THE CYBER TOOLCHAIN

NO. 086

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 26, 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 232 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 49 tools

▾ // THE BRIEF

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

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

Three unrelated things stand out. ConductorOne put secrets under the same access plane it already governs, with inventory, zero-knowledge

vim ? help

sharing and tripwire decoys that fire when someone reads a credential they shouldn't. Tailscale's PAM beta brokers SSH, database and Kubernetes sessions against identity instead of standing bastion credentials. Hashcat 7.1.2 adds four attack modes, including per-account association cracking.

PROTECT

Grant a database or SSH session against an identity instead of a shared credential on a bastion

Tailscale's PAM beta puts SSH, database and Kubernetes access behind identity-aware connectors, which is the thing jump hosts and shared service accounts exist to fake; ConductorOne's secrets suite adds inventory and zero-knowledge sharing so credentials stop living in tickets and chat, plus decoy secrets that alert on access. Aembit covers the workload side with SAML-bound dynamic credentials and in-cluster mTLS issuance, removing the long-lived Kubernetes secret.

tailscale · ConductorOne · Aembit

PROTECT

Confine what an AI agent's tool calls can actually touch on the host

nono brokers filesystem, network and credential access per command and carries that brokering across chained tool calls, so an agent that gets prompt-injected into running something destructive hits a policy instead of your shell. Arcjet pairs Rego policy guards with OTel traces and audit evidence, which is what you need when an auditor asks what the agent did; Runlayer tightens MCP launcher and Shadow AI controls over the same surface.

nono · Arcjet · Runlayer

Does **Opus 5** read better? 👍 👎

OFFENSIVE SECURITY

◆ Exploitation & C2

Bugcrowd Researcher Platform

Feedback ⓘ

changelog · Sep 26, 2026 CHANGELOG ↗ RANK

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

> Bugcrowd Researcher Platform added Submission Collaborator, letting multiple researchers work together on a single bug submission.

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

01 **Submission Collaborator for shared submissions** NEW

33

Adds Submission Collaborator, allowing multiple researchers to collaborate on a single bug submission.

– Describes the feature but no mechanism, limits, or how to invoke it.

changelog-20260926-f01e5362

> changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE
>
> Bugcrowd Researcher Platform adds Submission Collaborator, enabling researchers to collaborate on submissions.

WAS THIS USEFUL? 👍 👎

adscan

Feedback ⓘ v13.1.0 NOTES ↗ RANK

Free Active Directory pentesting tool for Linux, macOS and Windows. Automates AD and LDAP enumeration, Kerberoasting, ASREPRoast, password spraying, ADCS/ESC1, DCSync, RBCD, gMSA, shadow credentials, NTLM relay and credential dumping across 104 techniques, mapping BloodHound-compatible attack paths to Domain Admin.

> ADscan v13.1.0 sharpens the honesty and prioritization of its AD attack-path reporting, adding tri-state ADCS outcome labeling, an exposure-ranked accounts view, Entra Connect Tier-0 classification, attack-path deduplication, and a reclassification of Shadow Credentials as a persistence indicator.

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

01 **Tri-state ADCS finding outcomes** IMPROVED 75

ADCS findings now report exactly one of three outcomes: validated (CA reachable and escalation completed end-to-end), data gap (weakness detected but CA unreachable, with the next step named), or directory hygiene (decommissioned CA still registered). An ESC8 finding additionally names the web-enrollment transport (HTTP or HTTPS) and the observed channel-binding state.

— Names ESC8, transport states and channel-binding, but no command surface v13.1.0

02 **Shadow Credentials reclassified as persistence indicator**

IMPROVED 75

Shadow Credentials (pre-existing msDS-KeyCredentialLink entries) are now presented as a persistence indicator rather than an active attack step. The operator panel groups and caps objects, surfaces Tier-0 assets first, and defaults the client finding to Medium severity, rising to higher severity only when the affected object is itself Tier-0.

— Names the attribute, severity default and escalation condition v13.1.0

03 **Most Exposed Accounts ranking block** NEW 60

A new Most Exposed Accounts block appears under the exposure headline in the report and platform, ranking accounts by number of compromise paths that reach them (and how many were validated end-to-end), to help prioritize Protected Users membership, tiering, credential rotation, and LAPS.

— Clear UI location and purpose, no exact navigation path given v13.1.0

THINNER COVERAGE BELOW

04 **Entra Connect sync-account Tier-0 classification** NEW 55

ADscan now identifies Microsoft Entra Connect / MSOL_* directory-synchronization accounts holding directory-replication rights and classifies them as Tier-0 control-plane assets in the report, coverage report, and platform.

— Names MSOL_* accounts and Tier-0 classification, no direct action step v13.1.0

05 **Attack-path fan-out deduplication** IMPROVED 55

Attack-path listings in the exposure report now group fan-out routes that reach the same target via the same technique from different footholds into a single finding that names every foothold, reducing dozens of near-duplicate rows to a handful of findings.

— Describes the deduplication mechanism but no config or command v13.1.0

v13.1.0 NOTES ↗ LATEST STABLE

ADscan v13.1.0 adds honest ADCS outcome labeling, a Most Exposed Accounts block, and Entra Connect sync-account Tier-0 classification.

└──▷ GET THIS VERSION

```
$ git clone --branch v13.1.0 https://github.com/ADScanPro/adscan.git
# already have the repo? check out this version:
$ git checkout v13.1.0
```

ADCS findings now report exactly one of three outcomes: validated (CA reachable and escalation completed end-to-end), data gap (weakness detected but CA unreachable, with the next step named), or directory hygiene (decommissioned CA still registered); an ESC8 finding additionally names the web-enrollment transport (HTTP or HTTPS) and the observed channel-binding state.

Adds a Most Exposed Accounts block under the exposure headline in the report and platform, ranking accounts by number of compromise paths that reach them (and how many were validated end-to-end), to prioritize Protected Users membership, tiering, credential rotation, and LAPS.

Identifies Microsoft Entra Connect / MSOL_* directory-synchronization accounts holding directory-replication rights and classifies them as Tier-0 control-plane assets in the report, coverage report, and platform.

Attack-path listings in the exposure report now group fan-out routes that reach the same target via the same technique from different footholds into a single finding that names every foothold, reducing dozens of near-duplicate rows to a handful of findings.

Shadow Credentials (pre-existing msDS-KeyCredentialLink entries) are now presented as a persistence indicator rather than an active attack step; the operator panel groups and caps objects, surfaces Tier-0 assets first, and defaults the client

> finding to Medium severity (rising to higher severity only when the affected object is itself Tier-0).

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

EFF Rayhunter

Feedback ⓘ v0.13.0 NOTES ↗ RANK

A tool that detects cell site simulators (IMSI catchers) on a cheap mobile hotspot.

> Rayhunter v0.13.0 ships a graphical installer, adds NDJSON to ZIP exports, introduces a no-NAS analyzer, and sharpens its IMSI-catcher detection heuristics.

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

01 **Improved IMSI-catcher detection heuristics** IMPROVED 60

Reads the home PLMN from the SIM card to improve accuracy of the `imsi_requested` heuristic, and re-enables the `RRCConnectionRelease` classifier for detection coverage.

– Names both heuristics but no usage guidance v0.13.0

THINNER COVERAGE BELOW

02 **Graphical installer GUI** NEW 50

Adds a graphical installer-gui frontend (MVP) so users can install Rayhunter without the command line, and exposes CLI help text and default values inside the installer GUI.

– Names the GUI but no install steps or commands v0.13.0

03 **NDJSON in daemon ZIP export** IMPROVED 50

The daemon's ZIP download now includes NDJSON alongside the existing files.

– Names format and location but no schema detail v0.13.0

04 **No-NAS analyzer and alert** NEW 45

Adds a no-NAS analyzer and a corresponding UI alert to flag when NAS (Non-Access Stratum) traffic is absent.

– Explains what it detects but not how to enable it v0.13.0

05 **Clock drift warning control** IMPROVED 35

Adds an option to disable or autofix clock drift warnings in the daemon.

– No exact flag or config key named v0.13.0

v0.13.0 NOTES ↗ LATEST STABLE

Rayhunter v0.13.0 adds a graphical installer, NDJSON in ZIP exports, a no-NAS analyzer, and smarter IMSI heuristics.

└──▷ **GET THIS VERSION**

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

Includes NDJSON alongside existing files in the ZIP download from the daemon.

Adds a graphical installer-gui frontend (MVP) so users can install Rayhunter without the command line.

Exposes CLI help text and default values inside the installer GUI.

Adds an option to disable or autofix clock drift warnings in the daemon.

Adds a no-NAS analyzer and a corresponding UI alert to flag when NAS (Non-Access Stratum) traffic is absent.

+2 more

WAS THIS USEFUL? 👍 👎

ESP32 Marauder

Feedback ⓘ v1.17.0 NOTES ↗ RANK

A suite of WiFi/Bluetooth offensive and defensive tools for the ESP32.

> ESP32 Marauder v1.17.0 adds multiple saved WiFi profiles with scan-first selection and configurable wardrive geofences.

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

01 **Multiple saved WiFi profiles with scan-first selection** NEW

35

The device can now store multiple WiFi profiles and, when scanning, select from these stored network profiles rather than a single saved network.

— Describes behavior but no config keys or UI path v1.17.0

02 **Configurable wardrive geofences** NEW

25

Adds configurable geofences to restrict or define geographic boundaries for wardriving sessions.

— Names the feature but no mechanism or configuration detail v1.17.0

v1.17.0 NOTES ↗ LATEST STABLE

ESP32 Marauder v1.17.0 adds multiple saved WiFi profiles with scan-first selection and configurable wardrive geofences.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v1.17.0
https://github.com/justcallmekoko/ESP32Marauder.git
# already have the repo? check out this version:
$ git checkout v1.17.0
```

Adds support for multiple saved WiFi profiles with scan-first selection, allowing the device to scan and choose from stored network profiles.

Adds configurable wardrive geofences to restrict or define geographic boundaries for wardriving sessions.

WAS THIS USEFUL? 👍 👎

Hak5 WiFi Pineapple Pager

Feedback ⓘ 1.1.2-stable

NOTES ↗ RANK

A pocket sized WiFi Pineapple built for those on the go.

> WiFi Pineapple Pager 1.1.x adds payload-driven display takeover, new DuckyScript prompt commands, a WPA3-effective CSA deauth method with a selection workflow, PineAP recon improvements, and battery/power handling changes.

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

01 UI_TAKEOVER and UI_RELEASE DuckyScript commands NEW

80

Adds `UI_TAKEOVER` and `UI_RELEASE` DuckyScript commands to let a payload or application take over the framebuffer (display) and consume input (button presses) without shutting down the pineapplepager service.

— Named commands with clear mechanism and effect 1.1.2-stable

02 QR and image display DuckyScript commands NEW

60

Adds `PROMPT_QR` DuckyScript command to generate and display a QR code with a message to the user, and `PROMPT_IMAGE` to show an image on the screen.

— Named commands but no further mechanism given 1.1.1-stable

03 CSA deauth method and mode selection workflow NEW

60

Adds a CSA (Channel Switch Announcement) Deauthentication method, effective against some clients even when connected via WPA3, along with a deauth mode selection workflow letting operators choose between deauth methods.

— Names method and workflow, no exact UI path 1.1.1-stable

THINNER COVERAGE BELOW

04 Battery and power button handling improvements IMPROVED

50

Adds a safe shutdown procedure triggered on last-percent battery detection instead of allowing improper shutdown, and changes power button behavior to single-press sleep and wake rather than requiring selection from the power menu.

— Describes trigger and behavior change, no config exposed 1.1.1-stable

05 PineAP recon enhancements: MFP detection and censor mode NEW

Adds MFP (Management Frame Protection) detection to PineAP recon result models, and a censor mode toggle to the recon settings menu.

— Two thin one-line additions to recon settings 1.1.1-stable

1.1.2-stable NOTES ↗ LATEST STABLE

WiFi Pineapple Pager 1.1.2 adds `UI_TAKEOVER` and `UI_RELEASE` DuckyScript commands for payload-driven framebuffer control.

Adds `UI_TAKEOVER` and `UI_RELEASE` DuckyScript commands to let a payload or application take over the framebuffer (display) and consume input (button presses) without shutting down the pineapplepager service.

1.1.1-stable NOTES ↗ STABLE

WiFi Pineapple Pager 1.1.1 adds QR/image DuckyScript commands, CSA deauth for WPA3 clients, and MFP detection in PineAP recon.

Adds `PROMPT_QR` DuckyScript command to generate and display a QR code with a message to the user.

Adds `PROMPT_IMAGE` DuckyScript command to show an image on the screen.

Adds MFP (Management Frame Protection) detection to PineAP recon result models.

Adds CSA (Channel Switch Announcement) Deauthentication method, effective against some clients even when connected via WPA3.

Adds deauth mode selection workflow, allowing operators to choose between deauth methods.

+3 more

WAS THIS USEFUL? 👍 👎

DarkFlippers Unleashed Firmware

Feedback ⓘ

unlshd-093 NOTES ↗ RANK

Flipper Zero Unleashed Firmware, an unlocked alternative firmware.

> Unleashed unlshd-093 expands SubGHz protocol coverage with Nice O-Code and Security+ 2.0 keypad support (each backed by a dedicated recovery/PIN app), adds new KeeLoq manufacturers and the Superrollo roller-shutter protocol, reworks the Desktop menu with a second control page and selectable style themes, and lets users save signals directly from the Infrared Universal Remote and MIFARE Classic keys to a shared dictionary.

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

01 **Desktop menu customization: second page and style themes** NEW 80

Adds a second page to the Desktop up-button menu, accessible via left/right, with screen brightness, volume, and vibro controls. Adds `Settings -> Desktop -> Menu Style` to choose from Default, PS4, Wii, DSi, C64, Compact, CoverFlow, Vertical, and Grid main menu layouts, each shipped as its own plugin.

— Exact settings path and options given unlshd-093

02 **Nice O-Code SubGHz protocol support** NEW 76

Adds SubGHz Nice O-Code protocol support (52, 72 bit, Dynamic) with decode, save, and emulate, plus a new Nice O-Code app that recovers the installer code from 4+ captures.

— Names app and recovery mechanism, no exact menu path unlshd-093

03 **Security+ 2.0 keypad protocol support** NEW 76

Adds SubGHz Security+ 2.0 keypad support (86 bit, Dynamic) with decode, save, and emulate keypad frames, plus a new Security+ PIN app to test and change the keypad PIN in file.

— Names app and behaviour, no exact menu path unlshd-093

04 **Save signals from Infrared Universal Remote brute force** NEW

70

Adds the ability to save a working signal straight from the Infrared Universal Remote — pause brute force and press down to keep the signal as a new remote or append it to an existing one.

— Exact interaction steps given unlshd-093

05 **Save Keys to Dictionary for MIFARE Classic** NEW

70

Adds a Save Keys to Dictionary action on NFC read or saved MIFARE Classic cards, making keys found by any attack available to future reads and to NFC Magic.

— Named action, no exact menu navigation unlshd-093

06 **Superrollo protocol and new KeeLoq manufacturers** NEW

60

Adds SubGHz Superrollo (GW60) roller-shutter protocol (KeeLoq HCS361, 67-bit rolling code with CRC), including Add Manually support, plus new KeeLoq manufacturers JCM Tech Gen2, Stagnoli, Telcoma, SEA, and two additional Wisniowski variants.

— Lists protocols and manufacturers but no usage steps unlshd-093

unlshd-093 NOTES ↗ LATEST STABLE

Unleashed unlshd-093 adds Nice O-Code and Security+ 2.0 keypad support, new menu styles, and IR signal saving from Universal Remote.

└──▷ GET THIS VERSION

```
$ git clone --branch unlshd-093
https://github.com/DarkFlippers/unleashed-firmware.git
# already have the repo? check out this version:
$ git checkout unlshd-093
```

Adds SubGHz Nice O-Code protocol support (52, 72 bits, Dynamic) — decode, save, and emulate — plus a new Nice O-Code app that recovers the installer code from 4+ captures.

Adds SubGHz Security+ 2.0 keypad support (86 bits, Dynamic) — decode, save, and emulate keypad frames — plus a new Security+ PIN app to test and change the keypad PIN in file.

Adds SubGHz Superrollo (GW60) roller-shutter protocol (KeeLoq HCS361, 67-bit rolling code with CRC, including Add Manually support).

> Adds new KeeLoq manufacturers: JCM Tech Gen2, Stagnoli, Telcoma, SEA, and two additional Wisniowski variants.
>
> Adds a second page to the Desktop up-button menu, accessible via left/right, with screen brightness, volume, and vibro controls.
>
> +3 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

ZeroPath

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

ZeroPath is an AI-native SAST platform that finds 2x more real vulnerabilities with 75% fewer false positives.

> ZeroPath aligned its Issues API and UI to include dependency findings by default, a change integrations must account for.

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

01 **Dependency findings included by default in Issues API and UI** 71

BREAKING

The `issues.list`, `issues.search`, and MCP `issues_list` tool now include dependency findings when `detectionTypes` is omitted, aligning API defaults with the full open-findings set; integrations relying on the previous code-only default must now pass `detectionTypes` explicitly to preserve prior behavior. The Issues page `All` tab now also counts and lists dependency findings alongside code findings, matching the dashboard's `Total Open` card and repositories list.

— Names exact APIs, param, and UI elements; migration note included.

changelog-20260926-96035cdb

└──▷ BREAKING ON UPGRADE

! The `issues.list` and `issues.search` APIs now return dependency findings in addition to code findings when `detectionTypes` is omitted. Integrations that relied on the previous code-only default must pass `detectionTypes` explicitly to preserve their prior behavior.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

ZeroPath Issues API now includes dependency findings by default; Issues page `All` tab unified with dashboard totals.

The `issues.list`, `issues.search`, and MCP `issues_list` tool now include dependency findings when `detectionTypes` is omitted, aligning API defaults with the full open-findings set.

The Issues page `All` tab now counts and lists dependency findings alongside code findings, matching the dashboard's `Total Open` card and repositories list.

└──▷ BREAKING ON UPGRADE

! The `issues.list` and `issues.search` APIs now return dependency findings in addition to code findings when `detectionTypes` is omitted. Integrations that

relied on the previous code-only default must pass `detectionTypes` explicitly to preserve their prior behavior.

WAS THIS USEFUL? 👍 👎

StackHawk

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

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

> StackHawk shipped a new way to disable session verification with a required fallback check, clearer CLI version-support errors with quieter deprecation logging, and a fix preserving key HTTP headers on rescan replays.

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

01 Session verification disable flag in stackhawk.yml NEW

90

Adds `disabled: true` inside `app.authentication.sessionVerification` in `stackhawk.yml` to explicitly turn off session verification while still requiring a `loggedInIndicator` / `loggedOutIndicator` fallback check. Session verification can also be tuned with `pollFrequency`, `activeScanPollFrequency`, and `maxUnauthenticatedRequestPercent` to control how often HawkScan re-authenticates if it drifts into an unauthenticated state during an active scan.

Enable mid-scan session verification so HawkScan re-authenticates automatically if it drifts into an unauthenticated state during an active attack, keeping finding coverage high.

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

— Full config keys and behavior with a runnable yaml example. changelog-20260926-a9739752

02 Clearer CLI version-support errors and quieter logging

IMPROVED

80

HawkScan now reports a clear, non-zero-exit error — including the minimum supported version and upgrade location — when the platform rejects a scan due to an out-of-support CLI version, and logs deprecation advisories once per scan instead of on every platform call. CLI installs can be pinned and localized in CI using the `HAWK_INSTALL_VERSION` and `HAWK_INSTALL_DIR` environment variables.

Pin a specific HawkScan CLI version in CI so the pipeline is reproducible, and install it to a project-local directory instead of the user home.

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

— Behavior explained with a runnable pinned-install command. changelog-20260926-a9739752

03 **Preserved HTTP headers on rescan replay** IMPROVED 75

Entity and transport headers — Content-Type, Content-Length, Accept, Accept-Encoding, Host, User-Agent, Referer — are no longer stripped from replayed spider requests during rescans, even when listed in `authentication.testPath.requestHeaders` or `app.redact.headers`.

— Names exact headers and config fields but no usage example given.

changelog-20260926-a9739752

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

StackHawk: Adds `disabled: true` field inside `app.authentication.sessionVerification` to explicitly turn off session verification while still requiring a `loggedInIndicator` / `loggedOutIndicator` as a fallback check.

└──▷ **USE IT**

Enable mid-scan session verification so HawkScan re-authenticates automatically if it drifts into an unauthenticated state during an active attack, keeping finding coverage high.

stackhawk.yml

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

Pin a specific HawkScan CLI version in CI so the pipeline is reproducible, and install it to a project-local directory instead of the user home.

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

Reports a clear, non-zero-exit error — including the minimum supported version and upgrade location — when the platform rejects a scan due to an out-of-support CLI

version, and logs deprecation advisories once per scan instead of on every platform call.

Entity and transport headers — `Content-Type`, `Content-Length`, `Accept`, `Accept-Encoding`, `Host`, `User-Agent`, `Referer` — are no longer stripped from replayed spider requests during rescans, even when listed in `authentication.testPath.requestHeaders` or `app.redact.headers`.

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

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

> Semgrep added CSV export for Agentic Workflows issues and a navigation banner pointing to the latest run of a workflow.

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

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

Semgrep added CSV export for Agentic Workflows issues, available from the Issues page and through the public API.

— Names export format and two access surfaces, no endpoint detail

changelog-20260926-0de72f2b

02 **Banner linking to newest workflow run** IMPROVED 35

A banner now appears on older Agentic Workflows runs that links to the newest run of the same workflow on that branch.

— Describes behaviour but no UI path or config changelog-20260926-0de72f2b

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Semgrep: Adds CSV export for Agentic Workflows issues from the Issues page and through the public API.

Adds a banner on older Agentic Workflows runs that links to the newest run of the same workflow on that branch.

WAS THIS USEFUL? 👍 👎

GitHub

Feedback ⓘ

Enterprise managed settings in-product validator (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 shipped a wave of Copilot expansion this window — new AI models, JetBrains and VS Code agent session upgrades, and richer Slack/Teams integration — alongside platform changes including a managed-settings validator, pull-request review-stage timing in the usage metrics API, GA Relates to issue relationships, private saved views for issues, and a new default-enablement policy governing future Copilot features.

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

01 In-product validator for Copilot managed settings NEW

90

GitHub Copilot enterprise managed settings now include an in-product validator in the Copilot settings validation section of the enterprise AI controls page that catches malformed JSON, unsupported configurations, and invalid team mappings across `copilot/managed-settings.json`, `copilot/team-mappings.json`, and referenced team settings files.

After editing Copilot managed settings, confirm your configuration is valid before policies take effect.

> 📍 1. In your GitHub Enterprise account, go to Settings › Copilot › Enterprise AI controls. 2. Scroll to the Copilot settings validation section. 3. Review any listed issues — each shows the affected file (`copilot/managed-settings.json` or `copilot/team-mappings.json`) and the exact JSON path. 4. Fix the identified errors and commit to the default branch of your `.github-private` repository. 5. Reload the Agents page and re-check the validator to confirm no issues remain.

— Names exact files, UI path, and remediation steps

Enterprise managed settings in-product …

02 Copilot in Slack and Microsoft Teams gains context, model, and issue-detection features

NEW

90

Copilot in Slack and Microsoft Teams now supports switching models mid-conversation, with the choice persisting for the rest of the thread; setting per-channel model defaults

in Slack; setting default owners and repositories in Slack so Copilot knows where to create work without per-message prompting; duplicate-issue detection that checks for similar existing issues before creating a new one and includes direct links to created work and back to the originating conversation; sharing files, attachments, and message links as context in Slack; and using inline images, forwarded-message context, and channel and thread history as context in Microsoft Teams.

.

Route all Copilot-created issues and PRs to the right repo by default, so teammates don't have to specify it on every request in a shared Slack channel.

> 📍 In Slack, open the GitHub app settings and set `Default owner` and `Default repository` for your workspace, then mention @GitHub in any channel to create work without specifying the repo each time.

Give Copilot richer context in a Teams thread by forwarding a message from another channel before asking it to draft an issue.

> 📍 In Microsoft Teams, forward the relevant message into the current thread, then mention @GitHub and ask it to create an issue — Copilot will use the forwarded-message context and thread history in its response.

— Enumerates every named addition with steps for two

Updates to GitHub Copilot for Slack and… | GitHub Copilot weekly releases — Septem…

---

03 **Pull request review stage timings in usage metrics API** NEW

85

The enterprise and organization Copilot usage metrics API adds a `pull_request_review_times` array to each `repos-1-day` row, reporting `median_minutes_ready_to_first_review`, `p90_minutes_ready_to_first_review`, `median_minutes_first_to_final_review`, `p90_minutes_first_to_final_review`, `median_minutes_final_review_to_merge`, and `p90_minutes_final_review_to_merge`, all attributed to the day the pull request merged. Each entry also includes `authored_by`, `reviewed_by` (both human in this release), and `total_merged`, and is available to enterprise owners, billing managers, organization owners, and custom roles with the `View Copilot Metrics Permission` when the usage metrics policy is enabled; it returns `[]` on days with no qualifying pull requests.

— Full field list and permissions, no example request shown

Usage metrics API adds pull request rev…

---

04 **VS Code Copilot session management improvements** NEW

70

VS Code 1.139 enables running Copilot agents in Dev Containers on SSH, Tunnel, and WSL hosts using the remote project's tools and dependencies; adds a Compact View for session lists with filters to hide empty groups and the ability to rename sessions and chats directly in the list; and adds a preview layout option to choose between separate chat tabs or a single active-chat view within a session.

— Names VS Code version and specific UI features GitHub Copilot weekly releases — Septem…

05 **JetBrains Copilot agent session improvements** NEW 65

Copilot agent sessions in JetBrains add assisted approvals (public preview), automatically approving low-risk tool calls while prompting for higher-risk actions; message-editing to redirect agent sessions, where Copilot rewinds both the conversation and file changes before sending the replacement request; and shared organization and enterprise skills plus organization-managed custom instructions in local and agent sessions.

— Explains mechanism for each JetBrains addition GitHub Copilot weekly releases — Septem…

THINNER COVERAGE BELOW

06 **Private saved views for repository issues** NEW 55

Users can now create and save personalized private filtered views on repository issues pages (e.g., by assignee, label, or milestone) to return to quickly.

▶ Watch video

Save a personal filtered view of open issues assigned to you on a repository issues page so you can return to it quickly.

📍 1. Navigate to your repository's Issues tab. 2. Apply your desired filters (e.g., assignee, label, milestone). 3. Click Save view and select Private to store the view visible only to you.

— Includes UI steps but no config or API detail Private saved views for repository issu…

07 **New AI models added across Copilot plans** NEW 55

Copilot adds `Claude Opus 5.5` to Pro+, Max, Business, and Enterprise plans; GPT-6 Sol to Pro+, Max, Business, and Enterprise plans; GPT-6 Luna to Pro, Pro+, Max, Business, and Enterprise plans; and `Grok 4.7` to Pro, Pro+, Max, Business, and Enterprise plans.

— Names each model and plan tier, no usage steps GitHub Copilot weekly releases — Septem…

08 **'Relates to' issue relationship reaches general availability**

IMPROVED 45

The `Relates to` issue relationship is now generally available across the REST API, GraphQL API, webhooks, timeline events, issues search, and projects search.

— Lists surfaces but no endpoints or fields given

`Private saved views for repository issu…`

09 **OpenTelemetry support in the Copilot app** NEW 45

The GitHub Copilot app adds OpenTelemetry support to track agent activity in existing monitoring tools, configured through enterprise-managed settings.

— Names OpenTelemetry but no config keys given `GitHub Copilot weekly releases — Septem…`

└──▷ **BREAKING ON UPGRADE**

! Starting October 22, any eligible generally available Copilot feature left `Unconfigured` will automatically follow the global default policy — meaning features that were previously unconfigured (and therefore inaccessible) may become enabled enterprise-wide if the default is set to `Enabled`.

`Enterprise managed settings in-product validator (2026-09-25)` BLOG ↗

LATEST STABLE

GitHub Copilot enterprise managed settings now include an in-product validator that catches malformed JSON, bad team mappings, and unsupported configs.

└──▷ **HOW TO FIND IT**

After editing Copilot managed settings, confirm your configuration is valid before policies take effect.

📍 1. In your GitHub Enterprise account, go to Settings › Copilot › Enterprise AI controls. 2. Scroll to the `Copilot settings validation` section. 3. Review any listed issues — each shows the affected file (`copilot/managed-settings.json` or `copilot/team-mappings.json`) and the exact JSON path. 4. Fix the identified errors and commit to the default branch of your `.github-private` repository. 5. Reload the Agents page and re-check the validator to confirm no issues remain.

`Usage metrics API adds pull request review stages (2026-09-25)` BLOG ↗

STABLE

Copilot usage metrics API gains `pull_request_review_times` array with per-stage PR review durations at median and p90.

Adds `pull_request_review_times` array to each `repos-1-day` row in the enterprise and organization Copilot usage metrics API, reporting `median_minutes_ready_to_first_review`, `p90_minutes_ready_to_first_review`, `median_minutes_first_to_final_review`,

`p90_minutes_first_to_final_review`, `median_minutes_final_review_to_merge`, and `p90_minutes_final_review_to_merge` — all durations in minutes, attributed to the day the pull request merged.

Each `pull_request_review_times` entry includes `authored_by`, `reviewed_by` (both human in this release), and `total_merged` to identify which pull requests contribute to each timing row.

Available to enterprise owners, billing managers, organization owners, and custom roles with the View Copilot Metrics Permission when the Copilot usage metrics policy is enabled; returns `[]` on days with no qualifying pull requests.

Private saved views for repository issues and “Relates to” issue relationship is gener

BLOG ↗ STABLE

GitHub adds private saved views for repository issues and GA Relates to issue relationships across APIs, webhooks, and search.

└──▷ **HOW TO FIND IT**

Save a personal filtered view of open issues assigned to you on a repository issues page so you can return to it quickly.

📍 1. Navigate to your repository's Issues tab. 2. Apply your desired filters (e.g., assignee, label, milestone). 3. Click Save view and select Private to store the view visible only to you.

Adds Relates to issue relationship support to the REST API, GraphQL API, webhooks, timeline events, issues search, and projects search.

Adds private saved views for repository issues pages, letting users create and save personalized issue filters.

GitHub Copilot weekly releases — September 21 (2026-09-25) BLOG ↗ STABLE

Copilot gains Claude Opus 5.5, GPT-6, Grok 4.7, local sandboxing with OpenTelemetry, and new Slack/JetBrains/VS Code capabilities.

Adds local sandboxing in the GitHub Copilot app (public preview) to limit agents' access to files, networks, and credentials.

Adds OpenTelemetry support in the GitHub Copilot app to track agent activity in existing monitoring tools, configured through enterprise-managed settings.

Adds Claude Opus 5.5 model to Copilot Pro+, Max, Business, and Enterprise plans.

Adds GPT-6 Sol model to Copilot Pro+, Max, Business, and Enterprise plans.

Adds GPT-6 Luna model to Copilot Pro, Pro+, Max, Business, and Enterprise plans.

+11 more

Updates to GitHub Copilot for Slack and Microsoft Teams (2026-09-25) BLOG ↗

STABLE

GitHub Copilot in Slack and Teams gains file/image context, model switching, and default owner/repo settings.

└──▷ **HOW TO FIND IT**

Route all Copilot-created issues and PRs to the right repo by default, so teammates don't have to specify it on every request in a shared Slack channel.

> 📍 In Slack, open the GitHub app settings and set `Default owner` and `Default repository` for your workspace, then mention @GitHub in any channel to create work without specifying the repo each time.

Give Copilot richer context in a Teams thread by forwarding a message from another channel before asking it to draft an issue.

> 📍 In Microsoft Teams, forward the relevant message into the current thread, then mention @GitHub and ask it to create an issue — Copilot will use the forwarded-message context and thread history in its response.

Adds support for setting default owners and repositories in Slack so Copilot knows where to create work without per-message prompting.

Supports inline images, forwarded-message context, and channel and thread history as context in Microsoft Teams.

Enables switching models mid-conversation for the next message, with the choice persisting throughout that conversation.

Copilot now checks for similar existing issues before creating a new one, and includes direct links to created work and back to the originating conversation.

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ

Client logs now show the full request chain behind a package download (2026-09-25)

BLOG ↗ RANK

A fully-managed, enterprise-scale solution to control, secure, and distribute software packages and containers.

> Cloudsmith client logs now capture format-specific metadata across all supported package formats, giving visibility into the full request chain behind a download, not just the direct download itself.

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

01 **Full request chain visibility in client logs** IMPROVED

50

Client logs in the Cloudsmith web app now capture format-specific metadata across every supported package format, exposing package metadata fetches alongside direct package downloads for full traffic visibility.

.
.
.
.

— Explains what changed but no specific UI path or config.

Client logs now show the full request c…

> Client logs now show the full request chain behind a package download (2026-09-25)
>
> BLOG ↗ LATEST STABLE
>
> Cloudsmith client logs now capture format-specific metadata for all package formats, revealing full request chains beyond direct downloads.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

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

> Veracode's CLI gained a large batch of new commands (`veracode fix sca` , `veracode package discover` , `veracode auth login`) and packaging/scanning flags this window, while the platform added a Custom PDF REST API, expanded Reporting API fields, Manual Penetration Testing flaw matching, and multi-team DAST targets, alongside a breaking removal of the legacy healthcheck endpoint.

└──▷ **WHAT SHIPPED · 21 FEATURES** most completely described first ? what's the number?

01 veracode package command flags and exit codes IMPROVED

85

`veracode package` gains `--exclude` to prevent specified Java packages or JAR files from being scanned, `--skip-packagers` to exclude specific artifacts (e.g. unsupported SQL) from packaging results without affecting other languages, and `--strict` which returns exit code 4 for any build failures during packaging; new exit codes 5, 6, and 7 were also introduced to improve CI/CD pipeline automation and error handling.

Package a Java project excluding a specific JAR that causes false positives or is unsupported.

```
$ veracode package --exclude com.example:legacy-lib ./myapp
```

Enable strict mode in CI so any build failure during packaging immediately surfaces as exit code 4 rather than a silent partial result.

```
$ veracode package --strict ./myapp
```

— Names exact flags and exit codes with runnable examples CLI updates (2026-09-23)

02 Reporting API filters and fields IMPROVED

80

Adds `most_recent_scan` filter to the Reporting API Scans report (returns only the most recent scan per asset when `true`), `custom_severity` filter to the Reporting API Findings report, `collections_list` field to the Finding Reporting API response, `library_instance_ref` field to uniquely identify the library instance associated with a finding, `data_refresh_date` field indicating the most recent Analytics data refresh

time, `action_type` filter for the Audit Log report, and a new Audit Log report to the Reporting REST API for querying detailed user activity information.

— Names every field and filter added to the Reporting API `Platform updates (2026-09-24)`

03 veracode scan enhancements: policy, vulnerability data, and results

IMPROVED

75

`veracode scan` now supports applying a custom policy defined in the Veracode Platform and viewing policy evaluation for individual findings, specifying a custom location from which vulnerability databases can be downloaded, parsing Dockerfile heredoc syntax, categorizing secret-type findings as Critical, prioritizing vulnerability results from the NVD database when available, creating a separate asset per repository in ADO workflows, and schema enhancements for container and IaC scan results.

— Names seven distinct scan behaviors, no runnable flags given `CLI updates (2026-09-23)`

04 Package language and toolchain support expansion IMPROVED

75

Adds `VERACODE_PACKAGE_GOLANG_GENERATE` environment variable to enable client-side CGO code generation when packaging Go applications, `SRCLR_MAKE_WINDOWS` environment variable to enable preprocessing C/C++ source files with Make on Windows using a native Windows gcc or clang compiler, support for Go v1.24 in `veracode package`, support for the Poetry package management tool and `requirements-dev.txt` in Python autopackaging, support for preprocessing C/C++ source code on Linux via `veracode package` (replacing the existing Linux autopackaging solution), and support for preprocessor header files (PCH) and processing all Makefiles in a directory tree.

— Names all env vars and language versions, no runnable examples `CLI updates (2026-09-23)`

05 Healthcheck endpoint removed BREAKING

75

The Veracode Healthcheck endpoint `/healthcheck/status` has been removed and now returns 404; use the `v2/principal` endpoint in the Identity REST API instead.

— Names old and new endpoints with clear migration path `Platform updates (2026-09-24)`

06 SCA dependency fix generation via veracode fix sca NEW

70

Adds `veracode fix sca` command to generate fixes for vulnerable dependencies identified in SCA scans; currently Limited Availability.

Generate SCA-based dependency fixes for a JavaScript project after running an SCA scan.

```
$ veracode fix sca
```

— Named command with runnable example, but no mechanism detail CLI updates (2026-09-23)

07 **veracode package discover toolchain detection** NEW

70

Introduces `veracode package discover` command to detect toolchains required to build scannable artifacts, including detection for Java projects using Gradle and JavaScript projects using npm or Yarn, plus new exit codes for `veracode package discover` to support automation.

— Names command and target ecosystems, no example given CLI updates (2026-09-23)

08 **SBOM generation via veracode sca** NEW

65

Adds SBOM generation via `veracode sca` agent-based scans, with output in CycloneDX 1.6 and SPDX 2.3 formats.

— Names command and both output formats CLI updates (2026-09-23)

09 **Custom PDF REST API and reporting content** NEW

60

Adds a Custom PDF REST API to programmatically generate custom PDF reports; Customizable PDF now also includes SCA Component License Risk policy compliance in the Executive Summary and Detailed SCA Findings by Component sections.

— Names the API and specific report sections it added Platform updates (2026-09-24)

10 **Veracode Analytics platform upgrades** IMPROVED

60

Adds a Firewall Usage Dashboard and Firewall Usage Explore for metrics on Veracode Firewall, adds Collections fields to the Applications explore and new Collections visualizations to the Policy Compliance Overview dashboard, upgrades Veracode Analytics to Looker 26.6 (adding finer widget size granularity, persistent filtering, and tabbed dashboards), and reduces Veracode Analytics data refresh time from four hours to two hours.

— Names dashboards, Looker version, and refresh time change Platform updates (2026-09-24)

11 **Manual Penetration Testing flaw matching** NEW

60

Adds Manual Penetration Testing (MPT) Flaw Matching to correlate findings across tests, providing consistent flaw IDs, status values, and inclusion in remediation grace periods.

— Explains mechanism but no UI/API entry point given Platform updates (2026-09-24)

12 **Multi-team DAST targets** NEW

60

Enables assigning multiple teams to a DAST target when creating a web application or API specification, so more than one team can share ownership of scan findings.

Assign two teams to a DAST web application target so both AppSec and DevOps teams share ownership of scan findings.

> 📍 In the Veracode Platform, go to DAST › Targets › Create Web Application. Under `Teams`, select multiple teams from the dropdown (e.g., `AppSec` and `DevOps`), then save the target.

— Includes UI navigation steps to use the feature `changelog-20260926-ee875fb6`

THINNER COVERAGE BELOW

13 **`veracode auth login command`** NEW 55

Introduces `veracode auth login` command to onboard new users via OAuth2.

— Named command but no example or further mechanism `CLI updates (2026-09-23)`

14 **`OAuth Client Credentials authentication`** NEW 55

Adds OAuth Client Credentials authentication support for Veracode XML APIs, in addition to existing REST API support, and adds support for creating and managing OAuth 2.0 Client credentials on the API Credentials page in the Veracode Platform.

— Names both surfaces gaining OAuth support `Platform updates (2026-09-24)`

15 **`Container and IaC scan results in the Platform`** NEW 55

Adds container and IaC scan results from CLI scans to the Veracode Platform, including severity counts, filtering, scan history, and mitigation guidance.

— Describes what surfaces but not how to view them `Platform updates (2026-09-24)`

16 **`Identity and SSO enhancements`** IMPROVED 55

Enables mapping multiple domains to a single identity provider (IdP) for organizations using SP-initiated SAML SSO, and allows disabling Veracode-defined roles for users who use custom roles, via the Identity API.

— Names Identity API and SSO scenario, no exact call shown `Platform updates (2026-09-24)`

17 **`Security policy standard updates`** IMPROVED 50

Updates the Auto-Update OWASP rule and Auto-Update CWE Top 25 rule to enforce OWASP Top 10:2025 and MITRE CWE Top 25 for 2025 standards, and adds six new built-in security policies to the Veracode Platform.

— Names the specific standards updated and policy count `Platform updates (2026-09-24)`

18 **GitHub Workflow Integration policy parameter for IaC scans**

IMPROVED

50

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

— Names parameter and integration version, thin on mechanism changelog-20260926-ee875fb6

19 **IDE and CLI Usage dashboard, now in Europe** NEW

45

Introduces the IDE and CLI Usage Explore and Dashboard in Veracode Analytics to visualize an organization's IDE and CLI plugin usage, and makes it available in the European Region.

— Names the dashboard and its regional rollout changelog-20260926-ee875fb6

Platform updates (2026-09-24)

20 **DryRun Security integration** NEW

45

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

— Describes what it does, no setup steps given changelog-20260926-ee875fb6

21 **veracode dynamic scan target UUID support** IMPROVED

40

Adds `veracode dynamic scan` support for target UUIDs when triggering scans; webhook UUIDs remain supported.

— Bare capability name with minimal mechanism CLI updates (2026-09-23)

└──▷ **BREAKING ON UPGRADE**

! The Veracode Healthcheck endpoint `/healthcheck/status` has been removed and now returns 404; use the `v2/principal` endpoint in the Identity REST API instead.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Veracode adds multi-team DAST targets, DryRun Security integration, iOS 27 and Go 1.27 static analysis, and European region IDE/CLI dashboards.

└──▷ **HOW TO FIND IT**

Assign two teams to a DAST web application target so both AppSec and DevOps teams share ownership of scan findings.

📍 In the Veracode Platform, go to DAST › Targets › Create Web Application. Under Teams, select multiple teams from the dropdown (e.g., AppSec and DevOps), then save the target.

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

Enables assigning multiple teams to a DAST target when creating a web application or API specification.

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

Makes the IDE and CLI Usage Explore and Dashboard available in the European Region.

Adds support for Jira 11 (including Java 21 compatibility) in Veracode Integration for Jira Server v5.0.2.

`Platform updates (2026-09-24)` **BLOG ↗** STABLE

Veracode adds Custom PDF REST API, OAuth for XML APIs, IDE/CLI Usage Analytics, and new Reporting API fields and filters

Adds a Custom PDF REST API to programmatically generate custom PDF reports.

Adds OAuth Client Credentials authentication support for Veracode XML APIs, in addition to the existing REST API support.

Adds `most_recent_scan` filter to the Reporting API Scans report; when set to `true`, returns only the most recent scan per asset.

Adds `custom_severity` filter to the Reporting API Findings report.

Adds `collections_list` field to the Finding Reporting API response, returning a list of collection names associated with the application.

`+17 more`

└─▷ **BREAKING ON UPGRADE**

! The Veracode Healthcheck endpoint `/healthcheck/status` has been removed and now returns 404; use the `v2/principal` endpoint in the Identity REST API instead.

`CLI updates (2026-09-23)` **BLOG ↗** STABLE

Veracode CLI gains SCA fix generation, new package discovery, auth login, SBOM export, and a raft of new flags across scan and package commands.

└─▷ **TRY IT**

Generate SCA-based dependency fixes for a JavaScript project after running an SCA scan.

```
$ veracode fix sca
```

Package a Java project excluding a specific JAR that causes false positives or is unsupported.

```
$ veracode package --exclude com.example:legacy-lib ./myapp
```

Enable strict mode in CI so any build failure during packaging immediately surfaces as exit code 4 rather than a silent partial result.

```
$ veracode package --strict ./myapp
```

Adds `veracode fix sca` command to generate fixes for vulnerable dependencies identified in SCA scans (Limited Availability).

Adds `--exclude` flag to `veracode package` to prevent specified Java packages or JAR files from being scanned.

Adds `--skip-packagers` flag to `veracode package` to exclude specific artifacts (e.g. unsupported SQL) from packaging results without affecting other languages.

Adds `VERACODE_PACKAGE_GOLANG_GENERATE` environment variable to enable client-side CGO code generation when packaging Go applications.

Adds `SRCLR_MAKE_WINDOWS` environment variable to enable preprocessing C/C++ source files with Make on Windows using a native Windows gcc or clang compiler.

+21 more

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Self-Hosted

Feedback ⓘ

changelog · Sep 26, 2026 CHANGELOG ↗ RANK

GitGuardian Self-Hosted is an on-premise secret detection tool that scans repositories and prevents sensitive credentials from being committed.

> GitGuardian Self-Hosted 2026.9 adds GitGuardian Bridge for scanning isolated network sources, honeytoken-based endpoint protection, expanded AI Hooks coverage, and ggscout-based vaulted-secret classification, alongside a ClickHouse chart restructuring, a Jira/Confluence PAT renewal removal, and an Argo CD proxy-secret change that all require action before upgrading.

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

01 **ClickHouse chart restructured under first-party key** BREAKING

95

ClickHouse is now rendered under `clickhouse.*` in the GitGuardian chart instead of the vendored `clickhouse-server` subchart, adding structured `clickhouse.objectStorage` and `clickhouse.backup` config keys; `clickhouse.serverConfig` keys switch to camelCase. Any `clickhouse-server.*` values from 2026.8 must be moved to `clickhouse.*` before upgrading — the chart refuses to render while a `clickhouse-server` block is present.

— Exact keys named with required migration steps. changelog-20260926-851756fd

2026.9 (2026-09-25)

02 **Custom PostgreSQL schema for in-app analytics** NEW

90

Adds `inAppAnalytics.postgresql.applicationSchema`, `useSeparateAnalyticsSchema`, and `analyticsSchema` Helm/config keys so in-app analytics can write to a custom PostgreSQL schema instead of `public`, for databases that forbid the `public` schema or cannot grant broad privileges to the GitGuardian user.

Route in-app analytics writes to a dedicated PostgreSQL schema when the GitGuardian database user cannot create objects in `public`.

```yaml
inAppAnalytics:
  postgresql:
    applicationSchema: gganalytics
useSeparateAnalyticsSchema: true
analyticsSchema: gganalytics
```

Write in-app analytics to a non-public PostgreSQL schema when the database forbids the `public` schema or cannot grant broad privileges to the GitGuardian user.

```yaml
inAppAnalytics:
  postgresql:
    applicationSchema: gitguardian_analytics
useSeparateAnalyticsSchema: true
analyticsSchema: gitguardian_analytics
```

— Exact config keys plus a runnable YAML example. 2026.9 (2026-09-25) changelog-20260926-851756fd

03 **Argo CD proxy secrets no longer auto-created** BREAKING

90

Argo CD installations that do not use an outbound proxy will see `1 resources require pruning` and stay `OutOfSync` after upgrade, because the chart no longer creates the `gim-proxy` and `gim-migration-proxy` Secrets when no proxy URL is set. Resolve by running `argocd app sync <app> --prune` or enabling automated pruning.

— Exact secret names and the fix command are given. changelog-20260926-851756fd 2026.9 (2026-09-25)

04 **GitGuardian Bridge for isolated network scanning** NEW

85

Adds `GGBridge.bridges` Helm values key to declare GitGuardian Bridge servers, enabling scanning of sources in isolated networks (e.g. a GitLab instance behind a firewall) without opening inbound connections.

— Names exact Helm key and concrete use case. 2026.9 (2026-09-25)

05 **Support bundle pod scheduling controls** NEW

85

Adds `replicated.supportBundle.nodeSelector`, `tolerations`, and `affinity` Helm keys to schedule the support bundle collection pod onto specific or tainted nodes

so it does not compete with workload pods.

Schedule the support bundle collection pod onto a tainted node so it does not compete with production workloads.

```yaml
replicated:
  supportBundle:
    nodeSelector:
      node-role: ops
    tolerations:
      - key: ops
        operator: Exists
        effect: NoSchedule
    affinity: {}
```

Pin the support bundle pod to a dedicated tainted node so bundle collection does not compete with workload pods.

```yaml
replicated:
  supportBundle:
    nodeSelector:
      node-role: support
    tolerations:
      - key: dedicated
        operator: Equal
        value: support
        effect: NoSchedule
    affinity: {}
```

— Exact Helm keys with a runnable scheduling example. 2026.9 (2026-09-25)

changelog-20260926-851756fd

06 **Vaulted-secret classification via ggscout** NEW 80

`ggscout` now runs the Secrets Detection Engine on every collected value, enabling NHI inventory filtering by detector, detector type, secret family, category, and provider; requires `ggscout` 0.32.0+ — earlier versions outside the cluster do not classify collected secrets with the Secrets Detection Engine.

— Names filtering fields and the required minimum version. 2026.9 (2026-09-25)

changelog-20260926-851756fd

07 **Jira/Confluence PAT renewal removed** BREAKING 80

The automatic Personal Access Token renewal system for Jira Data Center and Confluence Data Center integrations is removed. Before upgrading, manually update the token for every Jira Data Center and Confluence Data Center integration (sources and notifiers) with a long-lived PAT via the integration settings — an integration still relying on renewal will stop working when its token expires.

— Exact migration steps required before upgrade. changelog-20260926-851756fd

2026.9 (2026-09-25)

08 **AI Hooks extended to Codex and Mistral Vibe** IMPROVED

70

AI Hooks now blocks secrets inside Codex and Mistral Vibe, alongside existing Claude Code, Cursor, Copilot CLI, and VS Code support, and adds a new inventory of AI agents and MCP servers per endpoint; requires `ggshield` 1.54.0+.

— Names every supported tool and the version requirement. 2026.9 (2026-09-25)

09 **ClickHouse health and monitoring** NEW

70

Adds ClickHouse health coverage to the Admin Area Health check page, admin banner, and notification emails, including HTTP reachability, metadata volume fill level against a configurable threshold, and backup freshness from the backup sidecar. Also adds Prometheus metrics with recommended thresholds and alerts for ClickHouse and the backup sidecar, and includes their logs and collectors in the support bundle.

— Names monitoring surfaces but not exact metric names. 2026.9 (2026-09-25)

10 **Honeytoken protection for endpoints** NEW

65

`ggshield` now plants a decoy AWS credential in credential-harvesting malware read paths on endpoints, with fleet coverage tracked in the Endpoints dashboard; requires Endpoint Protection and ClickHouse.

— Names dashboard and requirements but no exact command. 2026.9 (2026-09-25)

11 **GitHub app installation approval requirement** IMPROVED

60

GitHub app installations made outside the dashboard (e.g. from the app's public GitHub page) now queue for Owner or Manager approval on the GitHub / GitHub Enterprise Server settings page instead of automatically joining the monitored perimeter.

— Names exact settings page and approval roles. 2026.9 (2026-09-25)

12 **New secret detectors and analyzers** NEW

60

Secrets Detection Engine v2.170 adds detectors for Firebase ID Token, Coze Access Token, Keycloak Access Token, Keycloak Refresh Token, Fish Audio API Key, and Azure Custom Vision Training Key, plus an Apify Token analyzer; v2.171 adds Akamai API Credentials and Docker Credentials analyzers.

— Names every new detector and analyzer with engine version. 2026.9 (2026-09-25)

13 **Custom MCP gateway OAuth redirect URIs** NEW 60

Adds `mcpServer.oauth.extraRedirectUris` Helm key to register a custom MCP gateway as an OAuth callback.

— Names the Helm key but gives no example. 2026.9 (2026-09-25)

THINNER COVERAGE BELOW

14 **Jira ticket comments cover more incident events** IMPROVED 50

Linked Jira ticket comments now report additional incident lifecycle events: reassignment, regression, validity changes, access grants and revocations, public sharing, and feedback.

— Lists event types but no config or UI detail. changelog-20260926-851756fd

15 **Validity-check mode for air-gapped networks** NEW 45

Adds a configurable validity-check mode that omits a detector's default public host and checks secrets only against custom hosts, supporting air-gapped networks.

— Describes behavior without naming the config key. 2026.9 (2026-09-25)

16 **Per-team incident permissions for promoted members** NEW 35

When promoting a user to Member, GitGuardian Self-Hosted now supports setting per-team incident permissions for that member.

— Named only in a summary, no mechanism or UI path given. changelog-20260926-851756fd

BREAKING ON UPGRADE

! The automatic Personal Access Token renewal system for Jira Data Center and Confluence Data Center integrations is removed. Before upgrading, manually update the token for every Jira Data Center and Confluence Data Center integration (sources and notifiers) with a long-lived PAT via the integration settings. An integration still relying on renewal will stop working when its token expires.

! ClickHouse is now rendered under `clickhouse.*` in the GitGuardian chart instead of the vendored `clickhouse-server` subchart. Any `clickhouse-server.*` values from 2026.8 must be moved to `clickhouse.*` before upgrading; the chart refuses to render while a `clickhouse-server` block is present. Object storage and backup settings move to `clickhouse.objectStorage` and `clickhouse.backup`, and `clickhouse.serverConfig` keys switch to camelCase.

- ! Argo CD installations that do not use an outbound proxy will see `1 resources require pruning` and stay `OutOfSync` after upgrade, because the chart no longer creates the `gim-proxy` and `gim-migration-proxy` Secrets when no proxy URL is set. Resolve by running `argocd app sync <app> --prune` or enabling automated pruning.
- ! ggscout outside the cluster must be upgraded to 0.32.0 or later; earlier versions do not classify collected secrets with the Secrets Detection Engine.
- ! Jira Data Center and Confluence Data Center: the automatic PAT renewal system is removed. Every Jira Data Center and Confluence Data Center integration (sources and notifiers) must be updated with a long-lived PAT in the integration settings before upgrading, or the integration will stop working when its current token expires.
- ! ClickHouse (Helm installations): `clickhouse-server.*` subchart values are replaced by `clickhouse.*`. Any `clickhouse-server.*` block must be migrated to `clickhouse.*` before upgrading — the chart refuses to render while a `clickhouse-server` block is present. Object storage and backup settings move to `clickhouse.objectStorage` and `clickhouse.backup`; `clickhouse.serverConfig` keys switch to camelCase.
- ! Argo CD installations: proxy Secrets `gim-proxy` and `gim-migration-proxy` are no longer created when no proxy URL is set, leaving the previously created Secrets behind. The sync ends with '1 resources require pruning' and the application stays OutOfSync until you prune (e.g. `argocd app sync <app> --prune`).
- ! `ggscout` minimum version moves to 0.32.0; earlier versions will not classify collected secrets with the Secrets Detection Engine.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

GitGuardian Self-Hosted: Adds per-team incident permissions when promoting a user to Member, and extends linked Jira ticket comments to cover reassignment, regression, validity, access grants and revocations, public sharing, and feedback events.

└──▷ **USE IT**

Write in-app analytics to a non-public PostgreSQL schema when the database forbids the `public` schema or cannot grant broad privileges to the GitGuardian user.

yaml

```yaml
inAppAnalytics:
  postgresql:
    applicationSchema: gitguardian_analytics
useSeparateAnalyticsSchema: true
analyticsSchema: gitguardian_analytics
```

Pin the support bundle pod to a dedicated tainted node so bundle collection does not compete with workload pods.

```yaml
replicated:
  supportBundle:
    nodeSelector:
      node-role: support
    tolerations:
      - key: dedicated
        operator: Equal
        value: support
        effect: NoSchedule
    affinity: {}
```

└──▷ **BREAKING ON UPGRADE**

- ! The automatic Personal Access Token renewal system for Jira Data Center and Confluence Data Center integrations is removed. Before upgrading, manually update the token for every Jira Data Center and Confluence Data Center integration (sources and notifiers) with a long-lived PAT via the integration settings. An integration still relying on renewal will stop working when its token expires.
- ! ClickHouse is now rendered under `clickhouse.*` in the GitGuardian chart instead of the vendored `clickhouse-server` subchart. Any `clickhouse-server.*` values from 2026.8 must be moved to `clickhouse.*` before upgrading; the chart refuses to render while a `clickhouse-server` block is present. Object storage and backup settings move to `clickhouse.objectStorage` and `clickhouse.backup`, and `clickhouse.serverConfig` keys switch to camelCase.
- ! Argo CD installations that do not use an outbound proxy will see `1 resources require pruning` and stay `OutOfSync` after upgrade, because the chart no longer creates the `gim-proxy` and `gim-migration-proxy` Secrets when no proxy URL is set. Resolve by running `argocd app sync <app> --prune` or enabling automated pruning.
- ! ggscout outside the cluster must be upgraded to 0.32.0 or later; earlier versions do not classify collected secrets with the Secrets Detection Engine.

2026.9 (2026-09-25) BLOG ↗ STABLE

GitGuardian Self-Hosted 2026.9 adds Bridge networking, honeytoken endpoint protection, AI hooks for Codex and Mistral Vibe, and vaulted-secret classification.

└──▷ **USE IT**

Route in-app analytics writes to a dedicated PostgreSQL schema when the GitGuardian database user cannot create objects in `public`.

```yaml
inAppAnalytics:
  postgresql:
    applicationSchema: gganalytics
useSeparateAnalyticsSchema: true
analyticsSchema: gganalytics
```

Schedule the support bundle collection pod onto a tainted node so it does not compete with production workloads.

```yaml
replicated:
  supportBundle:
    nodeSelector:
      node-role: ops
    tolerations:
      - key: ops
        operator: Exists
        effect: NoSchedule
    affinity: {}
```

Adds `GGBridge.bridges` Helm values key to declare GitGuardian Bridge servers, enabling scanning of sources in isolated networks (e.g. a GitLab instance behind a firewall) without opening inbound connections.

Adds `inAppAnalytics.postgresql.applicationSchema`, `useSeparateAnalyticsSchema`, and `analyticsSchema` Helm/config keys so in-app analytics can write to custom PostgreSQL schemas instead of `public`.

Adds `replicated.supportBundle.nodeSelector`, `tolerations`, and `affinity` Helm keys to schedule the support bundle pod on specific or tainted nodes.

Adds `mcpServer.oauth.extraRedirectUris` Helm key to register a custom MCP gateway as an OAuth callback.

Adds `clickhouse.objectStorage` and `clickhouse.backup` Helm keys (replacing `clickhouse-server.*` subchart values) as structured storage config under the new first-party `clickhouse.*` key.

+10 more

└─▷ **BREAKING ON UPGRADE**

! Jira Data Center and Confluence Data Center: the automatic PAT renewal system is removed. Every Jira Data Center and Confluence Data Center integration (sources and notifiers) must be updated with a long-lived PAT in the integration settings

before upgrading, or the integration will stop working when its current token expires.

! ClickHouse (Helm installations): `clickhouse-server.*` subchart values are replaced by `clickhouse.*`. Any `clickhouse-server.*` block must be migrated to `clickhouse.*` before upgrading — the chart refuses to render while a `clickhouse-server` block is present. Object storage and backup settings move to `clickhouse.objectStorage` and `clickhouse.backup`; `clickhouse.serverConfig` keys switch to camelCase.

! Argo CD installations: proxy Secrets `gim-proxy` and `gim-migration-proxy` are no longer created when no proxy URL is set, leaving the previously created Secrets behind. The sync ends with '1 resources require pruning' and the application stays OutOfSync until you prune (e.g. `argocd app sync <app> --prune`).

! `ggscout` minimum version moves to 0.32.0; earlier versions will not classify collected secrets with the Secrets Detection Engine.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Arcjet

Feedback ⓘ announced · Sep 17, 2026 BLOG ↗ RANK

Arcjet developer security blog.

> Arcjet launched agent runtime security for AI agents, combining OpenTelemetry-based observability, Rego/OPA-powered policy guards, native framework integrations, and audit evidence collection with flexible storage.

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

01 Arcjet Guards deterministic policy enforcement NEW

95

Arcjet Guards are Rego and Open Policy Agent-powered deterministic policies applied to tools, APIs, database calls, and input variables, configurable via the Arcjet Console, CLI, MCP, or API without code redeployment. A guard can be written as a Rego policy (e.g. denying an `injection-on-destructive-tool` when `input.signals.prompt_injection.incoming_message.detected` and `input.values.destructive` are true) and applied in JS agents by wrapping a tool with guardTool(arcjet, refundTool, { action, actor, inputs }) so the tool never executes on a DENY decision.

Guard a payment refund tool so the agent can never execute it when prompt injection is detected on a destructive action.

rego

```
package arcjet.guard

import rego.v1

deny contains "injection-on-destructive-tool" if {
  input.signals.prompt_injection.incoming_message.detected
  input.values.destructive
}
```

Wrap an existing refund tool with Arcjet Guards in a JS agent so the tool never runs on a DENY decision.

```javascript
const arcjet = launchArcjet({ key: process.env.ARCJET_KEY });

const refund = guardTool(arcjet, refundTool, {
  action: "order.refunded",
  actor: user.id,
  inputs: ({ amount }) => ({
    amount: policyInput.server.number(amount),
    role: policyInput.server.string(user.role),
  }),
});
```

— Includes runnable Rego and JS code with exact fields launch-20260917-2fe5f51a

THINNER COVERAGE BELOW

02 **Audit evidence collection with flexible storage** NEW

55

Arcjet collects audit evidence for every policy execution — including inputs, results, and session metadata — with options for cloud, single-tenant, or private VPC storage, and real-time log export to external detection systems.

— Describes evidence scope and storage options, no config surface

launch-20260917-2fe5f51a

03 **Agent activity observability via OpenTelemetry and Claude Compliance API**

NEW

48

Arcjet ingests agent activity through OpenTelemetry events without requiring code changes, enabling real-time visualization and analysis in the Arcjet Console. It also pulls agent activity data from the Claude Compliance API for teams using Claude.

— Names OTel and Claude Compliance API but no config keys launch-20260917-2fe5f51a

04 **Native integrations for agent frameworks** NEW

45

Arcjet adds native integrations for agent frameworks including Claude Agents SDK, Claude Managed Agents, OpenAI Agents SDK, LangChain, LangFuse, Strands, Mastra, and Microsoft's Agent Framework.

— Names all integrations but no setup mechanism given launch-20260917-2fe5f51a

announced · Sep 17, 2026 BLOG ↗

Arcjet launches agent runtime security with observe, enforce, and audit capabilities for AI agents.

└──▷ **USE IT**

Guard a payment refund tool so the agent can never execute it when prompt injection is detected on a destructive action.

```rego
package arcjet.guard

import rego.v1

deny contains "injection-on-destructive-tool" if {
  input.signals.prompt_injection.incoming_message.detected
  input.values.destructive
}
```

Wrap an existing refund tool with Arcjet Guards in a JS agent so the tool never runs on a DENY decision.

```javascript
const arcjet = launchArcjet({ key: process.env.ARCJET_KEY });

const refund = guardTool(arcjet, refundTool, {
  action: "order.refunded",
  actor: user.id,
  inputs: ({ amount }) => ({
    amount: policyInput.server.number(amount),
    role: policyInput.server.string(user.role),
  }),
});
```

Supports ingesting agent activity via OpenTelemetry events without code changes, enabling real-time visualization and analysis in the Arcjet Console.

Pulls agent activity data from the Claude Compliance API for teams using Claude.

Introduces Arcjet Guards — Rego and Open Policy Agent-powered deterministic policies applied to tools, APIs, database calls, and input variables, configurable via the Arcjet Console, CLI, MCP, or API without code redeployment.

Adds native integrations for agent frameworks including Claude Agents SDK, Claude Managed Agents, OpenAI Agents SDK, LangChain, LangFuse, Strands, Mastra, and Microsoft's Agent Framework.

Collects audit evidence for every policy execution — including inputs, results, and session metadata — with options for cloud, single-tenant, or private VPC storage, and

> real-time log export to external detection systems.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Sophos Email Security

Feedback ⓘ changelog · Sep 26, 2026

CHANGELOG ↗ RANK

Sophos Email Security protects organizations from email-based threats including phishing, malware, and spam through advanced threat detection and filtering.

> Sophos Email Security now lets administrators apply custom email protection policies directly to distribution lists and shared mailboxes.

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

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

NEW 45

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

— Explains what it does but no config path or UI steps given changelog-20260926-76ba23d7

> changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE
>
> Sophos Email Security now supports custom policies applied directly to distribution lists and shared mailboxes.

WAS THIS USEFUL? 👍 👎

Sophos Access Points

Feedback ⓘ changelog · Sep 26, 2026

CHANGELOG ↗ RANK

Sophos Access Points provides secure remote access and authentication management for protecting enterprise network resources.

> Sophos AP6 access points gain new wired network configuration options and deeper diagnostic logging for troubleshooting.

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

01 **Wired network settings for AP6 profiles** NEW 66

A new Network settings tab on access point profiles lets admins configure wired Ethernet port settings on AP6 access points, including Spanning Tree Protocol, IGMP snooping, and LLDP-MED, to control multicast traffic and support VoIP endpoint discovery.

Enable IGMP snooping and LLDP-MED on an AP6 access point profile to control multicast traffic and support VoIP endpoint discovery on wired ports.

> 📍 In Sophos Fusion, navigate to the access point profile, open the Network settings tab, and enable Spanning Tree Protocol, IGMP snooping, and LLDP-MED for the relevant Ethernet ports.

— Names UI tab and protocols with exact configuration steps changelog-20260926-b9e48cc9

THINNER COVERAGE BELOW

02 **Advanced diagnostic logging in Sophos Fusion** NEW 53

Advanced diagnostic logging can be enabled from Diagnostics > System logs in Sophos Fusion for deeper wireless network troubleshooting.

Capture advanced diagnostic logs when troubleshooting intermittent wireless connectivity issues on an AP6 deployment.

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

— Exact UI path given, but mechanism of logging unexplained changelog-20260926-b9e48cc9

03 **Port trunking for AP6 840E models** NEW 40

Port trunking configuration is added under the Network settings tab, available exclusively for AP6 840E access point models.

— Names exact model and location but no usage steps changelog-20260926-b9e48cc9

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Sophos AP6 gains wired network settings (STP, IGMP snooping, LLDP-MED, port trunking) and advanced diagnostic logging.

└──▷ **HOW TO FIND IT**

Enable IGMP snooping and LLDP-MED on an AP6 access point profile to control multicast traffic and support VoIP endpoint discovery on wired ports.

> 📍 In Sophos Fusion, navigate to the access point profile, open the Network settings tab, and enable Spanning Tree Protocol, IGMP snooping, and LLDP-MED for the relevant Ethernet ports.

Capture advanced diagnostic logs when troubleshooting intermittent wireless connectivity issues on an AP6 deployment.

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

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

Adds port trunking configuration under the Network settings tab, available exclusively for AP6 840E models.

Adds advanced diagnostic logging, enabled from Diagnostics > System logs in Sophos Fusion, for deeper wireless network troubleshooting.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Chainguard

Feedback ⓘ changelog · Sep 26, 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 ship this window is a new Agent Skills hardening service (`chainctl skills harden` / `status` / `list` / `describe` / `pull`) that scans and hardens uploaded skill artifacts; Guardener's automated migration tool now also covers GitHub Actions, and a sonar-scanner-cli default-user change breaks existing Kubernetes pipelines starting October 12.

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

01 **Agent Skills hardening service** NEW 95

Chainguard Agent Skills is a beta service (requiring a Skills entitlement and accepted Skills terms) that hardens uploaded AI/agent skill artifacts. Users upload originals to `uploads.cgr.dev/<your-org>/<skill>` and submit them via `chainctl skills harden`, tracking jobs with `chainctl skills status --wait --timeout`; hardening jobs are deduplicated server-side so the same org/user/skill/digest returns an existing job ID instead of creating a new one. Results publish to `skills.cgr.dev/<org-uidp>/users/<user-namespace>/<skill>@sha256:<digest>`, downloadable via `chainctl skills describe` and `chainctl skills pull`, and include a `HARDENING.md` report alongside `SKILL.md` in `./hardened/<skill-name>/` summarizing changes and open scanner findings; all skills across both registries can be listed with `chainctl skills list --source all --recursive`.

Submit a local skill for hardening in CI and capture the job ID so the pipeline can check status asynchronously without blocking on completion.

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

Harden a previously pushed artifact by digest to ensure you target an exact, immutable version rather than a mutable tag.

```
$ chainctl skills harden --group your-organization --name harden-demo \
  --digest 'sha256:<64-hex-digest-from-push>' --wait --timeout 30m
```

List all hardened and uploaded skills across both registries to confirm a job result is available before installing.

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

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

```
$ chainctl skills harden
```

Upload an agent skill artifact and submit it for hardening, then track the job until the hardened image is available in your org registry.

```
$ chainctl skills harden uploads.cgr.dev/<your-org>/<skill>
```

— Full commands, endpoints, file names and dedup mechanism all named `product docs`

02 **Guardener migration extended to GitHub Actions** IMPROVED

60

Guardener, previously used to migrate container images, now also covers GitHub Actions: it inventories Actions usage across an organization's repositories, maps them to hardened Chainguard equivalents, and opens pull requests pinned to a specific SHA rather than a mutable tag. It runs in two modes — an upfront pass that surfaces existing usage and opens migration PRs, and ongoing standardization that watches workflow files and suggests Chainguard equivalents as new upstream Actions appear.

— Mechanism and modes described but no runnable command shown `changelog-20260926-9002d47e`

THINNER COVERAGE BELOW

03 **26+ new container images** NEW 15

Chainguard adds 26+ new container images to its catalog in this release, per the release summary.

— Only a count is given, no image names or details changelog-20260926-9002d47e

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Chainguard adds EU Sovereign Artifacts beta, Guardener now covers GitHub Actions, and 26+ new container images.

└──▷ **USE IT**

Temporarily restore the old `sonar-scanner-cli` default-user behavior in Kubernetes after the October 12 breaking change takes effect.

yaml

```yaml
securityContext:
  runAsUser: 0
```

Redirect the Sonar scanner scratch directory to the old workspace path so pipelines reading `report-task.txt` continue to work without switching to root.

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

Guardener automated migration tool now covers GitHub Actions in addition to container images, inventorying Actions across an organization's repositories, mapping them to hardened Chainguard equivalents, and opening pull requests pinned to a specific SHA rather than a mutable tag.

Guardener runs in two modes: an upfront pass that surfaces existing Actions usage and opens migration pull requests, and ongoing standardization that watches workflow files and suggests Chainguard equivalents as new upstream Actions appear.

Adds tutorial for pulling packages from Chainguard package repositories through Sonatype Nexus as pull-through caches for apk packages.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Lacework FortiCNAPP

Feedback ⓘ snapshot-20260926 NOTES ↗

RANK

Lacework FortiCNAPP is a cloud-native application protection platform providing visibility, compliance, and threat detection across cloud infrastructure.

> FortiCNAPP's latest release focuses on risk scoring tied to inline firewall protection, broader OS vulnerability coverage, and new alert filtering, account rollback, and on-demand scanning capabilities.

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

01 **Risk score adjustment for inline firewall protection** NEW 40

Adds risk score adjustment for inline FortiGate or other next-generation firewall protection, allowing risk scores to reflect whether a workload is shielded by an inline NGFW.

– Names mechanism but no config surface or scope details snapshot-20260926

02 **Vulnerability assessment for COS and Ubuntu** NEW 40

Adds vulnerability assessment for Google Container-Optimized OS and for Ubuntu 25.04 through 26.10.

– Names specific OS versions but no mechanism detail snapshot-20260926

03 **Alert filtering and grouping improvements** NEW 35

Adds click-to-filter alerts by category and sub-category, and adds resource group support for alerts.

– Two thin alert UI additions grouped, no detail on mechanism snapshot-20260926

04 **On-demand container image scanning** NEW 30

Adds on-demand container image scanning.

– States capability but no command or config to invoke it snapshot-20260926

05 **Cloud account rollback and deployment history** NEW 25

Adds cloud account rollback and deployment history tracking.

– Bare description, no mechanism or UI path given snapshot-20260926

06 **New AWS service coverage** NEW 10

Adds new AWS service coverage.

— Generic statement with no named services snapshot-20260926

snapshot-20260926 NOTES ↗ LATEST STABLE

FortiCNAPP adds risk score adjustment for inline firewall protection, vulnerability assessment for COS and Ubuntu, and Polygraph v2 graph-based anomaly alerts.

Adds risk score adjustment for inline FortiGate or other next-generation firewall protection.

Adds vulnerability assessment for Google Container-Optimized OS.

Adds vulnerability assessment for Ubuntu 25.04 through 26.10.

Adds click-to-filter alerts by category and sub-category.

Adds cloud account rollback and deployment history.

+3 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

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

> Huntress added an org-level rollout control for deployments, made Microsoft Defender scan behavior configurable, and overhauled Unsecured Credentials detection with a preview/allowlist workflow, alongside new macOS visibility and detection additions.

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

01 Unsecured Credentials preview and allowlist IMPROVED

75

Unsecured Credential Incident Reports can now be previewed before enabling: in Managed EDR > Process Insights > Credential Files, admins can see which files would trigger a report without turning reporting on, then select already-reviewed files and add them to an allowlist before enabling reporting.

Before re-enabling Unsecured Credential Incident Reports on a noisy environment, preview exactly which files would trigger reports and allowlist already-reviewed ones.

> 📍 In the Huntress console, go to Managed EDR › Process Insights › Credential Files. Use the preview mode to see which files would generate an Incident Report without turning reporting on, then select reviewed files and add them to the allowlist before enabling reporting.

— Clear workflow and UI path, though endpoint/config key not named.

changelog-20260926-292ff4c5

THINNER COVERAGE BELOW

02 macOS Screen Sharing visibility NEW

10

Huntress added visibility into macOS Screen Sharing activity, per the release summary; no further mechanism or scope was described.

— Only named in the summary, no supporting detail. changelog-20260926-292ff4c5

03 macOS quarantine-strip detections NEW

10

New detections were added for macOS quarantine-strip activity, per the release summary; no further mechanism or scope was described.

— Only named in the summary, no supporting detail. changelog-20260926-292ff4c5

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Huntress adds Modified Deployments rollout control, Unsecured Credentials overhaul, macOS Screen Sharing visibility, and new macOS quarantine-strip detections

└──▷ HOW TO FIND IT

Before re-enabling Unsecured Credential Incident Reports on a noisy environment, preview exactly which files would trigger reports and allowlist already-reviewed ones.

> 📍 In the Huntress console, go to Managed EDR › Process Insights › Credential Files. Use the preview mode to see which files would generate an Incident Report without turning reporting on, then select reviewed files and add them to the allowlist before enabling reporting.

Enable Archive File Scanning and disable Network File Scanning for Microsoft Defender on endpoints where mapped-drive performance is a concern.

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

Adds `Modified Deployments` org-level setting, letting administrators build a custom control set from the Huntress policy library, select low/medium/high impact items, set deployment days, and approve new controls before they roll out — ideal for strict compliance or change-control environments.

Adds `Network File Scanning` and `Archive File Scanning` toggles for Microsoft Defender in `EDR > Managed Antivirus > Defender Configuration > Scans`, previously fixed to Huntress defaults.

WAS THIS USEFUL? 👍 👎

Automox

Feedback ⓘ

Improved macOS Vulnerability Severity Data (2026-09-25) BLOG ↗ RANK

Check out our full list of release notes, and stay up to date with what

> Automox extended KEV status, EPSS scores, and CVE severity data to macOS, so existing KEV- and EPSS-targeted policies now automatically apply to macOS devices without any configuration changes.

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

01 macOS vulnerability severity data in policies and reporting

50

BREAKING

KEV status, EPSS scores, and CVE severity data now cover macOS patches, extending existing KEV- and EPSS-targeted policies to automatically include macOS devices with no policy configuration changes required. Console views, Analytics, and reports are updated to reflect CVE counts and severity levels for macOS based on this expanded vulnerability data. As a result, severity-based policies may behave differently on macOS — patches may now fall in or out of scope depending on the severity levels a policy targets — and CVE counts/severity levels shown in console views, Analytics, and reports may change.

— Describes mechanism and impact but no specific UI paths or config keys.

Improved macOS Vulnerability Severity D…

└──▷ **BREAKING ON UPGRADE**

! Severity-based policies may behave differently on macOS — macOS patches may now fall in or out of scope depending on the severity levels a policy targets.

! Existing KEV- and EPSS-targeted policies will now include macOS devices automatically, without any configuration change.

! CVE counts and severity levels for macOS in console views, Analytics, and reports may change as vulnerability data has been expanded and refined.

Improved macOS Vulnerability Severity Data (2026-09-25) BLOG ↗ LATEST

STABLE

Automox extends KEV status, EPSS scores, and CVE severity data to macOS patches, enabling risk-based policy targeting across all platforms.

Extends existing KEV- and EPSS-targeted policies to automatically include macOS devices with no policy configuration changes required.

Updates console views, Analytics, and reports to reflect CVE counts and severity levels for macOS based on the expanded vulnerability data.

└──▷ **BREAKING ON UPGRADE**

- ! Severity-based policies may behave differently on macOS — macOS patches may now fall in or out of scope depending on the severity levels a policy targets.
- ! Existing KEV- and EPSS-targeted policies will now include macOS devices automatically, without any configuration change.
- ! CVE counts and severity levels for macOS in console views, Analytics, and reports may change as vulnerability data has been expanded and refined.

WAS THIS USEFUL? 👍 👎

Iru

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

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

> Iru's September 25th release adds MSP Billing with itemized invoicing, brings LAPS support to Mac devices, strengthens Compliance evidence collection, and adds macOS 27 support.

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

01 **LAPS for Mac as Library Item** NEW 48

Adds LAPS for Mac as a Library Item, extending local admin password rotation to Mac devices.

— Names feature and delivery mechanism (Library Item) but no exact steps

What's New in Iru | September 25th

02 **MSP Billing with itemized invoicing** NEW 38

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



— Describes what it does but no config or UI path given

What's New in Iru | September 25th

03 **macOS 27 support** NEW 22

Iru adds support for macOS 27.

— Only mentioned in summary, no further detail provided

What's New in Iru | September 25th

04 Compliance evidence collection enhancements IMPROVED

18

Enhances Iru Compliance evidence collection.

— Vague description with no specifics on what changed What's New in Iru | September 25th

> What's New in Iru | September 25th NOTES ↗ LATEST STABLE
>
> Iru adds MSP Billing, LAPS for Mac, compliance evidence enhancements, and macOS 27 support.
>
> Adds MSP Billing with itemized billing to increase visibility into MSP margins.
>
> Adds LAPS for Mac as a Library Item.
>
> Enhances Iru Compliance evidence collection.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ changelog · Sep 26, 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 more granular AI usage reporting, markdown conversion functions, live-updating cases and dashboards, and change control tracking for cases in the storyboard.

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

01 **Granular AI usage API dimensions** IMPROVED 40

The AI usage API gains new dimensions, enabling more granular reporting on AI consumption.

– Names the API but not new dimension fields changelog-20260926-c9c47f7f

02 **Live-updating cases and dashboards** IMPROVED 35

Cases and dashboards now stay up to date when viewed in background browser tabs.

– States behaviour but no mechanism changelog-20260926-c9c47f7f

03 **Markdown conversion functions** NEW 25

New functions are added for converting content to markdown.

– No function names or syntax given changelog-20260926-c9c47f7f

└──▷ **ALSO FROM THESE RELEASES**

. . . .

> changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE
>
> Tines: Adds new dimensions to the AI usage API, enabling more granular reporting on AI consumption.
>
> Adds new functions for converting content to markdown.
>
> Cases and dashboards now stay up to date when viewed in background browser tabs.
>
> Enables tracking cases with change control directly in the storyboard.

WAS THIS USEFUL?  

◆ Detection Engineering

Elastic Detection Rules

Feedback ⓘ dev-v3.0.0 NOTES ↗

RANK

Detection Rules is Elastic's framework for defining and managing security detection logic to identify threats and anomalies in system and network data.

> Detection Rules 3.0.0 replaces remote ES|QL validation with an offline validator and removes the old validation flag, adds XDR/SIEM taxonomy and rule-profiling tags, supports ES|QL sub-queries, and ships a large wave of new detection rules spanning Windows, Linux, Azure, AWS, Entra ID, macOS and Anthropic.

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

01 Offline ES|QL validation replaces remote validation

BREAKING

90

Rules are now validated offline at load time via `detection-rules-esql-py` 0.1.0, checking queries against `min_stack_version` and local schemas instead of posting them to an Elasticsearch cluster. The `--esql-remote-validation` flag has been removed entirely, so any workflow or script that passed it will fail; `DR_BYPASS_ESQL_KEEP_VALIDATION` and `DR_BYPASS_ESQL_METADATA_VALIDATION` remain but now only skip the `KEEP` and `FROM` metadata checks, meaning rules that previously loaded only because remote validation was skipped can now fail local syntax, schema, or type checks.

— Names package, env vars and removed flag with clear before/after dev-v3.0.0

02 New detection rules across Windows, Linux, Azure, AWS and SaaS

NEW

65

Ships a wave of new detection rules and one hunting query: Potential ClickFix Command via Windows Run Dialog; Linux Firewall Rule Creation or Modification; Potential Fileless Execution via Unusual memfd Create Call; Potential Fileless Execution via `O_TMPFILE`; three Linux file-related web server rules; Azure AKS Pod Exec Potential Reverse Shell; Newly Seen Commonly Abused Network Scanner; RMM Installation via Internet-Originated MSI; First Seen Network Flow Exporter; Azure WireServer Abuse Related Detections; Potential Data Exfiltration via AzCopy or Azure Storage Explorer; Entra ID Windows Hello for Business Key-Borrowing Detections; Entra ID device-bound PRT cookie replay from unusual IP; AWS SES Account Email Sending Enabled; the Java_Ghost Threat Actor Infrastructure and Identity Markers hunting query; SSH brute force detection rules for the macOS Security Events integration; and Anthropic compliance audit detection rules.

— Every rule named but no query logic or coverage detail given dev-v3.0.0

THINNER COVERAGE BELOW

03 **Rule profiling tags** NEW 45

Adds rule profiling tags Recommended, Noise, Performance, and Threat that can be applied to detection rules.

— Named tag values but no mechanism for how they're set or used dev-v3.0.0

04 **XDR/SIEM taxonomy tags (phase 1)** NEW 30

Adds XDR/SIEM taxonomy tags to rules, described as phase 1 of the effort.

— Named but no mechanism or usage detail given dev-v3.0.0

05 **ES|QL sub-query support in validation** NEW 25

Rule validation now supports ES|QL sub queries.

— Bare statement with no mechanism or example dev-v3.0.0

└──▷ BREAKING ON UPGRADE

! The `--esql-remote-validation` flag is removed; any workflow or script that passed it will now fail. ES|QL rules are validated offline at load time against local schemas and `min_stack_version`, and queries that previously loaded only because remote validation was skipped can now fail local syntax, schema, or type checks.

dev-v3.0.0 NOTES ↗ LATEST STABLE

dev-v3.0.0 ships offline ES|QL validation, XDR/SIEM taxonomy tags, ES|QL sub-query support, and a wave of new detection rules.

└──▷ GET THIS VERSION

```
$ git clone --branch dev-v3.0.0
https://github.com/elastic/detection-rules.git
# already have the repo? check out this version:
$ git checkout dev-v3.0.0
```

Replaces remote ES|QL validation with offline validation via `detection-rules-esql-py` 0.1.0, checking rules at load time against `min_stack_version` and local schemas; `DR_BYPASS_ESQL_KEEP_VALIDATION` and `DR_BYPASS_ESQL_METADATA_VALIDATION` still skip only the `KEEP` and `FROM` metadata checks.

Removes the `--esql-remote-validation` flag; queries no longer need to be posted to an Elasticsearch cluster for validation.

Adds XDR/SIEM taxonomy tags (phase 1) to rules.

Adds rule profiling tags: `Recommended`, `Noise`, `Performance`, `Threat`.

Supports ES|QL sub queries in rule validation.

`+17 more`

└─▷ **BREAKING ON UPGRADE**

! The `--esql-remote-validation` flag is removed; any workflow or script that passed it will now fail. ES|QL rules are validated offline at load time against local schemas and `min_stack_version`, and queries that previously loaded only because remote validation was skipped can now fail local syntax, schema, or type checks.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Wazuh

Feedback ⓘ v4.14.8 NOTES ↗ RANK

Wazuh - The Open Source Security Platform. Unified XDR and SIEM protection for endpoints and cloud workloads.

> Wazuh v4.14.8 extends FIM eBPF whodata alerts with additional audit identity fields for better attribution of file integrity events.

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

01 **Audit identity fields in FIM eBPF whodata alerts** IMPROVED

55

FIM whodata alerts generated by the eBPF provider now include the fields `audit_uid`, `audit_name`, `audit_gid`, and `audit_group_name`, adding audit identity context to file integrity monitoring events.

— Names exact fields but no usage example or mechanism detail v4.14.8

v4.14.8 NOTES ↗ LATEST STABLE

Wazuh v4.14.8 adds audit identity fields to FIM eBPF whodata alerts.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v4.14.8 https://github.com/wazuh/wazuh.git
# already have the repo? check out this version:
$ git checkout v4.14.8
```

WAS THIS USEFUL? 👍 👎

Databricks

Feedback ⓘ

Genie Code can now build intelligent document processing pipelines (2026-09-25)

BLOG ↗ RANK

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

> Databricks added a SQL function for time-series anomaly detection and gave Genie Code the ability to build full document-processing pipelines.

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

01 **Anomaly detection SQL function** NEW 65

The new `ai_detect_anomalies` SQL function identifies anomalous values in historical time series data, returning expected values, prediction bounds, and anomaly scores.

— Names exact function and outputs but no usage example. product docs

02 **Genie Code document processing pipelines** NEW 60

Genie Code can now orchestrate end-to-end intelligent document processing pipelines that handle parsing, classification, extraction, and validation against unstructured PDFs, contracts, or invoices, with iterative quality improvement evaluated against labeled data.

— Describes mechanism and scope but no concrete command or config to act on.

Genie Code can now build intelligent do…

> Genie Code can now build intelligent document processing pipelines (2026-09-25)
>
> BLOG ↗ LATEST STABLE
>
> Genie Code can now orchestrate end-to-end intelligent document processing pipelines for PDFs, contracts, and invoices.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Cyberbro

Feedback ⓘ v0.15.2 NOTES ↗ RANK

A simple application that extracts your IoCs from garbage input and checks their reputation using multiple CTI services.

> Cyberbro adds a Microsoft Tenant Check engine that determines whether a domain or email belongs to a registered Microsoft Entra ID tenant and whether it uses Office 365 mail.

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

01 **Microsoft Tenant Check engine** NEW 88

New `ms_tenant_check` engine queries `https://login.microsoftonline.com/<domain>/.well-known/openid-configuration` to confirm whether a domain is registered as a Microsoft Entra ID (Azure AD) tenant, returning `tenant_id`, `tenant_region_scope` (e.g. `EU`, `NA`), `mx_records`, `is_office365_mx`, `tenant_found`, and `domain` fields. It supports both FQDN and email observables (extracting the domain from addresses like `user@contoso.com`), checks Google DNS for MX records, detects `mail.protection.outlook.com` to flag Office 365 mail usage, and requires no API key.

Determine if a suspicious domain is a real Microsoft tenant and whether it uses Office 365 mail — useful when triaging phishing domains or validating vendor identity.

```
$ curl -s 'https://login.microsoftonline.com/contoso.com/.well-known/openid-configuration'
```

— Names endpoint, returned fields, and detection logic in full. v0.15.2

v0.15.2 NOTES ↗ LATEST STABLE

Cyberbro v0.15.2 adds Microsoft tenant check capability.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.15.2 https://github.com/stanfrbd/cyberbro.git
# already have the repo? check out this version:
$ git checkout v0.15.2
```

WAS THIS USEFUL?  

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ changelog · Sep 26, 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 focuses on integration data hygiene, cloud asset visibility, and internal performance, alongside a CycloneDX-based CBOM export for auditing cryptographic assets.

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

01 **Automatic retirement of expired integration data** IMPROVED 35

Software and Vulnerabilities collected from third-party integrations are now retired automatically when the source integration's data expires, keeping the inventory from holding stale findings.

– Describes behavior but no config or threshold details changelog-20260926-8988beed

02 **Email notifications for Issue assignments** NEW 30

Issue owners now receive an email with a direct link when they are assigned an Issue.

– Simple behavior, no settings or triggers named changelog-20260926-8988beed

03 **Expanded asset types in AWS, Azure, GCP imports** IMPROVED 30

AWS, Azure, and GCP imports now include additional asset types, broadening what cloud inventory data runZero pulls in.

– Names three cloud providers but not which asset types changelog-20260926-8988beed

04 **Scan and task performance optimizations** IMPROVED 30

runZero optimized MAC merging so large scan jobs process faster, and now processes directory and device data concurrently to improve task performance.

– Explains two mechanisms but no benchmarks or metrics changelog-20260926-8988beed

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

runZero: Software and Vulnerabilities collected from integrations are now retired when that integration data expires.

└──▷ **HOW TO FIND IT**

Download the Cryptographic Bill of Materials (CBOM) as a CycloneDX JSON file to audit cryptographic assets across the inventory.

> 📍 In the console, go to Reports, select `Cryptographic Bill of Materials (CBOM)`, and choose `Export as CycloneDX JSON`.

Issue owners now receive an email with a direct link when assigned an Issue.

AWS, Azure, and GCP imports now include additional asset types.

runZero optimized MAC merging so large scan jobs now process faster.

runZero now processes directory and device data concurrently, improving task performance.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Aembit

Feedback ⓘ changelog · Sep 26, 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 SAML-bound dynamic credentials, in-cluster mTLS certificate issuance for Kubernetes workloads, and several MCP-focused controls including a managed access token provider and tool-filtering audit logging.

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

01 SAML-bound dynamic claims for OAuth credentials NEW

85

The OAuth 2.0 Authorization Code Credential Provider can now read values from a SAML 2.0 response using `${saml.response.subject.nameId}`, `${saml.response.issuer}`, and `${saml.response.attributes['<name>']}` expressions, binding each issued credential to the individual SAML user's identity (e.g. their group membership).

Bind a SAML user's group membership into a dynamically issued credential so each user gets a credential scoped to their own identity.

```text
${saml.response.attributes['groups']}
```

— Names exact expression syntax and provider it applies to. changelog-20260926-58a29786

02 MCP Enterprise Managed Access Token Credential Provider NEW

70

Enabled in the Aembit Tenant UI, this provider obtains per-user MCP server access tokens through a corporate identity provider, letting users reach MCP servers behind the MCP Identity Gateway without a per-user consent prompt.

— Names the provider and UI location but no config detail. changelog-20260926-58a29786

03 MCP tool filtering decision logging NEW

60

MCP workload authorization events now record tool filtering decisions under `application.mcp.contentSecurity`, capturing which control decided the outcome, which rules matched, which tools were removed, and the return count.

— Names the event field but is purely informational. changelog-20260926-58a29786

04 **Multi-cluster attestation for Kubernetes Service Account Trust Provider**

IMPROVED 60

The Kubernetes Service Account Trust Provider can now accept multiple OIDC endpoints, one per cluster issuer URL, added via Add an Endpoint, enabling attestation of a single Client Workload running active-active across two or more clusters.

— Names the UI action but not underlying config schema. changelog-20260926-58a29786

└──▷ **BREAKING ON UPGRADE**

! Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret, because `X509Svid` writes a `kubernetes.io/tls` Secret while every other credential type writes a generic Secret.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Aembit: Adds `${saml.response.subject.nameId}`, `${saml.response.issuer}`, and `${saml.response.attributes['<name>']}` expressions so OAuth 2.0 Authorization Code Credential Provider dynamic claims can read values from a SAML 2.0 response, binding each issued credential to the individual SAML user's identity.

└──▷ **USE IT**

Issue mTLS certificates to a Kubernetes workload without the private key ever leaving the cluster — set `credentialType: X509Svid` on a refresh schedule and Secrets Operator handles CSR submission and auto-renewal.

yaml

```yaml
credentialType: X509Svid
```

Bind a SAML user's group membership into a dynamically issued credential so each user gets a credential scoped to their own identity.

text

```
${saml.response.attributes['groups']}
```

Adds `credentialType: X509Svid` to `AembitSecretRefreshSchedule` so Aembit Secrets Operator generates a private key in-cluster, obtains a signed certificate from Aembit, and writes both to a `kubernetes.io/tls` Secret (`tls.crt` and `tls.key`) for mutual TLS; the private key never leaves the cluster and certificates renew automatically.

Enables the MCP Enterprise Managed Access Token Credential Provider in the Aembit Tenant UI, obtaining per-user MCP server access tokens through a corporate

identity provider so users reach MCP servers behind the MCP Identity Gateway without a per-user consent prompt.

Records MCP tool filtering decisions — which control decided the outcome, which rules matched, which tools were removed, and the return count — under `application.mcp.contentSecurity` in MCP workload authorization events.

Allows a Kubernetes Service Account Trust Provider to accept multiple OIDC endpoints (one per cluster issuer URL, added via `Add an Endpoint`), enabling attestation of a single Client Workload running active-active across two or more clusters.

└──▷ **BREAKING ON UPGRADE**

! Switching an `AembitSecretRefreshSchedule` to or from `credentialType: X509Svid` deletes and recreates the target Kubernetes Secret, because `X509Svid` writes a `kubernetes.io/tls` Secret while every other credential type writes a generic Secret.

WAS THIS USEFUL? 👍 👎

Okta for AI Agents

Feedback ⓘ changelog · Sep 26, 2026

CHANGELOG ↗ RANK

> Okta for AI Agents replaced its on-prem SCIM agent and highlighted network zone controls that let admins block traffic by IP service category.

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

01 **Okta On-prem SCIM agent replaces SCIM Server agent** BREAKING

40

The Okta On-prem SCIM Server agent has been replaced by the Okta On-prem SCIM agent, reducing dependencies and enabling new features. Existing deployments relying on the old agent name or its dependencies will need to migrate to the new agent.

— Names the old and new agent but no migration steps given changelog-20260926-2fd4fb13

└──▷ **BREAKING ON UPGRADE**

! Okta On-prem SCIM Server agent has been replaced by Okta On-prem SCIM agent — existing deployments relying on the old agent name or its dependencies will need to migrate to the new agent.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Okta for AI Agents: Replaces Okta On-prem SCIM Server agent with Okta On-prem SCIM agent, reducing dependencies and enabling new features.

└──▷ **HOW TO FIND IT**

Block access attempts originating from known malware proxy networks at the global session policy level.

📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add the `MALWARE_PROXIES` IP service category to the zone, then reference that zone in Security › Authentication Policies (or the Global Session Policy) with an Access denied action.

Prevent analysts and contractors from authenticating through remote desktop sessions by excluding `REMOTE_DESKTOP` traffic in an app sign-in policy.

📍 In the Admin Console, go to Security › Networks, open or create an Enhanced Dynamic Network Zone, add the `REMOTE_DESKTOP` IP service category, then in Security › Authentication Policies select the relevant app policy and add a rule: 'If network zone is `REMOTE_DESKTOP` — Deny access'.

└──▷ **BREAKING ON UPGRADE**

! Okta On-prem SCIM Server agent has been replaced by Okta On-prem SCIM agent — existing deployments relying on the old agent name or its dependencies will need to migrate to the new agent.

WAS THIS USEFUL?  

iDenfy

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

Explore iDenfy documentation for identity verification (KYC), business verification (KYB), AML screening, fraud prevention, and more.

> iDenfy's most notable change this window is a breaking update to its iOS SDK, which now supports iOS 27 but raises the minimum deployment target to iOS 15, alongside smaller additions like Android landscape mode, PNG logo uploads, and field filtering on the KYB company endpoint.

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

01 iOS SDK 9.2.2 update with breaking deployment target change

55

BREAKING

iOS SDK 9.2.2 adds support for iOS 27, but the iOS minimum deployment target is raised to iOS 15, which is a breaking change for apps targeting older iOS versions.

— Version numbers given but no migration steps beyond the deployment target.

changelog-20260926-928497ee

02 PNG logo upload support NEW

25

Adds support for uploading logos in PNG format.

— No endpoint or UI path given for where to upload. changelog-20260926-928497ee

└──▷ **BREAKING ON UPGRADE**

! iOS minimum deployment target is raised to iOS 15 in iOS SDK 9.2.2.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

iDenfy: Adds PNG logo upload support.

└──▷ **TRY IT**

Retrieve only specific fields from the KYB company endpoint to reduce payload size and avoid exposing unnecessary data.

```
$ curl -X GET 'https://api.idenfy.com/kyb/company?
fields=name,registrationNumber,directors' -H 'Authorization: Bearer
<token>'
```

Adds landscape mode support in Android SDK 9.2.0.

Adds iOS 27 support in iOS SDK 9.2.2.

> └──▷ **BREAKING ON UPGRADE**
>
> ! iOS minimum deployment target is raised to iOS 15 in iOS SDK 9.2.2.

WAS THIS USEFUL?  

Passbolt

Feedback

changelog · Sep 26, 2026

CHANGELOG ↗

RANK

Manage and share passwords securely with Passbolt. Open source, audited and built for teams that need collaboration, compliance and control.

> Passbolt 5.16.0 introduces a Beta Offline Mode in its browser extension, letting users access and autofill selected credentials when the server is unreachable.

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

01 **Offline Mode (Beta) in browser extension** NEW 45

Introduces Offline Mode (Beta) in the browser extension, providing secure read-only access and autofill for selected credentials when the Passbolt server cannot be reached.



— Describes behaviour but no config, flags, or limits given changelog-20260926-fa63499d

> changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE
>
> Passbolt 5.16.0 adds Offline Mode (Beta) for secure read-only access and autofill when the server is unreachable.

WAS THIS USEFUL?

Didit

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

Didit provides identity verification, fraud prevention, and transaction-monitoring APIs, SDKs, and MCP tools.

> Didit added camera-control query parameters for its hosted verification URL, letting integrators lock and hide liveness/document camera selection outside the SDK.

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

01 Camera control parameters for verification URL NEW

83

Added `liveness_camera` and `liveness_camera_switch` query parameters to the verification URL, allowing camera selection (e.g. rear camera) and switcher visibility to be controlled when opening the flow directly without the SDK. Related SDK config options include `defaultLivenessCamera`, `showLivenessCameraSwitchButton`, `defaultDocumentCamera`, and `showDocumentCameraSwitchButton`.

Lock a kiosk or operator-facing device to the rear camera for liveness capture and hide the in-capture switcher so the user cannot switch lenses.

```json
{
  "defaultLivenessCamera": "back",
  "showLivenessCameraSwitchButton": false,
  "defaultDocumentCamera": "back",
  "showDocumentCameraSwitchButton": false
}
```

Open the verification URL directly with rear-camera liveness and the switcher disabled, without going through the SDK.

```
$ https://verify.didit.me/session/<session_id>?
liveness_camera=back&liveness_camera_switch=false
```

— Names exact query params, config keys, and a runnable URL example

changelog-20260926-32287544

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

> Didit: Adds `liveness_camera` and `liveness_camera_switch` query parameters to the verification URL for controlling camera selection and switcher visibility when

opening the flow directly without the SDK.

└──▷ USE IT

Lock a kiosk or operator-facing device to the rear camera for liveness capture and hide the in-capture switcher so the user cannot switch lenses.

```json
{
  "defaultLivenessCamera": "back",
  "showLivenessCameraSwitchButton": false,
  "defaultDocumentCamera": "back",
  "showDocumentCameraSwitchButton": false
}
```

Open the verification URL directly with rear-camera liveness and the switcher disabled, without going through the SDK.

```
$ https://verify.didit.me/session/<session_id>?
liveness_camera=back&liveness_camera_switch=false
```

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗

RANK

Govern workforce access, deploy AI agents, and get work done with C1.

> ConductorOne's headline addition is a serverless TypeScript Functions platform for custom automation, alongside a new Access Profiles system with JML-driven auto-enrollment/auto-unenrollment and anomaly safeguards, and a secrets-management suite spanning inventory, zero-knowledge sharing, workload federation, MCP authorization and tripwire decoys — plus breaking changes to the secrets revoke and task-status APIs.

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

01 Serverless Functions platform NEW

100

Serverless TypeScript Functions run custom automation logic, call external systems, and implement identity governance workflows. Functions retrieve API keys and secrets via functions.getConfig(), import external npm packages using the `npm:` prefix, and call pre-authenticated C1 APIs via the `@c1/functions-sdk` package. They can be triggered from automation workflows on user lifecycle events (created, updated, deactivated), access events (grant created, revoked), review events (completed, decisions made), and schedule triggers (daily, weekly, custom), run on-demand from the web UI via Platform > Functions with JSON input/output, or invoked programmatically via the C1 API. Each function executes in an isolated sandbox authenticated as a service principal with role-binding-based runtime permissions, a network allowlist restricting reachable external domains, and `console.log` execution logs streamed back to the UI.

Manually invoke a published function to test custom provisioning logic with a specific input payload.

> 📍 In the C1 web UI, go to Platform › Functions, select your function, click Run, provide input JSON, then click Run to execute and inspect the output JSON.

— Names every API surface, trigger type and runtime control product docs

02 Dynamic Access Control via JML membership automation NEW

100

A `JML` page on each access profile lets admins configure joiner rules (enrollment approval policy, optional `Maximum membership duration`) and leaver rules (revocation policy, post-unenrollment entitlement handling), plus `Membership automation` settings that select entitlements as the membership condition, define an `Excluded users` list, and toggle `Grant and revoke entitlements on membership change` to drive auto-enrollment

and auto-unenrollment. Leaver rules support three post-unenrollment actions — leave access as-is, revoke all entitlements granted by the profile, or revoke only `unjustified` entitlements not duplicated elsewhere — and both joiner and leaver rules support a 'bypass individual entitlement policies' option treating profile approval as implicit approval for contained entitlements. An anomaly-detection safeguard (`Pause automation when anomalies are detected`) halts any sync removing more than 20% of members within an hour for profiles with 100 or more members, with the cutoff adjustable via Removed members threshold (%); paused profiles show banner alerts and a Slack notification to C1 admins, with a `Review` / `Resume sync` workflow to inspect and confirm staged removals. A `Sync now` button triggers an immediate sync outside the default hourly cadence, a `pending members` count with drill-through surfaces users awaiting enrollment approval on the `Members` page, and an `Enroll users` button lets admins manually add users who don't meet the automation condition.

Configure a contractor access profile to auto-enroll any user who holds a `Contractor` entitlement, auto-approve enrollment, and immediately revoke unjustified access when they leave — with a safeguard that pauses if more than 10% of members would be removed in one hour.

> 📍 1. Open the access profile and go to its JML page. 2. In the Joiner section, click Edit, select your auto-approve policy for the enrollment entitlement, choose `Bypass each entitlement's approval policy` , set Maximum membership duration if needed, then click Save. 3. In the Leaver section, click Edit, select your revocation policy, set post-unenrollment action to `Revoke unjustified entitlements` , choose `Bypass each entitlement's revoke policy` , then click Save. 4. In the Membership automation area, click Edit, select the `Contractor` entitlement under Membership automation, enable 'Grant and revoke entitlements on membership change', set Removed members threshold (%) to 10 in the Anomaly detection section, then click Save. 5. Flip Enabled to On and click Save.

Force an immediate membership sync after adding a new entitlement to an access profile, rather than waiting up to an hour for the next scheduled run.

> 📍 On the access profile's JML page, click Sync now.

Resume a paused membership automation sync after reviewing staged removals that triggered the anomaly-detection safeguard.

> 📍 1. On the main Access profile page, locate the access profile showing a safeguard alert banner. 2. Open the profile and click Review in the banner. 3. Inspect the list of members staged for removal. 4. If the list looks correct, click Resume sync; otherwise click Go back and adjust the Membership automation settings.

— Every config key, threshold and workflow step named `product docs`

03 **Access Profiles for bundled entitlements** NEW 100

Access Profiles (Governance > Access profiles) group apps and entitlements into requestable bundles scoped by audience — all org members or only users holding specific entitlements — and can be requested as a whole via the Profiles tab on the New request form. The JML tab adds Maximum membership duration and Access change behavior settings — Leave access as-is , Revoke all entitlements , or Revoke unjustified entitlements — to cap enrollment length and control grants at expiry. Role mining recommendations with confidence scores (80%+ strong fit, 50–79% possible fit, 49% or lower poor fit) suggest entitlements for a profile based on org-wide access patterns. A Create requests for currently enrolled users checkbox generates access request tasks for all enrolled users when entitlements change, a Manage provisioning action on the profile's Apps tab (including a Request for all bulk option) manually requests missing entitlements for current holders, and entitlements can be added to one or more profiles directly from the entitlement's details page via the Access profiles field in the Access requests section.

Create a department-scoped access profile that only Engineering staff (users holding a specific entitlement) can see and request, keeping the catalog tidy for everyone else.

> 📍 In ConductorOne, go to Governance › Access profiles and click New profile . After naming it, open the Requests tab, click Edit in the Allow requests area, toggle on Allow requests , set Profile is visible to to Users with entitlements , and select the Engineering entitlement from the dropdown. Click Save.

Enforce time-limited access for contractors by capping how long they can hold a profile and automatically revoking grants that aren't covered by another active enrollment when it expires.

> 📍 Open the access profile's JML tab. Set Maximum membership duration to your desired cap (e.g. 90 days). Set Access change behavior to Revoke unjustified entitlements . Save.

Retroactively grant a newly added entitlement to all users already enrolled in a profile without waiting for their next renewal.

> 📍 On the access profile's Apps tab, locate the newly added entitlement, click the '...' (more actions) menu, select Manage provisioning , then click Request for all to create access request tasks for every enrolled user who lacks the entitlement.

— Every setting, tab and bulk action named with values product docs

04 **Secrets inventory tab** NEW 80

A Secrets tab on the Inventory dashboard tracks API tokens and service account keys across connected systems (Okta, GitHub, AWS, Snowflake, Datadog, and Google Cloud Platform), attributing each to an owning identity and alerting on recently-unused or orphaned keys.

Review unrotated or orphaned API keys across all connected SaaS systems before a compliance audit.

> 📍 In the ConductorOne console, go to Inventory › Secrets tab — review flagged secrets that have not been used recently, noting the attributed identity for each, then initiate revocation or rotation for any orphaned keys.

— Names integrations and alert triggers; UI-only path product docs

05 **Zero-knowledge secret sharing** NEW 75

Zero-knowledge secret sharing delivers any secret — passwords, API keys, JSON service-account files, `.env` blocks — as a one-time share to a named recipient inside or outside the org. Encryption happens in the sender's browser so C1 never sees plaintext, and access expires after a configured time or view count.

— Mechanism and expiry controls named, no exact UI path product docs

06 **Revoke inherited access** NEW 75

A Revoke inherited access toggle per app traces entitlement access through direct grants, nested groups, role hierarchies, and delegated provisioning, revoking upstream sources automatically rather than just the immediate entitlement. Approvers see a banner listing downstream effects before a revocation is committed, and the behavior applies to all revocation triggers, including denied access review campaigns.

Enable inherited-access revocation for an app so that removing an entitlement also removes the underlying group memberships or role assignments that granted it — preventing grants from being resurrected at the next sync.

> 📍 In the ConductorOne admin console, open the app, go to Settings, find the Revoke inherited access card, and enable the toggle. Once on, any revocation — manual, policy-driven, or from a denied access review campaign — will trace and remove upstream grants, and approvers will see a banner listing all downstream effects before confirming.

— Explains mechanism and triggers, lacks a named config surface product docs

07 **Workload federation for CI/CD** NEW 75

Workload federation lets CI/CD pipelines and workloads exchange a platform OIDC token (GitHub Actions, GitLab CI, HCP Terraform, SPIFFE, or any custom OIDC issuer)

for a short-lived C1 access token scoped to a service principal's roles, with no stored client secret required.

— Names supported issuers and token exchange mechanism product docs

THINNER COVERAGE BELOW

08 **Decoys tripwire credentials** NEW 55

Decoys are tripwire credentials shaped like real human, service-principal, connector, workload-federation, or long-lived-token credentials that raise a Critical finding automatically when used, serving as a high-confidence indicator of compromise.

— Names credential types and finding severity, no setup steps product docs

09 **Credential stores for provisioning-generated passwords** NEW 50

Credential stores hold provisioning-generated passwords created when C1 provisions a new account through a connector, decryptable by store owners or the account recipient.

— Explains mechanism but no UI or API path given product docs

10 **Enterprise-managed authorization for AI agents via MCP** NEW 50

Enterprise-managed authorization for AI agents reaching an MCP server exchanges the signed-in user's identity for a short-lived, entitlement-checked token scoped to that one server.

— Mechanism stated but no config or endpoint given product docs

11 **Group member count on group detail page** IMPROVED 30

Groups now display member count at the top of the group detail page.

— Bare UI addition, no further detail changelog-20260926-e7a9b8d7

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! A revoke criteria scoped only by a legacy resource-type name now fails the step instead of silently matching every app sharing that name — re-scope any automation still relying on name-only matching.

! `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 SDK methods named Revoke are renamed to Delete.

! `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` value.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

ConductorOne: Removes `GET /api/v1/tasks/{task_id}/spend-remedy-status`; JS/TS clients must handle spend-control revision values as strings and handle new `TYPE_SPEND_REMEDY = 7` task action type.

└──▷ **TRY IT**

Revoke a credential after deleting a shared secret — now a two-step operation since deletion no longer triggers revocation automatically.

```
$ 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>"}'
```

Automatically assign access findings to the right owner at classification time using the new Transformation rule step.

📍 In the C1 console, go to Governance › Transformation rules, open or create a rule, add a Set assignee step, choose the target user, and save — manual assignments on individual findings will still take precedence.

Renders entitlement request forms directly in the Slack modal when requesting or extending access through the C1 app for Slack.

Groups now display member count at the top of the group detail page.

└──▷ **BREAKING ON UPGRADE**

! A revoke criteria scoped only by a legacy resource-type name now fails the step instead of silently matching every app sharing that name — re-scope any automation still relying on name-only matching.

! `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 SDK methods named Revoke are renamed to Delete.

! `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` value.

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ changelog · Sep 26, 2026

CHANGELOG ↗ RANK

Discover every SaaS and AI asset, harden security posture, and scale identity governance.

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

01 **In-browser OAuth device-code phishing warning** NEW

60

A new in-browser nudge warns employees on OAuth device code authorization pages for Microsoft, Google, GitHub, GitLab, Salesforce, AWS, Docker, Zoom, Vercel, and OpenAI; every acknowledgment of the warning is recorded as a nudge response.

— Names ten providers and recording mechanism, no UI steps given

changelog-20260926-4265b231

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Nudge Security adds batched password-manager nudges and in-browser device-code phishing warnings across 10 OAuth providers.

└──▷ **HOW TO FIND IT**

Sweep up shadow password-manager use in one campaign — send every employee a single nudge that lists all their unapproved accounts at once.

📍 In the Nudge Security console, go to Accounts Inventory › Filter by Password Manager › select Everything not set as approved › multi-select results › Send nudge: Use approved password manager

Adds a batched Use approved password manager nudge: filter the accounts inventory by observed password manager (including not approved), multi-select, and each employee receives one consolidated nudge listing all their affected accounts.

Adds an in-browser nudge that warns employees on OAuth device code authorization pages for Microsoft, Google, GitHub, GitLab, Salesforce, AWS, Docker, Zoom, Vercel, and OpenAI — every acknowledgment is recorded as a nudge response.

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ changelog · Sep 26, 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 biggest moves this window were the launch of SaaS Management, a Password Vault in the User Portal, unified Patch Management dashboards, and Dynamic Device/User Groups, alongside new SentinelOne and Microsoft 365 integrations and an AI Assistant knowledge base — while deprecating the GET reactivate endpoint, older Dynamic Groups filter operators, weak LDAP cipher suites, and the Linux AppImage build.

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

01

Unified Patch Management dashboards NEW

95

New unified Patch Management dashboards live under `Device Management > Patch Management` with Overview, Windows, and Apple tabs; admins can filter the KB list by Global Severity (e.g. Critical), bulk-remediate missing patches via Install missing critical patches /'Install Now', and export the KB list or a full per-device Device Patch Report as CSV or JSON, with inline preview of up to 500 rows.

Identify Windows devices missing critical patches and trigger immediate remediation without waiting for the next policy window.

> 📍 In the Admin Portal, go to Patch Management › Windows, apply the Global Severity filter to Critical , review the KB list for missing patches, select affected devices, and choose Install Now under Bulk Remediation.

Export a full per-device Windows KB patch status report for audit or ticketing workflows.

> 📍 In the Admin Portal, go to Patch Management › Device Patch Report (Windows), preview up to 500 rows inline, then click Export to download the complete organization result set as CSV.

Identify Windows devices missing critical patches across the fleet and bulk-remediate them immediately from one view.

> 📍 In the Admin Portal, go to Device Management › Patch Management, open the Windows tab, use the Global Severity filter to select Critical , review the KB list for missing patches, then select affected devices and choose Install missing critical patches to remediate immediately.

— Names exact UI path, filters, export formats, and row limit.

Release Notes 2026 (2026-09-25)

02

SSO forced re-authentication controls NEW

85

JumpCloud Identity Provider now honors `ForceAuthn="true"` in SAML and `prompt=login` in OpenID Connect, prompting fresh credentials even during active sessions; forced re-authentication events are visible in Directory Insights by filtering event type `sso_auth` .

Audit which users triggered forced re-authentication events across your SAML and OIDC apps for a compliance review.

> 📍 In Directory Insights, filter events by type `sso_auth` to view all forced re-authentication events with associated user, application, and timestamp details.

— Names exact SAML/OIDC parameters and audit filter. Release Notes 2025 (2026-09-25)

03 **Dormant application access detection** NEW 80

Dormant Application Access Detection in `Access Risk > Configuration` now alerts only when a user signs back into an SSO app after a configured inactivity threshold (e.g. 90 days), rather than triggering on first-time provisioning access.

Adjust the dormant application inactivity threshold so Access Risk only alerts when a user returns to an SSO app after a meaningful gap — not on first provisioning.

> 📍 In the Admin Portal, go to Access Risk › Configuration and set the dormant application threshold to your desired inactivity period (e.g. 90 days).

Reduce Access Risk noise from newly provisioned SSO apps by tuning the dormant access inactivity threshold.

> 📍 In the Admin Portal, go to Access Risk › Configuration and set the dormant application threshold to your desired inactivity period; the factor will now trigger only on sign-in after inactivity, not on first-time access.

Tune dormant app detection to flag re-access after 90 days of inactivity without generating noise from newly provisioned users.

> 📍 In the Admin Portal, go to Access Risk › Configuration, locate the Dormant Application Access threshold setting, and set your inactivity period (e.g. 90 days). First-time logins will not trigger the factor; only return logins after the threshold will.

— Names config location and threshold behaviour precisely.

Release Notes 2026 (2026-09-25) changelog-20260926-b05ed387

---

04 **LDAP cipher suite restriction** BREAKING 80

LDAP connections using cipher suites relying on CBC mode or SHA-1/SHA-256 with RSA or ECDHE key exchange are no longer accepted; only `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and `TLS_RSA_WITH_AES_128_GCM_SHA256` are now permitted.

— Lists every permitted cipher suite explicitly. Release Notes 2025 (2026-09-25)

---

05 **macOS Ventura upgrade blocking policy** NEW 75

New macOS software management policies let admins delay or block unwanted upgrades to macOS 13 Ventura, applied via `Device Management > Policies > New Policy` selecting the Delay or Block MacOS Ventura policy for a target device group.

Block unauthorized macOS Ventura upgrades across managed devices before users trigger them unintentionally.

> 📍 In the Admin Portal, go to Device Management › Policies › New Policy, select macOS, and apply the Delay or Block MacOS Ventura software management policy to the target device group.

— Names exact policy and UI navigation path. Release Notes 2022 (2026-09-25)

---

06 **Push MFA user group exclusions** NEW 70

Admins can exclude a user group from Push MFA enforcement via `Security > MFA > JumpCloud Protect > User Group Exclusions`, letting excluded users authenticate with TOTP only.

Exclude a remote region user group from Push MFA enforcement so those users can authenticate using TOTP only via JumpCloud Protect.

> 📍 In the Admin Portal, go to Security › MFA › JumpCloud Protect, open User Group Exclusions, and add the target user group to exclude it from Push enforcement.

— Names exact UI path and fallback method. Release Notes 2025 (2026-09-25)

---

07 **Password Vault in User Portal** NEW 60

Adds Password Vault in the User Portal with bulk import via CSV or Excel, credential sharing with granular permissions, and personal/shared folder organization.

— Names import formats and sharing model but no config key.

Release Notes 2026 (2026-09-25)

THINNER COVERAGE BELOW

08 **SentinelOne integration via Commands** NEW 55

Adds a SentinelOne integration via JumpCloud Commands to download and install the SentinelOne Agent on macOS, Windows, and Linux devices.

— Names platforms and delivery mechanism, no exact command shown.

Release Notes 2022 (2026-09-25)

---

09 **Android Bluetooth sharing restrictions expansion** IMPROVED 50

The Android Bluetooth Restrictions policy now controls outbound Bluetooth file sharing, and the Work & Personal Usage policy adds personal-profile Bluetooth sharing restrictions for company-owned devices that have a work profile.

— Names the two policies but no config path or example. product docs

---

10 **MSP billing visibility** NEW 50

Adds MSP billing visibility showing committed user licenses, available and consumed counts, and a downloadable JumpCloud contract.

— Names metrics shown but no UI path. Release Notes 2023 (2026-09-25)

11 **Access Requests tracking and webhooks** NEW 50

Adds Access Requests, tracked as events in Directory Insights, with webhook support for real-time ticket creation in external systems.

— Names webhook capability but no endpoint or payload. Release Notes 2025 (2026-09-25)

12 **Remote Assist for Windows Server login screen** IMPROVED 45

Enables Remote Assist access to the initial Windows Server login screen, allowing admins to authenticate and interact before the desktop loads.

— Describes behaviour, no config detail given. Release Notes 2023 (2026-09-25)

13 **Interactive password complexity dialog** IMPROVED 40

Adds an interactive as-you-type password complexity dialog with a mobile-responsive, WCAG-compliant design for setting policy-compliant passwords.

— Describes UI behaviour but no config path or example. Release Notes 2022 (2026-09-25)

14 **Microsoft 365 security group integration** NEW 30

Adds Microsoft 365 integration support for M365/Azure AD security groups.

— One-line description with no mechanism or UI path. Release Notes 2022 (2026-09-25)

15 **Multiple IP addresses in RADIUS config** IMPROVED 30

Adds multiple IP address support to RADIUS server configuration.

— One-line addition with no configuration detail. Release Notes 2024 (2026-09-25)

16 **Android 15 support in MDM** NEW 15

Supports Android 15 in MDM.

— Bare one-line mention, no scope or mechanism. Release Notes 2024 (2026-09-25)

└──▷ **ALSO FROM THESE RELEASES**

Protect admin accounts against MFA fatigue by verifying Push Bombing Mitigation is active and reviewing blocked attempts in Directory Insights.

> 📍 In the Admin Portal, go to Security › MFA Configurations › JumpCloud Protect, confirm push bombing mitigation is enabled (default: max 1 notification per resource per 60 seconds), then review blocked events under Insights › Directory Insights.

Filter system users whose department is `Finance | Accounting` or `IT`, using a backslash to escape the literal pipe character inside the search term.

```
$ GET https://console.jumpcloud.com/api/systemusers?
filter=department:$in:Finance \| Accounting|IT
```

Enforce MFA for all MSP admins across the tenant from the Multi-Tenant Portal without touching individual admin accounts.

> 📍 In the Admin Portal, go to Account › Settings and enable `Global MFA Requirement`.

Subscribe to JumpCloud product release note and agent update notifications in your RSS reader to stay current without manual polling.

> 📍 1. Go to the JumpCloud Release Notes and Bug Fixes page. 2. Click `Subscribe to Release Notes and Bug Fixes RSS Feed`. 3. Open the resulting RSS URL in your RSS reader (most browsers have a built-in reader; Chrome users can install the `RSS Feed Reader` extension).

└──▷ **BREAKING ON UPGRADE**

- ! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; any script, automation, or integration using it must be updated to use `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.
- ! The `$le` and `$ge` memberQuery filter operators in the Dynamic Groups API are deprecated and replaced by `$lte` and `$gte`; API requests using the old operators must be updated.
- ! LDAP connections using cipher suites relying on CBC mode or SHA-1/SHA-256 with RSA or ECDHE key exchange are no longer accepted; only `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and `TLS_RSA_WITH_AES_128_GCM_SHA256` are now permitted.
- ! The `.appimage` Linux build is no longer supported as of v2.2.197 (April 18, 2024); users must switch to the RPM or Debian build.

> changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

JumpCloud: Adds Knowledge Base Uploads for AI Assistant, letting admins upload internal documentation in PDF, TXT, DOCX, and MD formats (10 MB max per file, 25 files max per org) to an org-scoped knowledge base the assistant references when answering questions.

└──▷ **HOW TO FIND IT**

Adjust the dormant application inactivity threshold so Access Risk only alerts when a user returns to an SSO app after a meaningful gap — not on first provisioning.

> 📍 In the Admin Portal, go to Access Risk › Configuration and set the dormant application threshold to your desired inactivity period (e.g. 90 days).

Upload a runbook or SOP so AI Assistant can answer org-specific questions without exposing data outside your org.

> 📍 In the Admin Portal, go to AI Assistant › Knowledge Base and upload your internal documentation file (PDF, TXT, DOCX, or MD; 10 MB max per file, 25 files max per org).

Match the OAuth client authentication method required by Amazon Bedrock AgentCore or another external platform for an AI agent.

> 📍 In the Admin Portal, open the agent › Authentication tab, and under client authentication type select `Client secret post` to send the client ID and secret in the request body instead of the authorization header.

`Release Notes 2022 (2026-09-25)` `BLOG ↗` `STABLE`

JumpCloud: Adds SentinelOne integration via JumpCloud Commands to download and install the SentinelOne Agent on macOS, Windows, and Linux devices.

└──▷ **HOW TO FIND IT**

Block unauthorized macOS Ventura upgrades across managed devices before users trigger them unintentionally.

> 📍 In the Admin Portal, go to Device Management › Policies › New Policy, select macOS, and apply the `Delay or Block MacOS Ventura` software management policy to the target device group.

Adds Microsoft 365 integration support for M365/Azure AD security groups.

Adds new macOS software management policies to delay or block unwanted upgrades to macOS 13 Ventura.

Adds interactive as-you-type password complexity dialog with mobile-responsive, WCAG-compliant design for setting policy-compliant passwords.

Release Notes 2023 (2026-09-25) BLOG ↗ STABLE

JumpCloud: Adds Dynamic Device Groups and Dynamic User Groups allowing admins to automate group membership via attribute-based rules, with an optional review step before changes take effect.

└──▷ **HOW TO FIND IT**

Automatically scope a device group to all Windows endpoints on a specific public IP range without manual membership management.

> 📍 In the Admin Portal, go to Device Groups › create or edit a group › enable Dynamic › add a rule for `Public IP Address` matching your corporate egress range, then optionally enable membership review before changes apply.

Protect admin accounts against MFA fatigue by verifying Push Bombing Mitigation is active and reviewing blocked attempts in Directory Insights.

> 📍 In the Admin Portal, go to Security › MFA Configurations › JumpCloud Protect, confirm push bombing mitigation is enabled (default: max 1 notification per resource per 60 seconds), then review blocked events under Insights › Directory Insights.

Adds MSP billing visibility for committed user licenses, available and consumed counts, and JumpCloud contract download.

Enables Remote Assist access to the initial Windows Server login screen, allowing admins to authenticate and interact before desktop load.

Release Notes 2024 (2026-09-25) BLOG ↗ STABLE

JumpCloud: Replaces deprecated `GET api/systemusers/reactivate/{id}` with `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` (the GET now returns 404).

└──▷ **TRY IT**

Filter system users whose department is `Finance | Accounting` or `IT`, using a backslash to escape the literal pipe character inside the search term.

```
$ GET https://console.jumpcloud.com/api/systemusers?
filter=department:$in:Finance \| Accounting|IT
```

Reactivate a user account using the new POST endpoint after migrating away from the deprecated GET request.

```
$ curl -X POST
'https://console.jumpcloud.com/systemusers/reactivate/{id}' -H 'x-
api-key: <api_key>' -H 'Content-Type: application/json' -d
'{"activation_email": "user@example.com"}'
```

Enforce MFA for all MSP admins across the tenant from the Multi-Tenant Portal without touching individual admin accounts.

> 📍 In the Admin Portal, go to Account › Settings and enable `Global MFA Requirement`.

Supports Android 15 in MDM.

Launches SaaS Management, enabling discovery and tracking of SaaS app usage via the JumpCloud Go browser extension, SSO integration, app categorization, and access restriction with customizable warning or blocking messaging.

Adds multiple IP address support to RADIUS server configuration.

> └──▷ **BREAKING ON UPGRADE**
>
> ! The `GET api/systemusers/reactivate/{id}` endpoint is deprecated and now returns a 404 error; any script, automation, or integration using it must be updated to use `POST https://console.jumpcloud.com/systemusers/reactivate/{id}` with `activation_email` in the request body.
>
> ! The `$le` and `$ge` memberQuery filter operators in the Dynamic Groups API are deprecated and replaced by `$lte` and `$gte`; API requests using the old operators must be updated.

`Release Notes 2025 (2026-09-25)` `BLOG ↗` `STABLE`

JumpCloud adds Access Requests, MCP Server, AI Commands Builder, Shadow AI Dashboard, Custom Admin Roles, and deeper SSO/MFA controls in 2025.

└──▷ **HOW TO FIND IT**

Audit which users triggered forced re-authentication events across your SAML and OIDC apps for a compliance review.

> 📍 In Directory Insights, filter events by type `sso_auth` to view all forced re-authentication events with associated user, application, and timestamp details.

Exclude a remote region user group from Push MFA enforcement so those users can authenticate using TOTP only via JumpCloud Protect.

📍 In the Admin Portal, go to Security › MFA › JumpCloud Protect, open `User Group Exclusions`, and add the target user group to exclude it from Push enforcement.

JumpCloud Identity Provider now honors `ForceAuthn="true"` in SAML and `prompt=login` in OpenID Connect, prompting fresh credentials even during active sessions.

Tracks Access Request events in Directory Insights and supports webhooks for real-time ticket creation in external systems.

└──▷ **BREAKING ON UPGRADE**

! LDAP connections using cipher suites relying on CBC mode or SHA-1/SHA-256 with RSA or ECDHE key exchange are no longer accepted; only `TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256`, `TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384`, `TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384`, and `TLS_RSA_WITH_AES_128_GCM_SHA256` are now permitted.

`Release Notes 2026 (2026-09-25)` **BLOG ↗** STABLE

JumpCloud: Adds `Client secret basic` and `Client secret post` client authentication type options on an agent's `Authentication` tab for OAuth credential configuration.

└──▷ **HOW TO FIND IT**

Tune dormant app detection to flag re-access after 90 days of inactivity without generating noise from newly provisioned users.

📍 In the Admin Portal, go to Access Risk › Configuration, locate the Dormant Application Access threshold setting, and set your inactivity period (e.g. 90 days). First-time logins will not trigger the factor; only return logins after the threshold will.

Identify Windows devices missing critical patches across the fleet and bulk-remediate them immediately from one view.

📍 In the Admin Portal, go to Device Management › Patch Management, open the Windows tab, use the Global Severity filter to select `Critical`, review the KB list for missing patches, then select affected devices and choose `Install missing critical patches` to remediate immediately.

Upload an internal runbook so AI Assistant can answer org-specific process questions without exposing it to external models.

📍 In the Admin Portal, open AI Assistant › Knowledge Base, click Upload, and select your runbook file (PDF, TXT, DOCX, or MD, up to 10 MB). The assistant will cite the source document when it draws on the uploaded content in responses.

Adds Password Vault in the User Portal with bulk import via CSV or Excel, credential sharing with granular permissions, and personal/shared folder organization.

Adds unified Patch Management dashboards under `Device Management > Patch Management` with Overview, Windows, and Apple tabs; supports CSV and JSON export of the KB list.

Enables Dormant Application Access Detection in `Access Risk > Configuration` to alert only when a user signs back into an SSO app after a configured inactivity threshold, not on first-time access.

What's New: Password Manager Desktop App (2026-09-25) BLOG ↗ STABLE

JumpCloud: Adds an ARM build for macOS devices with M1 and above Apple Silicon chips, available on the Downloads page with automatic update detection.

└──▷ **HOW TO FIND IT**

When rolling out Password Manager to a fleet without requiring users to have admin rights for future updates, deploy via the JumpCloud command template.

📍 In the JumpCloud Admin Portal, go to Commands and run the Install and Update The JumpCloud Password Manager App command template v.2.0.0 targeting Windows or macOS devices.

Introduces item sorting by nickname.

└──▷ **BREAKING ON UPGRADE**

! The `.appimage` Linux build is no longer supported as of v2.2.197 (April 18, 2024); users must switch to the RPM or Debian build.

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ

New Capability: Manual Task Support for Governance Groups (2026-09-25) BLOG ↗

RANK

SailPoint is an identity governance and access management platform that automates user provisioning, access reviews, and compliance across enterprise systems.

> SailPoint now lets manual provisioning tasks be routed to entire Governance Groups rather than a single owner, with full audit attribution and updated notifications.

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

01 Governance Groups for manual provisioning tasks NEW

70

Sources can now be assigned a Governance Group so manual provisioning tasks route to all group members instead of a single owner, with ISC recording which individual member completed each task. The Pending Manual Changes, Pending Task Daily Digest, and Task Reassignment notifications now reach all members of an assigned Governance Group, and group-assigned tasks appear as a single row listing all assigned members in the admin task view.

— Names notifications and view but no config/API path given.

New Capability: Manual Task Support for…

> New Capability: Manual Task Support for Governance Groups (2026-09-25) BLOG ↗
>
> LATEST STABLE
>
> SailPoint ISC now routes manual provisioning tasks to Governance Groups, letting any member complete work with full audit attribution.
>
> Enables assigning a Governance Group to a source so that manual provisioning tasks are routed to all group members instead of a single owner, with ISC recording which individual member completed each task.
>
> Extends existing task notifications — Pending Manual Changes, Pending Task Daily Digest, and Task Reassignment — to reach all members of an assigned Governance Group.
>
> Displays group-assigned manual tasks as a single row listing all assigned members in the admin task view.

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password SaaS Manager #1805 (2026-09-24)

BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password SaaS Manager's September release adds an AI Assistant beta for app risk questions, new API and integration surfaces including `GET /contracts` enrichment and Kolide EU region support, and a batch of Access Policy, Workflow and profile-management refinements.

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

01 **Main application field in contracts API** IMPROVED 60

`GET /contracts` API responses now include the main application associated with each contract.

— Named endpoint and field, but no example response shown

1Password SaaS Manager #1805 (2026-09-2…

THINNER COVERAGE BELOW

02 **Workflow step enhancements for triggers and people** IMPROVED 55

Workflows gain: triggers that fire a remediation workflow when another workflow step fails; the ability to create or update a person by email address within a Workflow; and an extension to the Set account groups workflow step to use values from custom text fields on a person.

— Names three specific workflow mechanisms without full config detail

1Password SaaS Manager #1805 (2026-09-2…

03 **Custom person field editing from profile** IMPROVED 50

Adds the ability to set or clear custom person fields, including number fields with a value of 0, directly from a person's profile.

— Names exact location and edge case handled 1Password SaaS Manager #1805 (2026-09-2…

04 **Temporary access policies and matching order** NEW 45

Access Policies now support time-limited grants via new temporary-access-only Access Policies, which no longer require setting an unlimited access duration option. When multiple policies match, SaaS Manager now prioritizes policies scoped to specific people and teams over broader ones.

— Explains behavior change but no config path shown

1Password SaaS Manager #1805 (2026-09-2…

05 **Automatic usage-data retention policy** NEW 45

SaaS Manager now automatically retains daily usage data for one year and monthly usage data for five years.

— Exact retention numbers given but no user action possible

1Password SaaS Manager #1805 (2026-09-2…

06 **Workflow Runs UI improvements** IMPROVED 40

Workflow Runs now show the external ID that triggered a workflow step, and gain row-selection checkboxes for easier multi-edit actions.

— Two concrete UI changes but no walkthrough given

1Password SaaS Manager #1805 (2026-09-2…

07 **BambooHR Division-to-team mapping** IMPROVED 40

The BambooHR integration now supports mapping Division values to the Team for Later app.

— Names the integration and field mapped, no setup steps

1Password SaaS Manager #1805 (2026-09-2…

08 **Team ownership of applications in App Catalog** NEW 35

Adds the ability to assign a team as the owner of an application in the App Catalog.

— Clear UI action but no navigation path detailed

1Password SaaS Manager #1805 (2026-09-2…

09 **AI Assistant beta for app risk questions** NEW 30

A new AI Assistant, in public beta, answers questions about app risk levels.

— States purpose but not scope, inputs, or access path

1Password SaaS Manager #1805 (2026-09-2…

10 **Time increment picker for AI Spend chart** IMPROVED 30

Adds a time increment picker to the AI Spend and Consumption chart for switching between detail levels.

— Names the chart and control but not the increments

1Password SaaS Manager #1805 (2026-09-2…

11 **Kolide EU region connection** NEW 25

SaaS Manager adds support for connecting to Kolide's EU region.

— Bare feature name, no setup steps given 1Password SaaS Manager #1805 (2026-09-2…

1Password SaaS Manager #1805 (2026-09-24) BLOG ↗ LATEST STABLE

1Password SaaS Manager adds AI Assistant beta, temporary access policies, Kolide EU region, and contract API enrichment.

Adds `GET /contracts` API responses that now include the main application associated with each contract.

Adds an AI Assistant (public beta) that answers questions about app risk levels.

Adds temporary-access-only Access Policies, allowing time-limited grants without requiring an unlimited access duration option.

Adds support for connecting to Kolide's EU region.

Adds the ability to assign a team as the owner of an application in the App Catalog.

+10 more

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ changelog · Sep 26, 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 OAuth client-onboarding controls for MCP-style public clients, new SMS delivery event logging, and eased enterprise SSO limits by allowing multiple SAML/OIDC connections and multiple Directory Sync feeds per Organization.

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

01 **OAuth client onboarding controls for MCP** NEW 83

OAuth application settings gain a Client onboarding section with a Publish CIMD support toggle, letting MCP and other public OAuth clients authenticate via an HTTPS URL as `client_id` without a pre-issued secret or Dynamic Client Registration. A companion Client admission control adds a Pre-registered clients only mode to restrict OAuth flows to reviewed clients.

Lock down your MCP server's OAuth flow to only pre-reviewed clients after enabling Client ID Metadata Document support.

> 📍 In the Clerk Dashboard, go to OAuth applications settings › Client onboarding, enable Publish CIMD support, then set Client admission to Pre-registered clients only and add trusted client ID URLs from the Applications tab.

— Names exact toggle and admission mode with dashboard path. changelog-20260926-ef501d89

THINNER COVERAGE BELOW

02 **Multi-directory sync per Organization** NEW 50

A single Organization can now receive Directory Sync feeds from multiple directories, with per-directory membership and deprovisioning tracked independently.

— Describes behavior but no config or API name given. changelog-20260926-ef501d89

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Clerk: Enables Publish CIMD support toggle under Client onboarding in OAuth application settings, letting MCP and other public OAuth clients authenticate via HTTPS URL as `client_id` without a pre-issued secret or Dynamic Client Registration.

└──▷ **HOW TO FIND IT**

Lock down your MCP server's OAuth flow to only pre-reviewed clients after enabling Client ID Metadata Document support.

> 📍 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.

Investigate a user-reported missing SMS verification code by filtering SMS logs to their phone number and checking the normalized failure reason.

> 📍 In the Clerk Dashboard, open the SMS logs view, filter by the destination phone number, and inspect the `sms.failed` or `sms.undeliverable` event's `reason` field (e.g. `invalid_phone_number` or `destination_unreachable`) to determine the remediation action.

Support an enterprise customer who uses two separate identity providers (e.g. post-acquisition) within a single Clerk Organization.

> 📍 In the Clerk Dashboard, go to SSO connections, select the Organization that already has a SAML or OIDC connection, and add a second connection with its own domain, protocol, and Directory Sync settings.

Adds `Client admission` control to OAuth application settings, with a `Pre-registered clients only` mode to restrict flows to reviewed clients.

Introduces SMS delivery events (`sms.accepted`, `sms.delivered`, `sms.failed`, `sms.undeliverable`, `sms.unconfirmed`) with a normalized `reason` field (e.g. `invalid_phone_number`, `destination_unreachable`, `rate_limited`) logged in a dedicated SMS logs view for all plans.

Allows Organizations to have more than one SAML or OIDC enterprise connection simultaneously, removing the previous `organization_already_has_sso_connection` limit, across the Clerk Dashboard and Backend API.

Enables a single Organization to receive Directory Sync feeds from multiple directories, with per-directory membership and deprovisioning tracked independently.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Twingate

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗

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: Expands each AI service entry in the AI Usage Overview to show matched domains and indicate whether activity is widespread or concentrated on specific devices.

└──▷ HOW TO FIND IT

Identify shadow AI tool usage across your fleet and export per-device data to drive policy discussions or license reviews.

> 📍 In the Twingate admin console, go to DNS Filtering › AI Usage Overview, review which AI services appear in DNS activity, expand a service to see its matched domains, then select Export CSV to download per-device activity data.

WAS THIS USEFUL? 👍 👎

DNSFilter

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

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

> DNSFilter's macOS Roaming Client 2.4.6 adds IPv6 support, DoH VPN fallback, and component mismatch detection while raising the minimum OS to macOS 13 Ventura and requiring a new root CA certificate; the dashboard gains a block-reason column, deployment access controls for restricted roles, and safeguards against accidental deletions.

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

01 IPv6 and DoH fallback in macOS Roaming Client 2.4.6 NEW

80

The macOS Roaming Client 2.4.6 adds IPv6 support, DoH VPN fallback, and component mismatch detection. DoH fallback can be enabled at install time via `ENABLE_DOH_FALLBACK=true` in `dns_agent.conf`, keeping DNS filtering enforced for remote workers whose corporate VPN blocks the DNS-over-TLS port.

Keep DNS filtering enforced for remote workers whose corporate VPN blocks the DNS-over-TLS port, by enabling DoH fallback at install time.

```env
ENABLE_DOH_FALLBACK=true
```

— Names exact config key and file, plus purpose changelog-20260926-4d627d69

02 macOS 13 and new root CA required for 2.4.6 BREAKING

80

The macOS Roaming Client now requires macOS 13 Ventura or later; devices on earlier macOS versions cannot run v2.4.6. It also checks for the new DNSFilter Root CA G1 certificate, and devices with the previous certificate profile will show a Fix Installation Issues prompt until the updated `Combined_Certificates.mobileconfig` (or branded/white-label equivalent) is pushed via MDM.

— Clear migration requirement with exact file to deploy changelog-20260926-4d627d69

03 Scoped deployment access controls for restricted roles NEW

65

Adds deployment access controls scoped to restricted roles: read-only and policy-only users now have Site Keys hidden, Relay deployment controls disabled, and Roaming

Client installer access restricted.

— Names roles and each restricted surface, UI-only path changelog-20260926-4d627d69

04 **'Blocked By' column in request grids** NEW 60

Adds a Blocked By column to the Active Requests and Request History grids, showing the reason each domain was blocked, consistent with the Query Log's Method column, and also visible in the approval drawer.

— Names grids and related Query Log column, no config key changelog-20260926-4d627d69

THINNER COVERAGE BELOW

05 **New columns in CyberSight Top Clients grids** NEW 55

Adds Last Logged In User as a default column and Friendly Name as an optional column in the CyberSight Top Clients View All grids on Activity Overview and AI Usage.

— Names exact fields and pages, but only UI path changelog-20260926-4d627d69

06 **Type-to-confirm on Roaming Client deletions** IMPROVED 40

Adds a type-to-confirm requirement for Roaming Client delete and cleanup deletion actions to prevent accidental mass deletions.

— States purpose but no exact UI steps changelog-20260926-4d627d69

07 **Direct link to Unblock Requests in admin emails** IMPROVED 40

Adds a direct link to the Unblock Requests page in admin notification emails, with the subject line including the requester's organization name.

— Simple prose description, no config or endpoint changelog-20260926-4d627d69

└──▷ BREAKING ON UPGRADE

! The macOS Roaming Client now requires macOS 13 Ventura or later; devices on earlier macOS versions cannot run v2.4.6.

! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices with the previous certificate profile will show a Fix Installation Issues prompt until the updated `Combined_Certificates.mobileconfig` (or branded/white-label equivalent) is pushed via MDM.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

DNSFilter macOS Roaming Client 2.4.6 adds IPv6 support, DoH VPN fallback, and component mismatch detection; dashboard gains Blocked By column and scoped deployment access controls.

└──▷ USE IT

Keep DNS filtering enforced for remote workers whose corporate VPN blocks the DNS-over-TLS port, by enabling DoH fallback at install time.

dns_agent.conf

```
ENABLE_DOH_FALLBACK=true
```

Adds a `Blocked By` column to the Active Requests and Request History grids, showing the reason each domain was blocked (consistent with the Query Log's Method column), also visible in the approval drawer.

Adds `Last Logged In User` as a default column and `Friendly Name` as an optional column in the CyberSight Top Clients `View All` grids on Activity Overview and AI Usage.

Adds type-to-confirm requirement for Roaming Client delete and cleanup deletion actions to prevent accidental mass deletions.

Adds deployment access controls scoped to restricted roles: read-only and policy-only users now have Site Keys hidden, Relay deployment controls disabled, and Roaming Client installer access restricted.

Adds direct link to the Unblock Requests page in admin notification emails, with subject line including the requester's organization name.

└──▷ **BREAKING ON UPGRADE**

- ! The macOS Roaming Client now requires macOS 13 Ventura or later; devices on earlier macOS versions cannot run v2.4.6.
- ! v2.4.6 checks for the new DNSFilter Root CA G1 certificate; devices with the previous certificate profile will show a `Fix Installation Issues` prompt until the updated `Combined_Certificates.mobileconfig` (or branded/white-label equivalent) is pushed via MDM.

WAS THIS USEFUL? 👍 👎

tailscale

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale shipped a device posture status detail view in the admin console, tailnet event webhooks, a container `log-mode` flag, and introduced Tailscale PAM (beta) with identity-aware connectors for SSH, databases, and Kubernetes.

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

01 Container log verbosity control via log-mode flag NEW

75

Adds a `log-mode` parameter to the container to control whether logging output is grouped, ungrouped, or quiet, e.g. `docker run tailscale/tailscale --log-mode=quiet` to suppress noise in CI pipelines or `--log-mode=ungrouped` for structured log aggregators.

Control container log verbosity — use `quiet` in CI pipelines to suppress noise, or `ungrouped` when forwarding to structured log aggregators.

```
$ docker run tailscale/tailscale --log-mode=quiet
```

— Names exact flag and values with a runnable command changelog-20260926-234c8f4a

02 Tailscale PAM (beta) with identity-aware connectors NEW

70

Introduces Tailscale PAM (beta), a Privileged Access Management layer combining VPN and PAM capabilities for identity-aware access to private tailnet resources. Adds PAM connectors — application-aware identity proxies that establish sessions between tailnet clients and services such as SSH, databases, and Kubernetes — with Ansible-based installation/configuration across one or more hosts, high-availability connector deployment to maintain connectivity during a single connector failure, and protocol-aware capabilities including session recording, policy validation, and secrets management.

— Names PAM, connectors, Ansible, HA, session recording but no config keys or commands

product docs

---

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

> tailscale: Adds a device posture status detail view in the admin console showing each posture's required assertions, per-machine values, pass/fail state, and a count of policy file rules requiring that posture.

└──▷ **TRY IT**

Control container log verbosity — use `quiet` in CI pipelines to suppress noise, or `ungrouped` when forwarding to structured log aggregators.

```
$ docker run tailscale/tailscale --log-mode=quiet
```

Supports subscribing to tailnet events and routing notifications through an integration or app via webhooks.

Adds a `log-mode` parameter to the container to control whether logging output is grouped, ungrouped, or quiet.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Ubiq Security

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗

RANK

Learn how Ubiq discovers sensitive data, maps identity-to-data access, and enforces protection policies across applications, databases, APIs, and AI.

> Ubiq Security shipped native Apache Airflow operators for column-level structured encryption, five new structured dataset types (Integer, Date, Date & Time, Token, and Generic String) spanning its language libraries, and expanded Access Graph tracing with IDP groups, alongside a new Grafana Loki SIEM destination and IDP-based authentication for the JavaScript library.

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

01 **Airflow operators for structured encryption** NEW 95

Ubiq adds `UbiqEncryptOperator` and `UbiqDecryptOperator` Apache Airflow operators (installed via `pip install ubiq_security`, placed in `plugins/operators/`) that perform structured encrypt and decrypt operations directly inside DAG tasks. Both operators take `credentials`, `data_task`, and `dataset_names` parameters for column-level encryption control over tabular data passed as a list of lists (compatible with `BigQueryGetDataOperator` output), with per-column dataset mapping via `dataset_names` — pass an empty string to leave a column unencrypted, and the array length must match row width.

Encrypt specific columns (SSN, Birthdate) from a BigQuery pull inside an Airflow DAG, then decrypt them in a downstream task — leaving other columns untouched.

```python
from operators.ubiq_encrypt_operator import UbiqEncryptOperator
from operators.ubiq_decrypt_operator import UbiqDecryptOperator
import ubiq_security as ubiq

creds = ubiq.configCredentials(config_file="ubiq-credentials",
profile="default")

get_data = BigQueryGetDataOperator(
    task_id="get_data",
    dataset_id=DATASET_NAME,
    table_id=TABLE_NAME,
    max_results=10,
    selected_fields="name,email,ssn,birthday",
)

enc_task = UbiqEncryptOperator(
    task_id="enc_task",
    credentials=creds,
    data_task=get_data.task_id,
    dataset_names=["", "", "SSN", "Birthdate"]
)

dec_task = UbiqDecryptOperator(
    task_id="ubiq_dec_task",
    credentials=creds,
    data_task=enc_task.task_id,
    dataset_names=["", "", "SSN", "Birthdate"]
)

get_data >> enc_task >> dec_task
```

— Full operator names, parameters, and a runnable DAG example are given.

changelog-20260926-d6c55ad6

02 New structured dataset types across libraries NEW

90

Ubiq datasets gain five new structured types: Integer (32-bit range -99,999,999 to 99,999,999 or 64-bit range -9,999,999,999,999,999 to 9,999,999,999,999,999, encrypting to the same integer type), Date (01/01/0001 to 11/28/2738, encrypting to a date with no time component), Date & Time (1653-02-10 06:13:21 to 2286-11-20 17:46:39 UTC, encrypting to a timestamp in the same column), Token (fixed-length base62 tokens of 64 characters for strings under 40 chars, or 128 characters for strings under 80 chars, accessible via `EncryptAsync`, `DecryptAsync`, and `encryptForSearchAsync`), and Generic String (for strings of unknown length/format,

with optional `Input Encoding` of Base32 or Base64 and an `Input Padding Character`). Typed encryption methods are exposed per language — `encryptInt`, `encryptLong`, `encryptDate` in Java and `CipherInt64`, `CipherDate`, `CipherDateTime` in Go — available from C/C++ 2.3.0.0, .NET 2.4.0, Go 2.3.0, Java 2.3.0, Node.js 2.4.0, PHP 2.1.0, and Python 2.4.0.

Encrypt a 64-bit integer field in place using the Go library — the result stays an integer, so no schema changes are needed.

```go
encrypted, err := client.CipherInt64(plaintextInt64)
```

Encrypt a Java long (64-bit) value with the typed dataset method so the ciphertext remains a valid long in your existing database column.

```java
long encrypted = ubiqEncrypt.encryptLong(plaintextLong);
```

Encrypt a 32-bit integer stored in a database column so its value is protected while remaining searchable across encrypted records.

```javascript
const credentials =
ubiq.UbiqFactory.readCredentialsFromFile('./credentials',
'default');
const configuration = ubiq.UbiqFactory.defaultConfiguration();
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

const int32 = await
ubiqEncryptDecrypt.EncryptNumberAsync('int32', 9, []);
const search32 = await
ubiqEncryptDecrypt.encryptNumberForSearchAsync('int32', 9, []);
```

Encrypt a DateTime field (e.g. a transaction timestamp) and generate all search-compatible ciphertexts so encrypted records remain queryable.

```javascript
const credentials =
ubiq.UbiqFactory.readCredentialsFromFile('./credentials',
'default');
const configuration = ubiq.UbiqFactory.defaultConfiguration();
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

const plainDateTime = new Date(new Date().setUTCMilliseconds(0));
const ct = await
ubiqEncryptDecrypt.EncryptDateTimeAsync('datetime',
plainDateTime, []);
const search = await
ubiqEncryptDecrypt.encryptDateTimeForSearchAsync('datetime',
plainDateTime, []);
```

— Exact ranges, method names, versions and code samples across languages. product docs

03 **Grafana Loki SIEM destination** NEW 65

Ubiq adds Grafana Loki as a SIEM destination so activity events can stream into an existing Loki stack alongside other integrations such as Splunk. Configure it from the Ubiq console under Integrations > SIEM by adding a new Grafana Loki connection with the Loki push endpoint and credentials; Ubiq support must enable the destination for the enterprise account.

Enable Grafana Loki as a SIEM destination to stream Ubiq activity events into your existing Loki stack.

> 📍 In the Ubiq console, go to Settings › Integrations › SIEM and add a new Grafana Loki connection. Contact Ubiq support to enable the Loki destination for your enterprise account.

Enable Grafana Loki as a SIEM destination for Ubiq activity events alongside an existing Splunk integration.

> 📍 In the Ubiq console, go to Integrations › SIEM, add a new Grafana Loki connection, enter your Loki push endpoint and credentials, then save. Contact Ubiq support to enable the feature for your enterprise account.

— Exact console path and setup steps given, gated by support enablement.

changelog-20260926-2f75a2ec changelog-20260926-d6c55ad6

04 **IDP credential authentication in JS library** NEW 60

Adds `createCredentialsWithIdp` factory method (ubiq.UbiqFactory.createCredentialsWithIdp(<username>, <password>, <host>)) to authenticate using an identity provider.

— Exact method signature given but no full usage example. product docs

THINNER COVERAGE BELOW

05 **buildStructuredAsync builder method in JS library** NEW

50

Adds `buildStructuredAsync` builder method on `ubiq.CryptographyBuilder` to construct structured encryption/decryption objects, alongside the existing `buildEncryptionAsync` and `buildDecryptionAsync` methods.

— Named method and class but no standalone example provided. product docs

06 **Access Graph identity and dataset tracing** IMPROVED

45

Access Graph now includes IDP groups as nodes and filters related nodes in both directions, enabling access tracing from either an identity or a dataset.

— Describes new behavior but no navigation path or example given.

changelog-20260926-2f75a2ec

07 **Security History attribution and Access Group field** BREAKING

40

Security History events are now attributed to the IDP identity that performed them, and event details show the Access Group field in place of the retired Dataset Group field.

— Names the retired field but gives no migration guidance. changelog-20260926-2f75a2ec

08 **Go library validates input character set before encryption**

IMPROVED

35

The Go library now validates input against the dataset character set before encryption, returning a clear error instead of failing silently.

— States behavior change but no code sample or error format shown.

changelog-20260926-2f75a2ec

09 **Dataset creation wizard for structured configuration (C library)**

NEW

30

Adds dataset creation wizard support for advanced structured dataset configuration in the C 2.3.0.0 library.

— Bare mention of the wizard with no mechanism or steps. changelog-20260926-d6c55ad6

10 **Node.js worker thread object lifecycle guidance** IMPROVED

30

Documents multithreaded environment behavior: Ubiq objects must be created and destroyed within each worker thread because Node.js workers serialize/copy objects rather than sharing them.

— Clarifies a constraint but offers no code example. product docs

└──▷ ALSO FROM THESE RELEASES

Inspect the key number behind 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("Key number: %d\n", keyNum)
```

Add the Ubiq Rust SDK to a new project to start encrypting structured data with key zeroization on drop.

toml

```
# Cargo.toml
[dependencies]
ubiq-rust = "1.0.0"
```

Add the Ubiq Rust SDK to a project to encrypt structured data without decrypting to retrieve the key number.

toml

```
# Cargo.toml
[dependencies]
ubiq-rust = "1.0.0"
```

Inspect the key number behind a Go ciphertext for key-rotation auditing without performing a full decryption.

```go
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Key number: %d\n", keyNum)
```

Tag usage records with structured metadata (e.g. tenant ID and environment) so analysts can filter billing or audit data by application context.

```javascript
const ubiqEncryptDecrypt = await (new ubiq.CryptographyBuilder())
  .withCredentialsObject(credentials)
  .withConfigurationObject(configuration)
  .buildStructuredAsync();

ubiqEncryptDecrypt.addReportingUserDefinedMetadata('{"tenant":"acme","
env":"prod"}');
```

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Ubiq Security: Adds Ubiq `encrypt` and `decrypt` operators for Apache Airflow DAGs to protect data in pipeline tasks.

└──▷ **USE IT**

Inspect the key number behind 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("Key number: %d\n", keyNum)
```

Add the Ubiq Rust SDK to a new project to start encrypting structured data with key zeroization on drop.

Cargo.toml

```
# Cargo.toml
[dependencies]
ubiq-rust = "1.0.0"
```

Enable Grafana Loki as a SIEM destination to stream Ubiq activity events into your existing Loki stack.

> 📍 In the Ubiq console, go to Settings › Integrations › SIEM and add a new Grafana Loki connection. Contact Ubiq support to enable the Loki destination for your enterprise account.

Adds dataset creation wizard support for advanced structured dataset configuration in the C 2.3.0.0 library.

changelog · Sep 26, 2026 CHANGELOG ↗ STABLE

Ubiq Security: Access Graph now includes IDP groups and filters related nodes in both directions, enabling access tracing from either an identity or a dataset.

└──▷ **USE IT**

Add the Ubiq Rust SDK to a project to encrypt structured data without decrypting to retrieve the key number.

Cargo.toml

```
# Cargo.toml
[dependencies]
ubiq-rust = "1.0.0"
```

Inspect the key number behind a Go ciphertext for key-rotation auditing without performing a full decryption.

go

```
keyNum, err := ubiq.GetKeyNumber(ciphertext)
if err != nil {
    log.Fatal(err)
}
fmt.Printf("Key number: %d\n", keyNum)
```

Enable Grafana Loki as a SIEM destination for Ubiq activity events alongside an existing Splunk integration.

📍 In the Ubiq console, go to Integrations › SIEM, add a new Grafana Loki connection, enter your Loki push endpoint and credentials, then save. Contact Ubiq support to enable the feature for your enterprise account.

Go library now validates input against the dataset character set before encryption, returning a clear error instead of failing silently.

Security History events are now attributed to the IDP identity that performed them, and event details show the Access Group field in place of the retired Dataset Group field.

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ changelog · Sep 26, 2026

CHANGELOG ↗ RANK

DigiCert CertCentral is a centralized platform for managing, deploying, and monitoring SSL/TLS certificates across organizations.

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

01 Persistent DNS TXT automatic domain revalidation NEW

80

CertCentral now revalidates domains using the Persistent DNS TXT Record DCV method automatically every five days, eliminating the need to manually resubmit domains before the nine-day revalidation window expires. To use it, go to Domains, select a domain, choose the Persistent DNS TXT Record DCV method, place the persistent URI TXT record in your DNS zone, and submit for validation — once validated, leaving the TXT record in place keeps the domain continuously validated.

Set up Persistent DNS TXT automatic domain revalidation in CertCentral so domains stay continuously validated without manual intervention.

> 📍 In CertCentral, go to Domains, select the domain you want to prevalidate, choose the Persistent DNS TXT Record DCV method, place the persistent URI TXT record in your DNS zone, and submit for validation. Once validated, leave the TXT record in place — CertCentral will revalidate the domain automatically every five days.

— Names exact mechanism, interval and UI steps changelog-20260926-5b329386

└──▷ **ALSO FROM THESE RELEASES**

. . . .

└──▷ **BREAKING ON UPGRADE**

! All new TLS certificate orders, renewals, reissues, and duplicates will be issued from `DigiCert TLS RSA4096 Root G5` or `DigiCert TLS ECC P384 Root G5` by default as of October 15, 2026 — customers managing trust stores, hard-coding CA certificate trust, or pinning root/intermediate CA certificates must update their configurations or face service disruptions.

! MPIC enforcement as of December 1, 2026 requires that DNS queries and website traffic be reachable from six remote network perspectives across at least two RIR regions — environments that restrict traffic by IP address, region, query volume, or rate may block validation and must be reconfigured.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

DigiCert CertCentral: Adds Persistent DNS TXT automatic domain revalidation: CertCentral revalidates domains via the Persistent DNS TXT DCV method every five days, eliminating the need to manually resubmit domains before the nine-day revalidation window expires.

└──▷ **HOW TO FIND IT**

Set up Persistent DNS TXT automatic domain revalidation in CertCentral so domains stay continuously validated without manual intervention.

📍 In CertCentral, go to Domains, select the domain you want to prevalidate, choose the `Persistent DNS TXT Record` DCV method, place the persistent URI TXT record in your DNS zone, and submit for validation. Once validated, leave the TXT record in place — CertCentral will revalidate the domain automatically every five days.

└──▷ **BREAKING ON UPGRADE**

! All new TLS certificate orders, renewals, reissues, and duplicates will be issued from `DigiCert TLS RSA4096 Root G5` or `DigiCert TLS ECC P384 Root G5` by default as of October 15, 2026 — customers managing trust stores, hard-coding CA certificate trust, or pinning root/intermediate CA certificates must update their configurations or face service disruptions.

! MPIC enforcement as of December 1, 2026 requires that DNS queries and website traffic be reachable from six remote network perspectives across at least two RIR regions — environments that restrict traffic by IP address, region, query volume, or rate may block validation and must be reconfigured.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Mondoo

Feedback ⓘ Mondoo 14.0 is out! NOTES ↗ RANK

Your security team is stretched thin. Critical CVEs go unpatched.

> Mondoo 14.0 ships a major MQL language upgrade, 500+ new resources, unified IaC scanning with OpenTofu support, new scan output formats, and a big cut in cloud API call volume.

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

01 **Unified IaC scanning with OpenTofu support** NEW 50

Adds a unified IaC provider that discovers and scans every kind of infrastructure-as-code in a repository, plus first-class OpenTofu scanning.

– Names the provider and OpenTofu but no command shown Mondoo 14.0 is out!

02 **500+ new MQL resources for endpoint and Windows surfaces** NEW 45

Adds 500+ new MQL resources covering Windows, IIS, MDM, and EDR surfaces, including deep Windows and IIS coverage and dedicated MDM and EDR resources for endpoint management.

– Names resource categories but not specific resource names Mondoo 14.0 is out!

03 **New scan output formats: OCSF, InSpec, CSV** NEW 45

Adds `OCSF`, `InSpec`, and `CSV` scan output formats.

– Names three formats but no flag to select them Mondoo 14.0 is out!

04 **Expanded OS detection precision** IMPROVED 45

Adds precise OS detection for 15 more operating systems, with deep Solaris, FreeBSD, and SUSE 16 support.

– Gives concrete numbers and named OSes but no usage steps Mondoo 14.0 is out!

05 **MQL strict mode and cross-version support** NEW 42

Introduces MQL strict mode for tighter query validation, and adds cross-version MQL support so queries work across different resource schema versions.

– Names two capabilities but no flag or command to invoke them Mondoo 14.0 is out!

06 **Reduced cloud scan API call volume** IMPROVED 40

Reduces cloud scan API call volume by up to 97%.

— Concrete percentage but no mechanism explained Mondoo 14.0 is out!

07 **Container scanning identifies hardened base images** NEW 30

Adds container scanning that identifies and names hardened base images.

— Describes capability but no mechanism or example given Mondoo 14.0 is out!

08 **Automatic client updates on Windows** NEW 25

Adds automatic client updates on Windows.

— Bare statement with no mechanism or config detail Mondoo 14.0 is out!

09 **Six new out-of-the-box policies** NEW 25

Adds six new out-of-the-box policies.

— No policy names or targets given Mondoo 14.0 is out!

└──▷ **ALSO FROM THESE RELEASES**

Mondoo 14.0 is out! NOTES ↗ LATEST STABLE

Mondoo 14.0 adds MQL strict mode, 500+ new resources, unified IaC scanning, OpenTofu support, OCSF/InSpec/CSV output, and up to 97% fewer cloud API calls.

Adds OCSF, InSpec, and CSV scan output formats.

Introduces MQL strict mode for tighter query validation.

Adds cross-version MQL support so queries work across different resource schema versions.

Adds 500+ new MQL resources covering Windows, IIS, MDM, and EDR surfaces.

Adds deep Windows and IIS coverage via new MQL resources.

+8 more

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ changelog · Sep 26, 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 added support for several new LLM providers, tightened MCP launcher and Shadow AI access controls, and expanded policy and Slack management UIs, while also adding a minimal-environment mode for local MCP servers and new Terraform ECS deployment inputs.

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

01 **Terraform ECS `backend_desired_count` input** NEW 85

Adds `backend_desired_count` input to override the backend service's initial `desired_count` without replacing the full `services_configurations` map; seeds only services never previously created, with Application Auto Scaling owning the live count afterwards.

— Names exact Terraform input and its scoped mechanism product docs

02 **OTLP telemetry ingestion with MDM-minted API keys** NEW 85

Adds write-only OTLP telemetry ingestion for AI clients including Claude Code and Glean; adds automatic API key minting via Settings → MDM configuration for AI Watch scan, hook, and OTLP submissions — keys with this scope cannot be created manually.

Enable OTLP telemetry collection from an AI client (e.g. Claude Code) by provisioning an MDM-managed API key — the only way to obtain the write-only OTLP scope.

📍 In the Runlayer console, go to Settings → MDM configuration and apply a configuration targeting your AI client fleet. Runlayer will automatically mint API keys with the OTLP Telemetry scope for each enrolled device.

Confirm which org or enrollment API key a managed device is actually using before troubleshooting an enrollment failure.

```
$ aiwatch config show
```

— Names exact UI path and scope, missing endpoint detail product docs

03 **New model support across providers** NEW 75

Adds Claude Opus 5.5 support for agents through the Anthropic API or AWS Bedrock, with adaptive thinking effort settings; adds GPT-6 Sol and GPT-6 Luna model support on the OpenAI provider; adds Gemini support through OpenAI-Compatible providers pointed at `generativelanguage.googleapis.com`, sending API keys in the format Google expects rather than as a Bearer token.

— Names models and endpoint but no exact config steps `changelog-20260926-0af8fe13`

04 **AI Watch Enforce blocks unauthorized MCP launchers** BREAKING

75

AI Watch Enforce now blocks MCP launchers whose `--host` does not point at your Runlayer deployment or that name an unknown or disabled connector; previously allowlisted launchers from the Shadow AI activity view that fall into these categories are also blocked and can no longer be allowlisted from that view.

— Names exact flag but gives no migration path `changelog-20260926-0af8fe13`

05 **Minimal environment isolation for local MCP servers** IMPROVED

75

`runlayer run` now starts local MCP servers with only the connector's configured variables plus proxy, TLS, and locale settings, instead of inheriting the full AI client environment.

— Names the exact command and scope of the isolation `1.28.811 (2026-09-25)`

06 **Policies UI: filtered views and per-entry management** IMPROVED

70

Admins get shareable, bookmark-able filtered Policies views, and connector notices now deep-link to Policies filtered to the relevant connector; Enforce/Protect policy lists for allowed and denied MCP URLs, local packages, and built-in servers now support per-entry management with search and pagination instead of a single bulk save, and Shadow AI overview can add a discovered MCP straight to the allowlist.

Lock down an access policy so it only applies when a client connected via device-code registration, not pre-registered or manually issued credentials.

json

```
{
  "condition": {
    "field": "meta.client.kind",
    "operator": "equals",
    "value": "dcr"
  }
}
```

— Describes UI paths and a policy condition example `1.28.811 (2026-09-25)` `changelog-20260926-0af8fe13`

THINNER COVERAGE BELOW

07 **Shadow AI access controls** NEW 55

Saved Shadow AI skill and plugin files are now limited to resource managers and security or analytics reviewers; organization-wide Shadow AI REST APIs and MCP tools now require organization-wide audit log access.

— Names roles and scope but no endpoint specifics `changelog-20260926-0af8fe13`

└──▷ ALSO FROM THESE RELEASES

Set per-role Aurora CPU alarms and an ACU ceiling alarm so a saturated writer or a single hot reader pages before the cluster is fully pinned.

```hcl
rds_writer_cpu_utilization_threshold = 75
rds_reader_cpu_utilization_threshold = 85
rds_writer_acu_utilization_threshold = 90
```

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")
```

Temporarily pause ECS task-replacement loops on a ToolGuard service where worker-recycle drain flips `/health` to 503 and every task owns a whole GPU.

```hcl
runlayer_tool_guard_health_check_grace_period_seconds = 2147483647
```

└──▷ BREAKING ON UPGRADE

! AI Watch Enforce now blocks MCP launchers whose `--host` does not point at your Runlayer deployment or that name an unknown or disabled connector; previously allowlisted launchers from the Shadow AI activity view that fall into these categories are also blocked and can no longer be allowlisted from that view.

! Slack channel monitors no longer post a reply for every run — silent runs are suppressed. The reply-behavior setting that previously controlled this is removed.

! AWS provider must be `>= 6.64.0, < 7.0.0`; upgrade any older AWS provider selection in your dependency lock file before running `terraform plan` after upgrading this module.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Runlayer: Adds Claude Opus 5.5 support for agents through the Anthropic API or AWS Bedrock, with adaptive thinking effort settings.

└──▷ **TRY IT**

Confirm which org or enrollment API key a managed device is actually using before troubleshooting an enrollment failure.

```
$ aiwatch config show
```

Create an agent via Runlayer MCP and skip Slack setup in one call, useful for headless or CI-provisioned agents that should not have a Slack presence.

```json
{
  "name": "data-pipeline-agent",
  "enable_channels": []
}
```

Lock down an access policy so it only applies when a client connected via device-code registration, not pre-registered or manually issued credentials.

```json
{
  "condition": {
    "field": "meta.client.kind",
    "operator": "equals",
    "value": "dcr"
  }
}
```

Adds GPT-6 Sol and GPT-6 Luna model support for agents on the OpenAI provider.

Adds channel-only mode for Slack agents so admins can turn off direct messages; users who DM the agent are redirected to mention it in a channel.

Adds queued agent messages so follow-ups sent while an agent is working run in order, with the ability to edit, reorder, or remove queued messages.

Adds shareable, bookmark-able filtered Policies views for admins, and connector notices now deep-link to Policies filtered to the relevant connector.

Adds Gemini support through OpenAI-Compatible providers pointed at `generativelanguage.googleapis.com`, sending API keys in the format Google expects rather than as a Bearer token.

+2 more

└──▷ **BREAKING ON UPGRADE**

! AI Watch Enforce now blocks MCP launchers whose `--host` does not point at your Runlayer deployment or that name an unknown or disabled connector; previously allowlisted launchers from the Shadow AI activity view that fall into these categories are also blocked and can no longer be allowlisted from that view.

! Slack channel monitors no longer post a reply for every run — silent runs are suppressed. The reply-behavior setting that previously controlled this is removed.

1.28.811 (2026-09-25) BLOG ↗ STABLE

Runlayer 1.28.811 adds per-entry MCP policy management with search/pagination and a minimal environment for local MCP servers.

Adds minimal environment isolation for local MCP servers: `runlayer run` now starts local MCP servers with only the connector's configured variables plus proxy, TLS, and locale settings, instead of inheriting the full AI client environment.

Adds per-entry management to Enforce/Protect policy lists for allowed and denied MCP URLs, local packages, and built-in servers, with search and pagination instead of a single bulk save — and lets Shadow AI overview add a discovered MCP straight to the allowlist.

Splits an agent's Slack settings into four sections (Replies, Channel access, Monitoring, and Sending) navigable from the left panel, with Channel monitoring showing channels and condition only while active.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NVIDIA OpenShell

Feedback ⓘ v0.1.0 NOTES ↗ RANK

OpenShell runs autonomous AI agents in policy-controlled sandboxes with isolated filesystem, network, and credential access.

> OpenShell's v0.1.0 release adds new sandbox exec and create flags, server-side sandbox templates, structured list output, and broadens authentication, egress, and middleware extension surfaces across the platform.

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

01 **`--no-login-shell` flag for sandbox exec** NEW 85

`openshell sandbox exec` gains a `--no-login-shell` flag that skips shell startup files (such as `.bashrc` or `.profile`) when entering a sandbox, avoiding startup-file side effects during debugging sessions.

Drop into a sandbox for debugging without sourcing .bashrc or .profile, so startup-file side-effects don't interfere with your session.

```
$ openshell sandbox exec my-sandbox --no-login-shell
```

— Named flag with runnable example command. v0.1.0

02 **`--from` flag for VM sandbox rootfs archives** NEW 85

VM driver sandbox creation accepts a `--from` flag pointing to a rootfs tar archive as the sandbox source, letting artifact pipelines that produce filesystem tarballs provision sandboxes directly instead of using container images.

Provision a VM-driver sandbox from a pre-built rootfs tar archive instead of a container image — useful when your artifact pipeline produces filesystem tarballs.

```
$ openshell sandbox create --name rootfs-sandbox --from
./rootfs.tar
```

— Named flag with runnable example command. v0.1.0

THINNER COVERAGE BELOW

03 **`defaults-without-telemetry` build feature alias** NEW 55

Adds a `defaults-without-telemetry` feature alias at build time, enabling telemetry-free builds without manually composing feature flags.

— Named build flag with clear purpose. v0.1.0

04 **Sandbox templates on the server** NEW 45

The server now supports sandbox templates, allowing reusable sandbox configuration to be defined once and referenced at creation time.

— Explains behaviour but no API or config name given. v0.1.0

05 **Middleware extension surface expansion** NEW 45

Middleware gains broadened HTTP header mutation authority, expanding the extension surface for interceptors, and a new HTTP response pre-return interface that provides additional extension hooks.

— Two thin middleware additions grouped, no code names given. v0.1.0

06 **Pinned MCP revision allowlists in policy** NEW 45

Policy now preserves exact MCP revision allowlists, enabling pinned MCP version enforcement.

— Names mechanism, no policy syntax shown. v0.1.0

07 **Configurable OCSF schema version** NEW 40

The OCSF schema version used for SIEM output is now configurable, supporting backward compatibility with SIEM integrations expecting older schema versions.

— Names the mechanism but no config key given. v0.1.0

08 **Corporate HTTP proxy egress for microVM sandboxes** NEW 40

microVM sandboxes now support corporate HTTP forward proxy egress.

— Names the surface, no config detail. v0.1.0

09 **EC and EdDSA JWKS validation** IMPROVED 40

The server's OIDC JWKS validation now supports EC and EdDSA keys.

— Names key types, no further mechanism. v0.1.0

10 **Helm chart split into gateway and workspace charts** IMPROVED 40

The Helm chart has been split into separate gateway and workspace charts.

— States the change but no migration guidance. v0.1.0

11 **Driver-owned sandbox authentication** IMPROVED 40

Provider profiles are now authoritative and drive sandbox authentication rather than being advisory, and sandbox authentication is delegated to compute drivers, enabling driver-owned auth flows.

— Architectural change described without config or API names. v0.1.0

12 **Structured output for list commands** IMPROVED 35

`openshell` list commands now support structured, machine-readable output.

— No exact flag or format named. v0.1.0

13 **Driver-advertised compute resource capabilities** NEW 30

Compute drivers now advertise resource capabilities, making sandbox resource metadata visible to the control plane.

— Vague on what metadata or how it's used. v0.1.0

14 **OpenShift support for e2e:kubernetes target** IMPROVED 30

The `e2e:kubernetes` test target now works transparently on OpenShift.

— Names the target but no further detail. v0.1.0

15 **Separated public and contributor skill workflows** IMPROVED 30

Public and contributor agent skill workflows are now separated within the `skills/` directory.

— Names directory but no workflow specifics. v0.1.0

16 **Codex Security release qualification in CI** NEW 25

The CI pipeline now includes Codex Security release qualification.

— Bare mention with no mechanism given. v0.1.0

v0.1.0 NOTES ↗ LATEST STABLE

OpenShell v0.1.0 adds sandbox templates, VM HTTP proxy egress, structured list output, OCSF schema versioning, and EC/EdDSA JWKS support.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.1.0 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.1.0
```

└──▷ **TRY IT**

Drop into a sandbox for debugging without sourcing .bashrc or .profile, so startup-file side-effects don't interfere with your session.

```
$ openshell sandbox exec my-sandbox --no-login-shell
```

Provision a VM-driver sandbox from a pre-built rootfs tar archive instead of a container image — useful when your artifact pipeline produces filesystem tarballs.

```
$ openshell sandbox create --name rootfs-sandbox --from ./rootfs.tar
```

Adds `--no-login-shell` flag to `openshell sandbox exec` to skip shell startup files when entering a sandbox.

Adds `--from` flag accepting a rootfs tar archive as the source for VM driver sandboxes.

Adds structured output to `openshell` list commands for machine-readable results.

Adds configurable OCSF schema version for SIEM backward compatibility.

Adds sandbox templates on the server, allowing reusable sandbox configuration to be defined once and referenced at creation time.

+13 more

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗ RANK

The free and flexible app for your private thoughts.

> Obsidian 1.14.2 adds a command to open files from outside the vault and improves Kanban error feedback, while moving font-size storage out of the synced appearance config.

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

01 **Font size moved out of synced appearance config** BREAKING

65

Base font size is removed from `appearance.json` and now stored in device local storage; any vault that intentionally synced font size via Obsidian Sync will no longer sync this setting after upgrading.

— Names exact config file and sync impact changelog-20260926-e3fd1cda

THINNER COVERAGE BELOW

02 **Kanban card-move error feedback** IMPROVED

40

Kanban now shows error feedback when moving a card fails, for example due to invalid note properties or a missing target folder.

— Describes behaviour but no config or command changelog-20260926-e3fd1cda

└──▷ **BREAKING ON UPGRADE**

! Base font size is removed from `appearance.json` and stored in device local storage — any vault where font size was intentionally synced via Obsidian Sync will no longer sync this setting after upgrading.

changelog · Sep 26, 2026 CHANGELOG ↗ LATEST STABLE

Obsidian 1.14.2 adds open-from-outside-vault, OS Open with integration, and kanban folder-column grouping by `file.folder`

└──▷ **HOW TO FIND IT**

Open a Markdown incident-report file that lives outside your vault — useful when a file is dropped in a shared folder your vault does not index.

📍 In Obsidian, open the Command Palette, search for 'Open file from outside the vault...', select it, then navigate to the file on your filesystem and confirm.

Adds the `Open file from outside the vault...` command to open any file on your computer in the current window (Desktop only).

Kanban: Adds error feedback when moving a card fails (e.g. invalid note properties or missing target folder).

> └─▷ **BREAKING ON UPGRADE**
>
> ! Base font size is removed from `appearance.json` and stored in device local storage — any vault where font size was intentionally synced via Obsidian Sync will no longer sync this setting after upgrading.

WAS THIS USEFUL? 👍 👎

joplin

Feedback ⓘ v3.7.21 NOTES ↗ RANK

Joplin - the privacy-focused note taking app with sync capabilities for Windows, macOS, Linux, Android and iOS.

> Joplin added x-callback-url support for desktop automation and expanded its AI Chat feature with persistent conversation history.

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

01 x-callback-url commands for external apps NEW 90

Joplin's desktop app adds two x-callback-url commands: `joplin://x-callback-url/getCurrentNote?x-success=<url>&x-error=<url>` returns the currently selected note to a calling application, and `joplin://x-callback-url/createNote?title=<title>&body=<body>&x-success=<url>&x-error=<url>` creates a note in the current notebook and returns it to the caller.

Trigger Joplin from Hook (or another x-callback-url app) to fetch the currently open note and hand its title and URL back to the caller.

```
$ joplin://x-callback-url/getCurrentNote?x-success=hook%3A%2F%2Fx-callback-url%2FsetCurrentNode&x-error=hook%3A%2F%2Fx-callback-url%2Ferror
```

Programmatically create a new note in the current Joplin notebook and receive a callback with the note URL once it is saved.

```
$ joplin://x-callback-url/createNote?title=Incident%20Report&body=Summary%20here&x-success=editorial%3A%2F%2Fworkflow-callback&x-error=editorial%3A%2F%2Fworkflow-error
```

— Exact command syntax and runnable examples given product docs

THINNER COVERAGE BELOW

02 Persistent AI Chat history and chat switcher IMPROVED 30

AI Chat conversations now persist across sessions, and a chat switcher lets users manage multiple conversations. A new setting also allows hiding the AI Chat toolbar button.

— Describes what changed but no config keys or steps v3.7.21

v3.7.21 NOTES ↗ LATEST STABLE

Joplin v3.7.21 adds a setting to hide the AI Chat toolbar button and persistent AI conversation history with a chat switcher.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.7.21 https://github.com/laurent22/joplin.git
# already have the repo? check out this version:
$ git checkout v3.7.21
```

Persists AI Chat conversation history across sessions and adds a chat switcher to manage multiple conversations.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ changelog · Sep 26, 2026 CHANGELOG ↗

RANK

Connect, protect, and build everywhere.

> Cloudflare's release brings MCP server portals to general availability, adds OpenTelemetry span instrumentation to Workers, and introduces event-driven Browser Run subscriptions via Queues alongside smaller Gateway, firewall, email and observability updates.

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

01 **Per-domain scope for email suppressions** NEW 85

Adds a `scope` field to the Email Sending suppressions API (`POST /client/v4/accounts/$ACCOUNT_ID/email/sending/suppressions`) with values `account` and `sending_domain` , enabling per-domain suppression granularity; `scope_type=sending_domain&scope_value=<domain>` filters list results.

Scope an email suppression to a single sending domain so a bounce on `mail.myappexample.com` does not block the parent domain.

```
$ curl
https://api.cloudflare.com/client/v4/accounts/$ACCOUNT_ID/email/s
ending/suppressions \
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --header "Content-Type: application/json" \
  --data '{"email": "user@example.net", "scope": {"type":
"sending_domain", "value": "mail.myappexample.com"}}'
```

— Endpoint, field values and filter params all named. changelog-20260926-4137165b

02 **Browser Run event subscriptions via Queues** NEW 80

Adds `wrangler queues subscription create <QUEUE_NAME> --source browserRun --events crawl.started,crawl.updated,crawl.finished` to subscribe to Browser Run crawl lifecycle events via Cloudflare Queues, enabling event-driven downstream processing without polling.

Subscribe to Browser Run crawl events so a downstream Worker triggers on job completion without polling.

```
$ npx wrangler queues subscription create my-crawl-events --
source browserRun --events
crawl.started,crawl.updated,crawl.finished
```

— Exact runnable command with named flags and events. changelog-20260926-4137165b

03 **MCP server portals reach general availability** NEW

75

Makes MCP server portals generally available, adding Gateway routing for HTTP logging and DLP scanning, Code Mode policies, static OAuth client credentials, session management, service token authentication, and Logpush support for SIEM export.

— Lists many concrete sub-features but no setup steps given. changelog-20260926-4137165b

04 **OpenTelemetry span APIs in Workers tracing** NEW

70

Adds tracing.startSpan(name), tracing.getActiveSpan(), span.recordException(exception), and span.setAttributes(attributes) to the Workers custom spans API, expanding OpenTelemetry span instrumentation support. Also adds setAttribute() and setAttributes() returning the span object so calls can be chained.

— All method names given but no usage example. changelog-20260926-4137165b

05 **Deployment annotations on Workers Metrics charts** IMPROVED

60

Adds release and gradual deployment annotations to Workers Metrics charts, with hover details showing previous/new versions, rollout duration, and per-step traffic percentage.

— Describes chart hover details but is UI-only. changelog-20260926-4137165b

06 **Traffic-destination Gateway policies** NEW

60

Adds Gateway network policies that can match on traffic destination, such as `net.offramp.type == "internet"` to apply a policy only to traffic exiting to the public Internet while leaving tunnel-routed traffic unaffected.

Apply a Gateway Network policy only to traffic exiting to the public Internet, leaving tunnel-routed traffic unaffected.

📍 `net.offramp.type` == internet

— Example expression given but no field documentation. changelog-20260926-4137165b

THINNER COVERAGE BELOW

07 **Managed Ruleset support for Unified Routing** NEW 30

Adds Managed Ruleset support for accounts using Unified Routing mode in Cloudflare Advanced Network Firewall.

— Bare mention with no configuration detail. changelog-20260926-4137165b

└──▷ **ALSO FROM THESE RELEASES**

changelog · Sep 26, 2026 CHANGELOG ↗ STABLE

Cloudflare adds Browser Run queue subscriptions, Workers OpenTelemetry span APIs, Gateway traffic-destination policies, and MCP portals GA

└──▷ **TRY IT**

Subscribe to Browser Run crawl events so a downstream Worker triggers on job completion without polling.

```
$ npx wrangler queues subscription create my-crawl-events --source
browserRun --events crawl.started,crawl.updated,crawl.finished
```

Scope an email suppression to a single sending domain so a bounce on `mail.myappexample.com` does not block the parent domain.

```
$ curl
https://api.cloudflare.com/client/v4/accounts/$ACCOUNT_ID/email/send
ing/suppressions \
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --header "Content-Type: application/json" \
  --data '{"email": "user@example.net", "scope": {"type":
"sending_domain", "value": "mail.myappexample.com"}}'
```

Apply a Gateway Network policy only to traffic exiting to the public Internet, leaving tunnel-routed traffic unaffected.

📍 `net.offramp.type` == internet

Adds `wrangler queues subscription create <QUEUE_NAME> --source browserRun --events crawl.started,crawl.updated,crawl.finished` to subscribe to Browser Run crawl lifecycle events via Cloudflare Queues, enabling event-driven downstream processing without polling.

Adds `scope` field to the Email Sending suppressions API (`POST /client/v4/accounts/$ACCOUNT_ID/email/sending/suppressions`) with values

`account` and `sending_domain`, enabling per-domain suppression granularity; `scope_type=sending_domain&scope_value=<domain>` filters list results.

Adds tracing.startSpan(name), tracing.getActiveSpan(), span.recordException(exception), and span.setAttributes(attributes) to the Workers custom spans API, expanding OpenTelemetry span instrumentation support.

Adds setAttribute() and setAttributes() returning the span object so calls can be chained in Workers tracing.

Makes MCP server portals generally available, adding Gateway routing for HTTP logging and DLP scanning, Code Mode policies, static OAuth client credentials, session management, service token authentication, and Logpush support for SIEM export.

+2 more

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 26, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue