THE CYBER TOOLCHAIN

NO. 104

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 12, 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 41 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 16 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. Pomerium records SSH sessions and reaches hosts behind NAT without a bastion; NodeZero pivots through … it finds mid-run so an external pentest reaches the second hop;

vim ? help

Cloudflare learns API schemas on demand, so you review proposed per-operation rules instead of waiting on background sampling.

PROTECT

Broker and record privileged SSH instead of handing out keys to a bastion

v0.32.0 adds SSH session recording and reverse tunneling, so admin sessions produce replayable evidence and reach hosts that cannot accept inbound connections. That removes the usual pattern of a jump box with standing credentials and no record of what was typed.

Pomerium

IDENTIFY

Reach the second hop on an external pentest without re-scoping the run

Network pivoting lets an external operation continue through hosts discovered mid-run, rather than stopping at the perimeter and filing the internal reach as an assumption. Rate limits on Password Spray and post-exploitation also cut the account lockouts that usually force testers to throttle by hand.

Horizon3 NodeZero

PROTECT

Get a schema on an API you never documented, and approve it operation by operation

On-demand schema learning builds a candidate schema when you ask for it and surfaces live per-operation review, so positive-security enforcement on a shadow or legacy endpoint no longer waits for enough background traffic to accumulate.

Cloudflare

DETECT

Parse untrusted Sigma rules without a YAML alias bomb taking the process down

v2.0.1 hardens YAML processing against alias expansion attacks — the billion-laughs class, where a tiny file inflates into gigabytes on load. Relevant to anyone ingesting community rule repos or tenant-supplied detections into a shared pipeline.

pysigma

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Mend.io Mend SAST

Feedback ⓘ changelog · Oct 12, 2026

CHANGELOG ↗ RANK

Mend SAST is a static application security testing service that scans source code for vulnerabilities and helps teams prioritize remediation.

> Mend SAST 26.9.2 adds configuration-driven entry point detection for contract-first applications and sharpens Python engine accuracy across eleven vulnerability classes.

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

01 **YAML/JSON entry point detection for contract-first apps** NEW

60

Mend SAST now supports YAML and JSON configuration files for entry point detection, automatically identifying API endpoints and mapping them to implementation classes. All generated entry point parameters are automatically marked as tainted, eliminating the need for manual custom framework configuration for contract-first applications.

— Names config formats and mechanism but no file path or flag changelog-20261012-b0e147dc

THINNER COVERAGE BELOW

02 **Python engine accuracy improvements across 11 CWEs** IMPROVED

55

Python engine detection accuracy was improved across 11 CWEs: `CWE-022` Path Traversal, `CWE-078` OS Command Injection, `CWE-079` XSS, `CWE-089` SQL Injection, `CWE-094` Code Injection, `CWE-209` Error Message Info Exposure, `CWE-295` Improper Cert Validation, `CWE-532` Sensitive Info in Log, `CWE-601` Open Redirect, `CWE-918` SSRF, and `CWE-943` NoSQL Injection.

— Lists exact CWEs but no detail on how accuracy changed changelog-20261012-b0e147dc

WAS THIS USEFUL? 👍 👎

Mend.io Mend SCA

Feedback ⓘ changelog · Oct 12, 2026

CHANGELOG ↗ RANK

Mend sca is a traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend SCA shipped workflow-based email notifications for findings and a dedicated zero-day alert channel, giving teams faster, more targeted awareness of new vulnerabilities outside scheduled scans.

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

01 **Configurable Send Email workflow action** NEW 60

A new Send Email workflow action notifies users of new findings — such as newly published CVEs — independent of scans, with configurable delivery frequencies. Workflow creators can directly select specific users or groups as recipients for these scan summary emails.

Get immediate notification when a newly published CVE affects your monitored packages, without waiting for a scheduled scan.

📍 In the Workflows editor, add a new action step, select Send Email, choose target users or groups as recipients, configure the delivery frequency, and save the workflow.

– Names the action and setup steps but no config file/API changelog-20261012-e411f675

THINNER COVERAGE BELOW

02 **Dedicated zero-day threat alerts** NEW 30

Adds dedicated zero-day alerts to provide immediate notification when the organization is vulnerable to emerging threats.

– States the capability but no mechanism or setup detail changelog-20261012-e411f675

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Levo

Feedback ⓘ changelog · Oct 12, 2026 CHANGELOG ↗

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

Levo adds in-product license status with expiry warnings, bulk API endpoint merging, audit-format Vulnerabilities PDF, and a create–validate–approve lifecycle for AI-authored tests.

Adds license status display in-product with expiry warnings, a grace period, and clean end-of-access enforcement.

Adds bulk-merge for duplicate API endpoints in the API Surface view.

Adds an audit-format output option for the Vulnerabilities PDF, suitable for handing to auditors or customer security teams.

Adds a full create–validate–approve authoring lifecycle for custom AI-authored tests, binding every run to the exact human-reviewed code.

Enables web application scans to finish testing within budget while recording every distinct finding.

UPDATED DETECTION TLS findings from the Satellite restored

└──▷ BREAKING ON UPGRADE

! License fields in the Workspace and Organization APIs have changed — any integration that reads those fields must be updated before upgrading to v1.2.135.7.

! The API Surface export format has changed — any integration that reads it must be updated before upgrading to v1.2.135.4.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

juice-shop

Feedback ⓘ v20.3.0 NOTES ↗ RANK

OWASP Juice Shop is an intentionally insecure web application that developers and security professionals use to practice finding and exploiting common vulnerabilities.

> Juice Shop v20.3.0 adds configurable alternate product image formats with AVIF support and introduces Farsi and Tamil language support.

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

01 Configurable alternate product images with AVIF support NEW

83

Adds an optional `alternateImages` configuration property that lets instances define alternative product image formats and densities, and uses it to serve responsive AVIF product images across 1x, 2x, and 3x display densities.

Serve modern AVIF images at multiple resolutions for a custom product by declaring `alternateImages` in your instance configuration.

yaml

```yaml
alternateImages:
  - format: avif
    density: 1x
    path: /assets/public/images/products/my-product.avif
  - format: avif
    density: 2x
    path: /assets/public/images/products/my-product@2x.avif
  - format: avif
    density: 3x
    path: /assets/public/images/products/my-product@3x.avif
```

— Exact config key and runnable YAML example provided v20.3.0

THINNER COVERAGE BELOW

02 Farsi and Tamil language support NEW

18

Adds support for the Farsi and Tamil languages.

— Only named languages, no mechanism or usage detail given v20.3.0

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Mend.io Mend Container

Feedback ⓘ

changelog · Oct 12, 2026 CHANGELOG ↗ RANK

Mend container is a traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend Container doubled its image scan size limit, allowing larger container images to be scanned across all registry integrations.

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

01 **Increased image scan size limit to 4 GB** IMPROVED 57

The image scan size limit was raised from 2 GB to 4 GB, enabling pulling, extraction, and scanning of larger container images across all registry integrations including ECR. Images still exceeding the new limit return a clear error message.

— Clear before/after limit named but no config/flag to act on changelog-20261012-67a80fd9

WAS THIS USEFUL? 👍 👎

Docker CLI

Feedback ⓘ changelog · Oct 12, 2026 CHANGELOG ↗

RANK

Docker cli is a docker Documentation is the official Docker library of resources, manuals, and guides to help you containerize applications.

> Docker 29.9.0 introduces a daemon flag to control lazy image pulling for the containerd image store and confirms/extends IPv6 support for Swarm service ports and bridge-network neighbour discovery.

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

01 **Lazy-pull control for containerd image store** NEW 85

Adds the `lazy-pull` daemon feature flag to the containerd image store, controlling whether `docker pull` skips downloading layer content already provided by the snapshotter. It is enabled by default for the `nydus`, `overlaybd`, `soci`, and `stargz` snapshotters.

— Names the flag and snapshotters but not its config file location

changelog-20261012-2f4f9c31

THINNER COVERAGE BELOW

02 **IPv6 support for Swarm service ports** IMPROVED 45

Published Swarm-service ports are now accessible at the host's IPv6 addresses when the userland proxy is enabled, confirmed as a tested and supported feature as of v29.9.0.

— Describes behaviour but gives no config or command to act on

changelog-20261012-2f4f9c31

03 **IPv6 Neighbour Discovery on isolated bridge networks** IMPROVED 45

Allows IPv6 Neighbour Discovery between containers on a bridge network with inter-container communication disabled, matching existing IPv4 behaviour.

— Explains the behaviour change but names no flag or setting changelog-20261012-2f4f9c31

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Detection Engineering

pysigma

Feedback ⓘ v2.0.1 NOTES ↗ RANK

pySigma is a Python library for parsing, converting, and manipulating Sigma detection rules into queries for various SIEM and log analysis platforms.

> pySigma v2.0.1 hardens YAML processing against alias-bomb attacks and relocates a core policy module, requiring an import path update for downstream code.

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

01 **`default_policy` moved to `sigma.policy`** BREAKING 60

`default_policy` has moved from its previous location to `sigma.policy` to re-establish the sigma namespace package; any code importing it from the old path will break.

— Names exact new import path and breaking impact. v2.0.1

THINNER COVERAGE BELOW

02 **Alias expansion checks against YAML alias bombs** NEW 33

Adds alias expansion checks to prevent YAML alias bombs in processing pipelines and validation.

— Explains what it does but not the mechanism or limits. v2.0.1

└──▷ **BREAKING ON UPGRADE**

! `default_policy` has moved from its previous location to `sigma.policy`; any code importing it from the old path will break.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Endpoint Management (UEM/RMM)

Iru

Feedback ⓘ changelog · Oct 12, 2026 CHANGELOG ↗ RANK

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

> Iru added a post-Liftoff processing option for Mac Library Items, letting admins control whether items run during or after device Liftoff.

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

01 **Post-Liftoff scheduling for Mac Library Items** NEW

50

Custom App, Custom Printer, and Custom Script Library Items for Mac now support a post-Liftoff processing option, letting admins choose whether these items run during or after Liftoff.

— Names the exact item types and timing option but no config path

changelog-20261012-9a55f2e3

WAS THIS USEFUL? 👍 👎

◆ Security Validation (BAS & Automated Pentest)

Horizon3 NodeZero

Feedback ⓘ changelog · Oct 12, 2026

CHANGELOG ↗ RANK

Horizon3 nodezero is a horizon3.ai enables organizations to continuously assess the security posture of their enterprise across many attack surfaces.

> NodeZero's latest release adds network pivoting for external pentests to reach hosts discovered mid-operation, lockout-reducing rate limits for Password Spray and post-exploitation, a new CrowdStrike Falcon Next-Gen SIEM integration, and incremental WebApp pentest and ServiceNow VR integration improvements.

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

01 **Network pivoting for external pentests** NEW 65

Adds network pivoting for external tests, enabling NodeZero to reach additional hosts discovered during an operation beyond the originally configured scope. Enabled in the NodeZero console by creating or editing an External pentest operation and turning on Network Pivoting under Attack Configuration.

Enable network pivoting on an external pentest to automatically enumerate and attack hosts discovered during the operation, beyond the originally scoped targets.

> 📍 In the NodeZero console, create or edit an External pentest operation, then under Attack Configuration enable Network Pivoting to allow the operation to pivot to additional hosts discovered at runtime.

— Names exact UI toggle and scope behavior, no API or limit detail

changelog-20261012-451f4bf9

02 **Rate-limiting controls for Password Spray and post-exploitation**

NEW 65

Adds new on/off and rate-limiting options for Password Spray, credential pivoting, and post-exploitation attack configuration to minimize account lockouts. Set via the on/off toggle and rate-limiting options under Password Spray in Attack Configuration before launching an operation.

Reduce account lockout risk during password spray campaigns by enabling rate-limiting controls in attack configuration before launching an operation.

📍 In the NodeZero console, open Attack Configuration for your operation and use the on/off toggle and rate-limiting options under `Password Spray` to set a safe request cadence before launching.

— Names the exact toggle and three covered attack areas changelog-20261012-451f4bf9

THINNER COVERAGE BELOW

03 **WebApp pentest login field, secrets, and feedback improvements**

IMPROVED 45

Adds custom login field configuration and support for longer secrets in WebApp pentest operations, and adds clearer feedback in WebApp pentest when a Runner or read-only state limits an action.

— Names three WebApp pentest changes but no UI path or mechanism

changelog-20261012-451f4bf9

04 **Integration Settings search, filtering, and Activity Log**

IMPROVED 35

Expands Integration Settings with broader search, filtering, and Activity Log details.

— Names three settings capabilities but no exact fields or UI path

changelog-20261012-451f4bf9

05 **ServiceNow VR integration matching fields** IMPROVED

30

Expands ServiceNow VR integration with additional matching fields for more-reliable interoperation.

— Names the integration but not which fields were added changelog-20261012-451f4bf9

06 **CrowdStrike Falcon Next-Gen SIEM integration** NEW 25

Adds a new integration with CrowdStrike Falcon Next-Gen SIEM.

— Bare announcement with no setup or data-flow detail changelog-20261012-451f4bf9

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

OSV-Scanner

Feedback ⓘ v2.7.0 NOTES ↗

OSV-Scanner is an open-source tool from Google that scans project dependencies against the OSV vulnerability database to identify known security vulnerabilities.

Adds `purl` to JSON output, shorter `--x-` flag aliases, CycloneDX severity/rating output, and `--max-file-size` enforcement for container image contents.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.7.0 https://github.com/google/osv-scanner.git
# already have the repo? check out this version:
$ git checkout v2.7.0
```

UPDATED SUPPORTED PLATFORM Adds `purl` (package URL) fields to JSON output when available.

Adds shorter `--x-` flag aliases as replacements for the deprecated `--experimental-` flags.

Includes vulnerability severity and rating in CycloneDX output when available.

Applies `--max-file-size` limit to container image contents and tags container image layers with Buildpacks lifecycle metadata.

Adds best-effort Maven property pre-interpolation for dependency versions and in-memory caching for Maven resolution.

NEW SUPPORTED PLATFORM Zig build.zig.zon extractor

NEW INTEGRATION Bitnami container SPDX SBOM scanning

NEW SUPPORTED PLATFORM MinimOS ecosystem mapping for APK extractor

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Pomerium

Feedback ⓘ changelog · Oct 12, 2026 CHANGELOG ↗ RANK

Pomerium is a protect private applications and services with identity-aware access that evaluates every request from users, services, and AI agents.

> Pomerium's v0.32.0 release focuses on cross-origin session handling, a faster Databroker query API, a hosted identity provider, MCP OAuth improvements, new control-plane debugging tools, and major SSH access upgrades including session recording and reverse tunneling.

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

01 **Cross-origin session cookies via additional hosts parameter** NEW 90

Adds the `pomerium_additional_hosts` query parameter at the route level to declare cross-origin authentication dependencies, so Pomerium sets session cookies across up to five domains during a single OAuth flow.

Declare that a frontend route depends on a backend API route so Pomerium sets session cookies on both domains during a single login — eliminating cross-origin fetch failures.

```
$ https://app.example.com/dashboard?
pomerium_additional_hosts=api.example.com
```

— Named config parameter with a runnable usage example changelog-20261012-68827fa7

02 **Native SSH Access enhancements** IMPROVED 75

Native SSH Access gains reverse tunneling, jump-host mode, an interactive routes portal, and per-key session revocation, and switches from the Device Code flow to the standard OAuth Authorization Code flow.

— Lists all named additions and the auth flow change changelog-20261012-68827fa7

03 **Databroker GetOptions query API** NEW 70

Adds a new `GetOptions` API on the Databroker with field-level indexing, additional filter operators, and `order by` support, replacing full-record scans with targeted lookups.

— Named API and mechanism but no example usage changelog-20261012-68827fa7

04 **SSH session recording and playback in Enterprise** NEW 65

Adds SSH session recording and playback for privileged SSH connections in Pomerium Enterprise, with full keystroke and terminal output capture reviewable from the Console.

— Describes mechanism and where to review, no exact command `changelog-20261012-68827fa7`

---

05 **MCP OAuth refresh tokens and dynamic client registration**

IMPROVED

60

Improves MCP support by adding OAuth refresh token issuance and honoring for MCP sessions, keeping long-running agents connected without repeated interactive logins, and by adding dynamic client registration via client ID metadata documents so MCP clients can register with Pomerium automatically.

— Names mechanisms but no concrete commands or endpoints `changelog-20261012-68827fa7`

THINNER COVERAGE BELOW

06 **Control-plane debugging and observability tooling** NEW

55

Adds a control-plane debug server with gRPC `channelz` support for inspecting live gRPC channels and connections, a Databroker browser debug view for browsing Databroker records directly to diagnose session, policy, or sync issues, and streaming health updates over gRPC so dashboards and the Enterprise Console can reflect component health in real time.

— Names three debug surfaces but no exact endpoints or paths `changelog-20261012-68827fa7`

07 **Hosted identity provider type** NEW

45

Adds a hosted identity provider type so Pomerium can manage its own OIDC provider, removing the requirement to stand up an external IdP before getting started.

— Explains benefit but no setup detail or config shown `changelog-20261012-68827fa7`

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

LoginRadius

Feedback ⓘ changelog · Oct 12, 2026 CHANGELOG ↗

RANK

LoginRadius: Privacy-compliant identity and access management platform. Offer Authentication and Authorization for your web and mobile apps in just a few clicks, with: SSO, MFA, social logins and more.

> LoginRadius's V3JS library gained per-component visibility controls and expanded Risk-Based Authentication Email OTP handling, including a new localization and error-mapping layer and a behavior change for suspicious sign-ins.

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

01 **Component-level config in V3JS** NEW 75

Adds `componentConfig` to the V3JS library to dynamically show or hide specific authentication components per flow, letting developers e.g. hide social login on a step-up auth screen without duplicating page templates.

Show or hide specific auth components per flow — e.g. hide social login on a step-up auth screen without duplicating page templates.

```json
componentConfig: {
  socialLogin: { show: false },
  emailOtp: { show: true }
}
```

— Named config key with a runnable example changelog-20261012-69bd079a

THINNER COVERAGE BELOW

02 **RBA Email OTP localization and error mapping** NEW 56

Adds `rba-email-otp` localization keys for Risk-Based Authentication Email OTP screens, and adds dedicated error-code mapping for Hosted Pages Email OTP verification.

Add translated strings for the RBA Email OTP challenge screens to support a localized deployment.

```json
"rba-email-otp": {
  "title": "Verify Your Identity",
  "instruction": "Enter the code sent to your email.",
  "submit": "Verify"
}
```

— Names localization key with example; error mapping left unspecified

changelog-20261012-69bd079a

03 **Mandatory email verification skip flag extended** IMPROVED

48

Extends `preventVerificationEmail` support to Mandatory Email Verification flows in the Registration Flow.

— Names flag and flow but no example changelog-20261012-69bd079a

04 **RBA Email OTP challenge replaces sign-in error** IMPROVED

40

Suspicious sign-ins now trigger an Email OTP challenge via Risk-Based Authentication rather than returning an error, changing the behavior users and integrators see on flagged logins.

— Describes behavior change but no config surface given changelog-20261012-69bd079a

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Third-Party & Vendor Risk (TPRM)

Risk Ledger

Feedback ⓘ changelog · Oct 12, 2026 CHANGELOG ↗

RANK

Risk Ledger is a cutting-edge platform for real-time supply chain security, which enables organisations to Defend-as-One because every link matters in cybersecurity.

> Risk Ledger adds centralized document storage with expiry reminders inside supplier profiles, replaces its community signals feature with a permissioned Community Review Sharing model, and gives users finer control over network map density and dormant-profile handling.

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

01 Document storage and expiry tracking in supplier profiles

NEW

70

A dedicated document storage space is added within any supplier profile, visible only to the host organisation, so desk-research files and supplier-sent packs sit alongside other supplier records. Uploaded documents now carry expiry-date tracking that triggers reminders 30 days before and on the day a document expires.

When a supplier is slow to engage, upload their documents to the supplier profile so evidence stays centrally accessible and you receive an expiry reminder before certificates lapse.

> 📍 In the supplier profile, go to the new Documents section › Upload document › set an expiry date to receive reminders 30 days before and on the day it expires.

— Names the module and exact reminder timing, with a usage walkthrough.

changelog-20261012-23ac9f3a

02 Community Review Sharing permissions

BREAKING

70

A new permission framework for Community Review Sharing lets federated organisations and selected communities control which data points are shared and with which peers, and lets a central team share reviews and review signals with peers so it can review suppliers on behalf of members. This replaces the previous community signals feature, which has been retired.

When a central security team wants to share completed supplier reviews with member organisations in a community, without exposing raw evidence or supplier answers.

> 📍 In the Community Review Sharing permission framework, select the community › choose which data points to share with peers › set access to approve

on request so peers can see that a review exists and receive only approved data points.

— Explains the replaced feature and the new sharing mechanism with steps.

changelog-20261012-23ac9f3a

---

03 **Network map density and display controls** IMPROVED 65

Network map display controls now let users choose how many nth parties to display, hide the wider Risk Ledger network, and toggle node animations on or off to keep large maps fast and focused.

When your network map has grown large and is slowing down, reduce clutter by limiting nth-party depth and disabling animations.

📍 In the network map view, open the display controls › set Adjust network map density to your desired nth-party depth › disable Show animations › enable Hide the wider Risk Ledger network to focus on your own supply chain.

— Concrete controls named with a step-by-step example. changelog-20261012-23ac9f3a

THINNER COVERAGE BELOW

04 **Suggest a contact for dormant supplier profiles** NEW 55

A Suggest a contact option appears on the supplier Overview page under Last sign in for suppliers inactive for a year or more, letting users nominate someone to take over a dormant profile.

— UI location named but no example or further mechanism given.

changelog-20261012-23ac9f3a

└──▷ **BREAKING ON UPGRADE**

! The previous community signals feature has been retired and replaced by Community Review Sharing.

WAS THIS USEFUL? 👍 👎

Conveyor

Feedback ⓘ changelog · Oct 12, 2026 CHANGELOG ↗ RANK

Conveyor delivers AI for high-stakes customer requests like security reviews, vendor assessments, and requests for proposal (RFP).

> Conveyor's ConveyorAI Cmd+K search now shows its work, streaming live reasoning steps and offering a post-hoc audit panel for how answers were assembled.

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

01 Live reasoning display in Cmd+K agentic search NEW

65

ConveyorAI Cmd+K agentic search now shows live step-by-step reasoning as it works, displaying matched sources, narrowing logic, and answer drafting in real time.

Audit an AI-generated answer to a vendor security question — useful when validating that the agent pulled from the right knowledge sources before sending a response.

> 📍 1. Press Cmd+K to open ConveyorAI search. 2. Type your security question and select the agentic search mode. 3. Watch live steps appear: knowledge search, source matching, narrowing, and answer drafting. 4. Click How this answer was put together to open the full reasoning audit panel.

— Names the surface and flow but no technical mechanism detail.

changelog-20261012-2c72fcf5

THINNER COVERAGE BELOW

02 Answer audit panel in Cmd+K NEW

58

A new How this answer was put together audit panel in Cmd+K lets users review after the fact which knowledge the agent used and how it reached its answer.

Audit an AI-generated answer to a vendor security question — useful when validating that the agent pulled from the right knowledge sources before sending a response.

> 📍 1. Press Cmd+K to open ConveyorAI search. 2. Type your security question and select the agentic search mode. 3. Watch live steps appear: knowledge search, source matching, narrowing, and answer drafting. 4. Click How this answer was put together to open the full reasoning audit panel.

— Names exact panel label and purpose, clear UI action to try it.

changelog-20261012-2c72fcf5

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Security Posture & Discovery

Mend.io Mend AI

Feedback ⓘ changelog · Oct 12, 2026

CHANGELOG ↗

Mend ai is a traditional application security misses AI risk. Standalone AI security tools miss your code.

UPDATED DETECTION Mend AI Runtime Protection | Guardrails multi-language support

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ

API Shield - Run schema learning on demand (2026-10-12) BLOG ↗ RANK

Cloudflare is a connectivity cloud platform delivering internet security, performance, and developer infrastructure services.

> Cloudflare shipped on-demand schema learning for API Shield with live per-operation review, alongside smaller updates spanning MCP server skills, R2 analytics, a new multimodal Workers AI model, and faster Markdown conversion for agents.

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

01 **On-demand schema learning in API Shield** NEW 93

API Shield adds an Ad-hoc learn schema option in Web Assets > Operations > More options, letting you trigger a zone-wide schema learning run immediately instead of waiting for the weekly schedule. The Schema learning dialog shows current learning status and the last completed run for the zone, while the Learned schema view shows per-operation results — both update live so operations can be reviewed while the zone-wide run continues.

After sending representative traffic to a newly deployed API endpoint, immediately trigger schema learning rather than waiting up to a week for the scheduled run.

📍 1. In the Cloudflare dashboard, go to Web Assets › Operations. 2. Open the More options menu on the operations table. 3. Select Ad-hoc learn schema. 4. In the Schema learning dialog, select Run learning. 5. Monitor progress in the dialog; review per-operation results in the Learned schema view as they complete.

— Exact UI path, dialog names, and step-by-step usage example given

API Shield - Run schema learning on dem…

THINNER COVERAGE BELOW

02 **Clef-omni multimodal model on Workers AI** NEW 40

Clef-omni, a new multimodal decision model accepting text, image, audio, and video input, is now available on Workers AI.

— Names model and input types but no usage path given product docs

03 **R2 bandwidth analytics by location** NEW 35

R2 bandwidth analytics now break down traffic by the Cloudflare location that served each request.

— Describes the breakdown but no navigation or metric detail product docs

04 **Faster conversion in Markdown for Agents** IMPROVED 32

Markdown for Agents now uses in-process streaming conversion and supports larger HTML pages.

— States the change but omits size limits or mechanism detail product docs

05 **Cloudflare skills via API MCP server** NEW 30

The Cloudflare API MCP server now serves Cloudflare skills through the Skills over MCP extension.

— Names the extension but no usage detail or mechanism product docs

WAS THIS USEFUL?  

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 12, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue