit that forced a simplified matching algorithm.

— States before/after and the concrete file-count limit removed.

changelog-20260925-e3fd1cda

THINNER COVERAGE BELOW

04 **OS-level file association for Markdown** NEW 45

Obsidian now registers in the OS Open with menu and can be set as the default app for Markdown files; this requires the latest installer.

— Clear mechanism but no exact setup steps given. changelog-20260925-e3fd1cda

05 **Extra large iOS 27 widgets** NEW 40

iOS 27 adds support for extra large Quick Capture and View Note home-screen widgets.

— Names widget types and OS version but no setup detail. changelog-20260925-e3fd1cda

06 **Localized links to Obsidian website** IMPROVED 20

All in-app links now route to the localized version of the Obsidian website.

— Bare statement with no mechanism or scope detail. changelog-20260925-e3fd1cda

└──▷ BREAKING ON UPGRADE

! Base font size is no longer stored in `appearance.json` and will not sync via Obsidian Sync; each device now holds its own value in local storage, so synced vaults will lose a shared font-size setting on upgrade.

changelog · Sep 25, 2026 CHANGELOG ↗ LATEST STABLE

Obsidian: Adds the `Open file from outside the vault...` command, letting you pick and view any file on your computer without adding it to your vault (Desktop only).

└──▷ HOW TO FIND IT

Open a Markdown incident-report file that lives outside your vault — useful when triaging files dropped into a shared folder that is not part of any vault.

📍 In Obsidian, open the Command Palette and run 'Open file from outside the vault...', then navigate to the target file in the system file picker.

Obsidian now registers in the OS Open with menu and can be set as the default app for Markdown files (requires latest installer).

Kanban: cards can now be moved between folder columns when the view is grouped by `file.folder`.

Kanban: creating a note in a column now places it in the corresponding folder when the view is grouped by `file.folder`.

iOS 27: adds support for extra large Quick Capture and View Note home-screen widgets.

Fuzzy matching in editor link suggestions (`[[`) and the quick switcher is now always used, removing the 10,000-file vault limit that previously forced a simplified

algorithm.

+2 more

└─▷ **BREAKING ON UPGRADE**

! Base font size is no longer stored in `appearance.json` and will not sync via Obsidian Sync; each device now holds its own value in local storage, so synced vaults will lose a shared font-size setting on upgrade.

WAS THIS USEFUL? 👍 👎

OpenSSL

Feedback ⓘ announced · Sep 23, 2026 BLOG ↗ RANK

General purpose TLS and crypto library

> OpenSSL 4.1 Beta adds DTLS 1.3 and RFC 8701 GREASE support, hardware-optimized post-quantum crypto operations, and a handful of breaking changes to Configure options, Windows build targets, and the tsget utility.

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

01 **Configure drops no-ecdsa/no-ecdh in favor of no-ec** BREAKING 70

Removes the `no-ecdsa` and `no-ecdh` options from Configure ; scripts or build systems using them will break. Use `no-ec` instead to disable elliptic curve support.

— Names exact flags and the migration path launch-20260923-e530ee45

02 **Hardware-optimized ML-DSA, ML-KEM and AES-CBC operations** IMPROVED 64

Adds optimized ML-DSA and ML-KEM NTT operations on `ppc64le` ; optimized ML-DSA operations on `s390x` and `x86_64` ; AVX-512-optimized SHAKE x4 for ML-DSA on `x86_64` ; and AVX-512 and VAES optimizations for AES-CBC decryption on `x86_64` .

— Names algorithms and target architectures precisely launch-20260923-e530ee45

03 **DTLS 1.3 support (RFC 9147)** NEW 62

Adds DTLS 1.3 support per RFC 9147, with usage details documented in the ossl-guide-dtlsv13(7) manual page. Also adds DTLS support in the SSL listener API.

— Names RFC and manual page but no code sample launch-20260923-e530ee45

04 **tsget requires Net::Curl::Easy** BREAKING 62

Changes the `tsget` utility to use `Net::Curl::Easy` (Net-Curl CPAN distribution) instead of the abandoned `WWW::Curl::Easy` ; install `Net::Curl::Easy` before upgrading.

— Names old and new dependency with upgrade action launch-20260923-e530ee45

THINNER COVERAGE BELOW

05 **New MSVC 2013 build targets** NEW 50

Adds `VC-WIN32-MSVC2013` and `VC-WIN64A-MSVC2013` build targets for C99 bridging on MSVC 2013.

— Names exact target strings, no further detail launch-20260923-e530ee45

06 **Windows Itanium and CE build targets removed** BREAKING

50

Removes the `VC-WIN64I` (Windows-on-Itanium) and `VC-CE` (Windows CE) build targets from Configurations.

— Names removed targets but no migration guidance launch-20260923-e530ee45

07 **RFC 8701 GREASE support** NEW

27

Adds support for RFC 8701 GREASE (Generate Random Extensions And Sustain Extensibility).

— Names the RFC but no usage detail launch-20260923-e530ee45

08 **Initial Elbrus2000 (e2k) architecture support** NEW

27

Adds initial support for the Elbrus2000 (`e2k`) architecture.

— Names architecture, no build or usage details launch-20260923-e530ee45

09 **IKEV2 KDF support** NEW

20

Adds support for IKEV2 KDF.

— Bare mention, no mechanism given launch-20260923-e530ee45

└──▷ **BREAKING ON UPGRADE**

! The `no-ecdsa` and `no-ecdh` options are removed from `Configure` ; scripts or build systems using them will break — use `no-ec` instead.

! The `VC-WIN64I` (Windows-on-Itanium) and `VC-CE` (Windows CE) build targets are removed from Configurations.

! `tsget` now requires `Net::Curl::Easy` (Net-Curl CPAN distribution); the previously used `WWW::Curl::Easy` is no longer supported — install `Net::Curl::Easy` before upgrading.

> announced · Sep 23, 2026 BLOG ↗
>
> OpenSSL 4.1 Beta adds DTLS 1.3, RFC 8701 GREASE, IKEV2 KDF, and ML-DSA/ML-KEM hardware optimizations
>
> Drops `no-ecdsa` and `no-ecdh` options from `Configure` ; use `no-ec` instead to disable elliptic curve support.
>
> Adds DTLS 1.3 support (RFC 9147); see the ossl-guide-dtlsv13(7) manual page for usage details.
>
> Adds DTLS support in the SSL listener API.

Adds support for RFC 8701 GREASE (Generate Random Extensions And Sustain Extensibility).

Adds support for IKEV2 KDF.

+4 more

└─▷ **BREAKING ON UPGRADE**

- ! The `no-ecdsa` and `no-ecdh` options are removed from Configure ; scripts or build systems using them will break — use `no-ec` instead.
- ! The `VC-WIN64I` (Windows-on-Itanium) and `VC-CE` (Windows CE) build targets are removed from Configurations.
- ! `tsget` now requires `Net::Curl::Easy` (Net-Curl CPAN distribution); the previously used `WWW::Curl::Easy` is no longer supported — install `Net::Curl::Easy` before upgrading.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ changelog · Sep 25, 2026 CHANGELOG ↗

RANK

Connect, protect, and build everywhere.

> Cloudflare's biggest ship this window is MCP server portals reaching general availability with a full set of logging, auth, and routing controls, alongside four new Workers AI models, a new Gateway Traffic Destination selector, DNSSEC sentinel support in 1.1.1.1, and smaller updates across R2, Mesh, Images, and Worker previews.

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

01 **DNSSEC root key sentinel support in 1.1.1.1** NEW 85

1.1.1.1 now supports RFC 8509 root key trust anchor sentinels, letting users diagnose DNSSEC root key rollovers by querying sentinel names such as `root-key-sentinel-is-ta-38696.dnstest.dev` to verify whether the resolver trusts a DNSSEC root key like KSK-2024 (key tag 38696).

— Gives exact sentinel query name to run changelog-20260925-1ce42832

02 **MCP server portals reach GA with new controls** NEW 80

MCP server portals are now generally available to all Cloudflare customers, providing a single endpoint for approved Model Context Protocol (MCP) servers with Access logging of tool, prompt, and resource activity. New capabilities include Logpush support to export portal activity to external storage or a SIEM, Static OAuth client credentials for providers without Dynamic Client Registration, Service token authentication for autonomous agents and machine-to-machine access, Code Mode policies to control how portals reduce tool definitions and token use, Gateway routing for HTTP logging and DLP scanning, and Session management for reconnecting servers and changing authorizations from the portal.

— Names every capability but omits exact API paths changelog-20260925-1ce42832

03 **New models added to Workers AI catalog** NEW 75

Workers AI adds four new models: `claude-opus-5.5` (Anthropic Claude Opus 5.5) with a 1 million token context window, 128K token max output, always-on adaptive thinking, and medium default effort level for long-running agentic coding and knowledge work; `gpt-6-luna` (OpenAI GPT-6 Luna), the most efficient GPT-6 variant for focused, high-volume tasks; `gpt-6-sol` (OpenAI GPT-6 Sol), the mid-tier GPT-6 variant for complex coding and agentic workflows; and `grok-4.7` (xAI), a frontier reasoning

model supporting text and image inputs, function calling, structured outputs, and configurable reasoning effort levels (`low` , `medium` , `high` , `xhigh`).

— Exact model IDs and specs given, no call example `product docs`

04 **Worker Previews for pull requests** NEW 75

Worker Previews give each pull request an isolated, production-like environment with its own code, configuration, URL, logs, errors, metrics, and traces, viewable from a Preview dropdown in the Worker dashboard.

Deploy a Worker Preview for an in-progress branch so reviewers can test against an isolated, production-like environment before merging.

```
$ npx wrangler preview
```

— Includes a runnable preview command `changelog-20260925-1ce42832`

05 **Traffic Destination selector for Gateway policies** NEW 65

Gateway Network (and HTTP) policies gain a Traffic Destination selector to match traffic based on its egress path, for example blocking traffic that exits via a specific destination or egress location. Configure it in the Cloudflare Zero Trust dashboard under Gateway > Policies > Network (or HTTP) by adding a selector and choosing Traffic Destination .

When you want a Gateway Network policy to match only traffic that exits via a specific destination type, use the new Traffic Destination selector to target that exit path.

> 📍 In the Cloudflare Zero Trust dashboard, go to Gateway › Policies › Network (or HTTP), create or edit a policy, add a selector, and choose Traffic Destination to match traffic based on how it exits Cloudflare.

Route Gateway Network policy decisions based on egress path — for example, block traffic that exits via a specific egress location.

> 📍 In the Cloudflare Zero Trust dashboard, go to Gateway › Policies › Network, create or edit a policy, add a selector, and choose Traffic Destination to match on how traffic exits Cloudflare.

— Exact UI navigation given but no API or flag `product docs`

06 **R2 bandwidth metrics page in dashboard** NEW 65

Adds a product-level Metrics page to the Cloudflare dashboard for R2, showing bandwidth usage split by object upload and download across all buckets or per bucket, also exposed via the GraphQL Analytics API.

— Names dashboard page and API surface, no query example changelog-20260925-1ce42832

07 **Durable Object name search limit raised to 128 characters**

IMPROVED 60

Durable Object name search in the Cloudflare dashboard now accepts up to 128 characters, up from 20, applying to the Metrics tab object filter and Data Studio object search.

— Exact before/after limits and locations named changelog-20260925-1ce42832

THINNER COVERAGE BELOW

08 **Guided onboarding for Cloudflare Mesh participants** NEW 45

Cloudflare Mesh adds guided onboarding for participants, including node deployment setup, client installers, organization name lookup, and a unified participant management table.

— Lists onboarding steps but no exact commands product docs

09 **Account-level transformation analytics for Images** NEW 25

Cloudflare Images adds account-level transformation analytics, giving visibility into image transformation usage across an account.

— Bare description, no metrics or navigation named product docs

changelog · Sep 25, 2026 CHANGELOG ↗ STABLE

Cloudflare: Adds Logpush support for MCP server portals, enabling export of portal activity to external storage or a SIEM system.

└──▷ TRY IT

Deploy a Worker Preview for an in-progress branch so reviewers can test against an isolated, production-like environment before merging.

```
$ npx wrangler preview
```

Adds Static OAuth client credentials support to MCP server portals for providers that do not support Dynamic Client Registration.

Adds `Service token authentication` to MCP server portals for autonomous agents and machine-to-machine access.

Adds `Code Mode policies` to MCP server portals to control how portals reduce tool definitions and token use.

MCP server portals are now generally available to all Cloudflare customers, providing a single endpoint for approved Model Context Protocol (MCP) servers with Access logging of tool, prompt, and resource activity.

Adds a product-level `Metrics` page to the Cloudflare dashboard for R2, showing bandwidth usage split by object upload and download across all buckets or per bucket, also exposed via the `GraphQL Analytics API`.

+5 more

WAS THIS USEFUL?  

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

vigolium

Feedback ⓘ v0.5.1 NOTES ↗ RANK

Vigolium - High-fidelity vulnerability scanner fusing agentic AI with native speed, modularity, and precision

> Vigolium v0.5.1 deepens its Anthropic integration with new reasoning-effort control, a reworked token-limit config, Claude 5 pricing, adaptive thinking, prompt caching, and clearer skill tool visibility.

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

01 **agent.olium.max_tokens config key** BREAKING 85

Adds `agent.olium.max_tokens` config key, defaulting to `0` (provider default, 64000 on Anthropic); oversized values are ignored. This is a breaking change: previously set oversized values are silently ignored on upgrade.

Cap token usage at the provider default rather than a hard limit, letting Anthropic apply its 64k ceiling automatically.

```yaml
agent.olium.max_tokens: 0
```

— Named key, default, and concrete provider ceiling given v0.5.1

02 **Reasoning effort control for Anthropic** NEW 75

Adds `reasoning_effort: max` for Anthropic, extending effort-level control to every provider that supports it.

Push Anthropic models to their highest reasoning depth for complex audit runs that benefit from extended chain-of-thought.

```yaml
reasoning_effort: max
```

— Exact config value given, mechanism and scope clear v0.5.1

THINNER COVERAGE BELOW

03 **Skill allowed-tools visibility** IMPROVED 45

Surfaces skill `allowed-tools` to the model, with warnings for tools the run does not expose.

— Names the field but no config path or example v0.5.1

04 **Conversation-tail prompt caching** NEW 40

Adds conversation-tail prompt caching on Anthropic, reducing input cost on long agent loops.

— Mechanism and benefit named but no config surface v0.5.1

05 **Claude 5 family pricing** NEW 30

Adds Claude 5 family pricing (Fable 5.1/5, Opus 5.5/5, Sonnet 5) in audit cost estimates.

— Names model tiers but no mechanism or usage detail v0.5.1

06 **Adaptive thinking for Claude 4.6+** NEW 20

Enables adaptive thinking for Claude 4.6+ models.

— Bare statement, no behavior or config explained v0.5.1

└──▷ BREAKING ON UPGRADE

! `agent.olium.max_tokens` now defaults to `0` (provider default); previously set oversized values are silently ignored on upgrade.

v0.5.1 NOTES ↗ LATEST STABLE

vigolium v0.5.1 adds Claude 5 pricing, `reasoning_effort: max`, adaptive thinking, and prompt caching for Anthropic models.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.5.1 https://github.com/vigolium/vigolium.git
# already have the repo? check out this version:
$ git checkout v0.5.1
```

└──▷ USE IT

Cap token usage at the provider default rather than a hard limit, letting Anthropic apply its 64k ceiling automatically.

```yaml
agent.olium.max_tokens: 0
```

Push Anthropic models to their highest reasoning depth for complex audit runs that benefit from extended chain-of-thought.

```yaml
reasoning_effort: max
```

Adds `reasoning_effort: max` for Anthropic, extending effort-level control to every provider that supports it.

Adds `agent.olium.max_tokens` config key, defaulting to `0` (provider default, 64000 on Anthropic); oversized values are ignored.

Adds Claude 5 family pricing (Fable 5.1/5, Opus 5.5/5, Sonnet 5) in audit cost estimates.

Enables adaptive thinking for Claude 4.6+ models.

Adds conversation-tail prompt caching on Anthropic, reducing input cost on long agent loops.

+1 more

└─▷ **BREAKING ON UPGRADE**

! `agent.olium.max_tokens` now defaults to `0` (provider default); previously set oversized values are silently ignored on upgrade.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 25, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue