THE CYBER TOOLCHAIN

NO. 092

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 2, 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 202 changes

Do you prefer **scroll**?

▸ // CONTENTS SHOW 44 of 45 tools

▾ // THE BRIEF HIDE

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

Three unrelated things stand out. Metasploit Pro adds a native Kerberos relay chain with AD CS ESC8 and IPv6 DNS coercion, so red teams can path defenders rarely test. Watcher ships a subdomain-takeover

vim ? help

engine that catches dangling DNS before someone else claims it. Snyk's Snyk API & Web brings DAST to running web and API targets.

PROTECT

Find the AD and DNS footholds before an attacker does

Metasploit Pro's native Kerberos authentication relay with AD CS ESC8 and IPv6 DNS-takeover coercion turns a chain that previously meant stitching together separate tools into one supported module, and SharpHound v2.17.0's collection of Active Directory Sites ACLs lets BloodHound draw site-level paths that were invisible to the graph before. Watcher v3.6.0's takeover detection engine, merged with dnstwist and Certificate Transparency results into one DNS threats table, means dangling CNAMEs get caught by the same review that covers lookalike domains.

Metasploit Pro · sharphound · Watcher

PROTECT

Scan the AI runtimes and data warehouses nobody had coverage for

Plerion now runs vulnerability, secret, and SBOM scanning inside Amazon Bedrock AgentCore runtimes with ten new AWS checks, and Upwind's Snowflake Security applies CIS Snowflake Benchmark posture checks while tracing access from database roles through to the backing cloud storage — the hop where warehouse grants quietly become object-store exposure. Snyk API & Web adds DAST against live web and API targets, covering the runtime behaviour that SCA and SAST never see.

Plerion · Upwind · Snyk

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

SpecterOps SharpHound

Feedback ⓘ v2.17.0 NOTES ↗ RANK

C# Data Collector for BloodHound

SharpHound v2.17.0 adds a new `Site` collection method for enumerating Active Directory Sites ACLs and site data.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.17.0
https://github.com/SpecterOps/SharpHound.git
# already have the repo? check out this version:
$ git checkout v2.17.0
```

NEW DETECTION Active Directory Sites ACL collection

WAS THIS USEFUL? 👍 👎

Nmap

Feedback ⓘ changes since 2026-09-11 CODE ↗

Nmap - the Network Mapper. Github mirror of official SVN repository.

Nmap gains NSE execution time and memory limits, plus upgraded libpcap 1.10.7, pcre2 10.48, Lua 5.4.9, and libdnet 1.18.2.

Upgrades bundled libraries to libpcap 1.10.7, pcre2 10.48, Lua 5.4.9, and libdnet 1.18.2.

Ncat in listen mode now forwards termination signals (SIGINT, SIGTERM, SIGHUP, SIGQUIT) only to child processes started for connections when `-k` and `cmdexec` are both in effect, instead of the entire process group.

WAS THIS USEFUL? 👍 👎

OpenOSINT

Feedback ⓘ v2.30.0 NOTES ↗ RANK

AI-powered OSINT agent with interactive REPL, MCP server, and CLI. 20 tools.

> OpenOSINT v2.30.0 adds a zero-install way to launch its browser Web UI directly via uvx.

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

01 **One-command Web UI launch via uvx** NEW 70

The `uvx openosint web` command launches the browser Web UI as a zero-install, one-command option that requires no prior `pip install` or virtual environment setup.

Spin up the OpenOSINT Web UI instantly on any machine with uvx — no pip install or virtual environment required.

```
$ uvx openosint web
```

— Exact runnable command given, but mechanism beyond launch unexplained v2.30.0

v2.30.0 NOTES ↗ LATEST STABLE

OpenOSINT v2.30.0 adds one-command launch of the browser Web UI via `uvx openosint web`.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.30.0
https://github.com/OpenOSINT/OpenOSINT.git
# already have the repo? check out this version:
$ git checkout v2.30.0
```

└──▷ **TRY IT**

Spin up the OpenOSINT Web UI instantly on any machine with uvx — no pip install or virtual environment required.

```
$ uvx openosint web
```

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

Great Scott Gadgets Cynthion Feedback ⓘ

sold out · Oct 1, 2026 STORE ↗

A USB test instrument.

Cynthion with aluminum enclosure (cynthion-enclosure) sold out at `shop.hak5.org` on 1 Oct 2026.

WAS THIS USEFUL? 👍 👎

Great Scott Gadgets GreatFET One Feedback ⓘ

sold out · Oct 1, 2026 STORE ↗

GreatFET firmware and host software.

GreatFET One (greatfet-one) sold out at `shop.hak5.org` on 1 Oct 2026.

WAS THIS USEFUL? 👍 👎

Great Scott Gadgets HackRF One Feedback ⓘ

sold out · Oct 1, 2026 STORE ↗

A low cost software radio platform.

HackRF One with ANT500 antenna (hackrf-one-ant500) sold out at `shop.hak5.org` on 1 Oct 2026.

WAS THIS USEFUL? 👍 👎

O.MG / Hak5 O.MG Cable

Feedback ⓘ sold out · Oct 1, 2026

STORE ↗

A hand-made USB cable with a WiFi-controlled implant hidden inside that injects keystrokes, so red teams can emulate sophisticated physical attacks.

Basic, USB-A to USB-C (Black TPE) (omg-cable-basic-usba-usbc-black) sold out at `shop.hak5.org` on 1 Oct 2026.

WAS THIS USEFUL? 👍 👎

Hak5 Shark Jack Display

Feedback ⓘ

sold out · Oct 1, 2026 STORE ↗

Jack into a network and instantly run advanced recon, exfiltration, attack and automation payloads.

Cyb3rpunk Clear, single (sjd-cyb3rpunk-single) sold out at `shop.hak5.org` on 1 Oct 2026.

Cyb3rpunk Clear, Pro Bundle (sjd-cyb3rpunk-pro) sold out at `shop.hak5.org` on 1 Oct 2026.

WAS THIS USEFUL? 👍 👎

Hak5 WiFi Pineapple

Feedback ⓘ sold out · Oct 1, 2026

STORE ↗

The leading rogue access point and WiFi pentest toolkit (Mark VII and Enterprise firmware).

Mark VII+AC (wifi-pineapple-mk7-ac) sold out at `shop.hak5.org` on 1 Oct 2026.

WAS THIS USEFUL? 👍 👎

Great Scott Gadgets YARD Stick One

Feedback ⓘ

sold out · Oct 1, 2026 STORE ↗

A USB sub-1 GHz digital wireless transceiver that runs rfcat firmware, used to test garage doors, remote keyless entry, sensors and other low-speed radio systems.

YARD Stick One with ANT500 antenna (yard-stick-one-ant500) sold out at `shop.hak5.org` on 1 Oct 2026.

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

Rapid7 Metasploit Pro

Feedback ⓘ

October 1, 2026: 5.1.0-2026093001 NOTES ↗ RANK

Metasploit Pro is a commercial penetration testing platform for identifying, exploiting, and validating security vulnerabilities in systems.

> Metasploit Pro 5.1.0 ships 41 new modules headlined by a native Kerberos authentication relay attack chain with AD CS ESC8 exploitation and IPv6 DNS-takeover coercion, alongside Windows AArch64 and MIPS payload support, an OPC-UA scanner, and a breaking change requiring authentication on JSON RPC by default.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 **Kerberos authentication relay attack chain** NEW 90

Introduces native Kerberos authentication relay capabilities to the framework's relay stack, including the new `esc8_kerberos` auxiliary module for exploiting AD CS Web Enrollment (ESC8). Adds `auxiliary/spoof/ipv6/ipv6_ra_dns_takeover` and `auxiliary/spoof/dhcp/dhcpv6_dns_takeover` modules to enable native IPv6 DNS-takeover coercion as part of the relay attack chain, `KerberosTicketTrace` support in `auxiliary/admin/kerberos/forge_ticket`, and `Rex::Proto::Gss::KerberosToken` as a shared low-level implementation for handling Kerberos tokens carried by GSS-API and SPNEGO.

– Full module names and mechanism given, but no runnable example.

October 1, 2026: 5.1.0-2026093001

02 **SessionType option for WinRM login scanner** IMPROVED 86

Adds `SessionType` option to `auxiliary/scanner/winrm/winrm_login` to create a cmd-backed shell, a PSRP-backed PowerShell session, or automatically fall back to PowerShell when cmd shell creation is denied.

On a WinRM login attempt, automatically fall back to a PowerShell session when cmd shell creation is denied — useful against hardened Windows targets.

```
$ use auxiliary/scanner/winrm/winrm_login
set RHOSTS 192.168.1.50
set USERNAME administrator
set PASSWORD P@ssw0rd
set SessionType powershell
run
```

— Exact flag, behavior, and runnable example provided. October 1, 2026: 5.1.0-2026093001

03 **OPC-UA server scanner module** NEW 75

Adds `auxiliary/scanner/scada/opcua_enum`, a scanner module that detects OPC-UA servers speaking the OPC-UA TCP binary transport (`opc.tcp://`).

Scan an ICS network segment for OPC-UA servers exposing the binary TCP transport to enumerate exposed industrial control endpoints.

```
$ use auxiliary/scanner/scada/opcua_enum
set RHOSTS 10.10.20.0/24
run
```

— Module path, protocol, and example given. October 1, 2026: 5.1.0-2026093001

04 **ZoomEye search module updated to v2 API** IMPROVED 65

Updates the ZoomEye search module to support the ZoomEye v2 API, adds a three-attempt retry limit, and caps results per page at 10.

— Concrete limits named, but no module path or example. October 1, 2026: 5.1.0-2026093001

05 **JSON RPC requires authentication by default** BREAKING 60

Updates JSON RPC support to now require authentication by default, either via the database with user credentials or an auth token; existing unauthenticated JSON RPC clients will be rejected on upgrade.

— Clear behavior change but no config key or migration steps.

October 1, 2026: 5.1.0-2026093001

06 **SPIP date-column SQL injection module** NEW 60

Adds `auxiliary/scanner/http/spip_annee_sqli`, which exploits a blind SQL injection in SPIP's date column escaping logic.

— Named module path but no usage example. October 1, 2026: 5.1.0-2026093001

THINNER COVERAGE BELOW

07 **Windows AArch64 payload and psexec support** NEW 55

Adds Windows AArch64 staged shell payloads and extends `windows/smb/psexec` to work with AArch64 targets.

— Names target module but not specific payload names. October 1, 2026: 5.1.0-2026093001

08 **msfconsole startup and CLI additions** IMPROVED 50

Adds `MSFLOGO=nyan-cat.txt` environment variable support to `msfconsole` for selecting the nyan cat banner, a new command-line flag to print the count of modules, and improves msfconsole start-up performance.

— Env var named but new flag and performance gain unspecified.

October 1, 2026: 5.1.0-2026093001

09 **Inductive Automation Ignition fingerprint module** NEW 45

Adds an auxiliary scanner module that fingerprints Inductive Automation Ignition gateways across all major version families by probing unauthenticated info endpoints.

— Names product and method but not module path. October 1, 2026: 5.1.0-2026093001

10 **Linux x64 sandbox evasion module** NEW 40

Adds a Linux x64 sandbox-evasion module that performs runtime environment checks and aborts execution when a likely sandbox or VM is detected.

— Explains mechanism but omits module path. October 1, 2026: 5.1.0-2026093001

11 **MIPS exec payload additions** IMPROVED 40

Adds a MIPS64 exec payload, updates the MIPSEL and MIPSBE exec payloads, and adds prepends for all three.

— Names architectures but not payload paths or usage. October 1, 2026: 5.1.0-2026093001

12 **Writable check in Windows post mixin** NEW 40

Adds support to the post mixin for Windows sessions to run the `.writable?` method.

— Names method but no usage context given. October 1, 2026: 5.1.0-2026093001

13 **Reverse handler port fingerprint scanner** NEW 35

Adds a scanner module that enumerates ports on a host to determine whether each is a Metasploit reverse handler and what kind of shell it expects.

— Describes function but omits module path. October 1, 2026: 5.1.0-2026093001

14 **Persistence modules switched to create_process** IMPROVED

35

Updates persistence modules to use `create_process` instead of `cmd_exec`.

— Names mechanism change but not affected modules. October 1, 2026: 5.1.0-2026093001

15 **Persistence module for Kate** NEW 25

Adds a multi-platform persistence module for Kate.

— No module path, platforms, or mechanism named. October 1, 2026: 5.1.0-2026093001

└──▷ **BREAKING ON UPGRADE**

! JSON RPC now requires authentication by default (via database credentials or auth token); existing unauthenticated JSON RPC clients will be rejected on upgrade.

October 1, 2026: 5.1.0-2026093001 NOTES ↗ LATEST STABLE

Metasploit Pro 5.1.0-2026093001 adds 41 new modules including OPC-UA, Kerberos relay, IPv6 DNS-takeover, and Windows AArch64 payloads.

└──▷ **TRY IT**

Scan an ICS network segment for OPC-UA servers exposing the binary TCP transport to enumerate exposed industrial control endpoints.

```
$ use auxiliary/scanner/scada/opcua_enum
set RHOSTS 10.10.20.0/24
run
```

On a WinRM login attempt, automatically fall back to a PowerShell session when cmd shell creation is denied — useful against hardened Windows targets.

```
$ use auxiliary/scanner/winrm/winrm_login
set RHOSTS 192.168.1.50
set USERNAME administrator
set PASSWORD P@ssw0rd
set SessionType powershell
run
```

Adds `auxiliary/scanner/scada/opcua_enum`, a scanner module that detects OPC-UA servers speaking the OPC-UA TCP binary transport (`opc.tcp://`).

Adds `SessionType` option to `auxiliary/scanner/winrm/winrm_login` to create a cmd-backed shell, a PSRP-backed PowerShell session, or automatically fall back to PowerShell when cmd shell creation is denied.

Adds `MSFLOGO=nyan-cat.txt` environment variable support to `msfconsole` for

selecting the nyan cat banner.

Adds a new command-line flag to `msfconsole` to print the count of modules.

Updates JSON RPC support to now require authentication by default, either via the database with user credentials or an auth token.

+16 more

└──▷ **BREAKING ON UPGRADE**

! JSON RPC now requires authentication by default (via database credentials or auth token); existing unauthenticated JSON RPC clients will be rejected on upgrade.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Snyk

Feedback ⓘ changelog · Oct 2, 2026 DOCS ↗ BLOG ↗ RANK

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

> Snyk shipped a packaged ADS agent installer for macOS and Windows that removes custom wrapper scripts, alongside Snyk API & Web, a new DAST scanning capability for web and API targets.

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

01 **Packaged ADS installer for macOS and Windows** NEW 85

Snyk now offers signed macOS `.pkg` and Windows `.msi` installers for the ADS (Agent Supply Chain Security) agent, downloadable from the Settings page in Evo, removing the need to wrap a downloaded file in a custom script. The installer registers a scheduled job on the host machine so Agent Supply Chain Security scans run on their own schedule without requiring MDM.

Deploy the ADS agent on a macOS developer machine without MDM by using the new signed installer and your Tenant ID and push key.

> 📍 In Evo, open Settings, choose your products, and select Save & Publish. Choose macOS and your architecture, then select Download ADS Installer. Install the .pkg, then run the installer with your Tenant ID and push key.

— Names exact installer formats and gives a full install walkthrough. 2026-10-02-ca10b44a

02 **Snyk API & Web DAST scanning** NEW 60

Introduces Snyk API & Web, a new dynamic application security testing (DAST) scanning capability supporting web targets and API targets defined via Postman Collection, RAML, or Bruno Collection in OpenCollection format. Authentication methods include basic auth, login form, login sequence, two-factor authentication, OpenAPI, Postman, GraphQL, and mutual TLS.

— Lists target formats and auth methods but no setup steps given. 2026-10-02-ca10b44a

changelog · Oct 2, 2026 DOCS ↗ BLOG ↗ LATEST STABLE

Snyk adds signed macOS .pkg and Windows .msi ADS installers plus a new Snyk API & Web DAST scanning capability.

└──▷ HOW TO FIND IT

Deploy the ADS agent on a macOS developer machine without MDM by using the new signed installer and your Tenant ID and push key.

📍 In Evo, open Settings, choose your products, and select Save & Publish. Choose macOS and your architecture, then select `Download ADS Installer`. Install the .pkg, then run the installer with your Tenant ID and push key.

Adds signed macOS `.pkg` and Windows `.msi` packaged installers for the ADS (Agent Supply Chain Security) agent, downloadable from the Settings page in Evo, eliminating the need to wrap a downloaded file in a custom script.

The ADS installer registers a scheduled job on the host machine so Agent Supply Chain Security scans run on their own schedule without requiring MDM.

Introduces Snyk API & Web, a new dynamic application security testing (DAST) scanning capability supporting web targets, API targets (Postman Collection, RAML, Bruno Collection in OpenCollection format), and authentication methods including basic auth, login form, login sequence, two-factor authentication, OpenAPI, Postman, GraphQL, and mutual TLS.

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

Confidential comments on repository security advisories (2026-10-02) BLOG ↗

RANK

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

> GitHub's biggest moves this window were on the Copilot agent side — computer use for desktop automation and code-defined dynamic workflows for multi-agent orchestration — alongside a push to harden private vulnerability reporting (structured forms, CWE enforcement, rate limits) and confidential advisory comments, plus a unified retention policy for Actions checks, runs, and statuses.

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

01 Structured, enforceable forms for private vulnerability reports

98

BREAKING

Maintainers can add `.github/VULNERABILITY_REPORT.yml` (placed in an org or account's `.github` repo to apply it org-wide) to define a custom structured report form, with fields supporting `min_length` validation. A new REST API endpoint returns the form a repository enforces and surfaces it in error responses when API submissions don't match it. The default form now requires summary, details, a proof of concept (minimum 150 characters), and impact; reporters can check an 'I used AI assistance to find or write up this report' box; repos can require a CWE before submission (Settings > Advanced Security > Private vulnerability reporting, enforceable via org/enterprise policy); and repos with a `SECURITY.md` now show a banner linking to it before submission.

Enforce a structured vulnerability report form with minimum detail requirements across all repositories in your organization.

```yaml
# .github/VULNERABILITY_REPORT.yml in your org's .github repo
name: Vulnerability Report
description: Report a security vulnerability
body:
  - type: textarea
    id: summary
    attributes:
      label: Summary
    validations:
      required: true
      min_length: 50
  - type: textarea
    id: details
    attributes:
      label: Details
    validations:
      required: true
      min_length: 100
  - type: textarea
    id: poc
    attributes:
      label: Proof of Concept
    validations:
      required: true
      min_length: 150
  - type: textarea
    id: impact
    attributes:
      label: Impact
    validations:
      required: true
```

Require reporters to assign a CWE before submitting, and enforce the setting across your organization via policy.

> 📍 In the repository, go to Settings › Advanced Security › Private vulnerability reporting and enable the option to require a CWE. Organization and enterprise owners can then enforce this setting via a policy applied across all repositories.

— Named config file, endpoint, and validation fields with a runnable example.

Structured forms for private vulnerabil…

---

02 **Computer use for desktop app automation in Copilot** NEW

Copilot gains computer use (public preview) letting it interact with desktop applications — clicking controls, entering text, pressing keys, scrolling, dragging, and navigating workflows — including legacy or GUI-only software with no API, CLI, or MCP integration. It's controlled in Copilot CLI via `/computer on`, `/computer show`, and `/computer off`, and in the GitHub Copilot app on macOS and Windows via Settings > Computer Use > Enable Computer Use, with per-app approval controls and the ability to review or reset always-allowed apps. On macOS it walks users through granting Accessibility and Screen Recording permissions, and organizations can disable computer use across managed devices.

▶ Watch video

Enable computer use in Copilot CLI so Copilot can automate a GUI-only desktop workflow — for example, moving data through a legacy expense-report app with no API.

```
$ /computer on
```

Check whether computer use is active in the current Copilot CLI session, or turn it off after finishing an automated workflow.

```
$ /computer show
/computer off
```

Enable Computer Use from the GitHub Copilot app UI on macOS or Windows when you prefer not to use the CLI.

📍 Open the GitHub Copilot app › Settings › Computer Use › turn on `Enable Computer Use`. On macOS, follow the prompted steps to grant Accessibility and Screen Recording permissions.

— Exact CLI commands, UI path, and permission steps all given.

GitHub Copilot can now interact with de…

03 Dynamic workflows for multi-agent orchestration in Copilot NEW

91

Dynamic workflows enable code-defined, multi-agent orchestration across Copilot CLI (via `--experimental` flag or `/experimental on` in an interactive session), the GitHub Copilot app (no setup required), and the GitHub Copilot SDK (programmatic access to Copilot's extensibility APIs). They support pause-and-resume checkpoints for reviewing intermediate findings before continuing a long or expensive run, and parallel sub-agent execution for tasks like reviewing many pull-request files or sweeping a large codebase across directories simultaneously. Built-in Copilot authoring guidance lets practitioners

ask Copilot to write a workflow for them, or ask 'What dynamic workflows are available?' to discover existing ones.

Enable dynamic workflows in Copilot CLI to orchestrate a multi-agent incident investigation across logs and telemetry.

```
$ gh copilot --experimental
```

Turn on experimental features mid-session and discover existing dynamic workflows without restarting the CLI.

```
$ /experimental on
What dynamic workflows are available?
```

Have Copilot author a reusable release-check workflow that pauses for human review before resuming.

> 📍 Ask Copilot: 'Create a dynamic workflow that runs release checks, asks one agent to assess failures, and pauses at a checkpoint for me to review before resuming.'

— Exact CLI flags and discovery commands given across CLI, app, and SDK.

Dynamic workflows in Copilot CLI and th…

---

04 Unified retention for checks, runs, and statuses in Actions

82

BREAKING

The existing Actions retention setting — labeled `Check, workflow run, status, artifact and log retention` in the UI — now also cleans up checks, workflow runs, and statuses created by GitHub Actions and third-party applications, in addition to artifacts and logs. Public repository retention for these is capped at a maximum of 90 days, configurable under Settings > Actions > General. Data older than the configured period is automatically deleted and cannot be restored once removed.

Review or update your repository's unified retention period so checks, runs, and statuses are kept as long as your audit or compliance workflow requires.

> 📍 In the repository, go to Settings › Actions › General, then locate 'Check, workflow run, status, artifact and log retention' and set the desired number of days (max 90 for public repositories).

— Exact setting label, path, and 90-day cap given.

Actions retention now covers checks, ru…

05 **Rate limits for private vulnerability reports** NEW 77

Private vulnerability reporting now enforces daily per-user rate limits, showing a `try again later` message to reporters who exceed them (comments on existing advisories are unaffected). Repository administrators can set a custom daily overall reporting limit and add trusted reporters to an allow list so they're never rate limited, both configurable under Settings > Advanced Security > Private vulnerability reporting.

Add a trusted security researcher to the allow list so their reports are never blocked by rate limits.

> 📍 1. Go to your repository on GitHub. 2. Click Settings › Advanced Security. 3. Click Settings next to `Private vulnerability reporting`. 4. Under the allow list section, add the trusted reporter's GitHub account. 5. Save changes.

Set a custom daily reporting cap for your repository to throttle bulk submissions while still accepting legitimate reports.

> 📍 1. Go to your repository on GitHub. 2. Click Settings › Advanced Security. 3. Click Settings next to `Private vulnerability reporting`. 4. Enter your desired daily overall reporting limit. 5. Save changes.

— Exact settings path and behavior given, no numeric default limit stated.

`Rate limits for private vulnerability r…`

---

06 **Confidential comments on security advisories** NEW 74

Repository security advisories now support confidential comments, visible only to maintainers with write access and returned via the GraphQL API (not the REST API). Confidential comments are marked in the advisory timeline, comment views are logged in the audit log, and access follows write permissions — losing write access removes the ability to read existing confidential comments. Post one by selecting `Confidential` below the comment box before submitting; the setting cannot be changed afterward.

Post a confidential internal note on an advisory — e.g. to record investigation details without exposing them to the reporter or invited collaborators.

> 📍 In the security advisory, type your comment, select `Confidential` below the comment box, then click Submit. The comment will appear with a confidential marker visible only to maintainers with write access.

— Full mechanism, permissions model and steps given, no API name beyond GraphQL/REST.

`Confidential comments on repository sec…`

---

07 **Dev Container support for agent sessions** NEW 69

Adds a `Use Dev Container` option in the Agents window to start agent sessions inside a project's Dev Container, working on local, SSH, Tunnel, and WSL hosts.

Run an agent session inside your project's Dev Container so it has access to the exact toolchain and dependencies defined in your devcontainer config.

> 📍 In the Agents window, open the menu for your local or remote folder and select `Use Dev Container` to launch the agent session inside the configured container (works on SSH, Tunnel, and WSL hosts).

— Named option and supported hosts with a runnable example.

GitHub Copilot in VS Code, September 20…

---

08 **Scheduled automations for agent tasks** NEW 68

Copilot in VS Code adds automations to schedule recurring agent tasks on an hourly, daily, or weekly cadence, or run them on demand from a template or custom prompt (preview), configured via `New Automation` in the Agents window.

▶ Watch video

Run a nightly dependency-audit prompt against your repo without manual intervention.

> 📍 In the Agents window, select `New Automation`, choose a daily schedule, and enter your prompt (or pick a template) to have Copilot run the task automatically each day.

— Cadence options and exact UI steps given. GitHub Copilot in VS Code, September 20…

---

09 **Agent merge for pull request completion** NEW 68

Agent merge (preview) lets the agent handle review feedback, failed checks, merge conflicts, and workflow reruns within an active session, enabled via `Enable agent merge` in the Agents window.

Hand off post-PR-creation chores to the agent so you can focus on the next task instead of babysitting CI.

> 📍 In the Agents window, open an active agent session, then select `Enable agent merge` to let the agent automatically respond to review feedback, resolve merge conflicts, rerun failed checks, and merge the pull request when ready.

— Clear scope and exact toggle to enable it. GitHub Copilot in VS Code, September 20…

THINNER COVERAGE BELOW

---

10 **Pull request creation from agent sessions** NEW 53

Copilot, Claude, or Codex agent sessions in the Agents window can now create pull requests directly, with title/description editing and draft or merge options.

— UI location and options named but no step-by-step example.

GitHub Copilot in VS Code, September 20…

11 Session and chat navigation improvements in Copilot for VS Code

NEW

52

Adds `Mark as Done` suggestions with automatic cleanup to keep completed sessions organized (preview); application badge support (preview) to surface new results, input requests, and PR checks on the dock, launcher, or taskbar; a hierarchical sessions list view so agents can branch new chats or sessions while navigating back to the source chat; an `Add Context` entry-point with URL-paste support in the new-session input to attach a GitHub issue or pull request to any chat; the ability to continue a general Copilot quick chat (started without a workspace) by attaching a local folder mid-conversation; and the ability to pick up a Codex conversation from the ChatGPT app inside VS Code to continue work on a codebase without copying context or files.

▶ Watch video

▶ Watch video

— Six named UX additions but no step-by-step usage given.

GitHub Copilot in VS Code, September 20…

12 HydraFusion automatic model selection in Copilot NEW

30

Adds a `HydraFusion` model-picker option (research preview) that automatically selects models and workflows for coding tasks.

— Named feature but described in a single thin sentence.

GitHub Copilot in VS Code, September 20…

└──▷ BREAKING ON UPGRADE

! If you add a custom `VULNERABILITY_REPORT.yml` form, reports submitted through the REST API must match it — API submissions that do not match will receive an error.

! Checks, workflow runs, and statuses older than the configured retention period are now automatically deleted — including those created by third-party applications. Data already removed before this change cannot be restored.

> Confidential comments on repository security advisories (2026-10-02) BLOG ↗
>
> LATEST STABLE
>
> GitHub security advisories gain confidential comments — visible only to write-access maintainers, logged in the audit log, and available via GraphQL.
>
> **└──▷ HOW TO FIND IT**
>
> Post a confidential internal note on an advisory — e.g. to record investigation details without exposing them to the reporter or invited collaborators.

📍 In the security advisory, type your comment, select `Confidential` below the comment box, then click Submit. The comment will appear with a confidential marker visible only to maintainers with write access.

Adds confidential comments to repository security advisories, accessible in the GraphQL API but not returned by the REST API.

Confidential comments are clearly marked in the advisory timeline and views are recorded in the audit log.

Access to confidential comments follows repository write permissions — if someone loses write access, they can no longer read existing confidential comments.

Post a confidential comment by selecting `Confidential` below the comment box before posting; the setting cannot be changed after posting.

`Rate limits for private vulnerability reports (2026-10-01)` BLOG ↗ STABLE

GitHub adds daily rate limits, custom caps, and allow-lists to private vulnerability reporting to block bulk submissions.

└──▷ HOW TO FIND IT

Add a trusted security researcher to the allow list so their reports are never blocked by rate limits.

📍 1. Go to your repository on GitHub. 2. Click Settings › Advanced Security. 3. Click Settings next to `Private vulnerability reporting`. 4. Under the allow list section, add the trusted reporter's GitHub account. 5. Save changes.

Set a custom daily reporting cap for your repository to throttle bulk submissions while still accepting legitimate reports.

📍 1. Go to your repository on GitHub. 2. Click Settings › Advanced Security. 3. Click Settings next to `Private vulnerability reporting`. 4. Enter your desired daily overall reporting limit. 5. Save changes.

Applies daily per-user rate limits to new private vulnerability reports, with reporters who hit the limit shown a `try again later` message; limits do not affect comments on existing advisories.

Lets repository administrators set a custom daily overall reporting limit for their repository via the Advanced Security settings.

Lets repository administrators add trusted reporters to an allow list so they are never rate limited, configurable under repository Settings > Advanced Security > Private vulnerability reporting.

Structured forms for private vulnerability reports (2026-10-01) BLOG ↗

STABLE

GitHub adds structured, customizable forms for private vulnerability reports, with CWE enforcement, AI-disclosure, and a new API endpoint.

└──▷ **USE IT**

Enforce a structured vulnerability report form with minimum detail requirements across all repositories in your organization.

.github/VULNERABILITY_REPORT.yml

```yaml
# .github/VULNERABILITY_REPORT.yml in your org's .github repo
name: Vulnerability Report
description: Report a security vulnerability
body:
  - type: textarea
    id: summary
    attributes:
      label: Summary
    validations:
      required: true
      min_length: 50
  - type: textarea
    id: details
    attributes:
      label: Details
    validations:
      required: true
      min_length: 100
  - type: textarea
    id: poc
    attributes:
      label: Proof of Concept
    validations:
      required: true
      min_length: 150
  - type: textarea
    id: impact
    attributes:
      label: Impact
    validations:
      required: true
```

Require reporters to assign a CWE before submitting, and enforce the setting across your organization via policy.

📍 In the repository, go to Settings › Advanced Security › Private vulnerability reporting and enable the option to require a CWE. Organization and enterprise owners can then enforce this setting via a policy applied across all repositories.

Adds `.github/VULNERABILITY_REPORT.yml` to configure a custom structured form for private vulnerability reports; place the file in your org or account's `.github` repository to apply it across all repos you own.

Form fields support `min_length` (issue form syntax) so maintainers can enforce a minimum level of detail per field.

Adds a new REST API endpoint that returns the form a repository enforces, surfaced in error responses when an API submission does not match the custom form.

Default structured form now requires four fields from reporters: summary, details, proof of concept (minimum 150 characters), and impact.

Enables a 'I used AI assistance to find or write up this report' disclosure checkbox for reporters.

+2 more

└──▷ BREAKING ON UPGRADE

! If you add a custom `VULNERABILITY_REPORT.yml` form, reports submitted through the REST API must match it — API submissions that do not match will receive an error.

`GitHub Copilot can now interact with desktop apps with computer use (2026-10-01)`

BLOG ↗ STABLE

GitHub Copilot gains computer use in public preview, letting it click, type, and navigate desktop apps on macOS and Windows.

└──▷ TRY IT

Enable computer use in Copilot CLI so Copilot can automate a GUI-only desktop workflow — for example, moving data through a legacy expense-report app with no API.

```
$ /computer on
```

Check whether computer use is active in the current Copilot CLI session, or turn it off after finishing an automated workflow.

```
$ /computer show
/computer off
```

Enable Computer Use from the GitHub Copilot app UI on macOS or Windows when you prefer not to use the CLI.

> 📍 Open the GitHub Copilot app › Settings › Computer Use › turn on `Enable Computer Use`. On macOS, follow the prompted steps to grant Accessibility and Screen Recording permissions.

Adds `/computer on` in Copilot CLI to enable computer use, `/computer show` to check its status, and `/computer off` to disable it.

Adds Computer Use public preview in GitHub Copilot CLI and the GitHub Copilot app on macOS and Windows, letting Copilot interact with desktop applications — clicking controls, entering text, pressing keys, scrolling, dragging, and navigating workflows — including legacy and GUI-only software with no API, CLI, or MCP integration.

Adds a Computer Use settings panel in the GitHub Copilot app (Settings > Computer Use > Enable Computer Use) with per-app approval controls and the ability to review or reset always-allowed apps.

On macOS, guides users through required Accessibility and Screen Recording permissions as part of the computer use setup flow.

Supports organization-managed settings to disable computer use across managed devices.

`GitHub Copilot in VS Code, September 2026 releases (2026-10-01)` **BLOG ↗**

STABLE

Copilot in VS Code gains scheduled automations, agent merge, PR creation from sessions, and HydraFusion multi-model orchestration.

└──▷ HOW TO FIND IT

Hand off post-PR-creation chores to the agent so you can focus on the next task instead of babysitting CI.

> 📍 In the Agents window, open an active agent session, then select `Enable agent merge` to let the agent automatically respond to review feedback, resolve merge conflicts, rerun failed checks, and merge the pull request when ready.

Run a nightly dependency-audit prompt against your repo without manual intervention.

> 📍 In the Agents window, select `New Automation`, choose a daily schedule, and enter your prompt (or pick a template) to have Copilot run the task automatically each day.

Run an agent session inside your project's Dev Container so it has access to the exact toolchain and dependencies defined in your devcontainer config.

📍 In the Agents window, open the menu for your local or remote folder and select `Use Dev Container` to launch the agent session inside the configured container (works on SSH, Tunnel, and WSL hosts).

Adds `HydraFusion` model-picker option (research preview) to automatically select models and workflows for coding tasks.

Adds automations to schedule recurring agent tasks on an hourly, daily, or weekly cadence, or run them on demand from a template or custom prompt (preview).

Adds agent merge in active sessions (preview) to let the agent handle review feedback, failed checks, merge conflicts, and workflow reruns.

Adds pull request creation directly from Copilot, Claude, or Codex agent sessions in the Agents window, with title/description editing, draft and merge options.

Adds `Mark as Done` suggestions and automatic cleanup to keep completed sessions organized (preview).

`+6 more`

`Actions retention now covers checks, runs, and statuses (2026-10-01)` **BLOG ↗**

STABLE

GitHub Actions retention now governs checks, workflow runs, and statuses alongside artifacts and logs.

└──▷ HOW TO FIND IT

Review or update your repository's unified retention period so checks, runs, and statuses are kept as long as your audit or compliance workflow requires.

📍 In the repository, go to Settings › Actions › General, then locate 'Check, workflow run, status, artifact and log retention' and set the desired number of days (max 90 for public repositories).

Extends the existing Actions retention setting — labeled `Check, workflow run, status, artifact and log retention` in the UI — to also clean up checks, workflow runs, and statuses created by GitHub Actions and third-party applications.

Caps public repository retention for checks, workflow runs, and statuses at a maximum of 90 days.

└──▷ BREAKING ON UPGRADE

! Checks, workflow runs, and statuses older than the configured retention period are now automatically deleted — including those created by third-party applications. Data already removed before this change cannot be restored.

Dynamic workflows in Copilot CLI and the Copilot app (2026-10-01) BLOG ↗

STABLE

GitHub Copilot gains dynamic workflows — code-defined, multi-agent orchestration now in CLI, app, and SDK.

└──▷ TRY IT

Enable dynamic workflows in Copilot CLI to orchestrate a multi-agent incident investigation across logs and telemetry.

```
$ gh copilot --experimental
```

Turn on experimental features mid-session and discover existing dynamic workflows without restarting the CLI.

```
$ /experimental on
What dynamic workflows are available?
```

Have Copilot author a reusable release-check workflow that pauses for human review before resuming.

📍 Ask Copilot: 'Create a dynamic workflow that runs release checks, asks one agent to assess failures, and pauses at a checkpoint for me to review before resuming.'

Enables dynamic workflows in the Copilot CLI via the `--experimental` command-line option or `/experimental on` in an interactive session, allowing code-defined multi-agent orchestration.

Adds dynamic workflows to the GitHub Copilot app with no setup required, supporting parallel agent execution, checkpoints, and structured result passing.

Introduces dynamic workflows in the GitHub Copilot SDK, giving extension authors programmatic access to Copilot's extensibility APIs for multi-step, multi-agent processes.

Supports pause-and-resume checkpoints in dynamic workflows, letting practitioners review intermediate findings before continuing a long or expensive run.

Supports parallel sub-agent execution within a dynamic workflow, enabling tasks like reviewing many pull-request files or sweeping a large codebase across directories simultaneously.

+1 more

WAS THIS USEFUL? 👍 👎

Cloudsmith Feedback ⓘ

Trace downloads through connected repositories in client log exports (2026-10-01)

BLOG ↗ RANK

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

> Cloudsmith added a new traceability field to client log exports, letting teams follow package downloads across connected repositories.

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

01 **Connected repository tracing in client log exports** NEW

55

Client log exports now include a `connected_repositories` field listing every connected repository a package was served through, allowing downloads to be traced back to their origin repository.

— Names the exact field but no usage example or endpoint.

Trace downloads through connected repos…

> Trace downloads through connected repositories in client log exports (2026-10-01)
>
> BLOG ↗ LATEST STABLE
>
> Client log exports gain a `connected_repositories` field to trace package downloads through the serving repository chain.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Invicti

Feedback ⓘ Release 20261001 (2026-10-01) BLOG ↗ RANK

Official documentation for Invicti AppSec Platform, Invicti Platform, and Invicti Enterprise — web application security testing, DAST, API security, and vulnerability management.

> Invicti's latest releases centered on agentic pentesting: a new AI credit usage ledger with billing changes, HTML report downloads, agent log archives, and broadened agentic detection and agent behavior across Recon, Octo, and time-based confirmation.

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

01 AI credit usage ledger and billing rules

BREAKING 72

Adds an AI Credit Usage section under the Pentests nav item showing every credit addition, deduction, pending reservation, and current available balance in a transaction ledger, plus real-time credit consumption on the Summary tab with zero credit charge for failed assessments and pro-rated charges for canceled ones. Credit values are now expressed as whole numbers instead of decimals, and unused credits now expire at the end of the subscription term.

Review pending and committed AI credits before starting a new agentic pentest to ensure sufficient balance.

> 📍 In the Invicti Platform, go to Pentests › AI Credit Usage to view the full transaction ledger, including pending reservations for in-progress assessments and your current available balance.

Review credit spend and committed balance before starting a new Agentic Pentest assessment to avoid surprises.

> 📍 In the platform UI, go to Pentests › AI Credit Usage to see every credit addition, deduction, and pending reservation alongside your current available balance.

— Names exact nav path and billing mechanics, including the breaking rounding/expiry change

Release 20261001 (2026-10-01)

THINNER COVERAGE BELOW

02 Agent scan log download and retention

NEW 55

Enables downloading scanner logs and raw agent traces as a single compressed archive per assessment, for investigating a finding or sharing diagnostics with Invicti Support,

and adds a configurable retention period and storage location for internal agent scan logs.

Download raw scanner logs and agent traces after an assessment to investigate a finding or share diagnostics with Invicti Support.

> 📍 In the Invicti Platform, open the completed assessment, then select the option to download the scanner logs and agent traces as a single compressed archive.

Download raw agent traces after an assessment completes to investigate a finding in depth or share diagnostics with Invicti Support.

> 📍 In the platform UI, open the completed assessment, then select `Download scanner logs and agent traces` to save a single compressed archive to your machine.

— Exact download steps given but retention config location/values unspecified

Release 20261001 (2026-10-01)

---

03 **Expanded agentic vulnerability detection techniques** IMPROVED

55

Expands agentic detection coverage to include Ruby deserialization vulnerabilities, additional JWT authentication bypass techniques, and path-traversal bypasses against simple blocklist filters, plus more accurate login-field identification in the credential brute-force agent.

— Names specific vuln classes and agents but no steps to invoke

Release 20261001 (2026-10-01)

---

04 **Sitemap agent visibility** IMPROVED

50

Adds agent columns to the CSV sitemap export, listing which agents tested each endpoint, and now surfaces findings from scan-wide agents at the specific path covered rather than at the site root.

— Names CSV export and behavior change, no exact navigation given

Release 20261001 (2026-10-01)

---

05 **Recon agent scope and source ingestion** IMPROVED

50

The Recon agent now respects excluded paths, preventing discovery or probing of excluded assessment paths, and ingests README files and architecture diagrams inside uploaded source code archives to broaden attack-surface discovery.

— Two concrete behavior changes named, no configuration steps given

Release 20261001 (2026-10-01)

---

06 **Scan CLI authentication error handling** IMPROVED

50

Scan CLI now distinguishes a rejected API token as a credential error, separate from target or scan errors, and automatically retries requests that receive a temporary authentication error.

– Names Scan CLI and specific error types, no exact command shown

Release 20261001 (2026-10-01)

07 **Octo scope enforcement and false-positive triage** IMPROVED

45

Octo now blocks attack-agent requests that fall outside the target scope and configuration settings, and automatically marks DAST findings as false positives when agentic analysis confirms they are false positives.

– Describes two concrete behavior changes but no usage path Release 20261001 (2026-10-01)

08 **Multi-round time-based vulnerability confirmation** IMPROVED

45

Time-based vulnerability confirmation now runs multiple rounds with varied delay values, including per-round timing evidence in finding reports.

– Explains mechanism but no concrete numbers or configuration

Release 20261001 (2026-10-01)

09 **Agentic Pentest management in platform UI** NEW

45

Adds management of Agentic Pentests directly from the platform UI, with findings flowing into the same pipeline as standard DAST scans.

– States integration with DAST pipeline but no navigation detail product docs

10 **HTML format for assessment reports** NEW

40

Enables downloading assessment reports in HTML format for browser viewing or stakeholder sharing.

Download an HTML assessment report to share agentic pentest results with stakeholders who need a portable, browser-ready document.

📍 In the Invicti Platform, open the completed assessment and choose the HTML format option to download the report, then open it in any browser or attach it to a stakeholder communication.

– Clear download steps given, but no further mechanism detail

Release 20261001 (2026-10-01)

11 **Included Agents field populates immediately** IMPROVED

30

The `Included Agents` field now shows configured agents immediately when an assessment starts, instead of appearing empty.

— Names the exact field but is a minor UI fix `product docs`

---

12 **Attack agent input prioritization** IMPROVED 25

Attack agents now prioritize the most relevant inputs first to maximize coverage within time and credit limits.

— Describes behavior change only, no concrete mechanism or scope

`Release 20261001 (2026-10-01)`

---

13 **Clearer API source upload size limit messaging** IMPROVED 15

Adds clearer messaging about maximum file size limits for API source uploads.

— Vague wording with no stated limit or location `product docs`

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! Unused credits now expire at the end of your subscription term.

! Credit values are now whole numbers; decimal credit values are no longer used.

`Release 20261001 (2026-10-01)` BLOG ↗ LATEST STABLE

Invicti adds AI credit ledger, HTML/log-archive downloads, and expanded agentic vuln detection in Release 20261001.

└──▷ HOW TO FIND IT

Review pending and committed AI credits before starting a new agentic pentest to ensure sufficient balance.

> 📍 In the Invicti Platform, go to Pentests › AI Credit Usage to view the full transaction ledger, including pending reservations for in-progress assessments and your current available balance.

Download raw scanner logs and agent traces after an assessment to investigate a finding or share diagnostics with Invicti Support.

> 📍 In the Invicti Platform, open the completed assessment, then select the option to download the scanner logs and agent traces as a single compressed archive.

Download an HTML assessment report to share agentic pentest results with stakeholders who need a portable, browser-ready document.

📍 In the Invicti Platform, open the completed assessment and choose the HTML format option to download the report, then open it in any browser or attach it to a stakeholder communication.

Adds an `AI Credit Usage` section under the Pentests nav item showing every credit addition, deduction, pending reservation, and current available balance in a transaction ledger.

Adds real-time credit consumption to the Summary tab, showing credits used per assessment, with zero credit charge for failed assessments and pro-rated charges for canceled ones.

Enables downloading assessment reports in HTML format for browser viewing or stakeholder sharing.

Enables downloading scanner logs and raw agent traces as a single compressed archive per assessment.

Adds configurable retention period and storage location for internal agent scan logs.

+13 more

WAS THIS USEFUL? 👍 👎

DataDome

Feedback ⓘ Fraud SDK Go v2.1.0 (2026-10-02) BLOG ↗

RANK

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

> DataDome shipped SDK-level updates: an optional event-time field in the Fraud SDK for Go, and Swift Package Manager support plus an iOS privacy manifest in the Flutter Dio SDK.

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

01 **Swift Package Manager support and iOS privacy manifest** NEW

70

Flutter Dio SDK 3.5.0 adds Swift Package Manager (SPM) support for iOS native dependencies alongside existing CocoaPods support; apps on Flutter 3.44+ with SPM enabled resolve the plugin via SPM automatically. It also adds an iOS privacy manifest file `PrivacyInfo.xcprivacy` to the plugin.

— Names file, version threshold and dependency mechanism

Flutter Dio SDK 3.5.0 (2026-10-02)

THINNER COVERAGE BELOW

02 **eventDate field in Fraud SDK Go** NEW

45

The Fraud SDK Go module v2.1.0 adds an optional `eventDate` field to the payload to specify the effective event time.

— Names exact field but no further mechanism detail Fraud SDK Go v2.1.0 (2026-10-02)

> Fraud SDK Go v2.1.0 (2026-10-02) BLOG ↗ LATEST STABLE
>
> DataDome Fraud SDK Go v2.1.0 adds optional `eventDate` field for setting effective event time in payloads.

> Flutter Dio SDK 3.5.0 (2026-10-02) BLOG ↗ STABLE
>
> DataDome Flutter Dio SDK 3.5.0 adds Swift Package Manager support for iOS and an iOS privacy manifest.
>
> Adds Swift Package Manager (SPM) support for iOS native dependencies alongside existing CocoaPods support; apps on Flutter 3.44+ with SPM enabled resolve the plugin via SPM automatically.
>
> Adds an iOS privacy manifest (`PrivacyInfo.xcprivacy`) to the plugin.

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

ALEAPP

Feedback ⓘ v2026.4.3 NOTES ↗ CODE ↗ RANK

Android Logs Events And Protobuf Parser

> ALEAPP v2026.4.3 adds new SwiftKey keyboard artifact parsing, extends Telegram decryption to layer 227, and upgrades the raw image staging engine to fix several file-content extraction bugs.

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

01 **qnxprobe 1.57 raw image staging fixes** IMPROVED 77

Updates `qnxprobe` to 1.57 for `-t raw` acquisitions, now correctly staging OneDrive Files On-Demand placeholders (instead of zeros), XPRESS-compressed files (via Windows Overlay Filter) as their decompressed content, and NTFS-compressed files with early-ending compression units at their full length.

Parse a raw Android disk image to extract SwiftKey language model artifacts, including Terms by App to see per-app typing history.

```
$ python aleapp.py -t raw -i <path_to_image.img> -o
<path_for_report_output>
```

— Names version, flag and three concrete bug fixes with a usable command v2026.4.3

02 **SwiftKey language model artifacts** NEW 72

Adds four new artifact parsers derived from `dynamic.lm` files: Models, Terms, Term Sequences, and Terms by App. Covers both the Samsung Keyboard path (`app_SwiftKey`) and the standalone SwiftKey app path, surfacing per-app typing history.

Parse a raw Android disk image to extract SwiftKey language model artifacts, including Terms by App to see per-app typing history.

```
$ python aleapp.py -t raw -i <path_to_image.img> -o
<path_for_report_output>
```

— Names exact artifact types and file paths, has a runnable example v2026.4.3

03 **Telegram layer 227 decryption support** IMPROVED 68

Supports Telegram layer 227, including sender rank and guest peer fields on older constructors `0x3ae56482` and `0x95ef6f2b`, seventeen older message classes

including secret chat ones, service message sender reporting, and table preservation when a single record cannot be walked.

— Detailed mechanism and constructors but no direct usage example v2026.4.3

v2026.4.3 NOTES ↗ CODE ↗ LATEST STABLE

ALEAPP v2026.4.3 adds SwiftKey language model artifacts, Telegram layer 227 support, and improved raw NTFS image staging

└──▷ GET THIS VERSION ▾

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

└──▷ TRY IT

Parse a raw Android disk image to extract SwiftKey language model artifacts, including Terms by App to see per-app typing history.

```
$ python aleapp.py -t raw -i <path_to_image.img> -o
<path_for_report_output>
```

Adds four SwiftKey language model artifacts from `dynamic.lm` files — Models, Terms, Term Sequences, and Terms by App — covering both the Samsung Keyboard (`app_SwiftKey`) and SwiftKey app paths.

Supports Telegram layer 227, including sender rank and guest peer fields on older constructors (`0x3ae56482`, `0x95ef6f2b`), seventeen older message classes (including secret chat ones), service message sender reporting, and table preservation when a single record cannot be walked.

Updates `qnxprobe` to 1.57 for `-t raw` acquisitions, now correctly staging OneDrive Files On-Demand placeholders (instead of zeros), XPRESS-compressed files (via Windows Overlay Filter) as their content, and NTFS-compressed files with early-ending compression units at their full length.

WAS THIS USEFUL? 👍 👎

iLEAPP

Feedback ⓘ v2026.4.4 NOTES ↗ CODE ↗ RANK

iOS Logs, Events, And Plist Parser

> iLEAPP v2026.4.4 adds new Safari history-tag and knowledgeC location-activity artifacts, extends sysdiagnose and Session message parsing, and upgrades raw acquisition support via qnxprobe 1.57 to cover more filesystem and compression formats.

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

01 **Raw acquisition parsing via qnxprobe 1.57** IMPROVED 90

Updates `-t raw` acquisition reading via qnxprobe 1.57: adds JFFS2 compressed with OpenWrt's LZMA, UBI/UBIFS with bit errors, two adjacent JFFS2 partitions read as two separate filesystems, U-Boot environment and Belkin NVRM stores listed as single-file volumes (`uboot-env.bin` / `nvram.bin`), correct staging of OneDrive Files On-Demand placeholders, XPRESS-compressed NTFS files, and NTFS-compressed files with early-ending compression units.

— Names flag, version, and exact filesystem/file names v2026.4.4

02 **Safari history tags artifact and new History columns** NEW 80

Adds a `Safari Browser - History Tags` artifact with one row per tag-to-history-item link and one row per unlinked tag, and adds `Tags` , `Tag Identifiers` , and `Profile Name` columns to the existing `Safari Browser - History` artifact, with the profile name read from `SafariTabs.db` .

— Names artifact, new columns and source database v2026.4.4

03 **knowledgeC Location Activity artifact** NEW 80

Adds a `knowledgeC Location Activity` artifact for the `/app/locationActivity` stream, surfacing place name, address, coordinates, activity UUID, bundle ID, and the other coordinate pairs in the compressed payload with per-pair distances.

— Names stream path and exact surfaced fields v2026.4.4

04 **Session messaging parsing improvements** IMPROVED 70

Extends Session contact parsing to databases written by older Session versions that lack the `lastKnownClientVersion` column, and links Session 2.14.2 attachments named by a hash of the download URL.

— Names column and version but no UI path v2026.4.4

05 **Sysdiagnose Wi-Fi status artifact and TCC parsing** NEW

Adds a `Sysdiagnose - Wifi Status` artifact and extends TCC app-permission parsing to read from a sysdiagnose archive packed inside an extraction.

— Names artifact and source but little mechanism detail `v2026.4.4`

`v2026.4.4` `NOTES ↗` `CODE ↗` **LATEST** `STABLE`

iLEAPP v2026.4.4 adds Safari history tags, knowledgeC location activity, sysdiagnose Wi-Fi status, and broader flash-dump parsing via qnxprobe 1.57.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.4.4
https://github.com/abrignoni/iLEAPP.git
# already have the repo? check out this version:
$ git checkout v2026.4.4
```

Adds a `Safari Browser - History Tags` artifact with one row per tag-to-history-item link and one row per unlinked tag, and adds `Tags`, `Tag Identifiers`, and `Profile Name` columns to the existing `Safari Browser - History` artifact (profile name read from `SafariTabs.db`).

Adds a `knowledgeC Location Activity` artifact for the `/app/locationActivity` stream, surfacing place name, address, coordinates, activity UUID, bundle ID, and the other coordinate pairs in the compressed payload with per-pair distances.

Adds a `Sysdiagnose - Wifi Status` artifact and extends TCC app-permission parsing to read from a sysdiagnose archive packed inside an extraction.

Extends Session contact parsing to databases written by older Session versions that lack the `lastKnownClientVersion` column, and links Session 2.14.2 attachments named by a hash of the download URL.

Updates `-t raw` acquisition reading via qnxprobe 1.57: adds JFFS2 compressed with OpenWrt's LZMA, UBI/UBIFS with bit errors, two adjacent JFFS2 partitions read as two separate filesystems, U-Boot environment and Belkin NVRM stores listed as single-file volumes (`uboot-env.bin` / `nvram.bin`), correct staging of OneDrive Files On-Demand placeholders, XPRESS-compressed NTFS files, and NTFS-compressed files with early-ending compression units.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ

October 1, 2026 - CycloneDX VEX menu, Files downloads, and SBOM export errors (2026-10

BLOG ↗ RANK

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

> Finite State Platform added CycloneDX VEX export support with a new public API endpoint, and introduced Version Rescan for re-running scans while preserving prior results.

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

01 **CycloneDX VEX downloads and API endpoint** NEW 81

Adds a `GET /public/v0/vex/cyclonedx/{projectVersionId}` endpoint to retrieve CycloneDX VEX data for a project version, plus two new options in the Generate SBOM menu on a Version: CycloneDX VEX (vulnerabilities only) and CycloneDX VEX (triaged only).

Pull a CycloneDX VEX document for a specific project version to feed into a downstream vulnerability management pipeline.

```
$ curl -X GET
'https://platform.finitestate.io/public/v0/vex/cyclonedx/{projec
tVersionId}' -H 'Authorization: Bearer <token>'
```

Download a triaged-only CycloneDX VEX file from the platform UI when you only want to share externally vetted, triaged vulnerability data.

📍 In the platform, open a Version, click Generate SBOM, select CycloneDX VEX (triaged only), and click Download.

— Names exact endpoint and both download options with usage examples

October 1, 2026 - CycloneDX VEX menu, F…

02 **Version Rescan** NEW 68

A new Rescan action on a Version confirms and checkpoints the current Version, starts a new Version under the original name, and re-runs the scan from the stored file; for binary scans, analysts can choose which analyses to run while prior results remain accessible on the checkpoint Version.

Re-run a binary scan on an existing Version after updating analysis settings, without losing the original scan results.

> 📍 In the Finite State Platform, navigate to the Version you want to re-scan, select `Rescan`, choose which analyses to run for the binary scan, then confirm — the current Version is checkpointed and a new Version scan starts under the original name.

— Explains mechanism and workflow but no API/config surface named

`September 30, 2026 - Version Rescan (20…`

`October 1, 2026 - CycloneDX VEX menu, Files downloads, and SBOM export errors (2026-`

`BLOG ↗` **LATEST** `STABLE`

Finite State adds CycloneDX VEX downloads to the Generate SBOM menu, backed by a new GET `/public`/v0/vex/cyclonedx/{projectVersionId} endpoint.

└──▷ TRY IT

Pull a CycloneDX VEX document for a specific project version to feed into a downstream vulnerability management pipeline.

```
$ curl -X GET
'https://platform.finitestate.io/public/v0/vex/cyclonedx/{projectVe
rsionId}' -H 'Authorization: Bearer <token>'
```

Download a triaged-only CycloneDX VEX file from the platform UI when you only want to share externally vetted, triaged vulnerability data.

> 📍 In the platform, open a Version, click `Generate SBOM`, select `CycloneDX VEX (triaged only)`, and click Download.

Adds `GET /public/v0/vex/cyclonedx/{projectVersionId}` endpoint to retrieve CycloneDX VEX data for a project version.

Adds `CycloneDX VEX (vulnerabilities only)` and `CycloneDX VEX (triaged only)` download options to the Generate SBOM menu on a Version.

`September 30, 2026 - Version Rescan (2026-09-30)` `BLOG ↗` `STABLE`

Finite State Platform adds Version Rescan, letting analysts re-run a scan from the stored file while checkpointing prior results.

└──▷ HOW TO FIND IT

Re-run a binary scan on an existing Version after updating analysis settings, without losing the original scan results.

> 📍 In the Finite State Platform, navigate to the Version you want to re-scan, select `Rescan`, choose which analyses to run for the binary scan, then confirm — the current Version is checkpointed and a new Version scan starts under the original name.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Acalvio Technologies Acalvio

Feedback ⓘ

announced · Sep 12, 2026 | PRESS ↗ | RANK

Acalvio is a deception technology platform that deploys decoys and honeypots to detect advanced threats and insider attacks.

> Acalvio shipped two new product lines in this window: an agentless ITDR SaaS offering (ShadowPlex Identity Protection and Advanced Threat Defense) for identity attack-surface mapping and honeytoken-based detection, and the broader ShadowPlex distributed deception platform with its FluidDeception decoy engine, Shadow Networks adversary engagement, and DeceptionFarms management layer.

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

01 EDR, SIEM, SOAR, NAC and Splunk Adaptive Response integrations

NEW

70

Acalvio integrates with EDR, SIEM, SOAR, and NAC solutions to automate deception deployment, detect decoy credential use, and trigger real-time automated responses. It also ships pre-built integration with major SIEM systems for threat intelligence delivery, including Adaptive Response in Splunk Enterprise Security, which lets responders dynamically deploy deceptions directly from Splunk.

— Names specific integration surfaces including a concrete Splunk feature

launch-20260912-2ca2065a launch-20260912-6708abca

02 ShadowPlex Identity Protection attack-surface mapping NEW

60

ShadowPlex Identity Protection maps the identity attack surface across Microsoft Active Directory, Azure Active Directory, and endpoint credential caches to surface misconfigurations, attack paths, cached credentials, and lateral movement links.

— Names mapped surfaces but no usage path given launch-20260912-2ca2065a

03 DeceptionFarms management via ProjectionPoints NEW

60

DeceptionFarms manage distributed deceptions across network topologies using `ProjectionPoints`, eliminating the need for trunk or span port access.

— Names concrete mechanism ProjectionPoints and prior limitation removed

launch-20260912-6708abca

THINNER COVERAGE BELOW

04 **ShadowPlex Advanced Threat Defense honey accounts and tokens** NEW 55

ShadowPlex Advanced Threat Defense automatically generates honey accounts blended into Active Directory and disperses honey tokens across the enterprise network to detect credential theft and lateral movement.

— Explains mechanism but no concrete setup steps launch-20260912-2ca2065a

05 **Agentless SaaS deployment on cloud marketplaces** NEW 50

ShadowPlex Identity Protection and Advanced Threat Defense deploy as an agentless SaaS solution available on the Amazon Web Services, Microsoft Azure, and Google Cloud Platform marketplaces.

— Names exact marketplaces, a clear starting point launch-20260912-2ca2065a

06 **FluidDeception decoy generation engine** NEW 45

FluidDeception is a data-science-driven engine that automatically generates realistic, non-fingerprintable decoys, lures, baits, and breadcrumbs blended into the real network.

— Describes mechanism but no usage guidance launch-20260912-6708abca

07 **ShadowPlex distributed deception platform launch** NEW 30

Acalvio introduces ShadowPlex, a distributed deception platform supporting on-premises, private cloud, and public cloud deployment for breach and insider-threat detection.

— High-level launch description with no mechanism launch-20260912-6708abca

08 **Shadow Networks adversary engagement** NEW 30

Shadow Networks provide adversary engagement to capture TTPs (Tools, Techniques, and Processes) and inform response.

— Brief functional description, no operational detail launch-20260912-6708abca

announced · Sep 12, 2026 PRESS ↗

Acalvio launches agentless SaaS-based ITDR solution combining identity attack surface mapping and Active Defense deception.

Adds ShadowPlex Identity Protection, which maps identity attack surface across Microsoft Active Directory, Azure Active Directory, and endpoint credential caches to surface misconfigurations, attack paths, cached credentials, and lateral movement links.

Adds ShadowPlex Advanced Threat Defense, which automatically generates honey accounts blended into Active Directory and disperses honey tokens across the enterprise

network to detect credential theft and lateral movement.

Integrates with EDR, SIEM, SOAR, and NAC solutions to automate deception deployment, detect decoy credential use, and trigger real-time automated responses.

Deploys as an agentless SaaS solution available on Amazon Web Services, Microsoft Azure, and Google Cloud Platform marketplaces.

announced · Sep 12, 2026 PRESS ↗

Acalvio launches ShadowPlex, a distributed deception platform with automated decoy deployment, Shadow Networks engagement, and Splunk ES integration.

Introduces ShadowPlex, a distributed deception platform supporting on-premises, private cloud, and public cloud deployment for breach and insider-threat detection.

Includes FluidDeception, a data-science-driven engine that automatically generates realistic, non-fingerprintable decoys, lures, baits, and breadcrumbs blended into the real network.

Provides Shadow Networks for adversary engagement to capture TTPs (Tools, Techniques, and Processes) and inform response.

Introduces DeceptionFarms to manage distributed deceptions across network topologies using ProjectionPoints, eliminating the need for trunk or span port access.

Ships pre-built integration with major SIEM systems for threat intelligence delivery, including Adaptive Response in Splunk Enterprise Security to dynamically deploy deceptions from Splunk.

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ Fastly Terraform Provider 9.8.0 (2026-10-01)

BLOG ↗ RANK

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

> Fastly's Terraform Provider 9.8.0 adds new resources for managing Domain Management routing configurations.

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

01 Terraform resources for Domain Management routing NEW

70

New `fastly_routing_config` and `fastly_routing_config_domain_link` Terraform resources let users manage Domain Management routing configs with automatic versioning, and link those configs to domains.

— Names exact resource types but no usage example

Fastly Terraform Provider 9.8.0 (2026-1…

> Fastly Terraform Provider 9.8.0 (2026-10-01) BLOG ↗ LATEST STABLE
>
> Fastly Terraform Provider 9.8.0 adds `fastly_routing_config` and `fastly_routing_config_domain_link` resources for Domain Management routing configs.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Cloud Custodian

Feedback ⓘ 0.9.53.0 NOTES ↗ CODE ↗ RANK

Rules engine for cloud security, cost optimization, and governance, DSL in yaml for policies to query, filter, and take actions on resources

> Cloud Custodian 0.9.53 adds a `--vars-file` flag for runtime policy variable injection, expands coverage with 23 new AWS/GCP/K8s resource types spanning IoT, Amazon Q Business, S3 Tables, S3 file systems and SSM, adds new Azure Machine Learning and network security group actions, and standardizes dynamic tagging across clouds, while removing the `rename-tag` action from `aws.cloudwatch-dashboard`.

└──▷ **WHAT SHIPPED · 22 FEATURES** most completely described first ? what's the number?

01 **Runtime variable injection via vars-file** NEW 80

Adds `--vars-file` parameter to `custodian run` to inject variables into policy expansion at runtime without modifying policy files, e.g. `custodian run --vars-file vars.json -s out policy.yml`.

Inject environment-specific variables (e.g. account IDs, tag values) into a policy at run time without editing the YAML file.

```
$ custodian run --vars-file vars.json -s out policy.yml
```

— Runnable command and clear mechanism given. 0.9.53.0

02 **Tag management for ops-item and ssm-managed-instance** NEW 68

Adds `auto-tag-user`, `copy-related-tag`, `mark-for-op`, `remove-tag`, and `tag` actions plus a `marked-for-op` filter to `aws.ops-item` and `aws.ssm-managed-instance`.

— Enumerates all actions and filter with target resources. 0.9.53.0

THINNER COVERAGE BELOW

03 **IoT resource coverage expansion** NEW 58

Adds `aws.iot-certificate`, `aws.iot-ota-update`, and `aws.iot-policy` resource types, expanding IoT coverage, plus an `iot-logging` filter on `aws.account` for account-level IoT logging inspection.

— Names resources and filter but no deeper behavior. 0.9.53.0

04 **API Gateway v2 route resource with auth filtering** NEW

57

Adds `aws.apigwv2-route` resource type with authorization type filtering for API Gateway v2 routes.

– Includes a filtering mechanism beyond bare name. 0.9.53.0

05 **SSM resource types with tagging support** NEW

55

Adds `aws.ssm-association`, `aws.ssm-maintenance-window`, and `aws.ssm-patch-baseline` resource types with tagging support.

– Names resource types and tagging support, scope unclear. 0.9.53.0

06 **Amazon Q Business resource types** NEW

53

Adds `aws.qbusiness-application`, `aws.qbusiness-data-accessor`, `aws.qbusiness-data-source`, `aws.qbusiness-index`, and `aws.qbusiness-plugin` resource types for Amazon Q Business.

– Lists resource types but no mechanism detail. 0.9.53.0

07 **Azure Machine Learning lifecycle actions** NEW

52

Adds `archive` action to `azure.machine-learning-data-container` and `cancel` action to `azure.machine-learning-job`.

– Names both actions and their targets. 0.9.53.0

08 **Standardized dynamic tagging across clouds** IMPROVED

52

Standardizes dynamic tagging support across AWS, Azure, and GCP, including `default-value` with no `key` pattern on Azure and GCP.

– Names config pattern but scope of standardization vague. 0.9.53.0

09 **Bedrock custom model deployments filter** NEW

50

Adds `deployments` filter to `aws.bedrock-custom-model` for filtering by model deployments.

– Names filter and its purpose. 0.9.53.0

10 **DSQL cluster consecutive backups filter** NEW

50

Adds `consecutive-aws-backups` filter to `aws.dsql-cluster`.

– Names filter, no further mechanism. 0.9.53.0

11 **Removed rename-tag action from CloudWatch dashboard** BREAKING

50

The `rename-tag` action is removed from `aws.cloudwatch-dashboard`.

— Clear removal named but no migration path given. 0.9.53.0

12 **K8s custom resource dynamic CRD discovery** NEW 45

Adds `k8s.custom-resource` resource type for dynamic CRD discovery.

— Names behavior (dynamic discovery) with resource type. 0.9.53.0

13 **GCP Workspace user resource for MFA reporting** NEW 43

Adds `gcp.workspace-user` resource type for Google Workspace MFA reporting.

— States purpose but no further mechanism. 0.9.53.0

14 **S3 file system resource types** NEW 42

Adds `aws.s3files-access-point`, `aws.s3files-file-system`, and `aws.s3files-mount-target` resource types for S3 file system resources.

— Bare resource type names, no mechanism. 0.9.53.0

15 **S3 Tables resource types** NEW 40

Adds `aws.s3-table` and `aws.s3-table-bucket` resource types for S3 Tables.

— Bare resource type names, no mechanism. 0.9.53.0

16 **Audit mode for GCP Dataproc clusters** NEW 40

Adds audit mode support to `gcp.dataproc-clusters`.

— Names feature, mechanism unspecified. 0.9.53.0

17 **Parent-scoped dimensions in Azure metric filter** IMPROVED 40

Extends Azure metric filter to support parent-scoped dimensions.

— Describes extension without usage detail. 0.9.53.0

18 **RDS global cluster resource type** NEW 35

Adds `aws.rds-global-cluster` resource type for managing RDS global clusters.

— Bare resource type name only. 0.9.53.0

19 **AIOps investigation group resource type** NEW 35

Adds `aws.aiops-investigation-group` resource type for managing AWS AIOps investigation groups.

— Bare resource type name only. 0.9.53.0

20 **EventBridge event connection resource type** NEW 35

Adds `aws.event-connection` resource type for EventBridge event connections.

– Bare resource type name only. `0.9.53.0`

21 **GCP Vertex AI metadata artifact resource type** NEW 35

Adds `gcp.vertex-ai-metadata-store-artifact` resource type.

– Bare resource type name only. `0.9.53.0`

22 **Remove-rules action for Azure NSGs** NEW 35

Adds `remove-rules` action to `azure.networksecuritygroup`.

– Bare action name only. `0.9.53.0`

└──▷ BREAKING ON UPGRADE

! The `rename-tag` action is removed from `aws.cloudwatch-dashboard`.

`0.9.53.0` NOTES ↗ CODE ↗ LATEST STABLE

Cloud Custodian 0.9.53 adds 23 new AWS/GCP/K8s resource types, a `--vars-file` CLI flag, and new Azure ML/NSG actions.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 0.9.53.0 https://github.com/cloud-
custodian/cloud-custodian.git
# already have the repo? check out this version:
$ git checkout 0.9.53.0
```

└──▷ TRY IT

Inject environment-specific variables (e.g. account IDs, tag values) into a policy at run time without editing the YAML file.

```
$ custodian run --vars-file vars.json -s out policy.yml
```

Adds `--vars-file` parameter to `custodian run` to inject variables into policy expansion at runtime without modifying policy files.

Adds `aws.aiops-investigation-group` resource type for managing AWS AIOps investigation groups.

Adds `aws.apigwv2-route` resource type with authorization type filtering for API Gateway v2 routes.

Adds `aws.event-connection` resource type for EventBridge event connections.

Adds `aws.iot-certificate`, `aws.iot-ota-update`, and `aws.iot-policy` resource types, expanding IoT coverage.

+18 more

└──▷ **BREAKING ON UPGRADE**

! The `rename-tag` action is removed from `aws.cloudwatch-dashboard`.

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ Secure SaaS - October 01, 2026 BLOG ↗

RANK

Sysdig Secure is a container security platform providing runtime threat detection, forensics, and compliance monitoring for cloud-native environments.

> Sysdig Secure introduces a Technical Preview Supply Chain Findings page that surfaces license and end-of-life issues across the pipeline, registry, and runtime.

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

01 Supply Chain Findings page (Technical Preview) NEW

82

Adds the `Attack Surface > Supply Chain Findings` page (Technical Preview) listing license and end-of-life findings across pipeline, registry, and runtime resources, with filtering by severity, stage, finding type, license category, or license name, and grouping by package, image, or runtime resource. A new Supply Chain Findings tab is also added to the Resource Details drawer and the Component Details drawer, surfacing the same findings in context.

Review license and end-of-life supply chain findings for runtime resources, grouped by image, to prioritize remediation in production.

> 📍 In Sysdig Secure, go to Attack Surface › Supply Chain Findings. Set the Stage filter to runtime, the Finding Type filter to end-of-life, then use the Group By control to group by image.

— Names exact UI path, filters, and groupings; Technical Preview scope noted

Secure SaaS - October 01, 2026

Secure SaaS - October 01, 2026 BLOG ↗ LATEST STABLE

Sysdig Secure adds a Supply Chain Findings page (Technical Preview) surfacing license and end-of-life issues across pipeline, registry, and runtime.

└──▷ HOW TO FIND IT

Review license and end-of-life supply chain findings for runtime resources, grouped by image, to prioritize remediation in production.

> 📍 In Sysdig Secure, go to Attack Surface › Supply Chain Findings. Set the Stage filter to runtime, the Finding Type filter to end-of-life, then use the Group By control to group by image.

> Adds the `Attack Surface > Supply Chain Findings` page (Technical Preview) listing license and end-of-life findings across pipeline, registry, and runtime resources, with filtering by severity, stage, finding type, license category, or license name, and grouping by package, image, or runtime resource.
>
> Adds a `Supply Chain Findings` tab to the Resource Details drawer and the Component Details drawer.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Upwind

Feedback ⓘ

Upwind Now Secures Snowflake from AI Down to the Storage Underneath It (2026-10-01)

BLOG ↗ RANK

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

> Upwind launches Snowflake Security, a new CNAPP module that inventories Snowflake assets and AI workloads, applies CIS Snowflake Benchmark posture checks, traces access graphs from database roles to backing cloud storage, and flags risky token policy changes in real time.

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

01 **Snowflake asset inventory in Data sources** NEW 81

Adds a Snowflake Security integration under Data sources in the Upwind console, covering account-level asset inventory across warehouses, compute pools, users, account roles, network policies, database roles, secrets, network rules, stages, and external volumes. Connect a Snowflake account via `Data sources > Snowflake` to begin asset discovery.

Connect a Snowflake account to start receiving posture findings and AI asset inventory in the Upwind console.

> 📍 In the Upwind console, go to Data sources › Snowflake and add your Snowflake account to begin asset discovery and posture checks.

— Names exact UI path and full list of inventoried asset types

Upwind Now Secures Snowflake from AI Do…

02 **Cross-environment graph traversal to cloud storage** NEW 67

Adds cross-environment graph traversal that resolves Snowflake stages and external volumes to backing object storage buckets, storage integrations to assumed cloud roles, and network policies to cloud network topology — tracing a database role grant all the way to the bucket it exposes.

— Explains traversal mechanism end-to-end but no direct action step

Upwind Now Secures Snowflake from AI Do…

THINNER COVERAGE BELOW

03 **AI asset inventory within Snowflake schemas** NEW 58

Adds AI asset inventory within Snowflake schemas, covering foundation and custom models, agents, Cortex Search services, dynamic tables, notebooks, and MCP servers, nested by schema with role and user resolution.

— Lists asset types but no navigation or command given

Upwind Now Secures Snowflake from AI Do…

04 **CIS Snowflake Benchmark posture checks** NEW 53

Adds posture checks mapped to the CIS Snowflake Benchmark, covering privileges on account and database roles, authentication and network posture, AI asset access, and storage behind stages and external volumes.

— Names the benchmark and check areas but no runnable step

Upwind Now Secures Snowflake from AI Do…

05 **Token policy change monitoring** NEW 47

Adds continuous monitoring for account-level token policy changes that remove the network-policy requirement from programmatic access tokens, reported at the moment the setting changes rather than at audit time.

— Describes the specific policy change detected but no usage path

Upwind Now Secures Snowflake from AI Do…

Upwind Now Secures Snowflake from AI Down to the Storage Underneath It (2026-10-01)

BLOG ↗ LATEST STABLE

Upwind adds Snowflake Security: asset inventory, CIS Snowflake Benchmark posture checks, and graph traversal from roles to cloud storage.

└──▷ HOW TO FIND IT

Connect a Snowflake account to start receiving posture findings and AI asset inventory in the Upwind console.

📍 In the Upwind console, go to Data sources › Snowflake and add your Snowflake account to begin asset discovery and posture checks.

Adds Snowflake Security integration under Data sources in the Upwind console, covering account-level asset inventory across warehouses, compute pools, users, account roles, network policies, database roles, secrets, network rules, stages, and external volumes.

Adds AI asset inventory within Snowflake schemas, covering foundation and custom models, agents, Cortex Search services, dynamic tables, notebooks, and MCP servers, nested by schema with role and user resolution.

> Adds posture checks mapped to the CIS Snowflake Benchmark, covering privileges on account and database roles, authentication and network posture, AI asset access, and storage behind stages and external volumes.
>
> Adds cross-environment graph traversal that resolves Snowflake stages and external volumes to backing object storage buckets, storage integrations to assumed cloud roles, and network policies to cloud network topology — tracing a database role grant all the way to the bucket it exposes.
>
> Adds continuous monitoring for account-level token policy changes that remove the network-policy requirement from programmatic access tokens, reported at the moment the setting changes rather than at audit time.

WAS THIS USEFUL? 👍 👎

Plerion

Feedback ⓘ AgentCore runtimes get workload scanning:

NOTES ↗ RANK

Plerion is a cloud security platform that discovers and manages vulnerabilities and misconfigurations across infrastructure.

> Plerion added vulnerability, secret, and SBOM scanning for Amazon Bedrock AgentCore runtimes, backed by ten new AWS security checks and a Bucket ABAC tagging path to extend coverage to code-zip runtimes.

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

01 **Code-zip AgentCore runtime scanning via Bucket ABAC** NEW

90

Code-zip AgentCore runtimes can now be scanned like Lambda functions by turning on Bucket ABAC for the S3 bucket holding the zip and tagging it `PlerionAccess: Granted`.

— Exact config steps with named setting and tag value given

AgentCore runtimes get workload scannin…

02 **AgentCore runtime workload scanning** NEW

63

Plerion now scans Amazon Bedrock AgentCore runtimes that run a container image for vulnerabilities, sensitive data, and SBOM, surfacing results on the runtime's asset page. Scanning is activated automatically when Lambda scanning is already enabled for the integration.

— Mechanism and trigger explained, but no exact UI path given

AgentCore runtimes get workload scannin…

THINNER COVERAGE BELOW

03 **Ten new AgentCore vulnerability and secret checks** NEW

58

Adds ten new AWS checks `PLERION-AWS-1185` through `PLERION-AWS-1194` covering low/medium/high/critical image vulnerabilities, low/medium/high/critical secret exposure, exploitable vulnerabilities, and known-exploited vulnerabilities for AgentCore runtime images.

— Named check IDs but automatic, no user action required

AgentCore runtimes get workload scannin…

AgentCore runtimes get workload scanning: NOTES ↗ LATEST STABLE

Plerion now scans Amazon Bedrock AgentCore runtime container images for vulnerabilities, secrets, and SBOM.

Adds vulnerability, sensitive-data, and SBOM scanning for Amazon Bedrock AgentCore runtimes that run a container image, with results on the runtime's asset page — activated automatically when Lambda scanning is on for the integration.

Adds ten new AWS checks (PLERION-AWS-1185 through PLERION-AWS-1194) covering low/medium/high/critical image vulnerabilities, low/medium/high/critical secret exposure, exploitable vulnerabilities, and known-exploited vulnerabilities for AgentCore runtime images.

Enables code-zip AgentCore runtimes to be scanned like Lambda functions by turning on Bucket ABAC for the S3 bucket holding the zip and tagging it `PlerionAccess: Granted`.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Anchore syft

Feedback ⓘ v1.54.0 NOTES ↗ RANK

CLI tool and library for generating a Software Bill of Materials from container images and filesystems

> Syft v1.54.0 adds Perl ecosystem support for SBOM generation and extends container scanning to daemonless Podman.

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

01 **Perl (CPAN) ecosystem cataloging** NEW 50

Syft adds a new package cataloger for the Perl ecosystem (CPAN), enabling SBOM generation for Perl projects, and also reports the Perl interpreter itself as a CPAN package in generated SBOMs.

— Names ecosystem and package type but no commands or formats v1.54.0

02 **Non-daemon Podman scanning support** NEW 40

Syft now supports scanning container images via non-daemon Podman, i.e. daemonless mode.

— States capability but no flags or usage steps v1.54.0

v1.54.0 NOTES ↗ LATEST STABLE

syft v1.54.0 adds a Perl ecosystem cataloger, reports the Perl interpreter as a CPAN package, and supports non-daemon Podman analysis.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.54.0 https://github.com/anchore/syft.git
# already have the repo? check out this version:
$ git checkout v1.54.0
```

Adds a new package cataloger for the Perl ecosystem (CPAN), enabling SBOM generation for Perl projects.

Reports the Perl interpreter itself as a CPAN package in generated SBOMs.

Supports scanning container images via non-daemon Podman (daemonless mode).

WAS THIS USEFUL? 👍 👎

Google distroless

Feedback ⓘ changes since 2026-09-10

CODE ↗

Language focused docker images, minus the operating system.

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

Adds a Dart example (`examples/dart/bin/hello.dart`) demonstrating a distroless container built with the `base-nossl` image.

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

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Threat Intelligence

Watcher

Feedback ⓘ v3.6.0 NOTES ↗ RANK

Watcher - Open Source AI-powered Cyber Threat Intelligence & Hunting Platform. Developed with Django & React JS.

Unified DNS Threats Monitored table via `/api/dns_finder/threats_monitored/`, per-row triage workflow, and a first Kubernetes Helm chart under `kubernetes/watcher/`.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.6.0 https://github.com/thalesgroup-cert/Watcher.git
# already have the repo? check out this version:
$ git checkout v3.6.0
```

NEW DETECTION New `/api/dns_finder/threats_monitored/` endpoint serves a unified table merging dnstwist typosquats, Certificate Transparency keyword hits, and Subdomain Takeover findings in one view, with a `Source` filter (sub-filters: Fuzzer, Corporate Keyword, Provider) to narrow by origin.

New Helm chart under `kubernetes/watcher/` deploys Watcher, SearxNG, and CertStream with an optional Bitnami MySQL subchart, Ingress, migration Job, and secrets; settings available in `kubernetes/watcher/values.yaml`.

Unified triage status (`Pending`, `Suspected`, `Confirmed`, `Resolved`, `False Positive`) is now editable per-row across all three threat sources from a dropdown with no confirmation modal.

Comments of up to 300 characters can be added to any finding via the Edit modal.

Corporate proxy configured via `HTTP_PROXY` / `HTTPS_PROXY` in `deployment/.env` is now honoured by the Watcher backend; add internal service hostnames to `NO_PROXY` in `deployment/.env` to bypass the proxy for them.

└──▷ BREAKING ON UPGRADE

! In `/api/dns_finder/alert/`, the `status` field is now a string (`pending`, `suspected`, `confirmed`, `resolved`, `false_positive`) instead of `true` / `false` — any external script or integration that reads or writes this field must be updated.

! The database schema and data migrations applied by `make migrate` are not reversible — back up the database with `make backup` before upgrading.

! The Watcher backend now reads `HTTP_PROXY` / `HTTPS_PROXY` from `deployment/.env` — if Watcher previously reached internal services (TheHive, MISP,

SMTP relay, etc.) directly, add their hostnames to `NO_PROXY` in `deployment/.env` before running `make up` or those calls will be routed through the corporate proxy.

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ 7.261002.0 NOTES ↗ RANK

Open Cyber Threat Intelligence Platform

> OpenCTI shipped a cluster-wide usage-metrics manager with Prometheus export, new SMTP connection/greeting timeout controls, and several UI improvements including chatbot quota visibility, 1-click investigation duplication, and a split alerts/triggers notifications page.

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

01 Cluster-wide usage metrics manager NEW

88

A dedicated usage-metrics manager computes cluster-wide platform metrics once per interval and shares results via Redis so every node reports a consistent value; enable it with `PLATFORM_USAGE_METRICS_MANAGER__ENABLED=true`. It is disabled by default (metrics then report as `null` and are not exported to Prometheus) and includes a new `s3_used_size` metric reporting total platform file storage (imports, exports, embedded, support, template-generated, and draft files) summed from indexed sizes rather than by listing the S3/MinIO bucket directly.

Enable cluster-wide usage metrics collection so all nodes report a consistent, platform-wide figure rather than per-node observations.

env

```
PLATFORM_USAGE_METRICS_MANAGER__ENABLED=true
```

— Names env var, metric, Redis sharing and Prometheus behaviour product docs

02 SMTP connection and greeting timeouts NEW

78

New config keys `smtp:connection_timeout` / `SMTP__CONNECTION_TIMEOUT` cap the TCP connection wait to the SMTP server, and `smtp:greeting_timeout` / `SMTP__GREETING_TIMEOUT` caps the wait for the server greeting after connection opens, both specified in milliseconds.

Prevent SMTP hangs in high-latency environments by capping both the TCP connect and greeting phases.

yaml

```yaml
smtp:
  connection_timeout: 5000
  greeting_timeout: 10000
```

— Exact config keys and env vars with units and example product docs

THINNER COVERAGE BELOW

03 **Chatbot quota and feedback visibility** IMPROVED 30

The embedded chatbot UI now shows quota, prompts, and persisted message feedback.

— Names what changed but no mechanism or location 7.261002.0

04 **Split alerts and triggers in notifications** IMPROVED 30

The notifications page now splits alerts and triggers into separate sections.

— Names the page and change but no further detail 7.261002.0

05 **Stored filter remapping on user merge** IMPROVED 30

When users are merged, user references held in stored filters are now remapped to the surviving user.

— Explains mechanism but no scope or commands given 7.261002.0

06 **1-click investigation duplication** NEW 20

Investigations can now be duplicated with a single click.

— Bare description, no UI path or limits given 7.261002.0

7.261002.0 NOTES ↗ LATEST STABLE

OpenCTI 7.261002.0 adds chatbot quota visibility, 1-click investigation duplication, and a unified alerts/triggers notifications page.

└—▷ GET THIS VERSION ▾

```
$ git clone --branch 7.261002.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.261002.0
```

Shows quota, prompts, and persisted message feedback in the embedded chatbot UI.

Enables 1-click duplication of investigations.

Splits alerts and triggers into separate sections within the notifications page.

Remaps user references held in stored filters when users are merged.

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

Cribl Stream

Feedback ⓘ changelog · Oct 1, 2026 BLOG ↗ RANK

Cribl Stream is a data processing platform that routes, transforms, and enriches data from multiple sources to destinations in real-time.

> Cribl Stream adds the Gigamon Network Observability App, bringing deep-packet-inspection network metadata from Gigamon into Cribl Search, Lake, and dashboards.

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

01 **Gigamon Network Observability App integration** NEW 65

A new app, available in the Cribl Marketplace, integrates Gigamon Application Metadata Intelligence (AMI) network-derived telemetry into Cribl workflows. It ships prebuilt parsing, normalization, dashboards, saved searches, and investigation workflows for this telemetry, and lets teams retain rich AMI datasets in Cribl Lake while selectively routing subsets to SIEM, observability, and other analytics platforms per use case. Telemetry surfaced through Cribl Search includes DNS activity, TLS sessions, HTTP transactions, database interactions, and file operations alongside firewall, identity, and cloud telemetry, and extends visibility to unmanaged devices (IoT, OT, appliances) via Gigamon's agent-independent deep packet inspection across Layer 4 through Layer 7 traffic.

— Describes app, data sources, and workflows but no config/API specifics

2026-10-01-bd595565

changelog · Oct 1, 2026 BLOG ↗ LATEST STABLE

Cribl launches the Gigamon Network Observability App, bringing deep-packet-inspection network metadata into Cribl Search, Lake, and dashboards.

Adds the Gigamon Network Observability App for Cribl, available in the Cribl Marketplace, integrating Gigamon Application Metadata Intelligence (AMI) network-derived telemetry into Cribl workflows.

Provides prebuilt parsing, normalization, dashboards, saved searches, and investigation workflows designed around Gigamon AMI network telemetry inside Cribl.

Enables retention of rich AMI datasets in Cribl Lake while selectively routing subsets to SIEM, observability, and other analytics platforms based on per-use-case needs.

Surfaces Gigamon network-derived telemetry — covering DNS activity, TLS sessions, HTTP transactions, database interactions, and file operations — through Cribl Search alongside firewall, identity, and cloud telemetry.

> Extends network-derived telemetry visibility to unmanaged devices such as IoT, OT, and appliances by using Gigamon's agent-independent deep packet inspection across Layer 4 through Layer 7 traffic.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks Feedback ⓘ

Dashboard enhancements and bug fixes (2026-10-01) BLOG ↗ STABLE RANK

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

Databricks AI/BI dashboards gain SQL dataset materialization (Beta), multi-table data source selection, and Genie Code image widget authoring.

Extends dashboard materialization (Beta) to cover SQL datasets in addition to local metric views.

Allows selecting multiple tables simultaneously in the Add data source picker when building a dashboard data model.

Promotes numeric window measures to general availability in the low-code metric view UI.

Embedded dashboards can now show, collapse, or hide global filters by default.

Scheduled emails now preserve attachments that fit within the size limit instead of dropping all attachments when the combined image and PDF exceed it.

+3 more

UPDATED INTEGRATION Salesforce Data 360 file sharing connector enabled by default for compliance security profile workspaces

▾ 5 more releases in this issue · 2026-09-30 → 2026-10-01

Genie One enhancements (2026-10-01) BLOG ↗ STABLE

Genie One gains .txt file uploads, streaming responses, chat download, and personalized sample questions.

└──▷ HOW TO FIND IT

Download a Genie One chat session to share or archive a conversation with your team.

📍 In Genie One, open the sidebar and select Manage chats, then choose the chat you want and select Download.

Supports uploading plaintext `.txt` files to a Genie One conversation, in addition to previously supported file types.

Streams Genie One responses in real time as they are generated, so answers appear progressively rather than all at once.

Enables downloading a Genie One chat via the Manage chats sidebar.

Displays personalized sample questions beneath the search bar on the Genie One homepage.

`LTAP Direct Writes is generally available (2026-10-01)` BLOG ↗ STABLE

Databricks Lakebase: LTAP Direct Writes is GA, accelerating initial loads and full refreshes by bypassing the compute endpoint.

.

└──▷ **HOW TO FIND IT**

Opt in a new synced table to LTAP Direct Writes to avoid large initial loads competing with live queries on your compute endpoint.

📍 When creating a synced table in the Lakebase UI: go to Synchronization settings and select `LTAP Direct Writes`. This option is available for Lakebase projects running Postgres 16, 17, or 18.

`Schedule automatic data refresh in the Databricks Excel Add-in (Public Preview) (202`

BLOG ↗ LATEST STABLE

Databricks Excel Add-in gains scheduled automatic data refresh (Public Preview) with interval and calendar options.

`Control spreadsheet writeback with a workspace admin setting (2026-09-30)`

BLOG ↗ STABLE

Workspace admins can now gate Excel and Google Sheets writeback via the `Allow spreadsheet writeback` security setting.

└──▷ **BREAKING ON UPGRADE**

! Excel and Google Sheets writeback now requires a workspace admin to explicitly enable the `Allow spreadsheet writeback` security setting — existing writeback workflows will stop working until the setting is enabled.

`Salesforce Data 360 file sharing will soon be available by default for workspaces wi`

BLOG ↗ STABLE

Salesforce Data 360 file sharing connector becomes available by default for compliance security profile workspaces, enabling federated queries on Databricks compute.

WAS THIS USEFUL? 👍 👎

Exaforce

Feedback ⓘ announced · Oct 1, 2026 PRESS ↗ RANK

Exaforce runs AI Exabots that detect, triage, investigate and respond to threats across cloud, SaaS, identity, endpoint and code on a real-time knowledge graph.

Exaforce launches an agentic managed SOC service in Japan with autonomous Exabot AI agents for triage, investigation, automated response, and natural-language data querying.

NEW SUPPORTED PLATFORM Launches CTC-MSS AI SOC Service powered by Exaforce, a managed security service built on the Exaforce Agentic SOC platform and delivered by ITOCHU Techno-Solutions for enterprise customers in Japan.

Introduces Exabots — task-specific AI agents that autonomously handle detection, triage, investigation, and response, closing false positives with supporting evidence so only fully investigated, verified threats escalate to customer teams.

Enables automated response actions (such as resetting a password or revoking a session) configurable by ITOCHU Techno-Solutions per customer, rather than only sending notifications.

Provides natural-language investigation, allowing customer teams to query their own security data without learning complex query languages.

Operates the Exaforce real-time Data Platform as a SIEM, correlating identity, device/asset, log, configuration, and metadata sources into a real-time context graph for AI agent use.

+1 more

NEW INTEGRATION SaaS platform detection coverage

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ Iru now supports Windows 11 26H2 NOTES ↗ LATEST

STABLE RANK

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

Iru adds Windows 11 26H2 enrollment support and a new Managed OS Library Item targeting that release.

Supports enrolling devices running Windows 11 26H2 into Iru so teams can onboard new and upgraded hardware immediately.

Adds a new Managed OS Library Item for Windows 11 26H2, allowing IT teams to set 26H2 as the target version and keep devices on a consistent OS.

NEW INTEGRATION New Windows Auto Apps added to Iru Endpoint catalog

▾ 2 more releases in this issue · 2026-10-01 → 2026-10-02

> Auto App: Microsoft 365 Copilot NOTES ↗ STABLE
>
> Iru adds Microsoft 365 Copilot as a new Auto App in its catalog.

> New Windows Auto Apps NOTES ↗ STABLE

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Tenable Nessus

Feedback ⓘ Tenable Nessus10.12.5 (2026-10-01)

NOTES ↗ RANK

Tenable Nessus is a vulnerability scanner that identifies security weaknesses, misconfigurations, and compliance issues across networks and systems.

Manual plugin updates now require a valid Tenable signature file before archives are extracted or loaded; applies to Nessus and Nessus Manager only.

Manual plugin updates for Tenable Nessus and Tenable Nessus Manager now require a signature file to process plugin updates (does not affect Tenable Security Center plugin updates, Nessus Manager agent plugin updates, or Nessus application updates).

NEW SUPPORTED PLATFORM Alma Linux 10 scanner platform support

└──▷ BREAKING ON UPGRADE

! Manual plugin updates for Tenable Nessus and Tenable Nessus Manager now require a valid signature file — plugin archives without one will no longer be extracted or loaded. Tenable Security Center plugin updates, Nessus Manager agent plugin updates, and Nessus application updates are not affected.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Axonius

Feedback ⓘ Axonius Release Notes 9.0.8 (2026-09-24) BLOG ↗

RANK

Axonius is a cybersecurity platform that discovers, manages, and protects all devices and software across an organization's IT infrastructure.

> Axonius 9.0.8 introduces cross-environment Network Connections visibility, six new adapters (ARCON PAM, CyberArk Certificate Manager - SaaS, LinkShadow, ONESOURCE, Tencent Cloud EdgeOne, Zadara zCompute), ServiceNow ticketing for Recommended Actions, and expanded CVE filtering in Static Analysis Enrichment, alongside several breaking adapter changes (Exabeam Cloud, A10 Control, Bishop Fox, DefectDojo, CrowdStrike Falcon) that require reconfiguration.

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

01 Expanded CVE filtering in Static Analysis Enrichment NEW

90

Adds three filters to Static Analysis Enrichment Settings: Exclude CVEs without specific affected versions (default On for new customers, Off for existing), `Exclude CVEs published on or before [date]` (default date 2015-12-31 for new customers), and Exclude devices matching this query to remove matched devices from static analysis scope entirely.

Exclude aged and version-unbound CVEs from static analysis to reduce noise across your device inventory — navigate to Static Analysis Enrichment Settings and enable the two date and version filters.

> 📍 In the Axonius console, go to Settings › Exposures › Static Analysis Enrichment Settings, enable Exclude CVEs without specific affected versions and 'Exclude CVEs published on or before [date]', set the date cutoff (e.g., 2015-12-31), then save.

— Named config keys, defaults, and walkthrough steps given

Axonius Release Notes 9.0.8 (2026-09-24)

02 **Alerts as Incidents settings group for CrowdStrike Falcon**

80

BREAKING

Adds Alerts as Incidents settings group to the CrowdStrike Falcon adapter, replacing the standalone Fetch Alerts as Incidents toggle with a group that includes a fetch toggle and an Alert statuses to fetch selector; existing adapter configurations referencing the old standalone setting must be reconfigured.

— Named setting group and required migration Axonius Release Notes 9.0.8 (2026-09-24)

03 **Network Connections page and device tab** NEW

80

Introduces a new Network Connections page and per-device Network Connections tab for visibility into communication patterns across IT and OT environments, including protocols, ports, inbound/outbound connections, and cumulative connection counts.

Investigate unexpected lateral movement between OT and IT devices by reviewing the new Network Connections page for risky cross-segment traffic.

📍 In the Axonius console, go to Cyber-Physical Assets › Network Connections to see the environment-wide connection map, then open a specific device's asset profile and select the Network Connections tab to trace inbound and outbound sessions by protocol and port.

— Exact navigation path given in example Axonius Release Notes 9.0.8 (2026-09-24)

04 **Bishop Fox adapter stops creating Device assets from DNS**

80

BREAKING

The Bishop Fox adapter no longer creates Device assets from DNS records — the Devices from Domains, Devices from Subdomains, and Devices from Targets settings have been removed; existing connections are automatically upgraded to the corresponding URLs from Domains, URLs from Subdomains, or URLs from Targets settings.

— Names removed and replacement settings Axonius Release Notes 9.0.8 (2026-09-24)

05 **Checkmarx One adapter scan and application data** NEW

75

Adds Fetch Scan Ids, Fetch SAST Results, and Fetch API Security Results advanced settings to the Checkmarx One adapter for optional retrieval of SAST results and API Security risks from the last completed scan per project; the adapter now also fetches Business Applications.

— Named advanced settings, no example Axonius Release Notes 9.0.8 (2026-09-24)

06 **Required /api/v2 suffix for A10 Control adapter** BREAKING

75

The A10 Control adapter now requires the `/api/v2` suffix in the Host Name or IP Address connection setting (e.g., `https://a10control.example.com/api/v2`); existing connections using the old format will break.

— Exact required URL format given Axonius Release Notes 9.0.8 (2026-09-24)

---

07 **Application settings and user data in Akamai adapter** NEW

74

Adds Fetch Application Settings advanced setting to the Akamai Application Security adapter to retrieve application settings including account lock state and registration completion, and adds support for fetching Users assets including email, role, team, lock state, and registration state.

— Named setting and fetched user fields Axonius Release Notes 9.0.8 (2026-09-24)

---

08 **OAuth 2.0 client credentials for file-based adapters** NEW

74

Adds support for OAuth 2.0 client credentials authentication to Custom Files and other file-based adapters when using URL/FTP as the file source, sending an `Authorization: Bearer` header automatically before each fetch cycle.

— Names auth mechanism and header used Axonius Release Notes 9.0.8 (2026-09-24)

---

09 **Expanded asset coverage in Cisco Meraki adapter** NEW

73

Adds support in Cisco Meraki adapter for fetching client SaaS application data, Meraki organizations as Accounts assets, networks and VLANs as Networks assets, and appliance NAT rules (port forwarding, 1:1 NAT, 1:Many NAT) as Cisco Meraki NAT Rule Firewall assets.

— Names each new asset type and mechanism Axonius Release Notes 9.0.8 (2026-09-24)

---

10 **Project Resources enrichment in Digital Ocean adapter** NEW

71

Adds Project Resources advanced setting to the Digital Ocean adapter to enrich Droplets, Load Balancers, and Kubernetes Clusters assets with their DigitalOcean Project ID and Project Name.

— Named setting and enriched fields Axonius Release Notes 9.0.8 (2026-09-24)

---

11 **Rule-based enrichment limit and removed settings in Exabeam Cloud**

71

BREAKING

Adds option to limit rule-event enrichment to specific Rule IDs in the Exabeam Cloud adapter. The `Maximum devices per API call` and `Number of days to` advanced settings have been removed — existing configurations using these settings must be updated.

— Names removed settings requiring migration Axonius Release Notes 9.0.8 (2026-09-24)

---

12 **ServiceNow destination for Recommended Actions tickets** NEW

70

Adds ServiceNow as a supported destination in the Recommended Actions ticket creation template flow, enabling automatic routing to assigned remediation owners via the same out-of-the-box template mechanism used for Jira.

Restrict Recommended Actions ticket creation to ServiceNow so remediation owners receive auto-routed tickets without manual intervention.

> 📍 In the Axonius console, go to Exposures › Recommended Actions, open a template ticket flow, select ServiceNow as the ticketing destination, map the remediation owner to the configured ServiceNow connection, then save the template.

— Exact UI steps for configuration given Axonius Release Notes 9.0.8 (2026-09-24)

---

13 **Six new adapters across PAM, certificates, NDR, tax and cloud**

NEW

70

Adds ARCON PAM adapter (Users and Devices), CyberArk Certificate Manager - SaaS adapter (Certificates, Devices, Users), LinkShadow adapter (Devices and Users), ONESOURCE adapter (Users and Groups), Tencent Cloud EdgeOne adapter (Domains & URLs, Compute Services, Load Balancers, Network/Firewall Rules), and Zadara zCompute adapter (Devices, Compute Images, Disks, Network/Firewall Rules).

— Names each adapter and asset types fetched Axonius Release Notes 9.0.8 (2026-09-24)

---

14 **DefectDojo findings reclassified as Application Resources**

70

BREAKING

DefectDojo findings are now parsed as Application Resources instead of Application Services — existing queries, dashboards, and reports targeting Application Services must be updated to target Application Resources.

— Clear migration guidance given Axonius Release Notes 9.0.8 (2026-09-24)

---

15 **Reduced default worker count for Network Routes enrichment**

68

BREAKING

The default number of parallel workers for Network Routes enrichment is reduced from 10 to 2; environments relying on the previous default throughput will need to explicitly set a higher value (up to 15) if needed.

– Exact before/after numbers and remediation given

Axonius Release Notes 9.0.8 (2026-09-24)

16 **Fetch pending status in Adapter Fetch History** IMPROVED

65

Adds a Fetch pending status and Queued At column to Adapter Fetch History for connections queued and waiting for an available worker, with Terminate Fetch now also cancelling connections in Fetch pending status.

– Named status and column, no example Axonius Release Notes 9.0.8 (2026-09-24)

17 **Multiple CyberArk Vault instances in Enterprise Password Manager**

NEW

65

Adds support for configuring multiple CyberArk Vault instances in the Enterprise Password Manager, each with its own domain, port, Application ID, and API Prefix settings.

– Names config fields, no example Axonius Release Notes 9.0.8 (2026-09-24)

18 **AD/LDAP static role support in HashiCorp Vault integration** NEW

62

Adds `Active Directory / LDAP - Static Role` secret engine type support to HashiCorp Vault integration for retrieving current passwords for LDAP static roles.

– Named secret engine type, no example Axonius Release Notes 9.0.8 (2026-09-24)

19 **Device tag fetching in Cisco Catalyst Center adapter** NEW

62

Adds option to fetch tags assigned to network devices in the Cisco Catalyst Center adapter, supporting `Category: Value` format as key-value tags and unformatted tag names as key-only tags.

– Names tag format support Axonius Release Notes 9.0.8 (2026-09-24)

20 **Column Views role permissions in Adapters area** NEW

60

Adds `Column Views - View` and `Column Views - Add` role permissions to control access to column views in the Adapters area; predefined Admin and Owner roles have both enabled by default after upgrade.

– Named permissions but no usage example Axonius Release Notes 9.0.8 (2026-09-24)

21 **Email security grades in BitSight adapter** NEW 60

Adds option to fetch email security grades (DMARC Grade, SPF Grade, DKIM Grade) for domains in the BitSight Security Ratings adapter.

— Names exact grade fields Axonius Release Notes 9.0.8 (2026-09-24)

THINNER COVERAGE BELOW

22 **Application Settings fetch in Dynatrace adapter** NEW 59

Adds option to fetch Dynatrace Application Settings, including IP allowlist, user account status, group permission, and platform-token status.

— Names fetched fields, no example Axonius Release Notes 9.0.8 (2026-09-24)

23 **Configurable email response link expiration** NEW 55

Adds configurable email response link expiration under General GUI Settings, with a default of 24 hours and a maximum of 120 hours.

— Exact defaults given, no example Axonius Release Notes 9.0.8 (2026-09-24)

24 **Optional credentials in Cisco ISE connection config** IMPROVED 48

Adds option to make Username and Password fields optional in Cisco ISE connection configuration when pxGrid is enabled without credentials.

— Thin description of config condition Axonius Release Notes 9.0.8 (2026-09-24)

25 **Invoices as Expenses in Bill adapter** NEW 48

Adds option to fetch Invoices as Expenses in the Bill adapter, in addition to the existing Bills as Expenses fetch.

— Thin, single setting description Axonius Release Notes 9.0.8 (2026-09-24)

26 **Vulnerability status filter in CrowdStrike Falcon Discover** NEW 42

Adds option to filter fetched vulnerabilities by status in the CrowdStrike Falcon Discover adapter.

— Thin, no exact field name given Axonius Release Notes 9.0.8 (2026-09-24)

└──▷ BREAKING ON UPGRADE

! The Maximum devices per API call and Number of days to advanced settings have been removed from the Exabeam Cloud adapter — existing configurations using these settings must be updated.

- ! The A10 Control adapter now requires the `/api/v2` suffix in the Host Name or IP Address connection setting (e.g., `https://a10control.example.com/api/v2`); existing connections using the old format will break.
- ! The Bishop Fox adapter no longer creates Device assets from DNS records — the Devices from Domains, Devices from Subdomains, and Devices from Targets settings have been removed; existing connections are automatically upgraded to the corresponding URLs from Domains, URLs from Subdomains, or URLs from Targets settings.
- ! DefectDojo findings are now parsed as Application Resources instead of Application Services — existing queries, dashboards, and reports targeting Application Services must be updated to target Application Resources.
- ! The default number of parallel workers for Network Routes enrichment is reduced from 10 to 2; environments relying on the previous default throughput will need to explicitly set a higher value (up to 15) if needed.
- ! The CrowdStrike Falcon standalone Fetch Alerts as Incidents advanced setting has been replaced by the Alerts as Incidents settings group — existing adapter configurations referencing the old standalone setting must be reconfigured.

Axonius Release Notes 9.0.8 (2026-09-24) BLOG ↗ LATEST STABLE

Axonius 9.0.8 adds Network Connections visibility, six new adapters, ServiceNow template ticketing, and expanded static analysis CVE filtering.

└──▷ HOW TO FIND IT

Exclude aged and version-unbound CVEs from static analysis to reduce noise across your device inventory — navigate to Static Analysis Enrichment Settings and enable the two date and version filters.

> 📍 In the Axonius console, go to Settings › Exposures › Static Analysis Enrichment Settings, enable Exclude CVEs without specific affected versions and 'Exclude CVEs published on or before [date]', set the date cutoff (e.g., 2015-12-31), then save.

Investigate unexpected lateral movement between OT and IT devices by reviewing the new Network Connections page for risky cross-segment traffic.

> 📍 In the Axonius console, go to Cyber-Physical Assets › Network Connections to see the environment-wide connection map, then open a specific device's asset profile and select the Network Connections tab to trace inbound and outbound sessions by protocol and port.

Restrict Recommended Actions ticket creation to ServiceNow so remediation owners receive auto-routed tickets without manual intervention.

> 📍 In the Axonius console, go to Exposures › Recommended Actions, open a template ticket flow, select ServiceNow as the ticketing destination, map the remediation owner to the configured ServiceNow connection, then save the template.

Adds three new filtering options to Static Analysis Enrichment Settings: `Exclude CVEs without specific affected versions` (default On for new customers, Off for existing), `Exclude CVEs published on or before [date]` (default date 2015-12-31 for new customers), and `Exclude devices matching this query` to remove matched devices from static analysis scope entirely.

Adds `Column Views - View` and `Column Views - Add` role permissions to control access to column views in the Adapters area; predefined Admin and Owner roles have both enabled by default after upgrade.

Adds a `Fetch pending` status and `Queued At` column to Adapter Fetch History for connections queued and waiting for an available worker, with `Terminate Fetch` now also cancelling connections in `Fetch pending` status.

Adds `Alerts as Incidents` settings group to the CrowdStrike Falcon adapter, replacing the standalone `Fetch Alerts as Incidents` toggle with a group that includes a fetch toggle and an `Alert statuses to fetch` selector.

Adds `Fetch Scan Ids`, `Fetch SAST Results`, and `Fetch API Security Results` advanced settings to the Checkmarx One adapter for optional retrieval of SAST results and API Security risks from the last completed scan per project.

+24 more

└──▷ BREAKING ON UPGRADE

- ! The `Maximum devices per API call` and `Number of days to` advanced settings have been removed from the Exabeam Cloud adapter — existing configurations using these settings must be updated.
- ! The A10 Control adapter now requires the `/api/v2` suffix in the Host Name or IP Address connection setting (e.g., `https://a10control.example.com/api/v2`); existing connections using the old format will break.
- ! The Bishop Fox adapter no longer creates Device assets from DNS records — the `Devices from Domains`, `Devices from Subdomains`, and `Devices from Targets` settings have been removed; existing connections are automatically upgraded to the corresponding `URLs from Domains`, `URLs from Subdomains`, or `URLs from Targets` settings.
- ! DefectDojo findings are now parsed as Application Resources instead of Application Services — existing queries, dashboards, and reports targeting Application Services must be updated to target Application Resources.
- ! The default number of parallel workers for Network Routes enrichment is reduced from 10 to 2; environments relying on the previous default throughput will need to explicitly set a higher value (up to 15) if needed.
- ! The CrowdStrike Falcon standalone `Fetch Alerts as Incidents` advanced setting has been replaced by the `Alerts as Incidents` settings group — existing adapter configurations referencing the old standalone setting must be reconfigured.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Didit

Feedback ⓘ snapshot-20261002 NOTES ↗ RANK

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

> Didit's biggest shipment this window is production digital ID wallet verification across six European schemes, paired with new KYB ownership-tier and AML-monitoring APIs, five new MCP tools, and a wide set of workflow, SDK, and document-coverage improvements — alongside a breaking webhook-signature change.

└──▷ **WHAT SHIPPED · 29 FEATURES** most completely described first ? what's the number?

01 **KYB ownership data tiers and endpoints** NEW 85

Adds `POST /v3/kyb/shareholders/` and `POST /v3/kyb/ubo/` endpoints for standalone business ownership data, with the response naming the tier requested and the tier delivered (falling back to Lite when ownership data is unavailable). Also introduces selectable KYB data tiers — Lite, Shareholders, or UBOs — per country in the KYB Registry Check step and in one-off checks, with price shown before selection.

— Endpoints and tiers named; pricing mechanism not detailed snapshot-20261002

02 **AML monitoring API with batch and dry-run support** NEW 85

Adds `POST /v3/users/aml-monitoring/` and `POST /v3/businesses/aml-monitoring/` to turn ongoing AML monitoring on or off for up to 1,000 profiles at a time, with a dry-run estimate mode and background job support for large batches.

— Endpoints, limit, and modes all named snapshot-20261002

03 **New Didit MCP tools** NEW 85

Adds `didit_session_explain_decision` to trace why a session ended Approved, Declined or In Review; `didit_session_webhooks` to list a session's webhook deliveries and optionally wait for the first one; `didit_session_webhook_resend` to resend a specific delivery; `didit_workflow_build_graph` to turn a plain feature list (with countries named exactly) into a complete, validated workflow graph; and `didit_workflow_get_kyb_registry_catalog` to return registry tiers and prices per country.

— Five tools named with exact purpose each snapshot-20261002

04 **Sandbox AML hit simulation** NEW 81

Adds a `sandbox_scenario` parameter on `POST /v3/aml/` to simulate an AML hit in sandbox without an end-user session, so teams can validate alert-handling workflows before going live.

Simulate an AML hit in sandbox to validate your alert-handling workflow before going live

```
$ curl -X POST https://api.didit.me/v3/aml/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"vendor_data": "test-user-42", "sandbox_scenario":
"hit"}'
```

— Endpoint, parameter, and runnable curl example given `snapshot-20261002`

---

05 **Camera control options for kiosk and operator flows** NEW

80

Adds `defaultDocumentCamera`, `defaultLivenessCamera`, `showDocumentCameraSwitchButton` and `showLivenessCameraSwitchButton` options to the JavaScript SDK for kiosk and operator-facing camera control, also available as query parameters `liveness_camera=back` and `liveness_camera_switch=false`.

— Exact flags and query params named `snapshot-20261002`

---

06 **Branded session-history PDF endpoint** NEW

73

Adds `GET /v3/users/{vendor_data}/generate-pdf/` to return a branded PDF of a user's 20 most recent completed sessions.

Retrieve a user's branded session history PDF via API to attach to a compliance record or audit package

```
$ curl -X GET
'https://api.didit.me/v3/users/<vendor_data>/generate-pdf/' \
  -H 'Authorization: Bearer <token>' \
  --output session-history.pdf
```

— Exact endpoint and usage example provided `snapshot-20261002`

---

07 **Digital ID wallet verification in production** NEW

68

Introduces digital ID wallet verification in production for MitID (Denmark), Finnish Trust Network (Finland), Smart-ID (Estonia, Latvia, Lithuania, Belgium), Mobile-ID (Estonia, Lithuania), and BankID Sweden.

— Schemes and countries named, no integration steps `snapshot-20261002`

---

08 **Terminal node visibility in decision API** NEW 65

Exposes `workflow_terminal_node_id` on the decision API with `include=events` to surface which branch path and terminal node ended a session.

— Field and query param named, no usage sample snapshot-20261002

09 **New ID Verification warnings for missing back side and covered fields**

NEW 65

Adds a `DOCUMENT_BACK_SIDE_NOT_AVAILABLE` informational warning to ID Verification results when only a front side could be collected, and a covered-critical-field warning that fires when a document's number or another critical field is obscured, configurable and available on the standalone ID Verification API.

— Two warning conditions named, no config detail snapshot-20261002

10 **Webhook signature verification migration** BREAKING 65

`X-Signature-Simple` is deprecated — webhooks must now be verified with `X-Signature-V2`.

— Exact header names given, no migration timeline snapshot-20261002

11 **Transaction submission in web SDK** NEW 64

Adds submitTransaction() to the web SDK for submitting transactions for monitoring using a scoped transaction token, with automatic device intelligence collection.

— Method and mechanism named, no code sample snapshot-20261002

12 **Skip liveness and face match after database lookup** NEW 62

Adds a per-country `skip_liveness_and_face_match` switch (on by default) to control whether liveness and face-match steps run after a database lookup.

— Exact config key and default named snapshot-20261002

THINNER COVERAGE BELOW

13 **Expanded document type and country coverage** NEW 58

Adds support for Police ID cards from Argentina, Brunei, Colombia, and the Philippines; Alberta driving licence and ID card (2026) with barcode validation; Moldova driving licence (2026); NYC Identification Card (2019); and back-side capture for Brazilian state driving licences, Thai driving licences, and Canadian passport Page B.

— Many documents named, no API change involved snapshot-20261002

14 **Document liveness scores in decision API** IMPROVED 57

Adds `document_liveness` scores (portrait replacement, printed copy, screen replay) to every decision API result, even when no warning fired.

— Field and sub-scores named, no example shown snapshot-20261002

15 **Contradiction detection in workflow editor** NEW 57

Adds contradiction detection to the workflow editor: branches whose condition contradicts their destination node get a warning badge on the canvas and in the graph-validation API.

— API referenced loosely, no endpoint given snapshot-20261002

16 **Document-subtype branch routing in workflow editor** NEW 55

Adds branch-on-document-subtype routing in the workflow editor to route sessions on exact document subtypes (e.g. national ID card vs. residence permit).

— Mechanism named with example, UI-only snapshot-20261002

17 **Liveness video analysis can decline sessions and flag virtualized devices**

IMPROVED 55

Enables liveness video analysis to decline a session (not only report), and flags liveness captures from emulated or virtualized devices as an actionable risk.

— Behavior change explained, no config name given snapshot-20261002

18 **Transaction Monitoring rule library additions** NEW 55

Adds iGaming rule presets (bonus abuse, chip dumping, and others) to the Transaction Monitoring rules library, and current-value velocity filters to Transaction Monitoring rule conditions.

— Presets and filter type named, no config keys snapshot-20261002

19 **Didit Copilot skills and undo** NEW 55

Adds Didit Copilot slash skills — typing "/" in the composer browses page-specific action templates — and undo for Copilot workflow edits, where a change card lists what changed and `Undo` reverts it.

— Two UI mechanisms named, grouped under Copilot snapshot-20261002

20 **Per-country source selection for non-document lookups** NEW

55

Adds source selection per country in the Countries tab for non-document lookups where multiple sources exist, including Argentina without selfie and Brazil with verified selfie.

— Examples named, no API or config key `snapshot-20261002`

21 **Face Search toggle per workflow step** NEW

53

Adds a `face_search_enabled` switch on the Rules tabs for Liveness and Age Estimation steps to disable Face Search per workflow.

— Config switch named; little mechanism beyond that `snapshot-20261002`

22 **Detected codes field for documents** NEW

53

Adds a `detected_codes` field to Proof of Address and Document AI results, returning QR codes and barcodes found on the document.

— Field named with scope, no example `snapshot-20261002`

23 **Same-face-different-document control** NEW

50

Adds a Same face, different document control to configure what happens when a face already approved on another identity document reappears.

— Control named, mechanism brief `snapshot-20261002`

24 **AML screening of persons named in Document AI output** NEW

50

Adds AML screening of any person named by an earlier Document AI step (e.g. from a tax certificate or articles of incorporation).

— Mechanism and example given, no API name `snapshot-20261002`

25 **Wallet-ownership check in web SDK** NEW

48

Adds wallet-ownership check support to the web SDK: when a transaction requires a wallet-ownership check, the SDK runs it and reports completion via callback.

— Mechanism described loosely, no API name `snapshot-20261002`

26 **On-demand wallet screening check** NEW

48

Adds wallet screening as a single on-demand console check (now labelled Wallet Verification & Screening) without requiring a full session.

— UI label named, limited mechanism `snapshot-20261002`

27 **SMS sender name customization** NEW

48

Adds SMS sender name customization via the new SMS tab in Customization, with live preview.

— UI path named, minor feature `snapshot-20261002`

28 **Expired-document resubmission emails** NEW 45

Adds expired-document resubmission emails that specify the document expired, customizable via the Texts tab.

— UI path named, thin mechanism `snapshot-20261002`

29 **Additional SDK modal languages** NEW 38

Adds four additional modal languages to the web SDK: Bosnian, Kazakh, Kyrgyz, and Albanian.

— Thin addition, just a language list `snapshot-20261002`

└──▷ **BREAKING ON UPGRADE**

! `X-Signature-Simple` is deprecated — webhooks must now be verified with `X-Signature-V2`

`snapshot-20261002` NOTES ↗ LATEST STABLE

Didit adds digital ID wallets, KYB data tiers, AML monitoring APIs, MCP tools, and expanded workflow controls

└──▷ **TRY IT**

Retrieve a user's branded session history PDF via API to attach to a compliance record or audit package

```
$ curl -X GET
'https://api.didit.me/v3/users/<vendor_data>/generate-pdf/' \
  -H 'Authorization: Bearer <token>' \
  --output session-history.pdf
```

Simulate an AML hit in sandbox to validate your alert-handling workflow before going live

```
$ curl -X POST https://api.didit.me/v3/aml/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"vendor_data": "test-user-42", "sandbox_scenario": "hit"}'
```

Adds `POST /v3/kyb/shareholders/` and `POST /v3/kyb/ubo/` endpoints for standalone business ownership data; response names the tier requested and the tier delivered, with Lite as fallback when ownership data is unavailable

Adds `POST /v3/users/aml-monitoring/` and `POST /v3/businesses/aml-monitoring/` to turn ongoing AML monitoring on or off for up to 1,000 profiles at a time, with a dry-run estimate mode and background job support for large batches

Adds `GET /v3/users/{vendor_data}/generate-pdf/` to return a branded PDF of a user's 20 most recent completed sessions

Adds `sandbox_scenario` parameter on `POST /v3/aml/` to simulate an AML hit in sandbox without an end-user session

Adds `face_search_enabled` switch on the Rules tabs for Liveness and Age Estimation steps to disable Face Search per workflow

+33 more

└─▷ **BREAKING ON UPGRADE**

! `X-Signature-Simple` is deprecated — webhooks must now be verified with `X-Signature-V2`

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback

Roles Based Access Release Notes 2025 (Preview) (2026-10-01) BLOG ↗ RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> JumpCloud expanded its AI Gateway with a dozen new third-party integrations and overhauled custom admin role permissions, adding CRUD-level granularity while renaming and restructuring several role scopes.

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

01 **RBAC scope renaming and reorganization** BREAKING 78

JumpCloud renamed several custom role scopes: Policy Management became Access Management, System Management became Devices, and Authentication Policies became Conditional Access Policies and IP Lists. Identity Management was split to host only user-related scopes, with Directory Management and Group Management moved into a new Directory and Groups Management category; SaaS Applications and Asset Management moved under a new SaaS and Asset Management category; RADIUS Authentication moved under Access Management; Application Templates were folded into the Applications scope; and all User Portal-related APIs and operations moved under the Multi-Factor Authentication scope. Existing custom roles or documentation referencing the old scope names must be updated, and role assignments tied to the moved scopes should be reviewed.

— Names every renamed/moved scope and migration impact, no exact config path

Roles Based Access Release Notes 2025 (…

02 **CRUD permissions for Identity and Application Management** NEW 73

Custom admin roles now support full CRUD operations plus a No Access permission level for the Identity Management and Application Management scopes, enabling more granular separation between read-only and full-edit admin access for users and applications. Admins can configure this via Roles > Custom Roles > (select role) > Permissions, setting Identity Management and Application Management independently.

Configure a custom admin role with granular CRUD vs. read-only separation for user management and application management after the new scope split.

📍 In the JumpCloud Admin Portal, go to Roles › Custom Roles › (select or create a role) › Permissions, then set Identity Management to the desired CRUD level and Application Management to No Access or a specific CRUD level as appropriate for that admin.

— Explains mechanism and gives UI steps, but no API/flag names

Roles Based Access Release Notes 2025 (…

THINNER COVERAGE BELOW

03 **AI Gateway third-party integrations** NEW 45

JumpCloud's AI Gateway gained 12 new configurable integrations: Airtable, Attio, BigQuery, Calendly, Databricks, Granola, Lucid, monday.com, PagerDuty, PayPal, Shopify, and Stripe.

— Names all 12 integrations but no setup detail product docs

└──▷ BREAKING ON UPGRADE

! The Policy Management scope is renamed to Access Management — any custom roles or documentation referencing Policy Management must be updated.

! The System Management scope is renamed to Devices — any custom roles or documentation referencing System Management must be updated.

! The Authentication Policies scope is renamed to Conditional Access Policies and IP Lists — any custom roles or documentation referencing Authentication Policies must be updated.

! Identity Management no longer bundles all identity-related scopes; Directory and Groups Management is now a separate scope requiring explicit assignment in custom roles.

! Directory Management and Group Management scopes have moved under Directory and Groups Management — existing role assignments may need to be reviewed.

! SaaS Applications and Asset Management scopes have moved under SaaS and Asset Management — existing role assignments may need to be reviewed.

! RADIUS Authentication has moved under Access Management — existing role assignments scoped to its previous location may need to be updated.

! User Portal APIs and operations are now under the Multi-Factor Authentication scope — roles that previously controlled User Portal access may need reassignment.

Roles Based Access Release Notes 2025 (Preview) (2026-10-01) BLOG ↗

LATEST STABLE

JumpCloud RBAC gains CRUD-level permissions for Identity Management and Application Management, plus major scope reorganization.

└──▷ HOW TO FIND IT

Configure a custom admin role with granular CRUD vs. read-only separation for user management and application management after the new scope split.

📍 In the JumpCloud Admin Portal, go to Roles › Custom Roles › (select or create a role) › Permissions, then set Identity Management to the desired CRUD level and Application Management to No Access or a specific CRUD level as appropriate for that admin.

Adds CRUD operations and `No Access` permission to the Identity Management and Application Management scopes, enabling more granular admin permissions for users and applications.

Renames `Policy Management` to `Access Management`, `System Management` to `Devices`, and `Authentication Policies` to `Conditional Access Policies and IP Lists`.

Splits Identity Management so it hosts user-related scopes only, with Directory and Groups Management now available as a separate scope.

Moves Directory Management and Group Management scopes under a unified `Directory and Groups Management` permissions category.

Moves SaaS Applications and Asset Management scopes under a new `SaaS and Asset Management` permissions category.

`+3 more`

└──▷ BREAKING ON UPGRADE

- ! The `Policy Management` scope is renamed to `Access Management` — any custom roles or documentation referencing `Policy Management` must be updated.
- ! The `System Management` scope is renamed to `Devices` — any custom roles or documentation referencing `System Management` must be updated.
- ! The `Authentication Policies` scope is renamed to `Conditional Access Policies and IP Lists` — any custom roles or documentation referencing `Authentication Policies` must be updated.
- ! Identity Management no longer bundles all identity-related scopes; Directory and Groups Management is now a separate scope requiring explicit assignment in custom roles.
- ! Directory Management and Group Management scopes have moved under `Directory and Groups Management` — existing role assignments may need to be reviewed.
- ! SaaS Applications and Asset Management scopes have moved under `SaaS and Asset Management` — existing role assignments may need to be reviewed.
- ! RADIUS Authentication has moved under Access Management — existing role assignments scoped to its previous location may need to be updated.
- ! User Portal APIs and operations are now under the Multi-Factor Authentication scope — roles that previously controlled User Portal access may need reassignment.

WAS THIS USEFUL? 👍 👎

Yubico Feedback ⓘ

Hardware root of trust in a digital world: Verifying identity across every layer of th

BLOG ↗ RANK

Yubico manufactures hardware security keys and authentication devices that provide strong two-factor and multi-factor authentication for secure access.

> Yubico shipped YubiKey 5.8 with improvements to passkey support and FIDO2 performance, and expanded its Works with YubiKey catalog with six new integrations.

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

01 **Six new Works with YubiKey integrations** NEW 40

The Works with YubiKey catalog adds six new integrations: T. Rowe Price, ISEC7 Mail, OpenConf, Kakuremi, Epovest, and ArkWarden.

— Names all six integrations but no setup detail given.

Hardware root of trust in a digital wor…

02 **YubiKey 5.8 passkey and FIDO2 improvements** IMPROVED 30

YubiKey 5.8 strengthens passkey support, enhances FIDO2 performance, and improves cryptographic operations.

— No mechanism or concrete metrics given for the improvements.

Hardware root of trust in a digital wor…

> Hardware root of trust in a digital world: Verifying identity across every layer of
>
> BLOG ↗ LATEST STABLE
>
> YubiKey 5.8 strengthens passkey support and FIDO2 performance; six new Works with YubiKey integrations added.
>
> YubiKey 5.8 strengthens passkey support, enhances FIDO2 performance, and improves cryptographic operations.
>
> Adds six new Works with YubiKey catalog integrations: T. Rowe Price, ISEC7 Mail, OpenConf, Kakuremi, Epovest, and ArkWarden.

WAS THIS USEFUL? 👍 👎

iDenfy

Feedback ⓘ snapshot-20261002 NOTES ↗ RANK

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

Adds MCP support for AI assistants with scoped consent, a new Explore products menu, iFrame camera-fail fallback to a new tab, and Shopify age-estimation billing.

NEW INTEGRATION Adds MCP integration so Claude and Claude Code can connect to iDenfy with a consent screen controlling data-area access.

Adds a new Explore products menu surfacing iDenfy products not yet active on the account.

UPDATED COMPLIANCE FRAMEWORK Adds iFrame fallback to a new browser tab when the camera fails during a verification flow.

NEW INTEGRATION Enables Shopify merchants to sign up for and pay for age estimation directly inside Shopify.

UPDATED SUPPORTED PLATFORM Brazil CPF checks — date-of-birth collection

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Sep 30th, 2026 NOTES ↗ RANK

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

> Compyl shipped the ability to cancel integration syncs mid-run, expanded its ServiceNow CMDB connector to new asset types, added endpoint validation to the GoDaddy integration, and made SSO provider changes take effect without a restart.

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

01 **Cancel a running or queued integration sync** NEW 80

Adds the ability to cancel an integration sync directly: a queued sync stops immediately, a running sync stops cleanly at its next saved step, and cancelled syncs are marked as Cancelled so a new sync can start right away. The cancellation record shows the identity of the person who requested it.

When a runaway or misconfigured integration sync is consuming resources, cancel it from the UI so a corrected sync can start immediately.

📍 In Compyl, go to Integrations › Sync History, locate the queued or running sync, and click Cancel Sync. The record will be marked as Cancelled and will show the identity of the person who requested the cancellation.

— Clear mechanism and exact UI steps given. Sep 30th, 2026

02 **ServiceNow CMDB connector adds Business Applications and Servers**

IMPROVED 70

Extends the ServiceNow CMDB connector to support Business Applications, Service Instances, Servers, and their relationships, including custom compliance fields like System ID and PCI/SOX tags.

— Names specific objects and fields but no setup steps. Sep 30th, 2026

03 **Endpoint validation for the GoDaddy integration** NEW 70

Adds endpoint validation to the GoDaddy integration, letting you individually review and toggle Domains, Orders, and Subscriptions during setup.

When setting up the GoDaddy integration, use endpoint validation to selectively enable only the asset types your programme needs.

📍 In Compyl, go to Integrations › GoDaddy › Setup, then use the endpoint validation screen to individually toggle Domains, Orders, and Subscriptions on or off before completing the connection.

— Names toggleable asset types with exact UI navigation. Sep 30th, 2026

THINNER COVERAGE BELOW

04 **Immediate SSO provider changes without restart** IMPROVED

40

Enables SSO provider changes to take effect immediately at the next sign-in, with no application restart required.

— Describes behaviour but no named surface or steps. Sep 30th, 2026

└──▷ ALSO FROM THESE RELEASES

Sep 30th, 2026 NOTES ↗ LATEST STABLE

Compyl adds integration sync cancellation, ServiceNow Business Application/Server support, and GoDaddy endpoint validation.

└──▷ HOW TO FIND IT

When a runaway or misconfigured integration sync is consuming resources, cancel it from the UI so a corrected sync can start immediately.

📍 In Compyl, go to Integrations › Sync History, locate the queued or running sync, and click Cancel Sync. The record will be marked as Cancelled and will show the identity of the person who requested the cancellation.

When setting up the GoDaddy integration, use endpoint validation to selectively enable only the asset types your programme needs.

📍 In Compyl, go to Integrations › GoDaddy › Setup, then use the endpoint validation screen to individually toggle Domains, Orders, and Subscriptions on or off before completing the connection.

Adds the ability to cancel an integration sync directly: a queued sync stops immediately, a running sync stops cleanly at its next saved step, and cancelled syncs are marked as Cancelled so a new sync can start right away.

Enables SSO provider changes to take effect immediately at the next sign-in, with no application restart required.

> Extends the ServiceNow CMDB connector to support Business Applications, Service Instances, Servers, and their relationships, including custom compliance fields like System ID and PCI/SOX tags.
>
> Adds endpoint validation to the GoDaddy integration, letting you individually review and toggle Domains, Orders, and Subscriptions during setup.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.28.979 (2026-10-02) BLOG ↗ RANK

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

> Runlayer's biggest window items are organization-wide LLM budget enforcement that now blocks agent runs and virtual keys, a finer-grained invisible-character threat taxonomy with per-category actions, and a wave of deployment configuration additions (self-hosted Rev decision model, OpenTelemetry sidecars, extraEnv injection) spanning both the Terraform ECS module and the Kubernetes Operator. Also shipped: PagerDuty-triggered agents, AgentGuard Monitor mode, Claude.ai chat imports via Compliance Access Keys, expanded MCP management tools, and a default-on agent browsing policy change.

└──▷ **WHAT SHIPPED · 25 FEATURES** most completely described first ? what's the number?

01 Invisible character detection split into per-category actions

85

BREAKING

Invisible character detection now splits into distinct categories — ANSI colour codes, Hidden characters, Control characters, Isolated zero-width characters, C1 control characters, and Encoded hidden text — each independently configurable as Allow, Alert, Mask, or Block globally and overridden per connector, client, local-tool family, or tool. ANSI colour codes is allowed by default so colourised output from tools like git, lefthook, and pytest stops triggering hidden-character events; three strict categories (Isolated zero-width characters, C1 control characters, Encoded hidden text, covering Unicode tags, invisible operators, unbalanced bidi isolates, and dense zero-width sequences) are opt-in and allowed by default. Concealing sequences, bidi overrides, and zero-width runs are consolidated under Hidden characters, which cannot be switched off. Incidents report the firing category, and terminal-noise events now group under `CONTROL_CHARS` instead of `HIDDEN_CHARS`, so prior dismissals do not transfer.

— Named categories and actions but no config example shown 1.28.944 (2026-10-01)

02 OpenTelemetry sidecar support for ECS and Kubernetes NEW

83

Runlayer adds OpenTelemetry sidecar support to both deployment surfaces. On ECS, the `otel_collector` block (with `export_endpoint`) forwards telemetry to a customer-owned OTLP gRPC endpoint, `runlayer_obs_otel_enabled` toggles it, `FLOW_TRACE_COLLECTION_ENABLED` (defaults to `runlayer_obs_otel_enabled`) decouples backend flow-metric generation from export transport, and `FLOW_TRACE_ADAPTIVE_SAMPLING_ENABLED` enables adaptive sampling on hot-path

flow traces as an independent cost lever. On Kubernetes, `observability.otlp.endpoints` configures the sidecar, and sidecar health metrics (export queue depth, send failures, receiver refusals) are exported only to endpoints named `runlayer-*` using `protocol: http`.

Attach an OTel Collector sidecar to backend and worker tasks to forward telemetry to a customer-owned OTLP gRPC endpoint, while decoupling flow-trace collection from export transport.

```hcl
otel_collector = {
  enabled         = true
  export_endpoint = "grpc://otel.internal.example.com:4317"
}
runlayer_obs_otel_enabled          = true
FLOW_TRACE_COLLECTION_ENABLED      = true
FLOW_TRACE_ADAPTIVE_SAMPLING_ENABLED = true
```

Configure the OTLP endpoint so the OpenTelemetry sidecar exports its health metrics only to the Runlayer-monitored endpoint.

```yaml
observability:
  otlp:
    endpoints:
      - name: runlayer-telemetry
        protocol: http
        url: http://<otlp-collector-host>:4318
```

Expose OTel collector sidecar health metrics so Runlayer can monitor your Kubernetes deployment's telemetry pipeline — endpoint name must start with `runlayer-` and use HTTP.

```yaml
endpoints:
  - name: runlayer-otel-health
    protocol: http
    port: 8888
```

— Named config blocks and env vars for both deployment paths `1.28.943 (2026-10-01)`

03 **Claude.ai chat import via Compliance Access Keys** NEW

Import Claude.ai chats by pasting an Anthropic `sk-ant-api01-…` Compliance Access Key (scoped `read:compliance_user_data`) into Runlayer under Settings → Claude.ai Web (expand `Claude.ai Web` under Web clients), continuing imports after Anthropic's September 22 listing change. Optional Anthropic organization ID filtering (UUID or `org_…` form) limits which organizations' chats are pulled. A beta extends import to Claude Science, Claude for Microsoft 365, and Claude in Chrome — sources Runlayer hooks cannot capture — using the same Compliance Access Key, requires UUID-form org IDs, and must be enabled by the Runlayer account team.

— Exact scope, key format and settings path given `1.28.979 (2026-10-02)`

04 **Expanded MCP management tools** NEW 81

Runlayer's MCP tool surface gains `rl_session_id` on `list_audit_logs` to pull a single session's activity without paging the full audit window; `update_agent_schedule`, `update_agent_trigger`, and `update_agent_webhook` to edit existing schedules, triggers, and webhooks in place; `group_by=actor` on `get_agent_token_usage`, which rolls everyone past the top 30 into an `Other actors` row so per-person spend always totals to the org; `list_agents` filtering by drafts, tags, disabled status, team, and All Workspace with sorting by name, last used, or run count; and a new `list_teams` tool.

Pull audit logs for a single agent session in MCP without paging the full audit window.

```
$ list_audit_logs(rl_session_id="sess_abc123")
```

Break down per-person token spend across the full org, including a catch-all row for users outside the top 30.

```
$ get_agent_token_usage(group_by="actor")
```

— Named tools and params with runnable call examples `1.28.943 (2026-10-01)` `1.28.979 (2026-10-02)`

05 **Terraform ECS module tagging and env-var scoping inputs** NEW 78

The Terraform ECS module gains `additional_tags` (string map, default `{}`) to apply deployment metadata tags to directly managed workload resources, with resource identity tags taking precedence and the `aws-apn-id` tag requiring the format `pc:<Marketplace product code>` without replacing a provider `default_tags` value; `additional_backend_only_env_vars` to scope non-sensitive environment variables exclusively to the backend service task (excluding worker, consumers, and backend migration task), e.g. `ECS_TASK_PROTECTION_ENABLED`; and `sandbox_egress_proxy_keep_nat_gateway` to retain the module-created AgentCore

consumer VPC NAT gateway in `sandbox_egress_mode = 'proxy_only'`, keeping proxy tasks behind a stable NAT Elastic IP.

Keep a stable NAT Elastic IP for sandbox egress in proxy_only mode so third-party allowlists don't break when proxy tasks are replaced.

```hcl
sandbox_egress_mode                  = "proxy_only"
sandbox_egress_proxy_keep_nat_gateway = true
```

Scope a per-service environment variable to the backend task only, without leaking it to the worker or consumers.

```hcl
additional_backend_only_env_vars = {
  ECS_TASK_PROTECTION_ENABLED = "true"
}
```

Tag all Runlayer-managed workload resources with environment and team metadata for cost attribution and audit.

```hcl
additional_tags = {
  environment = "production"
  team        = "platform"
}
```

— Exact input names, defaults and precedence rules given `product docs`

06 **Procedures now require an agent sandbox runtime** BREAKING

70

The `procedures_validator` Terraform input is deprecated and ignored: validation now runs in the agent sandbox and requires `enable_runlayer_agents` with the `AgentCore` runtime, with `terraform plan` warning on the old input before it stops working in the next major module version. Creating or validating a Procedure without an agent runtime now returns `503` instead of succeeding.

— Deprecated input and exact failure code named, migration path clear

`1.28.979 (2026-10-02)`

07 **extraEnv injection for Kubernetes Operator and CRD** NEW

65

The Helm `extraEnv` config key appends environment variables to every backend-image workload, rejecting (rather than silently overriding) any name the operator already sets, while `RunlayerInstance.spec.extraEnv` does the same at the CRD level for backend, worker, migration job, and consumer workloads for settings with no dedicated field.

Inject deployment-wide environment variables into every backend-image workload without risking silent overrides of operator-managed names.

yaml

```yaml
extraEnv:
  - name: MY_CUSTOM_VAR
    value: "my-value"
  - name: ANOTHER_VAR
    value: "another-value"
```

Inject a custom environment variable (e.g. an internal API base URL) into all Runlayer backend workloads when the CRD has no dedicated field for it.

yaml

```yaml
spec:
  extraEnv:
    - name: INTERNAL_API_BASE
      value: "https://internal.example.com"
```

— Two exact config surfaces named with runnable examples `1.28.943 (2026-10-01)`

---

08 **Interactive chat cards and artifact previews, including Slack/Teams**

NEW

65

Agents can request approvals or connector sign-ins through interactive inline decision cards answered directly in the chat, with the conversation continuing from the choice. Agents can also display private artifact previews (generated reports or visuals) visible only to the recipient, which expire unless saved via Save artifact. Both card types and previews are now available in Slack and Microsoft Teams with inline button responses that keep the agent continuing in the same thread.

— Mechanism and platforms named, no config surface `1.28.979 (2026-10-02)`

---

09 **File attachment support in playground and Slack** NEW

65

The agent playground supports attaching images, PDFs, Office documents, and text files via paperclip, paste, or drag and drop. Slack agents additionally support `.zip`, `.mp4`, and `.mov` attachments up to 100 MB, which land in the agent's workspace for unpacking or inspection.

— File types and size limit named, UI entry point given 1.28.979 (2026-10-02)

10 **Agent failure analysis workflow** NEW 60

Agent runs can be marked Pass or Fail with a note from the run drawer; Analyze failures on the Evals tab groups failing notes into categories and lets a user launch an offline eval from any category.

— Named UI entry points give a clear starting place 1.28.979 (2026-10-02)

11 **Skill Usage & edits panel** NEW 60

A new panel shows skill uses, unique users, and edits over 7, 30, or 90 days with edit markers on the usage chart, full edit history with diff links, and top users.

— Named metrics and time windows, UI location implied 1.28.979 (2026-10-02)

12 **Agent browsing enabled by default via network policy** BREAKING 60

Web browsing is now enabled by default for all agents, with access governed entirely by the agent's network access policy rather than a per-agent toggle. The per-agent Web browsing toggle is removed, along with the `enable_web_fetch` API field and Terraform field.

— Removed fields named, migration path implied 1.28.943 (2026-10-01)

THINNER COVERAGE BELOW

13 **Instruction-change tracking in agent chat** NEW 55

When an agent changes its own instructions, the chat shows a diff with an Undo button to restore the previous version.

— Clear mechanism and UI control named 1.28.979 (2026-10-02)

14 **Org LLM budget enforcement extends to agent runs and virtual keys**

IMPROVED 50

The organization usage rule's total budget now counts spend from agent runs and virtual keys, not only devices, so the LLM gateway halts both when the budget is exhausted until an admin raises it.

— Clear mechanism but no exact config surface named 1.28.933 (2026-10-01)

15 **AgentGuard Monitor mode** NEW 50

AgentGuard can now run in Monitor mode, scoring sessions from devices in AI Watch Monitor mode and surfacing detections as alerts without ever blocking tool calls.

— Behaviour explained but no configuration path given 1.28.979 (2026-10-02)

16 **Claude Code subagent token usage tracking** NEW 50

Claude Code subagent token usage now appears on Sessions and the Usage page, attributed to the parent session and priced per model; it requires the updated Runlayer CLI / AI Watch hook.

— Named requirement but no exact command or version 1.28.979 (2026-10-02)

17 **Confusion matrix and labelling-confidence warning for eval scorers**

NEW 50

Comparative scorers on eval runs now show a per-scorer confusion matrix and a labelling-confidence warning, indicating how often each scorer's verdicts match human labels.

— Mechanism explained, no exact navigation given 1.28.943 (2026-10-01)

18 **PagerDuty triggers for agents** NEW 45

Agents can now start when a PagerDuty incident is triggered, acknowledged, reassigned, or has its priority changed.

— Trigger conditions named, no setup steps shown 1.28.979 (2026-10-02)

19 **Session status and failure drill-through in Analytics** IMPROVED

45

Open sessions now show Running while events arrive and Idle with the last-activity time once quiet. The Top erroring agents card in Analytics is clickable, opening the agent's Sessions tab filtered to failed runs, and Sessions gained a new Status filter.

— UI behaviour described, no exact navigation path 1.28.979 (2026-10-02)

20 **Analytics grouping by department and person** NEW 45

ROI artifacts can be grouped by department or individual, showing which outputs an agent run produced and allowing drill-into-spend.

— Feature named, mechanism thin 1.28.979 (2026-10-02)

21 **Duplicate GitHub skill URL detection** NEW 45

Importing a skill from a URL already present in the workspace flags it on the review step and offers Update from source instead of creating a second copy.

— Mechanism named, no exact UI path 1.28.979 (2026-10-02)

22 **SkillGuard engine for skill security scanning** IMPROVED

40

SkillGuard is now the engine behind skill security scanning, with findings that highlight the specific source lines the model flagged.

— Engine named but scope of scanning not detailed 1.28.943 (2026-10-01)

23 **Most/Least toggle on Top active accounts card** NEW 40

The Top active accounts card on Analytics → Overview gains a Most / Least toggle, letting admins surface users with the fewest actions in a selected range.

— Exact UI location named, limited mechanism 1.28.943 (2026-10-01)

24 **AI output analytics enabled for all workspaces** IMPROVED 35

The Output view in Analytics is now enabled for all workspaces without requiring admin activation.

— Thin rollout note, no new mechanism described 1.28.943 (2026-10-01)

25 **Model cost filtering by provider** IMPROVED 25

Detailed usage calculations can now be filtered by model provider, showing the matching row count.

— Thin one-line description, no location or example 1.28.979 (2026-10-02)

└──▷ BREAKING ON UPGRADE

! The `procedures_validator` input to the Terraform ECS module is deprecated and ignored; procedure validation now runs in the agent sandbox and requires `enable_runlayer_agents` with the `AgentCore` runtime. Setting `procedures_validator` causes `terraform plan` to warn; it will stop working in the next major module version.

! If you run both the local Claude Code hook and Anthropic inference hooks, reported Claude Code spend may decrease because turns previously counted once per feed are now deduplicated.

! Creating or validating a Procedure without an agent runtime (`with AgentCore`) now returns `503` instead of succeeding.

! Terminal-noise incidents are now grouped under `CONTROL_CHARS` instead of `HIDDEN_CHARS`; prior dismissals of those incidents do not transfer to the new grouping.

! The per-agent Web browsing toggle is removed; web access is now controlled entirely by the agent's network access policy.

! The `enable_web_fetch` API field and Terraform field are removed.

! OpenTelemetry collector sidecar health metrics are exported ONLY to endpoints named `runlayer-*` using `protocol: http`; endpoints with any other name never receive these metrics.

1.28.979 (2026-10-02) BLOG ↗ LATEST STABLE

Runlayer 1.28.979 adds PagerDuty triggers, AgentGuard Monitor mode, MCP session filtering, inline agent cards in Slack/Teams, and agent failure analysis.

└──▷ TRY IT

Pull audit logs for a single agent session in MCP without paging the full audit window.

```
$ list_audit_logs(rl_session_id="sess_abc123")
```

Break down per-person token spend across the full org, including a catch-all row for users outside the top 30.

```
$ get_agent_token_usage(group_by="actor")
```

Adds `rl_session_id` parameter to `list_audit_logs` MCP tool so agents can retrieve a single session's activity without paging the full audit window.

Adds `update_agent_schedule`, `update_agent_trigger`, and `update_agent_webhook` MCP tools to edit existing schedules, triggers, and webhooks in place without deleting and recreating them.

Adds `group_by=actor` support to `get_agent_token_usage` MCP tool, now returning an `Other actors` row for everyone past the top 30 so per-person spend always totals to the org.

Adds PagerDuty triggers: agents can now start when a PagerDuty incident is triggered, acknowledged, reassigned, or has its priority changed.

Adds AgentGuard in Monitor mode: scores sessions from devices in AI Watch Monitor mode and surfaces detections as alerts without ever blocking tool calls.

+15 more

└──▷ **BREAKING ON UPGRADE**

! The `procedures_validator` input to the Terraform ECS module is deprecated and ignored; procedure validation now runs in the agent sandbox and requires `enable_runlayer_agents` with the `AgentCore` runtime. Setting `procedures_validator` causes `terraform plan` to warn; it will stop working in the next major module version.

! If you run both the local Claude Code hook and Anthropic inference hooks, reported Claude Code spend may decrease because turns previously counted once per feed are now deduplicated.

1.28.944 (2026-10-01) BLOG ↗ STABLE

Runlayer 1.28.944 splits invisible character detection into per-category actions — Allow, Alert, Mask, or Block — with fine-grained overrides.

Splits invisible character detection into distinct categories — `ANSI colour codes`, `Hidden characters`, `Control characters`, `Isolated zero-width characters`, `C1 control characters`,

and `Encoded hidden text` — each with its own action (Allow, Alert, Mask, or Block) configurable globally and overridden per connector, client, local-tool family, or tool.

Adds an `ANSI colour codes` category (allowed by default) so colourised output from tools like git, lefthook, and pytest no longer triggers hidden-character events.

Adds three opt-in strict categories for connectors whose output should never carry invisible text: `Isolated zero-width characters`, `C1 control characters`, and `Encoded hidden text` (covering Unicode tags, invisible operators, unbalanced bidi isolates, and dense zero-width sequences), all allowed by default until explicitly enabled.

Concealing sequences (SGR conceal, background colours, OSC hyperlinks and clipboard writes), bidi overrides, and zero-width runs are consolidated under a `Hidden characters` category that cannot be switched off.

Incidents now report the specific category that fired and group terminal-noise events under `CONTROL_CHARS` instead of `HIDDEN_CHARS`, so prior dismissals of terminal-noise incidents do not transfer.

└──▷ BREAKING ON UPGRADE

! Terminal-noise incidents are now grouped under `CONTROL_CHARS` instead of `HIDDEN_CHARS`; prior dismissals of those incidents do not transfer to the new grouping.

`1.28.943 (2026-10-01)` `BLOG ↗` `STABLE`

Runlayer 1.28.943 adds Kubernetes extraEnv injection, OTel sidecar health export, SkillGuard security scoring, and expanded MCP agent filtering.

└──▷ USE IT

Inject a custom environment variable (e.g. an internal API base URL) into all Runlayer backend workloads when the CRD has no dedicated field for it.

yaml

```yaml
spec:
  extraEnv:
    - name: INTERNAL_API_BASE
      value: "https://internal.example.com"
```

Expose OTel collector sidecar health metrics so Runlayer can monitor your Kubernetes deployment's telemetry pipeline — endpoint name must start with `runlayer-` and use HTTP.

```yaml
endpoints:
  - name: runlayer-otel-health
    protocol: http
    port: 8888
```

Adds `RunlayerInstance.spec.extraEnv` to the Kubernetes CRD to append environment variables to every backend-image workload (backend, worker, migration job, consumers) for settings the CRD has no dedicated field for.

Adds OpenTelemetry collector sidecar health metric export to endpoints named `runlayer-*` (protocol: `http`) covering export queue depth, send failures, and receiver refusals — mirroring ECS telemetry monitoring for Kubernetes deployments.

Adds SkillGuard as the engine for skill security scanning, with findings that highlight the specific source lines the model flagged.

Expands `list_agents` in Runlayer MCP to filter by drafts, tags, disabled status, team, and All Workspace, and to sort by name, last used, or run count.

Adds a `list_teams` tool to Runlayer MCP.

+4 more

└─▷ **BREAKING ON UPGRADE**

- ! The per-agent Web browsing toggle is removed; web access is now controlled entirely by the agent's network access policy.
- ! The `enable_web_fetch` API field and Terraform field are removed.
- ! OpenTelemetry collector sidecar health metrics are exported ONLY to endpoints named `runlayer-*` using `protocol: http`; endpoints with any other name never receive these metrics.

1.28.933 (2026-10-01) BLOG ↗ STABLE

Organization LLM budget enforcement now extends to agent runs and virtual keys, blocking both at the gateway when the limit is spent.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Giskard

Feedback ⓘ v3.0.1 NOTES ↗ RANK

Giskard tests and evaluates LLM applications for performance, bias, and security issues.

> Giskard shipped two small developer-experience improvements: simpler generator configuration and an internal agent code-quality gate, alongside mention of new SOM support for LLM judges.

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

01 **Plain model strings for set_default_generator** IMPROVED 65

`set_default_generator` now accepts model strings directly (e.g. set_default_generator('openai/gpt-4o-mini')), so callers can specify a generator by name without wrapping it in a generator object.

Set the default generator for checks using a plain model string instead of a generator object.

```python
from giskard.checks import set_default_generator

set_default_generator('openai/gpt-4o-mini')
```

– Named function and runnable example given v3.0.1

THINNER COVERAGE BELOW

02 **Agent quality rails for code checks** NEW 40

Adds a per-file lint/typecheck gate plus a `uv` guard, enforcing code quality checks when agent hooks execute.

– Names mechanism but no usage path given giskard-core/v1.0.2

03 **SOM support for LLM judges** NEW 15

Both v3.0.1 and giskard-core v1.0.2 mention adding SOM (Self-Organizing Map) support for LLM judges, though no further mechanism or usage detail is given.

– Only named in summary, no detail or example giskard-core/v1.0.2 v3.0.1

> v3.0.1 NOTES ↗ LATEST STABLE
>
> Giskard v3.0.1 adds SOM support for judges and lets `set_default_generator` accept plain model strings.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.0.1 https://github.com/Giskard-AI/giskard-oss.git
# already have the repo? check out this version:
$ git checkout v3.0.1
```

└──▷ USE IT

Set the default generator for checks using a plain model string instead of a generator object.

```python
from giskard.checks import set_default_generator

set_default_generator('openai/gpt-4o-mini')
```

Allows `set_default_generator` to accept model strings directly, so callers can specify a generator by name without wrapping it in an object.

giskard-core/v1.0.2 NOTES ↗ STABLE

Giskard core v1.0.2 adds SOM (Self-Organizing Map) support for LLM judges and a per-file lint/typecheck agent quality gate.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch giskard-core/v1.0.2
https://github.com/Giskard-AI/giskard-oss.git
# already have the repo? check out this version:
$ git checkout giskard-core/v1.0.2
```

Adds agent quality rails — a per-file lint/typecheck gate plus `uv` guard — enforcing code quality checks on agent hook execution.

WAS THIS USEFUL? 👍 👎

Onyx Security

Feedback ⓘ announced · Oct 1, 2026 PRESS ↗

RANK

Onyx Security is a control plane that discovers, monitors and governs enterprise AI agents inline across browsers, coding assistants, cloud workloads and MCP servers.

> Onyx Security launched a native integration with Anthropic's inference hooks, extending its Secure AI Control Plane to enforce runtime policy directly inside Claude Enterprise without agents or proxies.

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

01 **Anthropic inference hooks integration for Claude Enterprise**

NEW 67

Onyx AI Guard connects to Anthropic Claude Enterprise inference hooks so every prompt, tool-call response, and uploaded-file text is evaluated against enterprise AI security policies before reaching the model, returning a real-time allow-or-deny verdict at the inference boundary with no agent, proxy, or per-device installation required. Coverage spans Claude, Claude Code, Cowork, and connected tools from a single integration, and the same policy layer already governing copilots, coding agents, SaaS AI applications, MCP servers, and AI gateways now governs Claude without a separate Claude-specific policy model.

— Describes mechanism and scope but no config/API names given `launch-20261001-a83126fd`

02 **Agent-aware policy enforcement at tool-call boundary** NEW

60

Onyx runtime policies block sensitive or regulated data exposure, credentials and secrets, intellectual property leakage, prompt injection and jailbreak attempts, and risky content from agent tool results or connector responses. Enforcement goes beyond prompt inspection by evaluating agent identity, permissions, session context, and the action being attempted at the tool-call boundary, allowing Onyx to block destructive operations, mask sensitive information, and restrict tools by policy.

— Explains enforcement logic in detail but no concrete config surface named `launch-20261001-a83126fd`

> announced · Oct 1, 2026 PRESS ↗
>
> Onyx AI Guard integrates with Anthropic inference hooks to enforce enterprise AI security policy natively inside Claude Enterprise.
>
> Connects Onyx AI Guard to Anthropic Claude Enterprise inference hooks, so every prompt, tool-call response, and uploaded-file text is evaluated against enterprise AI

security policies before reaching the model.

Returns a real-time allow-or-deny verdict at the inference boundary with no agent, proxy, or per-device installation required — coverage extends across Claude, Claude Code, Cowork, and connected tools from a single integration.

Enforces Onyx runtime policies to block sensitive or regulated data exposure, credentials and secrets, intellectual property, prompt injection and jailbreak attempts, and risky content from agent tool results or connector responses.

Extends Onyx's Secure AI Control Plane into Claude's native inference path, so the same policy layer governing copilots, coding agents, SaaS AI applications, MCP servers, and AI gateways also governs Claude without a separate Claude-specific policy model.

Supports agent-aware enforcement beyond prompt inspection — evaluates agent identity, permissions, session context, and the action being attempted at the tool-call boundary, enabling Onyx to block destructive operations, mask sensitive information, and restrict tools by policy.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare Feedback ⓘ

WAF - Attack Signature Detection is now available in Early Access (2026-09-07)

BLOG ↗ RANK

Connect, protect, and build everywhere.

Attack Signature Detection (Early Access) adds passive WAF signature matching with review in Security Analytics and scoped Security Rules integration.

NEW DETECTION Adds Attack Signature Detection (Early Access), which evaluates requests against Cloudflare attack signatures and records matches without applying a mitigation action, so you can investigate detected traffic before deciding how to respond.

Surfaces matched signature references, categories, confidence levels, and request outcomes in Security Analytics > Attack Analysis for pre-mitigation review.

Enables use of Attack Signature Detection fields in Security Rules, combinable with request properties such as hostname, path, and HTTP method, to apply scoped mitigation after investigation.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 2, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue