THE CYBER TOOLCHAIN

NO. 097

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 8, 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 76 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 17 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. Cloudflare lets you write CASB detections in Rego, so a data-exposure pattern your SaaS vendor never ed a check for becomes a finding you own. GitHub's Copilot agent

vim ? help

sessions now run sandboxed locally by default at GA. And CVE MCP Server's `triage_cve` collapses NVD, EPSS and KEV into one call with an SSVC decision.

PROTECT

Get one answer on whether a CVE actually matters, instead of three tabs

`triage_cve` fuses NVD severity, EPSS exploit probability and CISA KEV into a composite score with an SSVC v2 decision in a single call, replacing the manual cross-referencing that eats most of a triage shift. Codex Security pairs a CLI and SDK with a cloud GitHub scanner that validates findings before proposing fixes, which is the step that separates a usable queue from a wall of maybes.

CVE MCP Server · OpenAI Codex Security

DETECT

Write your own SaaS exposure checks rather than waiting on the vendor's catalog

Cloudflare CASB custom finding types are authored in Rego, so a misconfiguration specific to how your org uses a SaaS app becomes a first-class finding instead of a spreadsheet exception. Suricata's new `tcp.session` keyword matches on TCP lifecycle phase, giving rule authors a way to scope detections to handshake, established or teardown rather than inferring state from flowbits gymnastics.

Cloudflare · Suricata

PROTECT

Let coding agents run without handing them the host

GitHub's local sandboxing for agentic Copilot sessions is GA, so an agent executing commands during a task is confined rather than operating with the developer's full ambient access — plus a purpose-built model for leaked secret detection feeding push protection. Portkey's Lakera Guard integration screens prompts and responses inline and redacts PII before it reaches either the model or the user, which is the control that keeps a chat transcript from becoming a disclosure incident.

GitHub · Portkey

OFFENSIVE SECURITY

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ snapshot-20261008 NOTES ↗

Horizon3.ai enables organizations to continuously assess the security posture of their enterprise across many attack surfaces.

NEW INTEGRATION CrowdStrike Falcon Next-Gen SIEM integration

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

SpecterOps BloodHound

Feedback ⓘ v9.8.0 NOTES ↗

RANK

Six Degrees of Domain Admin

> BloodHound v9.8.0 adds ZIP bundle ingestion, AD site visualization, alerts retry, and expands saved queries with extension and category metadata alongside several API client and UI additions.

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

01 **Extension and category metadata for saved queries** NEW

70

Saved queries gain `extension_id` and `category` fields in API responses, support extension and category assignment on query upsert, and can now be organized and filtered by extension and category in the UI.

– Names API fields but no endpoint paths given v9.8.0

THINNER COVERAGE BELOW

02 **Alerts retry endpoint** NEW

55

Adds alerts retry capability via the `retry` endpoint.

– Names the endpoint, no path or method shown v9.8.0

03 **ZIP bundle ingestion for extensions** NEW

48

The extension upload API now supports ZIP bundle ingestion.

– Names surface, lacks size limits or usage detail v9.8.0

04 **Collection plan and collector job client APIs** NEW

45

Adds collection plan API client methods and collection plan schedule client methods, plus collector job type client support and collector job secret client APIs.

– Lists surfaces but no method signatures or endpoints v9.8.0

05 **Cross-platform flag on findings** NEW

43

Adds an `is_cross_platform` type field to the Findings Table for cross-platform finding classification.

– Names field but not how classification is used v9.8.0

06 **AD site visualization in graph explorer** NEW — 40

Adds visualisation of Active Directory site data in the graph explorer.

— Describes feature without rendering mechanism v9.8.0

07 **Database management page restricted to admins** BREAKING — 40

Restricts the database management page to admin-only access.

— States change but not prior access model v9.8.0

08 **Minor UI and instrumentation polish** IMPROVED — 38

Distinguishes duplicate node names using distinguished name in the graph UI, updates empty value displays in the entity panel, makes the navigation sidebar auto-expand and auto-contract at the responsive breakpoint, adds icons for the findings table in the shared UI component library, and adds a `dataAnalytics` prop on MultiSelect triggers to support Pendo tracking.

— Several thin UI tweaks across unrelated areas v9.8.0

09 **Custom Claim Key Parameter for client auth** NEW — 32

Adds a Custom Claim Key Parameter for client authentication configuration.

— No config key name or mechanism given v9.8.0

10 **Custom deny ACE count on AD objects** NEW — 30

Adds a `custom deny aces count` property on AD objects.

— Names property only, no usage context v9.8.0

WAS THIS USEFUL?  

APPLICATION SECURITY

◆ SAST / SCA / Code Security

GitHub

Feedback ⓘ Claude Haiku 5.5 in GitHub Copilot (2026-10-07)

BLOG ↗ RANK

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

> GitHub rolled out several Copilot platform upgrades: a purpose-built AI model for leaked secret detection with push protection and new billing controls, GA local sandboxing for agentic Copilot sessions, local model discovery in Copilot CLI, and general availability of Claude Haiku 5.5 as a selectable Copilot model.

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

01 **Purpose-built AI model for secret detection** NEW 90

GitHub introduced a fine-tuned model for leaked secret detection, extending AI-detected secret scanning to the Copilot `/security-review` command (flagging unstructured credentials in supported Copilot CLI and Copilot App sessions) and to a new AI push protection check that scans for unstructured credentials at push time, before they enter repository history (private preview, GitHub Enterprise Cloud and GitHub Teams with GHSP or GHAS). Existing AI-detected Password alert scans are automatically upgraded to the new model at no extra charge for GHSP/GHAS customers, AI-detected secret alerts come to GitHub Enterprise Server 3.23 in public preview, and a new Secret Protection AI Credits SKU appears under AI usage insights with opt-in SKU-level budget controls configurable at `Billing and licensing > Budgets and alerts` (Product: Advanced Security, SKU: Secret Protection AI Credits).

Run an AI-powered security review of active changes — including the new secret-classifier checks — before committing or opening a pull request.

```
$ /security-review
```

Cap spending on AI push protection and secret-classifier checks before opting into the public preview.

> 📍 In the GitHub UI: Billing and licensing › Budgets and alerts › New budget › SKU-level budget › Product: Advanced Security › SKU: Secret Protection AI Credits › Enable 'Stop usage when budget limit is reached'.

— Names exact command, endpoint path, SKU and billing flow

Purpose-built model for leaked secret d…

02 **Local model discovery in Copilot CLI** NEW 80

GitHub Copilot CLI (v1.0.94-0+) adds a `/model` command that discovers models from a running local Ollama instance alongside configured and cloud models, letting users add a discovered model for the current session via `Add and use for this session` or `Add without switching` without restarting the CLI. The `/model` picker surfaces provider connection failures with explanations, offline mode can be enabled explicitly via the `COPILOT_OFFLINE=true` environment variable (kept separate from local model selection), and local model providers can also be configured in the GitHub Copilot app at Settings > Model providers.

Discover and add a local Ollama model to your Copilot CLI session without restarting — useful when you want to test a locally hosted model mid-workflow.

```
$ /model
```

Add a supported local model provider in the GitHub Copilot app so it appears in the model picker alongside cloud models.

> 📍 In the GitHub Copilot app, go to Settings › Model providers and add your local Ollama provider.

— Names exact command, env var and UI path with runnable example

`Discover local models in GitHub Copilot…`

03 **GA local sandboxing for Copilot agents (MXC)** NEW 60

Local sandboxing for GitHub Copilot, powered by Microsoft eXecution Container (MXC), is now generally available in GitHub Copilot CLI, the GitHub Copilot app, and VS Code Agent Host sessions, restricting agentic tool and command access to the filesystem, network, and credentials. Per-policy controls govern which files/directories agent commands can read or modify and whether they can reach the internet, local networks, Git credentials, or GitHub CLI credentials; sandboxing extends to local MCP and language servers where supported, and enterprise-managed settings can require sandboxing with policies developers cannot weaken.

— Explains mechanism and scope but gives no runnable example

`Local sandboxing for GitHub Copilot now…`

THINNER COVERAGE BELOW

04 **Claude Haiku 5.5 in Copilot model picker** NEW 55

Claude Haiku 5.5 is now a selectable model in GitHub Copilot's model picker across Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot cloud agent, GitHub Copilot app, github.com, GitHub Mobile (iOS/Android), JetBrains IDEs, Xcode, and Eclipse. Copilot Enterprise and Business administrators can manage access via the

model policy in Copilot settings, with automatic enablement under the default model enablement policy unless explicitly disabled.

Select Claude Haiku 5.5 for fast, token-efficient coding tasks like quick edits or terminal work in VS Code.

> 📍 In GitHub Copilot, open the model picker and select `Claude Haiku 5.5` from the available models list.

— Names surfaces and admin policy but no deeper mechanism

Claude Haiku 5.5 in GitHub Copilot (202…

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ

October 7, 2026 - Project audit trail and Product Reachability (2026-10-07)

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 added Project-level scoping across the Audit trail, surfacing project identity in both the UI and the public API, alongside a stricter UUID validation rule on audit filter parameters.

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

01 Project object on audit event API responses IMPROVED

90

`GET /public/v0/audit`, `GET /public/v0/audit/{id}`, `GET /public/v0/audit-events`, `GET /public/v0/audit-events/{id}`, and `GET /public/v0/projects/{projectId}/audit` now return a `project` object (with `id` and `name`) in their responses whenever the event is associated with a Project.

Retrieve audit events scoped to a specific project and confirm the new `project` object is present in each result, including its id and name.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/public/v0/projects/<projectId>/audit' | jq '.[].project'
```

— Five endpoints and exact response field named, with a runnable curl example.

October 7, 2026 - Project on audit even…

02 Stricter UUID validation on audit filter parameters BREAKING

80

The `id`, `userId`, and `projectId` filter parameters (e.g. `?userId=`) on `GET /public/v0/audit`, `GET /public/v0/audit-events`, and `GET /public/v0/projects/{projectId}/audit` now accept only hyphenated UUIDs; hyphen-less and brace-wrapped UUIDs that were previously accepted now return HTTP 400 with an error naming the field and value.

— Exact endpoints, params and failure mode named but no migration example given.

October 7, 2026 - Project on audit even…

03 **Project filtering and export in Audit trail** NEW 63

Adds a `Project` column, filter, and CSV field to the Audit trail so events can be sliced and exported by Project, and adds a dedicated Audit trail page on each Project for viewing that Project's events in isolation.

Export the Audit trail filtered to a specific Project to feed into a compliance report or SIEM.

> 📍 In the Finite State Platform, go to Audit trail, apply the Project filter to select your target project, then use the CSV export to download the filtered events.

Review all activity scoped to a single Project without sifting through organisation-wide audit events.

> 📍 In the Finite State Platform, open the target Project, then navigate to its Audit trail page to see only events belonging to that Project.

— Names UI elements and export field but only a UI navigation path, no exact screen name.

October 7, 2026 - Project audit trail a…

└──▷ **BREAKING ON UPGRADE**

! The `id`, `userId`, and `projectId` filter parameters (e.g. `?userId=`) on `GET /public/v0/audit`, `GET /public/v0/audit-events`, and `GET /public/v0/projects/{projectId}/audit` now accept only hyphenated UUIDs; hyphen-less and brace-wrapped UUIDs previously accepted will now return HTTP 400 with an error naming the field and value.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ

Sharper Reliability answers in Well-Architected workloads: NOTES ↗ RANK

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

> Plerion overhauled Reliability pillar coverage in Well-Architected workloads this cycle, adding live-resource deletion/termination protection checks, 39 new live checks spanning backups, encryption, and failover across AWS services, while retiring 33 outdated checks and reclassifying 15 that were miscategorized as security checks.

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

01 Deletion and termination protection checks for live AWS accounts

NEW 80

New checks `PLERION-AWS-1197` (RDS DB instance deletion protection) and `PLERION-AWS-1198` (EC2 instance termination protection) now run as live-resource checks answering REL 9 and REL 11 in Well-Architected workloads. Existing checks `PLERION-AWS-311` (RDS cluster deletion protection) and `PLERION-AWS-614` (load balancer deletion protection) are extended to run against live AWS accounts in addition to infrastructure-as-code templates.

– Names exact check IDs and REL questions but no UI path

Sharper Reliability answers in Well-Arc…

02 Reliability pillar check cleanup removes retired and misclassified checks

70

BREAKING

Removes 33 retired checks and 15 misclassified security checks (such as IMDSv2 on Auto Scaling groups) from the Reliability pillar, raising AutoFill answer quality across 8 of 13 Reliability questions. This is a breaking change: if you use AutoFill, Reliability pillar answers in Well-Architected workloads will change on the next daily run as these checks are removed and the 39 new live checks are added.

– Explains before/after impact but no migration steps

Sharper Reliability answers in Well-Arc…

03 39 new live Reliability checks across AWS services NEW

65

Adds 39 live checks to the Reliability pillar covering backups for RDS, Aurora, DocumentDB, Neptune, ElastiCache, OpenSearch, and Redshift; snapshot encryption and sharing; RDS Multi-AZ; DynamoDB global tables; CloudFront origin failover; VPN tunnel redundancy; and SQS dead-letter queues.

— Lists covered services but omits individual check IDs

Sharper Reliability answers in Well-Arc…

└──▷ **BREAKING ON UPGRADE**

! If you use AutoFill, Reliability pillar answers in Well-Architected workloads will change on the next daily run: 33 retired checks and 15 security checks have been removed, and 39 new live checks have been added.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks

Feedback ⓘ

Google Workspace connector is available in Beta (2026-10-09) BLOG ↗ RANK

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

> Databricks shipped a breaking change to `CREATE OR REPLACE VIEW` that preserves grants, tags, and comments in-place by default starting December 2026, alongside new OAuth custom identity claims for per-user row-level security through a single service principal, default `remote_query` access for compliance security profile workspaces, chart-data extraction in `ai_parse_document`, and new priority and reserved throughput options for DeepSeek V4.1 Flash.

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

01 CREATE OR REPLACE VIEW preserves grants, tags, comments by default

95

BREAKING

Starting December 2, 2026, `CREATE OR REPLACE VIEW` updates views in-place instead of dropping and re-creating them, preserving grants, tags, and comments by default. Use `CREATE OR REPLACE VIEW <catalog>.<schema>.<view> NOCOPY *` to explicitly drop grants, tags, and comments as before (governed tags are still retained); dropping a column with a governed tag still referenced by an ABAC policy without removing the tag first raises a `CANNOT_DROP_TAGGED_COLUMN` error.

Explicitly drop grants, tags, and comments when replacing a view — necessary after December 2, 2026 when the default changes to in-place preservation.

```
$ CREATE OR REPLACE VIEW main.analytics.my_view NOCOPY * AS
SELECT id, name FROM main.raw.source_table;
```

— Exact syntax, error code, and runnable migration example given

CREATE OR REPLACE VIEW will soon preser…

02 OAuth custom identity claims for service principal tokens

NEW

95

Adds a `custom_claim` parameter to OAuth token requests using the `client_credentials` or `urn:ietf:params:oauth:grant-type:token-exchange` grant types, letting applications attach a custom identity claim to a service principal token. The new current_oauth_custom_identity_claim() SQL function reads that claim at query time (returning NULL when absent), and the Genie Conversation API, SQL Statement Execution API, and Databricks JDBC Driver all honor it, enabling per-user data isolation through a single service principal via dynamic view row filters.

Request a service principal token with a custom identity claim so downstream SQL queries can filter rows to only that end user's data.

```
$ curl -X POST https://<workspace-host>/oidc/v1/token \
  -H 'Content-Type: application/x-www-form-urlencoded' \
  -d 'grant_type=client_credentials' \
  -d 'client_id=<sp-client-id>' \
  -d 'client_secret=<sp-client-secret>' \
  -d 'scope=all-apis' \
  -d 'custom_claim=alice@example.com'
```

Create a dynamic view that isolates each end user's rows by comparing a column to the custom identity claim on the token — write the filter to match no rows when the claim is NULL.

```sql
CREATE OR REPLACE VIEW user_orders_view AS
SELECT * FROM orders
WHERE user_email = current_oauth_custom_identity_claim()
  AND current_oauth_custom_identity_claim() IS NOT NULL;
```

— Full parameter, function name, and runnable curl/SQL examples provided

OAuth custom identity claims for servic…

THINNER COVERAGE BELOW

03 **remote_query enabled by default for compliance security profile workspaces**

IMPROVED 50

The `remote_query` table-valued function will be enabled by default for workspaces with the compliance security profile enabled starting November 2026, allowing queries on external databases in their native SQL syntax using Unity Catalog connection credentials.

— Names function and mechanism but gives no usage example

The remote_query function will soon be …

04 **Chart data extraction in ai_parse_document** IMPROVED

50

`ai_parse_document` now extracts chart values, series, legends, and axes as structured JSON, enabling search and agents to use chart data in answers.

— Describes output fields but includes no usage example

ai_parse_document now extracts chart da…

05 **Priority throughput options for DeepSeek V4.1 Flash** NEW

45

DeepSeek V4.1 Flash now supports priority pay-per-token serving to reduce latency for latency-sensitive traffic, and reserved provisioned throughput for dedicated, predictable capacity.

— Names the two serving modes but no setup steps or endpoint

DeepSeek V4.1 Flash supports priority p…

NEW INTEGRATION Google Workspace connector (Beta) in Databricks Lakeflow Connect

└──▷ **BREAKING ON UPGRADE**

! Starting December 2, 2026, `CREATE OR REPLACE VIEW` no longer drops and re-creates the view — it updates in-place, preserving grants, tags, and comments. Workflows that relied on `CREATE OR REPLACE VIEW` to reset grants or tags will need to use `CREATE OR REPLACE VIEW <catalog>.<schema>.<view> NOCOPY *` explicitly, or remove governed tags from columns before dropping them.

> Google Workspace connector is available in Beta (2026-10-09) BLOG ↗ LATEST
>
> STABLE

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Patch Management

N-able N-central

Feedback ⓘ

New in N-zo: Asset Activity, Asset Tags, and Asset Selection (2026-10-06)

BLOG ↗ RANK

N-able N-central is a remote monitoring and management platform enabling IT service providers to manage and support clients' IT infrastructure remotely.

> N-able N-central's headline change this window is the 2026.4 GA release, which brings native dark mode and themes, a PowerShell 7 automation object, a new mobile app, and an Ecoverse navigation section, while Vulnerability Management gained Saved Views with a charts toggle and N-zo's AI assistant picked up asset selection, activity querying, and tag analysis.

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

01 Saved Views and chart toggle in Vulnerability Management NEW

85

The All Detections and By Vulnerability tabs in Vulnerability Management gain Saved Views, letting users save filters, sorting, search, and column settings (order, visibility, width, and pinning) as named views accessible from a Views menu in each tab's toolbar; views are stored per user account and accessible from any device, and a tab can be set to load a chosen view as its default. A `Charts` switch in the Saved Views menu shows or hides the visual filter charts on both tabs, with chart state saved as part of each named view.

Save a daily critical-severity triage layout so you never rebuild filters and column order at the start of each session.

> 📍 In Vulnerability Management, open the All Detections tab. Apply your filters (e.g. Critical severity, sorted by exploitability), arrange columns, then open the Views menu in the toolbar and choose `Save view`. Name it `Critical Servers – Immediate Patching` and mark it as your default view.

Keep a compact, chart-free layout for day-to-day triage and a chart-on layout for weekly risk reviews — without reconfiguring each time.

> 📍 In the Views menu, toggle the Charts switch off, then save the current state as `Grid Only – Daily Triage`. Create a second view with Charts toggled on and save it as `Charts – Weekly Review`. Switch between them from the Views menu.

— Names exact UI elements, menu, and persisted settings

Vulnerability Management: Saved Views a…

02 **Asset Activity querying in N-zo** NEW 75

N-zo can answer plain-language questions about device logins (successful and failed), user logoffs, workstation locks/unlocks, application and system crashes, and unexpected shutdowns, with the underlying data collected on a 30-minute cycle.

Scope a troubleshooting query to the right device when similar names exist across customers, then investigate recent crash activity — all without leaving the chat.

📍 Steps: 1. In the N-zo chat, type '@' and begin typing the device name. 2. Select the correct asset from the search suggestions that appear. 3. Ask: 'Show all system crashes and unexpected shutdowns from the last 7 days.'

— Lists concrete data types and collection cadence

New in N-zo: Asset Activity, Asset Tags…

03 **Asset Tags analysis in N-zo** NEW 75

N-zo supports plain-language queries across static and dynamic tags, tag categories, tag usage, and the devices associated with them, including multi-condition queries that combine tags with other N-zo data, e.g. 'Have any assets tagged VIP had a blue screen in the last 14 days?'

Cross-reference tag-based device groupings with security-relevant activity to quickly surface at-risk endpoints during an incident review.

📍 Steps: 1. In the N-zo chat, ask: 'Have any of my assets tagged VIP had a blue screen in the last 14 days?' 2. Review the returned device list and activity summary. 3. Follow up with: 'Compare reboots and shutdowns on those devices this week.'

— Concrete example query shows real multi-condition mechanism

New in N-zo: Asset Activity, Asset Tags…

04 **Run PowerShell 7 Script object in Automation Manager** NEW 75

Automation Manager adds a `Run PowerShell 7 Script` object, letting automation policies use a modern PowerShell 7 runtime on 64-bit Windows alongside existing PowerShell 5.1 policies.

Use the new `Run PowerShell 7 Script` object in an automation policy to run modern PowerShell on Windows endpoints without touching existing PS 5.1 policies.

📍 In Automation Manager, open or create an automation policy, drag the `Run PowerShell 7 Script` object onto the canvas, paste your PS7-compatible script, and

save. Deploy the policy to 64-bit Windows devices as usual.

— Names exact object and runtime scope with usage steps

Announcing the GA of N-central 2026.4 (…

05 **N-central Mobile app for iOS and Android** NEW 70

A new N-central Mobile app supports device search, Take Control remote sessions, N-able Login with MFA, and Entra ID sign-in, and requires N-central 2026.4 or later.

Enable technicians to launch remote sessions from the field using the new mobile app with MFA-protected N-able Login.

📍 1. Install the N-central Mobile app from the iOS App Store or Google Play. 2. Sign in with N-able Login (MFA or Entra ID). 3. Search for the target device. 4. Tap Take Control to open a remote session.

— Lists concrete capabilities and version requirement

Announcing the GA of N-central 2026.4 (…

06 **Migration from Integrated Backup to Cove Data Protection** NEW 60

Eligible devices can now be migrated from N-able Integrated Backup to Cove Data Protection without reinstalling Backup Manager, with backup monitoring automatically transitioned.

— Names migration path and before/after but no exact steps

Announcing the GA of N-central 2026.4 (…

THINNER COVERAGE BELOW

07 **@ mention asset selector in N-zo chat** NEW 55

N-zo chat adds `@` mention syntax to invoke the Asset Selector, surfacing search suggestions so technicians can confirm the correct endpoint before running a query or action, reducing mistakes when similar device names appear across multiple customers.

Scope a troubleshooting query to the right device when similar names exist across customers, then investigate recent crash activity — all without leaving the chat.

📍 Steps: 1. In the N-zo chat, type '@' and begin typing the device name. 2. Select the correct asset from the search suggestions that appear. 3. Ask: 'Show all system crashes and unexpected shutdowns from the last 7 days.'

— Names exact syntax and purpose but no deeper mechanism

New in N-zo: Asset Activity, Asset Tags…

08 **Dark mode, themes, and contrast options** NEW 55

N-central adds a native Dark Mode and Light Mode toggle, four built-in themes (Modern Classic, Coastal Sky, Lavender Mist, Jade Garden), and seven contrast modes for accessibility and readability.

— Names themes and modes but no navigation path given

Announcing the GA of N-central 2026.4 (…

09 **Script editor enhancements in Automation Manager** IMPROVED 50

The Automation Manager script editor gains syntax highlighting, line numbers, parameter auto-complete, word wrap, zoom controls, and current-line highlighting.

— Lists UI improvements but no usage path or example

Announcing the GA of N-central 2026.4 (…

10 **Ecoverse section in left navigation** NEW 40

A dedicated Ecoverse section is added to the left navigation for discovering and enabling cloud-powered capabilities via N-able Login.

— Bare description of new nav area, no further detail

Announcing the GA of N-central 2026.4 (…

11 **Pre-upgrade readiness checks and progress visibility** IMPROVED 40

N-central now runs pre-upgrade readiness checks before the server goes offline and provides enhanced upgrade progress visibility, aiming to reduce unexpected downtime during upgrades.

— Describes benefit but no mechanism or UI detail

Announcing the GA of N-central 2026.4 (…

WAS THIS USEFUL?  

ITarian

Feedback ⓘ

Introducing our new ITarian AI Native MSP platform. — Log In Today (2026-10-05)

BLOG ↗ RANK

Streamline IT operations with ITarian's all-in-one IT management software. Automate tasks, enhance security, and boost efficiency.

> ITarian introduced a new AI-native MSP platform that automatically diagnoses tickets and recommends remediation playbooks, now available as region-specific instances for US and Europe users.

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

01 **AI-native ticket diagnosis and playbook recommendation** NEW

40

The new AI-native MSP platform reads device history and past incidents when a ticket is opened, automatically builds a diagnosis, and recommends a remediation playbook before a technician engages, with execution occurring within authority.

— Describes behavior but no concrete flags, commands, or UI path

Introducing our new ITarian AI Native M…

02 **Region-specific US and Europe platform instances** NEW

30

The platform offers region-specific instances for the US and Europe, and existing users automatically carry over their data, devices, and history with no migration required.

— Names regions and no-migration claim but no setup steps

Introducing our new ITarian AI Native M…

NEW SUPPORTED PLATFORM Linux device monitoring via Custom Script monitors

> Release Notes: ITarian Platform- October 01, 2026 (2026-09-29) BLOG ↗
>
> STABLE
>
> Device File List now shows script files by default; breadcrumb preserves group context when opening a device.
>
> Device File List tab now shows non-executable files (e.g. PowerShell and command scripts) by default, with the Show with non-executable file(s) filter pre-checked.
>
> Breadcrumb now preserves group context when opening a device, showing the full customer and group navigation path.

└──▷ **BREAKING ON UPGRADE**

! The agent version option `Latest Stable` is renamed to `Public` for Communication Client, Client Security, and Xcitium Client - EXDR; existing version settings are unaffected but any documentation or scripts referencing the label `Latest Stable` will see a different name in the UI.

WAS THIS USEFUL? 👍 👎

Patch My PC

Feedback ⓘ

Patch My PC Publisher Production Version 3.1.1.0 Released (2026-09-30) BLOG ↗

RANK

Simplify IT with Patch My PC's endpoint management software. Automate patching, app updates, reporting, and analytics to save time and boost security.

> Patch My PC's Publisher shipped new content distribution controls, WSUS republishing, UI improvements to package info, and Publisher usability additions like a dependency filter and an assignment-overwrite warning across versions 3.1.1.0 and 3.1.0.462.

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

01 **Dependency filter in Publisher app list** NEW 57

Adds a filter in the Publisher to show only apps that have dependencies, letting admins quickly isolate dependency-linked applications to audit or troubleshoot dependency-driven deployments. Enabled by opening the app list and toggling the 'show apps that have dependencies' filter.

Quickly isolate apps with dependencies in the Publisher to audit or troubleshoot dependency-driven deployments.

> 📍 In the Publisher, open the app list and enable the 'show apps that have dependencies' filter to display only dependency-linked applications.

— Includes exact UI steps to use the filter Patch My PC Publisher Preview Version 3…

02 **Content distribution overrides at multiple levels** NEW 43

Publisher 3.1.1.0 adds content distribution overrides configurable at the Product, Vendor, and All Vendors level.

— Names configuration levels but not the UI path or mechanism

Patch My PC Publisher Production Versio…

03 **Version and Release Date columns in Show Package Info**

IMPROVED 35

Adds Version and Release Date columns to the Show Package Info view.

— Names exact columns and view but no further mechanism

Patch My PC Publisher Production Versio…

04 **Warning for 'All Vendors' assignment overwrites** NEW

30

Warns administrators when All Vendors assignment changes will overwrite product-level assignments.

— Describes the warning but not trigger conditions or UI location

Patch My PC Publisher Preview Version 3…

05 **Republishing for metadata-only WSUS products** NEW 25

Adds republishing support for metadata-only WSUS products.

— Bare statement with no detail on how republishing works

Patch My PC Publisher Production Versio…

WAS THIS USEFUL? 👍 👎

AutoPkg

Feedback ⓘ v3.0.0 NOTES ↗ RANK

Automating packaging and software distribution on macOS.

> AutoPkg 3.0 is a major hardening release: it adds a cached recipe map and a new clear-cache verb for faster day-to-day use, expands the audit command with new checks and CI-friendly flags, and tightens code-signature verification and path confinement across core processors in ways that can break existing recipes.

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

01 **Recipe-map caching for faster resolution** NEW 95

Adds `autopkg generate-recipe-map` command and an on-disk JSON cache at `~/Library/AutoPkg/recipe_map.json` that indexes every recipe and override by identifier and shortname, speeding up recipe resolution on systems with many repos; cache location is overridable with the `RECIPE_MAP_PATH` preference or `AUTOPKG_RECIPE_MAP_PATH` env var, and can be disabled with `DISABLE_RECIPE_MAP` preference or `AUTOPKG_DISABLE_RECIPE_MAP=1`.

Force-rebuild the recipe map in a CI workspace that writes to a custom path instead of ~/Library/AutoPkg.

```
$ AUTOPKG_RECIPE_MAP_PATH=/tmp/ci-workspace/recipe_map.json
autopkg generate-recipe-map
```

— Names command, cache path, and override env vars. v3.0.0

02 **URLDownloader hashing, header testing, and metadata caching** IMPROVED 90

Adds `COMPUTE_HASHES` input variable (default `False`) to URLDownloader to compute MD5, SHA1, and SHA256 of downloaded files as output variables for later processors; adds `HEADERS_TO_TEST` input variable and unifies change-detection/cache-metadata handling between `URLDownloader` and `URLDownloaderPython`; download metadata (ETag, Last-Modified, file size) is now persisted to an `.info.json` sidecar file alongside each download, improving reliability on filesystems/network volumes lacking xattr support.

— Names variables, sidecar file, and mechanism precisely. v3.0.0

03 **Audit CLI output and gating controls** NEW 90

Adds `audit --json` output mode writing a JSON array with one object per audited recipe, each finding tagged with a check name and severity (`error` , `warning` , `info`), ordered most-severe first (`--json` and `--plist` are mutually exclusive); adds `--fail-on error|warning|info` to return exit code 1 when a finding meets or exceeds the given severity, enabling CI gating; adds `--only-check` and `--skip-check` comma-separated flags to filter which checks run across all output modes; adds `--list-checks` to print available check names and severities.

Gate a CI pipeline on any audit finding at warning severity or above, using JSON output for machine parsing.

```
$ autopkg audit --json --fail-on warning MyApp.munki.recipe
```

— Exact flags and exit-code behavior shown with CI example. v3.0.0

04

New audit security checks NEW

85

Expands `audit` to recognize `URLDownloaderPython` as a downloader (flagging missing `CodeSignatureVerifier`) and `AppPkgCreator` / `ChocolateyPackager` as artifact creators (surfacing preceding modification processors); flags path-traversal markers, suspicious privileged install/copy paths, `PkgRootCreator` parent-directory references, unsafe DMG pseudo-paths, insecure `ftp:` URLs alongside `http:` , and `ChocolateyPackager` recipes using weak checksum algorithms (`md5` or `sha1`); flags top-level `Input` values that look like hard-coded credentials based on key names such as `password` , `secret` , `api_key` , `token` , `client_secret` , `access_key` , `private_key` , `credential` , `bearer` and variants, while ignoring variable references like `%CLIENT_SECRET%` and empty values.

— Lists every processor and check name covered. v3.0.0

05

Stricter code signature verification requirements BREAKING

85

`CodeSignatureVerifier` 's `strict_verification` now defaults to `True` (passing `--strict` to codesign), breaking recipes for apps that fail strict verification unless `strict_verification` is set to `False` ; it now requires a `requirement` for app bundles and `expected_authority_names` for installer packages, failing recipes that omit these instead of accepting any valid Developer ID; typoed versions of the `requirement` and `expected_authority_names` keys now error instead of warn; `CodeSignatureVerifier` and `SignToolVerifier` always warn when verification is disabled via `DISABLE_CODE_SIGNATURE_VERIFICATION` , including an up-front warning from `autopkg run` when set globally via environment variable, recipe list, or `-k` / `--key` ; `PkgCreator` 's `scripts` directory files are now included in recipe override trust information, so preinstall/postinstall script changes trigger trust verification failures.

— States exact defaults, keys, and failure conditions. v3.0.0

06 **New clear-cache verb** NEW 80

Adds new `autopkg clear-cache` verb to remove cached files for a named recipe or `all` cached files; supports `--dry-run` to preview deletions and `-v` / `-vv` for varying verbosity.

Preview what cached files would be removed for a recipe before actually deleting them.

```
$ autopkg clear-cache Firefox.munki --dry-run
```

— Exact command and flags given with a runnable example. v3.0.0

07 **Path confinement across download and package processors** 75

BREAKING

URLDownloader strips Content-Disposition filenames to their base name; DMG-relative paths are confined to the mount point; `PkgExtractor` confines `IFPkgFlagDefaultLocation` to the extraction root; `PkgRootCreator` confines `pkgdirs` to the pkgroot; `autopkginstalld` validates package paths before invoking the system installer; `RECIPE_CACHE_DIR` is now confined to `CACHE_DIR`, rejecting a recipe Identifier that would place it outside `CACHE_DIR` via an absolute path or `..` components.

— Names each processor and confinement rule affected. v3.0.0

08 **Automatic copy of external --pkg files into cache** IMPROVED 65

Packages passed via `--pkg` that sit outside the recipe cache are now automatically copied into the recipe's download folder before use; on APFS the copy is instant and uses no extra space.

— Explains behavior and APFS optimization, no flag example. v3.0.0

09 **Stricter boolean input validation in URLDownloader** BREAKING 60

URLDownloader and `URLDownloaderPython` now reject any boolean input variable value other than `true` / `false`, `yes` / `no`, `on` / `off`, `1` / `0`, or empty string, erroring on any other value instead of silently accepting it.

— Names accepted values but gives no migration guidance. v3.0.0

THINNER COVERAGE BELOW

10 Bundled Python 3.11.9 and split dependency files IMPROVED

55

Bundles Python 3.11.9 (security support through October 2027) and splits Python requirements into `requirements.in` (direct dependencies) and `requirements.txt` lockfile.

— Names versions and files, no usage detail given. v3.0.0

---

11 Stricter shared-processor reference resolution BREAKING

50

A fully-qualified shared-processor reference (`RECIPE_ID/Processor`) now fails with an `Unknown processor` error when the named recipe or its processor file cannot be found under the active search directories.

— Names reference format but offers no remediation detail. v3.0.0

└──▷ BREAKING ON UPGRADE

! CodeSignatureVerifier's `strict_verification` now defaults to `True` (passes `--strict` to codesign), breaking recipes for apps that fail strict verification; set `strict_verification` to `False` in affected recipes as a workaround.

! CodeSignatureVerifier now requires a `requirement` for app bundles and `expected_authority_names` for installer packages; recipes that omit these arguments now fail instead of passing with any valid Developer ID.

! Typoed versions of the `requirement` and `expected_authority_names` keys now result in an error instead of a warning.

! URLDownloader and URLDownloaderPython now reject any boolean input variable value other than `true` / `false` , `yes` / `no` , `on` / `off` , `1` / `0` , or empty string — any other value is an error.

! A fully-qualified shared-processor reference (`RECIPE_ID/Processor`) now fails with an `Unknown processor` error when the named recipe or its processor file cannot be found under the active search directories.

! `RECIPE_CACHE_DIR` is now confined to `CACHE_DIR` ; a recipe `Identifier` that would place it outside `CACHE_DIR` (via an absolute path or `..` components) is rejected.

WAS THIS USEFUL?  

Microsoft Intune

Feedback ⓘ

Week of September 28, 2026 (Service release 2609) NOTES ↗

Microsoft Intune is a cloud-based unified endpoint management solution built on Zero Trust. Empower IT to manage, secure and govern devices, apps and endpoints.

Intune 2609 adds GCC High support for Cloud PKI and Remote Help, auto-launches Defender during Android Enterprise enrollment, and raises iOS minimum to 18.

Enables the `Device features tour` Setup Assistant skip key in Automated Device Enrollment profiles to hide the features tour pane on iPhone and iPad during enrollment.

NEW SUPPORTED PLATFORM UPDATED SUPPORTED PLATFORM NEW INTEGRATION Supports assignment filters for Android Enterprise and Android Open Source Project (AOSP) Settings Catalog policies, allowing include/exclude rules based on device properties.

Automatically launches Microsoft Defender for Endpoint during out-of-box setup for corporate-owned Android Enterprise fully managed (COBO) and work-profile (COPE) devices when configured via `Endpoint security > Defender for Endpoint`.

Supports upgrading an existing Android Enterprise Gmail-account connection to a managed Google domain via `Devices > Enrollment > Android > Prerequisites > Managed Google Play`.

Prevents users from bypassing Managed Home Screen (MHS) authentication when accessing protected activities in MAM-integrated apps on Android Enterprise corporate-owned dedicated devices with shared device mode.

+3 more

UPDATED SUPPORTED PLATFORM Minimum supported iOS/iPadOS version raised to 18

└─▷ **BREAKING ON UPGRADE**

! iOS/iPadOS 17 and earlier are no longer supported for standard device-management, Company Portal, and app-protection scenarios — Intune now requires iOS/iPadOS 18 or later.

! The previous single device page in the Intune admin center is no longer available; all admins now see the new single device page when navigating to `Devices > All devices`.

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ October 15, 2026

NOTES ↗ RANK

DigiCert CertCentral is a centralized platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across organizations.

> CertCentral now issues TLS certificates from new G5 RSA/ECC root hierarchies by default and correspondingly trimmed the extra self-signed root from the Apache certificate download bundle.

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

01 Apache certificate bundle drops extra self-signed root

75

BREAKING

The '.crt (best for Apache / Linux)' download produces a `DigiCertCA.crt` file that now contains two certificates (the intermediate CA and the cross-signed root) instead of three, as part of the move to G5 issuance. Scripts or automation that strip the extra self-signed root from this file, or that assume a specific certificate count in `DigiCertCA.crt`, will break after October 15, 2026.

Preview the updated Apache bundle content for a G5-issued order directly in CertCentral before updating automation.

> 📍 In CertCentral, go to Certificates › Orders, select the order number of a TLS certificate issued from a G5 intermediate CA, then in the Download certificate as menu select '.crt (best for Apache / Linux)' and inspect `DigiCertCA.crt` to confirm it contains two certificates.

— Exact file name and before/after cert count, plus a UI check path October 15, 2026

02 Default TLS issuance moves to G5 RSA/ECC roots

BREAKING

73

CertCentral switches default TLS certificate issuance — for new orders, renewals, reissues, and duplicates — to `DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5` for all DV, OV, and EV certificates across the DigiCert, GeoTrust, Thawte, RapidSSL, and Encryption Everywhere brands. Trust stores that do not already include these two roots must add and trust them or newly issued certificates will fail validation, and any policy that hard-codes or pins root/intermediate CA trust must be removed since issuance now moves to the G5 hierarchies.

— Names exact root certs and affected brands/types, but no CLI/API step October 15, 2026

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! Scripts or automation that remove the extra self-signed root from the Apache download file, or that expect a specific number of certificates in `DigiCertCA.crt`, will break after October 15, 2026: the file now contains two certificates (intermediate CA and cross-signed root) instead of three.

! Trust stores that do not already include `DigiCert TLS RSA4096 Root G5` and `DigiCert TLS ECC P384 Root G5` must add and trust those roots, or newly issued certificates will fail validation.

! Hard-coded or pinned root/intermediate CA certificate trust will break as issuance moves to G5 hierarchies — policies requiring hard-coding or pinning must be removed.

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

1Password

Feedback ⓘ 1Password SaaS Manager #2001 (2026-10-07)

BLOG ↗

All the latest releases from the 1Password team

SaaS Manager adds a Budget threshold crossed workflow trigger, multi-year budgets, team app ownership, a new Access tab, and a self-serve AI Consumption view.

Adds a Budget threshold crossed workflow trigger that fires when a budget crosses a configured threshold or reaches 100%, and surfaces application, budget, current spend, budget amount, and budget period details.

Adds an Access tab to application pages, consolidating the existing Access Level table, edit actions, and delete actions into a dedicated view.

Supports budgets that span more than one year.

Enables assigning teams as application owners in the App Catalog, with ownership kept current as team membership changes.

Allows team member searches in App Catalog to include people outside the default active employee and contractor filters.

+2 more

UPDATED INTEGRATION BambooHR Later custom field mapping

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

Feedback ⓘ Tailscale v1.104.1 (2026-10-07) BLOG ↗

RANK

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale v1.104.1 adds client-side netmap caching for faster startup and offline resilience, expands device pinning to macOS, iOS, tvOS and Android, brings Services into the macOS app and Taildrop URL/text sharing, while dropping support for macOS Monterey 12 and iOS 15.

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

01 **DisableReauthWindow macOS policy** NEW 70

Adds the `DisableReauthWindow` system policy on macOS to block the re-authentication window from appearing when a user's device key has expired.

Prevent the re-authentication pop-up from interrupting users on managed macOS fleets when a device key expires.

```text
DisableReauthWindow
```

– Names exact config key but lacks deployment/rollout detail

Tailscale v1.104.1 (2026-10-07)

02 **Dropped macOS Monterey and iOS 15 support** BREAKING 62

Drops support for macOS Monterey 12, requiring macOS Ventura 13 or later as v1.104 is built with Go 1.27, and drops support for iOS 15, requiring iOS 16 or later.

– States exact minimum versions readers must upgrade to Tailscale v1.104.1 (2026-10-07)

03 **Device pinning across platforms** NEW 60

Adds the ability to pin Devices, Exit Nodes, and Services to the top of their respective lists in the macOS app window, with pins also appearing in a dedicated Pinned page, and adds device pinning to iOS, tvOS, and Android.

– Names all pinnable item types and the new Pinned page Tailscale v1.104.1 (2026-10-07)

THINNER COVERAGE BELOW

04 **Disable remote logging via Bug Report menu** NEW 51

Adds the ability to disable remote client logging via the Bug Report menu on iOS and tvOS.

— Names exact UI path but no detail on scope of logging Tailscale v1.104.1 (2026-10-07)

05 **Client-side netmap caching** NEW 48

Adds client-side netmap caching on Linux, Windows, and macOS to improve start-up speed and maintain connectivity when the Tailscale control plane is unreachable.

— Explains purpose and platforms but no mechanism or limits

Tailscale v1.104.1 (2026-10-07)

06 **Taildrop URL and text sharing** IMPROVED 45

Extends Taildrop on macOS, iOS, and Android to support sharing URLs and text selections in addition to files.

— Describes the before/after capability but no limits given

Tailscale v1.104.1 (2026-10-07)

07 **Tailscale Services in macOS app** NEW 37

Adds Tailscale Services and Tailscale PAM Services to the macOS app window.

— Names the feature but not what it lets users do Tailscale v1.104.1 (2026-10-07)

08 **iOS widgets on macOS Tahoe** NEW 30

Adds Tailscale iOS widgets to macOS Tahoe 26 and later.

— Bare statement with OS version but no further detail Tailscale v1.104.1 (2026-10-07)

09 **Fast user switching in iOS share sheet** NEW 25

Adds fast user switching in the iOS share sheet.

— One-line mention with no further mechanism Tailscale v1.104.1 (2026-10-07)

└──▷ **BREAKING ON UPGRADE**

! macOS Monterey 12 is no longer supported; Tailscale v1.104 requires macOS Ventura 13 or later (built with Go 1.27).

! iOS 15 is no longer supported; Tailscale v1.104 requires iOS 16 or later.

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.28.1078 (2026-10-07) BLOG ↗ RANK

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

> Runlayer shipped a new API endpoint and matching MCP tools for agent usage reporting, alongside UI improvements to artifact sorting, failed-run drilldowns, and faster AI Watch Monitor scanning.

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

01 **Agent usage reporting API and MCP tools** NEW 90

Adds `GET /api/v1/agents/metrics/usage` to pull every agent's creator, consumers, sharing, and token and cost usage into a BI tool or data warehouse, plus `get_agent_usage_report` and `list_agent_usage_consumers` MCP tools in Runlayer MCP so AI clients can query agent usage, sharing, and token data directly.

Pull per-agent token and cost usage into a data warehouse or BI dashboard for chargeback or governance reporting.

```
$ curl -H 'Authorization: Bearer <token>'
'https://app.runlayer.com/api/v1/agents/metrics/usage'
```

— Named endpoint, two MCP tools, fields, and a runnable curl example

1.28.1078 (2026-10-07)

THINNER COVERAGE BELOW

02 **Faster AI Watch Monitor mode** IMPROVED 50

Improves AI Watch Monitor mode so tool-call hooks receive their answer before output scanning runs, removing scanner latency from every tool-call round-trip in coding-agent sessions.

— Mechanism described but no flag or setting to invoke it 1.28.1078 (2026-10-07)

03 **Artifact sort controls** NEW 40

Adds sorting to the Artifacts page by newest, oldest, most recently updated, or least recently updated.

— UI sort options named but no API or config 1.28.1078 (2026-10-07)

04 **Failed-runs drilldown on Monitor tab** NEW 40

Adds a clickable failed-run count under Success rate on an agent's Monitor tab that opens Sessions filtered to failed runs.

— Clear UI navigation path but no further mechanism `1.28.1078 (2026-10-07)`

WAS THIS USEFUL?  

GOVERN

◆ AI/LLM Security

ai-safe2-framework

Feedback ⓘ CLI_v1.0.1 NOTES ↗ RANK

Ai-Safe2-Framework provides governance, risk, and compliance controls for securing agentic AI systems and non-human identities.

> AI SAFE² CLI v1.0.1 ships a major hardening pass across skill gating, MCP scanning and scoring, wrap-mode blocking, and project-level gating, while also publishing the package to PyPI for the first time and changing a finding ID that breaks existing automation.

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

01 MCP scoring tightened against self-attestation and hostile schemas

83

BREAKING

`safe2 score mcp` / `safe2 gate mcp URL` no longer let self-attestation (`.well-known/mcp-security.json`) or unvalidated OAuth metadata contribute to the score; hostile schema content now caps the score, withholds the badge, and causes `safe2 gate mcp` to exit 1.

— Names exact file, flags and exit behaviour change CLI_v1.0.1

02 Wrap-mode blocking withholds poisoned output and schema rug-pulls

NEW

81

`safe2 mcp wrap-stdio --block` now withholds poisoned tool output rather than delivering it to the client, and the new `safe2 mcp wrap-proxy --pin-schema` withholds a rug-pulled catalog when the tool schema changes unexpectedly.

— Two exact flags named with their blocking behaviour CLI_v1.0.1

03 MCP scanner catches new sinks and evasion tricks IMPROVED

73

`safe2 scan mcp` / `safe2 gate mcp PATH` now detect Python and TypeScript sinks previously missed, obfuscated execution, and evasion via file renaming (e.g. renaming a vulnerable file to `test_server.py`).

— Specific commands and a concrete evasion example given CLI_v1.0.1

04 Skill trust gating hardened with TG-013 to TG-024 IMPROVED

72

`safe2 gate skill` gains rules TG-013 to TG-024, rejecting hostile skills including paraphrased overrides, Unicode tags, homoglyphs, env exfiltration, persistence, and trigger hijack.

— Names exact rule range and attack types but no mechanism detail CLI_v1.0.1

05 **First PyPI release of ai-safe2** NEW 70

Publishes `ai-safe2==1.0.1` to PyPI for the first time (1.0.0 shipped on GitHub only), installable via `pip install 'ai-safe2[all]==1.0.1'`.

— Exact install command given, runnable immediately CLI_v1.0.1

06 **Project gate caps score on CRITICAL finding** IMPROVED 68

`safe2 gate` (project gate) now has a CRITICAL finding cap the overall score, preventing a project using `shell=True` and eval() on LLM output from clearing Tier 2.

— Concrete example given but limited to one command CLI_v1.0.1

07 **Finding ID change and score invalidation** BREAKING 65

DEP findings now cite `MCP-4` instead of `MCP-3`, matching the v3.1 control map, so any automation matching on the old finding ID will stop matching; scores and gate results from 1.0.0 or 0.9.x are no longer reliable and gates must be re-run under 1.0.1 scoring logic.

— Names old/new finding IDs but no migration tooling given CLI_v1.0.1

└──▷ **BREAKING ON UPGRADE**

! DEP findings now cite `MCP-4` instead of `MCP-3`, matching the v3.1 control map — any automation that matches on the old finding ID will stop matching.

! Scores and gate results from 1.0.0 or 0.9.x are no longer reliable; gates must be re-run to get valid results under 1.0.1 scoring logic.

WAS THIS USEFUL? 👍 👎

Aurascape

Feedback ⓘ

September New Innovations Recap: More Precise Protection for AI Work (2026-10-05)

BLOG ↗

Aurascape is the AI-native security platform that empowers enterprises with real-time visibility, intent-based controls, and protection for every AI interaction—securing innovation without slowing it down.

Aurascape September 2026 adds prompt masking in Protection policy, classification management APIs, SSO audit logging, and richer on-demand reporting.

Adds sensitive data masking as a mode of the `redact` action in Protection policy, replacing sensitive values in prompts before they reach the LLM and restoring real values in the response the user reads.

New APIs automate management of data patterns, data categories, and data combination lists for data classification configurations, and can retrieve the supporting objects those configurations reference.

Adds administrator SSO activity to Audit Logs, recording successful and failed login attempts with timestamp, status, and source IP address.

On-demand reports now include a `Unique User Counts by Application` metric and a `Top Protection Policy Hits` bar chart alongside the existing application breakdown.

Conversation records for coding sessions now separate `system prompt`, `user prompt`, and `endpoint details` into individual fields.

+2 more

UPDATED INTEGRATION Granular model-level access policy for Claude, Anthropic API, ChatGPT, OpenAI API, Google Vertex AI API, and Amazon Bedrock

UPDATED INTEGRATION Expanded MCP conversation record detail via Zero-Bypass MCP Gateway

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe Feedback on this issue