

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 19, 2026 · EVERY DAY

EDITIONS **tail** | grep | head | ls | diff | uniq

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 55 changes](#)Do you prefer **scroll**?[▶ // CONTENTS](#) [SHOW](#)

24 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. Zenity now checks every Claude Enterprise tool call before it runs, with taint tracking for injected

vim ? help

pts, so an agent is stopped rather than audited after the fact.

Artemis moves detections into pull requests, giving rule changes review and rollback. Lasso's AI-BOM maps which agents inherit a vulnerable component.

PROTECT

Stop an agent's tool call before it executes, not after

Zenity enforces at the pre-execution point for Claude Enterprise, classifying read versus write boundaries and tracking prompt-injection taint through to the call, so a poisoned instruction fails closed instead of showing up in a log review. Operant applies the same shape to agentic payments across x402, MPP and AP2 with spending limits and scope policies, and Arcade narrows exposure by handing an agent only the tools it needs over MCP rather than the whole catalog.

[Zenity](#) · [Operant AI](#) · [Arcade](#)

DETECT

Review and roll back a detection rule like code

Detections live as pull requests in the team's own version control instead of only in the vendor portal, so a rule change gets peer review, a diff and a revert path. The usual failure — someone edits a rule in a web UI at 2am and nobody can reconstruct what it looked like before — stops being possible.

[Artemis Security](#)

IDENTIFY

Trace a vulnerable AI component to every agent that depends on it

The AI-BOM inventories agentic components and follows the dependency graph through to each agent consuming them. When an advisory lands on a model, library or MCP server, the blast-radius question is answered from the inventory rather than by asking teams which agents they built on what.

[Lasso Security](#)

RESPOND

Hand pentest findings and retest scheduling to a workflow engine

The Torq integration exposes NodeZero's full GraphQL API to workflow builders, so remediation verification can be scheduled automatically and findings pulled into existing runbooks. The manual loop of fixing something, then remembering to re-run the pentest to prove it, goes away.

OFFENSIVE SECURITY

PurpleAILAB Decepticon

[Feedback](#)

v1.5.2

[NOTES ↗](#)

RANK

Autonomous Hacking Agent for Red Team

Decepticon v1.5.2 streamlines engagement setup by auto-generating RoE directly from the CLI and adds a version-check safeguard to its update process.

←→ WHAT SHIPPED • 2 FEATURES most completely described first ? what's the number?

01 Auto-generated RoE for new CLI engagements NEW 78

Starting a new CLI engagement now only requires an engagement name and target web URL; Decepticon automatically generates the RoE and its matching approval and opens Red for the first request — no scope JSON file, manual import/approval command, or Soundwave planning session required.

Start a new engagement directly into Red without a pre-written scope file — supply a name and root URL and Decepticon handles RoE generation and approval automatically.

```
$ decepticon update --channel latest
decepticon
```

— Clear before/after and exact command, but no flag names for RoE config v1.5.2

02 Version check on update channel IMPROVED 70

Updates via `decepticon update --channel latest` now include a shipped-image version check to detect mismatched CLI image metadata at install time.

Start a new engagement directly into Red without a pre-written scope file — supply a name and root URL and Decepticon handles RoE generation and approval automatically.

```
$ decepticon update --channel latest
decepticon
```

— Names exact command and flag but mechanism of check is thin v1.5.2

WAS THIS USEFUL?



APPLICATION SECURITY

◆ Software Supply Chain Security

Endor Labs

[Feedback](#)[announced · Oct 16, 2026](#)[BLOG ↗](#)

Endor Labs is an application security platform built for AI-generated and human-written code, helping teams prioritize and fix real risk across the SDLC.

Endor Labs expands its SSCS platform with CI/CD Security: pipeline discovery, RSPM, and build integrity verification via artifact signing.

- › Adds Pipeline Discovery to map tools and workflows across all repositories, surfacing shadow engineering, security coverage gaps, and policy violations.
- › Adds Build Integrity Verification via artifact signing, attaching provenance information (source repository, branch, commit, and repository owner) to packages and containers to confirm they were produced by CI/CD pipelines and have not been tampered with.

COVERAGE

NEW detection

Repository Security Posture Management (RSPM) out-of-the-box policies

NEW compliance framework

CIS benchmark and SLSA/NIST SSDF compliance coverage via RSPM

All of Endor Labs's coverage changes →

WAS THIS USEFUL?



JFrog software supply chain solutions provide an end-to-end pipeline to control your binaries from build to production. Power your software updates to the edge.

JFrog introduced the Business Value Center, a real-time dashboard inside MyJFrog that quantifies ROI and security-scan coverage for SaaS customers.

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

01 Business Value Center dashboard in MyJFrog NEW 60

A new MyJFrog dashboard called JFrog Business Value Center surfaces real-time ROI, developer hours reclaimed, cost savings, and storage avoided across the software supply chain. It also tracks what share of repositories are actively scanned by JFrog Xray and what share of incoming open-source packages are blocked by JFrog Curation policies, giving security teams auditable coverage metrics on demand.

View your platform's ROI, hours saved, and scan-coverage metrics before a renewal or budget review.

Sign in to MyJFrog, then navigate to Business Value Center to see real-time ROI, developer hours reclaimed, storage savings, and the percentage of repositories currently scanned by Xray.

- Names dashboard, metrics, and navigation path but no API or config.

launch-20261015-aab63d61

WAS THIS USEFUL? thumbs up thumbs down

Mend.io Mend Platform

[Feedback](#)[announced · Oct 15, 2026](#)[BLOG ↗](#)

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

New [docker.io/sbx/mend-guardrails-kit](#) mixin routes AI agent traffic through Mend Guardrails inside Docker Sandboxes, adding content inspection without agent changes.

- › Adds [docker.io/sbx/mend-guardrails-kit](#) mixin kit that routes OpenAI-compatible agent traffic through [mend-guardrails-server](#) inside a Codex Docker Sandbox using [OPENAI_BASE_URL](#), enabling AI Guardrails policy enforcement without rebuilding the agent.
- › Adds [--kit-arg mend-guardrails.interceptTui=true](#) flag to extend Guardrails interception to Codex [TUI](#) model traffic (requires a host OpenAI Platform API key with Responses write permission).
- › Adds [mend-guard-text](#) command to explicitly evaluate [MCP](#) and tool-result text content against Mend AI Guardrails policy inside the sandbox.
- › Adds [mend-guardrails-selftest](#) built-in self-test command to validate that Guardrails enforcement is working correctly inside the running sandbox.

COVERAGE

[NEW](#) detection[Mend AI Guardrails — Prompt Injection and Secret Keys detection in Docker Sandboxes](#)[NEW](#) integration [Mend Guardrails SBX Kit integration with Docker Sandboxes / Codex](#)[All of Mend Platform's coverage changes →](#)

WAS THIS USEFUL?



ImmuniWeb

Feedback



announced · Oct 15, 2026

BLOG ↗

RANK

ImmuniWeb® AI Platform for Application Security Testing, Attack Surface Management & Dark Web Monitoring.

ImmuniWeb's Q3 2026 update adds new in-house AI models to Neuron for deeper attack simulation and testing, expands Discovery with AI-focused threat monitoring and dashboards, and adds Ping Identity SSO support.

←→ WHAT SHIPPED · 6 FEATURES most completely described first ? what's the number?

01 New AI models in ImmuniWeb Neuron for deeper testing NEW

70

ImmuniWeb® Neuron gains new in-house AI models for simulation of chained attacks, advanced business logic testing, real-time assistance to penetration testers, faster analysis of single-page applications, and reverse engineering of API schemas.

– Lists five concrete AI capabilities but no commands or UI path `launch-20261015-fa9a3cba`

THINNER COVERAGE BELOW

02 New monitoring capabilities in ImmuniWeb Discovery NEW

50

ImmuniWeb® Discovery adds monitoring of AI training data sets for leaked or unlawfully scraped data, a new dedicated incident type for stolen credit cards, and enhanced detection of exposed AI infrastructure.

– Names three distinct monitoring targets but no mechanism detail

`launch-20261015-fa9a3cba`

03 New dashboards and alerting in ImmuniWeb Discovery NEW

50

Discovery adds a new dedicated Third-Party Risk Management (TPRM) dashboard, customizable alerts for the Cyber Threat Intelligence (CTI) dashboard, and advanced search of past alerts on a dedicated dashboard.

– Names three dashboard/alert additions but no navigation detail `launch-20261015-fa9a3cba`

04 **Advanced search for AI misconfigurations in Neuron** NEW

40

ImmuniWeb® Neuron adds advanced search for AI-related misconfigurations or leaked tokens.

– Names the search target but not how it is used `launch-20261015-fa9a3cba`

05 **AI-based asset ownership attribution in Discovery** NEW

30

ImmuniWeb® Discovery adds new in-house AI models to better attribute IT asset ownership.

– Bare description of capability with no mechanism or scope `launch-20261015-fa9a3cba`

06 **Ping Identity SSO integration** NEW

25

Adds Ping Identity SSO to the existing DevSecOps and other integrations.

– Names the provider but no setup or scope detail `launch-20261015-fa9a3cba`

WAS THIS USEFUL?



REVERSE ENGINEERING, MALWARE & FORENSICS

CAPEv2

Feedback



changes since 2026-10-17

CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 gains structured JSON logging with `json_logs = on`, TTY-aware ANSI colors, and `main_task_id` propagation across all analysis logs.

↳ GET THIS VERSION

```
$ git clone --branch commits-2026-10-17
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-17
```

↳ USE IT

Enable structured JSON logging so CAPE emits machine-readable `*.json.log` files ingestible by Fluentd or Google Cloud Logging alongside human-readable logs.

```
custom/conf/logging.conf

json_logs = on
color_logs = off
```

Enable JSON logging without editing config files — useful in containerised or CI environments where the config is baked in.

```
$ CAPE_JSON_LOGS=1 python3 cuckoo.py
```

- › Adds `StructuredJsonFormatter` and `init_json_file_handler` in `lib/cuckoo/core/log.py` to emit side-by-side single-line JSON log files (`*.json.log`) compatible with Google Cloud Logging / Fluentd, enabled via `json_logs = on` in `logging.conf` or the `CAPE_JSON_LOGS=1` environment variable.
- › Adds `color_logs = auto/on/off` setting and `NO_COLOR` environment variable support to gate ANSI color sequences in `ConsoleHandler` and `ColoredFormatter` based on TTY detection, preventing raw escape sequences in systemd/syslog logs.
- › Adds contextvars-based task/machine/correlation context propagation via `set_log_task_context`, `log_task_context`, and `TaskContextFilter` across

`AnalysisManager` , `utils/process.py` , and `utils/gcp_pubsub_service.py` .

- > Propagates `main_task_id` (from `Task.get_attribute('main_task_id')`) into `AnalysisLogger` , `AnalysisManager.run` , `GuestManager` , and `Scheduler` logs, formatting log lines as Task #<id> (<main_task_id>) to match distributed task formatting.
- > Propagates `main_task_id` and `task_label` through `ResultServer.add_task` and `ResultServer.del_task` into `HandlerContext` , `ProtocolHandler` , `FileUpload` , `LogHandler` , and `BsonStore` log messages in both single-process and multiworker modes.

+1 more

WAS THIS USEFUL?



CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback



Lambda Web Functions are now in your inventory:

NOTES ↗

Plerion is a cloud security platform that detects misconfigurations, vulnerabilities, and threats across cloud environments and provides remediation guidance.

COVERAGE

NEW

integration [AWS Lambda Web Functions inventory support](#)

[All of Plerion's coverage changes →](#)

WAS THIS USEFUL?



DETECTION & RESPONSE

Elastic Security Labs

[Feedback](#)[announced · Oct 16, 2026](#)[BLOG ↗](#)[RANK](#)

Elastic security labs is a powered by Elasticsearch. Ready to deploy your data?

Elastic Security shipped a native integration with Google Distributed Cloud (GDC) air-gapped, bundling unified [SIEM](#) and [EDR](#), curated Mandiant threat intelligence, and a [SecOps](#) Foundation layer for sovereign and disconnected deployments.

←→ [WHAT SHIPPED · 5 FEATURES](#) most completely described first [? what's the number?](#)

01 SIEMv2 on Google Distributed Cloud air-gapped NEW

55

Ingests and correlates infrastructure, network, and audit telemetry with pre-built [MITRE ATT&CK](#)-mapped detection rules and behavioral analytics in a fully disconnected environment.

– Names telemetry types and MITRE ATT&CK mapping but no exact rules or commands.

[launch-20261016-7f65f739](#)

02 EDRv2 on Google Distributed Cloud air-gapped NEW

53

Provides kernel-level behavioral protection via a single unified agent to detect memory exploits, polymorphic malware, and [lateral movement](#) without external network access.

– Explains mechanism and threat types but no deployment or config detail.

[launch-20261016-7f65f739](#)

03 SecOps Foundation layer for sovereign deployments NEW

39

Adds built-in case management, incident response plans, runbooks, and standard operating procedures for sovereign and regulated deployments on Google Distributed Cloud.

– Lists components but no interface or workflow detail.


[launch-20261016-7f65f739](#)

04 Elastic Security integration with Google Distributed Cloud air-gapped

NEW

35

Embeds unified SIEM and EDR natively into Google Distributed Cloud (GDC) air-gapped infrastructure, without custom log forwarders or external connectors, for sovereign environments.



– Describes the integration's purpose but no setup steps or config.

launch-20261016-7f65f739

05 Bundled Mandiant threat intelligence for air-gapped use NEW

34

Curated Mandiant threat intelligence — including indicators of compromise and adversary TTPs — made available for use in completely air-gapped environments.

– States content included but not how it's accessed or updated. launch-20261016-7f65f739

WAS THIS USEFUL?



Artemis security is a traditional SIEMs weren't built for today's world. Artemis closes that gap with an AI-native Protection Platform designed to dramatically reduce MTTD and MTTR for modern threats.

Artemis Security launched Detection as Code, letting teams manage detections as pull requests in their own version control system instead of only inside the Artemis portal.

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

01 Detection as Code via Git repositories NEW 91

Detection as Code is now GA, enabled from [Settings → Detectors → Detection as Code](#) to connect a repository and manage detections as pull requests. It supports GitHub (including GitHub Enterprise Server), GitLab, Bitbucket Cloud, and Cursor Origin as detection repository hosts, with a native Cursor Origin marketplace integration authorized via the Artemis app for Origin with admin-approved write access. Every new or tuned detection — adaptive, intelligence, user-created, and AI Mode detectors — is pushed to the repository as a pull request with an explanation and a link back to the detection in Artemis; existing branch protection rules and required checks are respected so detections stay unchanged until a pull request merges. Engineers can edit detections in their own editor and merge via the standard pull request flow, which Artemis picks up automatically, and when a repository is first connected Artemis opens an initial pull request organizing existing detections by data source.

Connect a GitHub repository so Artemis opens pull requests for every new or tuned detection your organization owns.

 In the Artemis portal, go to [Settings → Detectors → Detection as Code](#), select your repository host (e.g. GitHub), choose a repository and branch, and click Connect. Artemis will open an initial pull request adding your existing detections organized by data source.

– Names exact menu path, supported hosts, and PR behavior with a runnable setup example.

`launch-20261015-9a0b7309`

WAS THIS USEFUL?



Trend Vision One

Feedback



Automated tagging rules now have VDI type available as a condition (2026-10-16)

BLOG ↗

RANK

Trend Vision One is Trend Micro's unified cybersecurity platform that consolidates threat detection, response, and risk management across endpoints, cloud, email, and networks.

Trend Vision One's Tag Management now lets automated tagging rules target virtual desktops by persistence state.

←→ WHAT SHIPPED • 1 FEATURE most completely described first ? what's the number?

01 VDI type condition in automated tagging rules NEW

83

Tag Management's automated tagging rules gain a **VDI type** condition supporting **Persistent** and **Non-persistent** values for the **Devices** asset type, letting rules identify and tag VDI devices by persistence state. This condition can only be combined with **Asset Group** or **CREM exception list** platform tag conditions, which prevents duplicate-tag rule creation.

Tag all non-persistent VDI devices automatically so they can be scoped into ephemeral-device policies or excluded from persistent-host alerts.

📌 In the console, go to Service Management → Tag Management → Automated Tagging Rules, create or edit a rule, set the asset type to **Devices**, add a condition using **VDI type**, and select **Any of Non-persistent**. Combine only with an **Asset Group** or **CREM exception list** platform tag condition, then save.

– Names exact condition, values and UI path; lacks broader rollout context

Automated tagging rules now have VDI ty...

WAS THIS USEFUL?



XFA

Feedback



announced · Sep 22, 2026

BLOG ↗

RANK

XFA launched a free, no-time-limit plan that delivers device security checks and warnings at every sign-in without requiring a paid subscription.

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

01 Free no-time-limit plan with sign-in device checks NEW

60

XFA introduced a free, no-time-limit plan available at dashboard.xfa.tech/signup that delivers device security warnings at every sign-in, without requiring a paid subscription, IT team, or budget.

Sign up for the free plan to start checking devices at sign-in with no IT team or budget required.

📌 Open dashboard.xfa.tech /signup and create a free account to enable user-driven device security checks with warnings at every sign-in.

– Names signup URL and behaviour but no mechanism detail. [launch-20260922-474f216d](#)

WAS THIS USEFUL?



EXPOSURE & VULNERABILITY MANAGEMENT

Horizon3 NodeZero

Feedback



announced · Oct 15, 2026

BLOG ↗

RANK

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

NodeZero launches a Torq integration that exposes its full GraphQL API to workflow builders, letting Torq schedule automated remediation verification and pull pentest data directly from NodeZero.

↳ WHAT SHIPPED · 1 FEATURE most completely described first ? what's the number?

01 Torq integration with GraphQL API and verification workflows

NEW

85

NodeZero adds a Torq integration that retrieves pentests, weaknesses, attack paths, hosts, and credentials via a **Horizon3 API key** and gateway URL, and exposes a **GraphQL query** step in Torq that opens the full Horizon3 API beyond the pre-built action steps. Two API key permission levels are supported: a **User role** key that can retrieve results and schedule NodeZero verification tests, and a **Read-only** key that can retrieve results but cannot start tests. Torq workflows can schedule **1-Click Verify** retests of specific weaknesses, reusing the original pentest configuration, and retrieve resolution status per weakness to drive automated ticket closure or re-escalation.

After Torq remediates a credential exposure, use a User role API key to schedule a 1-Click Verify retest and confirm the weakness is no longer exploitable before closing the ticket.

📌 In Torq, add a **Schedule NodeZero Verification Test** step using a User role Horizon3 API key and the gateway URL for your region, select **1-Click Verify** against the original pentest ID, then add a **Get Weakness Resolution Status** step — if status is **resolved**, close the finding; otherwise route it back for review.

— Names API key types, GraphQL step and 1-Click Verify with a worked example.

launch-20261015-5ab0f342

WAS THIS USEFUL?



PDQ Deploy & Connect

Feedback



announced · Oct 15, 2026

BLOG ↗

RANK

PDQ deploy & connect is a ready to save time & automate your deployments? Learn how you can update software, deploy custom scripts, & make configuration changes in minutes with PDQ Deploy.

PDQ added two live, non-disruptive remote administration tools — an interactive Remote Shell and a Live File Explorer — alongside an expanded AI assistant and new dashboard customization options.

↳ WHAT SHIPPED · 4 FEATURES most completely described first ? what's the number?

01 Remote Shell for live command execution NEW 75

An interactive terminal inside PDQ for running commands and reviewing output in real time on managed Windows and macOS devices, without taking over the user's screen.

When a deployment fails and you need to follow the evidence interactively — run commands, review output, and decide the next step — without disrupting the end user.

📌 In PDQ, open the target device record, select **Remote Shell**, then run commands interactively in the live terminal (e.g. `Get-EventLog -LogName Application -Newest 20`) and review output in real time.

— Names mechanism and includes a runnable usage example. `launch-20261015-bbc8db4f`

02 Live File Explorer for remote file access NEW 75

Lets administrators browse, search, upload, download, and delete files on a live device file system without launching a remote desktop session, transferring one file at a time within an applicable file-size limit.

When you need to retrieve a log file or drop a script onto a device without launching a remote desktop session or interrupting the user.

📌 In PDQ, open the target device record, select **Live File Explorer**, navigate to the relevant directory, and use Upload or Download to transfer the file — one file at a time, within the applicable file-size limit.

— Clear mechanism plus step-by-step example with a stated limit. `launch-20261015-bbc8db4f`

THINNER COVERAGE BELOW

03 AI assistant expanded with deployment context IMPROVED

The AI assistant (open early access) now uses package, deployment, and device context to explain deployment results, break down output, and suggest targeted diagnostics, with admin approval required before any action runs on a device.

– Explains behavior but gives no concrete path to try it. `launch-20261015-bbc8db4f`

04 Dashboard customization and new widgets NEW

Administrators can now show/hide, reorder, and resize dashboard widgets, and the dashboard gains a most-vulnerable-devices widget, a configurable low-disk-space threshold, and a more compact OS patching visualization.

– Names all specific widgets but no exact navigation given. `launch-20261015-bbc8db4f`

WAS THIS USEFUL?



IDENTITY, ACCESS & DATA PROTECTION

Firezone

Feedback



Android 1.5.16 (2026-10-19)

BLOG ↗

RANK

Replace your VPN with Firezone, an open-source zero trust access platform built on WireGuard®. Connect users to anything, anywhere.

Firezone's Android client 1.5.16 ships a unified Compose-based settings UI, device pool grouping in the resource list, and several packet-processing performance improvements.

←→ WHAT SHIPPED • 3 FEATURES most completely described first ? what's the number?

01 Device pools in resource list IMPROVED 40

The resource list now shows device pools instead of a flat view of all connected devices.

– Describes what changed but not how pools are formed Android 1.5.16 (2026-10-19)

02 Packet-processing performance improvements IMPROVED 40

Android client packet processing was sped up in three ways: improved performance when flow logging is enabled, faster peer and route lookups, and reduced allocations and packet moves during network writes.

– Names the three optimization areas but no metrics or benchmarks

Android 1.5.16 (2026-10-19)

03 Unified Compose-based settings UI IMPROVED 25

The Android app now uses a unified Compose-based settings UI, replacing the previous settings interface.

– Named only in summary, no mechanism or screens detailed Android 1.5.16 (2026-10-19)

WAS THIS USEFUL?



◆ Identity & Access Management

Ory

[Feedback](#)

Ory Network, Ory Hydra, Ory Kratos, Ory Polis, Ory Elements v26.4.3 released

[NOTES](#)[RANK](#)

Discover the latest features, improvements, and updates to Ory.

Ory Network v26.4.3 adds configurable one-time code length and character set in Ory Kratos.

↳ **WHAT SHIPPED** · **1 FEATURE** most completely described first [? what's the number?](#)

01 **Configurable one-time code format in Kratos** NEW 37

Ory Network v26.4.3 supports configuring the length and characters of one-time codes issued by Ory Kratos Identities.

– Names the capability but no config key or exact syntax given.

Ory Network, Ory Hydra, Ory Kratos, Ory...

WAS THIS USEFUL?



Clerk

[Feedback](#)[announced · Oct 15, 2026](#)[BLOG ↗](#)[RANK](#)

Clerk is an authentication and user-management platform that provides pre-built UI components, APIs, and SDKs for adding sign-up, sign-in, and identity features to applications.

Clerk's MCP server now supports full account management: AI agents can authenticate via a remote endpoint, manage apps, instances, users and domains, walk through a guided production-deployment workflow, and preview changes with a dry-run mode — all protected by short-lived per-call scoped tokens.

↳ **WHAT SHIPPED · 5 FEATURES** most completely described first [? what's the number?](#)

01 Remote MCP server endpoint with account sign-in **NEW**

70

Adds `https://mcp.clerk.com/mcp` as a remote MCP server endpoint, letting AI clients authenticate with a Clerk account via a one-time consent approval and then read and modify workspace apps and instances within that same authenticated session.

Register the Clerk MCP server in your AI client so your agent can authenticate and manage Clerk resources.

text

`https://mcp.clerk.com/mcp`

— Names the endpoint and consent flow; readers can register it directly

`launch-20261015-6cb35cfd`

02 Production deployment workflow via MCP **NEW**

70

Adds a production-deployment workflow via MCP that creates a production instance on a given domain, copies development settings, and reports required DNS records and OAuth credential gaps.

Ask your agent to walk through production setup — it creates the instance, copies dev settings, and reports outstanding DNS and OAuth gaps.

📌 Take this app to production on example.com and tell me what's left.

— Describes mechanism and outcome with a usage example `launch-20261015-6cb35cfd`

03 MCP tools for apps, instances, users, and domains **NEW**

68

Adds MCP tools to read and create Clerk applications and to read or change instance settings (with sensitive values redacted); tools to list, filter, and create users within a Clerk instance; and tools to list an app's domains with their required DNS records and trigger DNS, certificate, and email verification re-checks.

– Names every tool surface but lacks exact tool call syntax `launch-20261015-6cb35cfd`

04 **Per-call scoped tokens for MCP tool calls** NEW 65

Enforces per-call scoped tokens expiring within 5 minutes — each MCP tool call receives only that tool's permissions, preventing cross-tool privilege escalation.

– Clear mechanism and limit, but no user action required `launch-20261015-6cb35cfd`

THINNER COVERAGE BELOW

05 **Dry-run mode for instance and user changes** NEW 55

Supports a dry-run mode for instance setting changes and user creation so agents can preview changes before applying them.

Preview a user creation in your development instance before committing, using the dry-run mode.

📌 Create a test user in my development instance with the email `ada@example.com` — dry run first, show me what would change.

– Simple behaviour described with a concrete usage example `launch-20261015-6cb35cfd`

↳ **ALSO FROM THESE RELEASES**

.

WAS THIS USEFUL?



Veeam Data Cloud

Feedback



changelog · Oct 15, 2026

BLOG ↗

Veeam Data Cloud is a SaaS backup and recovery service that protects Microsoft 365, Azure, and virtual machine data with cloud-hosted storage and management.

Veeam iResOps for Microsoft 365 (GA) adds AI Undo, AI-agent visibility, sensitive-data classification, ROT cleanup, and minimum viable recovery guidance via the DataAI Command Graph.

- › Launches **Veeam iResOps for Microsoft 365** as a generally available add-on to Veeam Data Cloud for Microsoft 365, powered by the **DataAI Command Graph** intelligence layer.
- › Adds AI Undo capability to automatically roll back AI-driven file modifications and deletions in Microsoft 365, identifying exactly which files an AI agent changed even across millions of files.
- › Adds ROT (redundant, obsolete, and trivial) data detection across Microsoft 365, with automatic quarantine of duplicate, near-duplicate, and ownerless files to reduce storage costs and compliance footprint.
- › Adds minimum viable recovery guidance via the **DataAI Command Graph**, linking production data context to backup coverage to identify gaps and prioritize recovery of business-critical data.
- › Opens **Veeam iResOps for Veeam Data Platform** and **Veeam iResOps for Okta** in Private Preview.

COVERAGE

NEW integrations

[Veeam iResOps for Microsoft 365 — Generally Available, Veeam iResOps for Veeam Data Platform — Private Preview](#)

NEW compliance framework

[Automated compliance checks for GDPR, HIPAA, SOX, DORA across Microsoft 365](#)

[All of Veeam Data Cloud's coverage changes →](#)

WAS THIS USEFUL?



Acronis Cyber Protect Cloud

Feedback



changelog · Oct 15, 2026

BLOG ↗

Acronis Cyber Protect Cloud is a cloud-based platform combining data backup, disaster recovery, cybersecurity, and endpoint management for managed service providers and businesses.

10 new integrations added to the Acronis Technology Ecosystem, spanning AI automation, remote support, documentation, password management, hardware and security visibility for MSPs.

COVERAGE

NEW

integrations

Discern Security, Flowgear, Iperius Remote Desktop, IT Glue, MSP Skills, NordLayer and 4 more

All of Acronis Cyber Protect Cloud's coverage changes →

WAS THIS USEFUL?



GOVERN

Operant AI

Feedback



announced · Oct 19, 2026

BLOG ↗

RANK

Operant AI discovers, detects and defends AI applications, agents, MCP servers, APIs and Kubernetes workloads at runtime.

Operant AI shipped Agentic Commerce Defense for its Agent Protector, a suite of new capabilities that evaluate and govern agentic payments across x402, MPP, and AP2 protocols in real time, including a semantic firewall, scope policies, spending limits, and a new dashboard.

↳ WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 Atomic spending limits across agents and networks NEW

80

Atomic spending limits can be applied by agent or service, MCP server, user, team, network, and asset, with configurable hourly, daily, weekly, or monthly windows; budget is reserved before forwarding each payment to prevent concurrent agents from double-spending the same allowance.

– Mechanism and scopes fully named, though no exact config key given.

launch-20261019-c9c4e63a

02 Real-time policy evaluation for agentic payments NEW

61

Agentic Commerce Defense adds transaction-level policy evaluation for agentic payments over the x402, MPP, and AP2 protocols, evaluating each payment before the money is forwarded.

– Names protocols and behavior but no config or command to invoke.

launch-20261019-c9c4e63a

03 Commercial scope policies restrict payment destinations NEW

60

Commercial scope policies restrict which MCP servers, payment networks, assets, and wallet destinations an agent may use, blocking payments to unapproved destinations regardless of credential validity or budget.

– Names policy scope dimensions but no config syntax shown.

launch-20261019-c9c4e63a

THINNER COVERAGE BELOW

04 Operant Semantic Firewall for payment sessions NEW

51

The new Operant Semantic Firewall evaluates paid tool activity against the session's original task, flagging behavior consistent with prompt injection, task drift, or payment loops even when payments are properly authorized and under budget.

– Clear mechanism described, but no setup or trigger details given.

launch-20261019-c9c4e63a

05 **Agentic Commerce Dashboard for payment visibility** NEW

43

The Agentic Commerce Dashboard provides visibility into payment requests, policy decisions, submissions, and reported outcomes, without storing payment proofs or holding customer signing keys.

– Describes dashboard purpose but not its exact location or use.

launch-20261019-c9c4e63a

WAS THIS USEFUL?  

Lasso Security

Feedback



announced · Oct 15, 2026

[BLOG](#)

RANK

Lasso Security protects employee GenAI use, LLM applications and AI agents with runtime guardrails, data-leak prevention and CPU-only guardrail models.

Lasso introduced an AI-BOM that inventories agentic AI components and traces vulnerabilities through to every dependent agent, integrated into the Discovery stage of its AI Security Platform.

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

01 AI-BOM for agentic component inventory NEW 78

The AI-BOM inventories AI building blocks — models, MCP servers and tools, vector databases, frameworks, and AI services — discovered from code, cloud, and runtime sources, recording first-seen and last-seen timestamps per component. Each component shows a **Used by** count of dependent agents and links through to the AI-SPM graph with the full exposure path highlighted; vulnerabilities are categorized against the OWASP Top 10 for Agentic Applications (e.g. **ASI05 Unexpected Code Execution** , **ASI04 Agentic Supply Chain Vulnerabilities**) with each finding naming the responsible code or configuration. Newly discovered components are flagged so AppSec can review them before they spread to additional agents, and the AI-BOM feeds into AI risk management, Automated AI Red Teaming, and runtime protection stages of Lasso's AI Security Platform.

– Names OWASP codes and platform stages but no UI path or command given

launch-20261015-355fc0bc

WAS THIS USEFUL?



Zenity

Feedback



announced · Oct 19, 2026

BLOG ↗

RANK

Zenity discovers, secures and governs enterprise AI agents from buildtime to runtime across SaaS, endpoint and cloud.

Zenity launched Universal Runtime Control for Claude Enterprise, adding pre-execution enforcement of every tool call plus read/write boundary classification, tool-provenance controls, prompt-injection taint tracking, and Compliance API support for the new denials.

↳ WHAT SHIPPED · 5 FEATURES most completely described first ? what's the number?

01 Pre-execution enforcement for Claude tool calls NEW

85

Adds pre-execution enforcement via Anthropic's `tool_call` Inference Hooks event, holding every Claude tool call for an allow-or-deny verdict. The `Validate tool calls` configuration routes supported Claude tool calls to the organization's Zenity security server for policy evaluation before Anthropic executes them. A single configuration covers claude.ai (web, desktop, mobile), Cowork, Claude Code, Claude for Microsoft 365, Claude in Chrome, Claude's built-in browser, and Claude Tag in Slack — no endpoint agent or browser extension required.

Enable pre-execution tool-call enforcement in Zenity so that every Claude tool call is held for a policy verdict before it runs — no endpoint agent needed.

📌 In the Zenity console, go to Claude Enterprise integration settings and enable `Validate tool calls`. Contact your Zenity account team to activate tool-call enforcement on your tenant.

— Names exact config, hook event, and setup steps; full mechanism given.

`launch-20261019-f198282d`

02 Tool-provenance access controls NEW

60

Adds tool-provenance controls using tool-call metadata to restrict access by provider type (platform, Claude application, third-party `MCP` server, or client tools such as `bash` and local MCP servers in Claude Code), with the ability to allow only Anthropic-verified servers.

— Names provider types and tools but no config steps shown.

`launch-20261019-f198282d`

03 Prompt-injection taint tracking NEW

60

Adds prompt-injection taint tracking: sessions where injected instructions are detected are marked tainted, and any subsequent tool call acting on those instructions is blocked at the pre-execution checkpoint, with the denial reason surfaced in Claude and logged in the Activity Feed.

– Clear mechanism and logging location, no setup command given. `launch-20261019-f198282d`

THINNER COVERAGE BELOW

04 **Read vs write boundaries for tool calls** IMPROVED 55

Extends Zenity Boundaries to classify tool calls as reads vs. writes/deletes/executes, allowing policies such as permitting CRM reads while blocking deletions — without relying on connector or tool names.

– Mechanism and example given but no exact policy syntax. `launch-20261019-f198282d`

05 **Compliance API support for pre-execution denials** IMPROVED 30

Extends Zenity's existing Compliance API integration to support audit and investigation of pre-execution denials.

– Brief mention with no endpoint or field names given. `launch-20261019-f198282d`

WAS THIS USEFUL?  

INFRASTRUCTURE

Cloudflare

[Feedback](#)

[Bots - Precursor adds Monitor mode \(2026-10-19\)](#)
[BLOG](#)
[RANK](#)

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

Cloudflare's biggest shipments this window are the beta launch of Cloud Workload Security for policing cloud API traffic and new Precursor Monitor/Arm modes with per-path Interstitial Rules for bot management. Alongside these, Cloudflare raised Tunnel route limits to 50,000 with paid Tunnels for Platforms, added a [cf docs search](#) CLI subcommand, CJK keyword search for AI Search, Cron Trigger local testing, configurable Workers Builds runtimes, [DLP](#) match confidence controls, custom Application Library apps, and a breaking kebab-case rename of CLI commands.

←→ [WHAT SHIPPED](#) • [10 FEATURES](#) most completely described first [? what's the number?](#)

01 Precursor Monitor mode and decoupled Interstitial Rules NEW

95

Adds a [Monitor](#) option to the [Precursor Status](#) dropdown (under [Security > Settings](#)) so traffic runs through Precursor detections without affecting bot score or Challenge clearance, alongside the existing [Arm](#) option that lets detections actively affect bot score and Challenge clearance in production. Decouples the Precursor Interstitial from Precursor mode: set a sitewide default with [Show Interstitial](#), then override per-path with [Interstitial Rules](#) (e.g. enforce an interstitial Challenge on [/checkout](#) while leaving it off elsewhere).

:

Safely evaluate Precursor detections on live traffic before they impact bot scores or block users.

📌 In the Cloudflare dashboard, go to Security > Settings, open the [Precursor Status](#) dropdown, and select [Monitor](#). Precursor detections will run but bot score and Challenge clearance remain unchanged.

Require an interstitial Challenge only on sensitive paths (e.g. checkout) while keeping the rest of the site unaffected.

📌 In Security > Settings, set [Show Interstitial](#) to off (sitewide default), then add an [Interstitial Rules](#) entry matching the path [/checkout](#) with the interstitial Challenge enabled.

02 **`cf docs search` subcommand for documentation lookup** NEW

85

The Cloudflare CLI (`1.0.0-beta.16`) adds a `cf docs search` subcommand that searches live Cloudflare documentation from the terminal or a coding agent, returning up to five ranked JSON results (title, description, URL) per query. Piping results through `jq` and requesting the `Accept: text/markdown` header fetches full documentation pages as Markdown.

Search Cloudflare docs from the terminal and retrieve the top result as Markdown — useful when working with a coding agent or in a shell-only environment.

```
$ url=$(cf docs search "rate limit requests with WAF" | jq -r '[0].url')
curl -H "Accept: text/markdown" "$url"
```

Search Cloudflare documentation without leaving the terminal during an incident or local dev session.

```
$ cf docs search 'tunnel route limit'
```

Search Cloudflare documentation without leaving the terminal during an incident or review.

```
$ cf docs search 'tunnel route limit'
```

– Exact command and flags given with runnable example product docs

03 **50,000 route limit and Tunnels for Platforms purchasing**

IMPROVED

80

Raises the default route limit for Cloudflare Tunnel (and Cloudflare Mesh) to 50,000 per account, up from 1,000, covering CIDR and hostname routes for `cloudflared`. Introduces Tunnels for Platforms, letting accounts purchase more than 1,000 tunnels at \$4.50 per tunnel per month via `Tunnels > Buy more tunnels` in the dashboard.

– Exact limits, price and dashboard path specified product docs

04 **Cron Trigger local testing in Local Explorer** NEW

75

Local Explorer can now invoke a Worker's `scheduled()` handler on demand during local development, including configured or ad hoc Cron Triggers with custom scheduled times, without needing to construct URLs manually.

Trigger a Worker's `scheduled()` handler on demand during local development without constructing URLs manually — useful for CI or agent-driven test harnesses.

```
$ curl -X POST "http://localhost:8787/cdn-cgi/local/explorer/api/local/scheduled?worker=my-worker" \
  -H 'Content-Type: application/json' \
  -d '{"cron":"0 * * * *","scheduled_time":1745856238000}'
```

Fire a Cron Trigger at an arbitrary scheduled time locally to verify time-sensitive job logic before deploying.

 In the Local Explorer UI, select your Worker, choose **Cron Triggers**, enter a custom scheduled time, and click **Invoke**.

– Runnable curl endpoint and UI steps given [product docs](#)

05 CJK keyword search via trigram tokenizer in AI Search NEW

75

Adds `keyword_tokenizer: trigram` support to AI Search keyword search, enabling Japanese, Chinese, and Korean natural-language queries. Changing the tokenizer triggers a full reindex.

Enable CJK keyword search for an AI Search index whose content is in Japanese, Chinese, or Korean — note that changing this setting triggers a full reindex.

```
yaml
keyword_tokenizer: trigram
```

– Exact config key and reindex caveat given [product docs](#)

06 Configurable default runtimes for Workers Builds IMPROVED

75

Workers Builds now supports environment-variable overrides for default runtimes:

`NODE_VERSION`, `PYTHON_VERSION`, `RUBY_VERSION`, `GO_VERSION`, `BUN_VERSION`, and `YARN_VERSION`. Bun 1.4.3 supports both `bun.lockb` (v1) and `bun.lock` (v2) lockfiles.

– Lists every override variable and lockfile versions [product docs](#)

07 Kebab-case command naming in Cloudflare CLI

Cloudflare CLI `1.0.0-beta.16` standardizes generated command names to kebab-case — for example, `cf accounts billing getCredits` is renamed to `cf accounts billing get-credits`. Any script or automation using the old camelCase command names will break.

– Exact before/after command names given [product docs](#)

08 Cloud Workload Security (CWS) beta NEW

60

Launches Cloud Workload Security in Beta, enabling request-level policy enforcement on cloud API traffic — including Rego checks and request filters — via a CWS proxy deployed in your environment that sits in front of requests before they reach the cloud provider. Requires a Zero Trust subscription.

– Describes mechanism but no config, endpoint or command shown [product docs](#)

THINNER COVERAGE BELOW

09 DLP match confidence controls in Workers NEW

30

Enables Cloudflare Workers to validate data loss prevention matches by adjusting confidence or rejecting matches from custom pattern entries.

– No API, flag or example given for the mechanism [product docs](#)

10 Custom applications in Application Library for Gateway policies

NEW

30

Supports custom applications in the Application Library, referenceable in Gateway DNS, Network, and HTTP policies.

– Only summary-level description, no steps or config given [product docs](#)

! —> BREAKING ON UPGRADE

! Cloudflare CLI `1.0.0-beta.16` standardizes generated command names to kebab-case — for example, `cf accounts billing getCredits` is renamed to `cf accounts billing get-credits`. Any script or automation using camelCase command names will break.

WAS THIS USEFUL?



OTHER / UNCATEGORIZED

vigolium

Feedback ⓘ v0.5.4 **NOTES** **RANK**

Vigolium - High-fidelity vulnerability scanner fusing agentic AI with native speed, modularity, and precision

Vigolium added automatic failover for out-of-band application security testing (OAST) when a custom server is unreachable.

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

01 **OAST fallback to default server** **NEW**

75

New `oast.fallback_to_default` config key (default `true`) automatically re-registers against `oast.pro` when the configured `server_url` does not respond, without forwarding the token.

– Names exact config key, default, and behavior v0.5.4

WAS THIS USEFUL?  

Arcade

Feedback



announced · Oct 15, 2026

BLOG ↗

RANK

Arcade is the actions runtime for enterprise AI agents [^]_^_^ one control point to enforce, execute, and govern every action, every time.

Arcade launched Tool Recommendation, letting agents dynamically pull only the tools they need over MCP instead of loading the full catalog, cutting input tokens by up to 91%, alongside policy enforcement and audit trails applied to every tool call.

↳ WHAT SHIPPED · 2 FEATURES most completely described first ? what's the number?

01 Tool Recommendation via on-demand MCP tool retrieval NEW

80

Arcade adds `Arcade_ListApps` , `Arcade_SelectTools` , and `Arcade_UseTool` as built-in tools exposed through the Arcade gateway, replacing full catalog loading with on-demand tool retrieval over MCP. Tool Recommendation uses purpose-built embedding and ranking models to match an agent's plain-language task description against the full tool catalog and return only the relevant tool schemas, cutting input tokens by up to 91%. It can be enabled per Arcade gateway with no SDK changes, framework-specific integration, or agent-side rebuilds, and works with any MCP client or framework; it is available in Arcade Cloud now and also deployable in customer cloud, on-premises, and air-gapped environments connected to existing identity providers and policy engines.

– Names exact built-ins and token savings but no setup command shown.

launch-20261015-8c478675

THINNER COVERAGE BELOW

02 Policy enforcement and audit trails for tool calls NEW

50

Pre- and post-tool-call policy enforcement hooks, PII screening, and DLP protection are applied to every recommendation and tool call, with audit trails exported via OpenTelemetry.

– Mechanism named but no config keys or endpoints given. launch-20261015-8c478675

WAS THIS USEFUL?



// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 19, 2026

Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ **Subscribe**

Feedback on this issue