THE CYBER TOOLCHAIN

NO. 054

# Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 23, 2026 · EVERY DAY



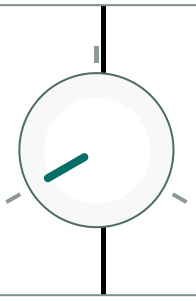
*The daily firehose — everything the toolchain shipped today, already filtered.*

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

**scroll**



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 27 tools

## OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

# Bugcrowd

Feedback ⓘ 1 RELEASE · 2026-08-23 CHANGELOG ↗

Bugcrowd is a crowdsourced security platform connecting organizations with ethical hackers to identify vulnerabilities through bug bounty programs.

> Bugcrowd introduced Savant Pathseeker 1.0 in Early Access, bringing an autonomous pentesting capability to the platform.

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

01 **Savant Pathseeker autonomous pentesting (Early Access)** NEW

25

Bugcrowd launched Savant Pathseeker 1.0 in Early Access, described as an autonomous pentesting capability available on the Bugcrowd platform.

— Names the capability but no mechanism, scope, or access details.

changelog-20260823-3027718a

WAS THIS USEFUL? 👍 👎

## rekono-kbx

Feedback ⓘ 1 RELEASE · 2026-08-22 NOTES ↗

Offensive security platform that automates attack surface discovery and vulnerability management

> Rekono 2.0 is a major overhaul: it adds finding triage and fixed-status tracking, multi-format reporting, an alerting system, pentesting notes, and a batch of enrichment integrations (VirusTotal, HackTricks, CVE Crowd, NVD), alongside a large set of breaking API, queue, broker and deployment changes.

└──▷ **WHAT SHIPPED · 20 FEATURES** most completely described first ? what's the number?

### 01 Settings API restructured for per-integration control

75

BREAKING

Replaces the `/api/system/` endpoint with `/api/settings/` plus one dedicated endpoint per integration, giving granular control over each integration; administrators can now enable and disable individual integrations through these endpoints.

Query the new per-integration settings endpoint to check which integrations are enabled before configuring enrichment.

```
$ curl -s -H 'Authorization: Bearer <token>' https://<rekono-host>/api/settings/
```

— Names old and new endpoints and gives a runnable query 2.0.0-1

### 02 Enrichment integrations for hosts and CVEs NEW

75

Adds integration with VirusTotal to retrieve reputation, malicious-analysis results, and WHOIS data for discovered hosts; HackTricks to surface related wiki resources alongside identified assets; and CVE Crowd to flag which detected CVEs are trending on social media. Adds support for NVD NIST API keys to raise rate limits and accelerate CVE enrichment, a monitor background job that keeps CVE-linked vulnerabilities up to date and purges stale executions, and domain and geolocation resolution for every discovered host.

— Names every integration and its concrete function 2.0.0-1

### 03 Deployment and environment requirement changes BREAKING

70

Removes Rekono Desktop (Electron apps and Debian package); Rekono is now deployed via Docker or from source only. Removes support for deploying under a custom web path

via `rootpath`. Raises required Python version to 3.11 and required Dirsearch version to 0.5.0, and moves TLS certificates into a Docker volume instead of the `docker/nginx/tls` directory.

— Names every removed artifact and each new version/path requirement `2.0.0-1`

04 **API field and operation removals** BREAKING 68

The host `address` field is renamed to `ip`; `scheduled_in` is removed from scans in favour of `scheduled_at` as the only scheduling field; the finding fields `first_seen`, `last_seen`, `detected_by`, and `is_active` are removed from the API; and the `enable` / `disable` operations on findings are removed.

— Lists exact renamed/removed fields and operations, no migration command `2.0.0-1`

05 **Execution configuration: headers, proxies, scope, and redirects**

NEW 65

Adds per-instance, per-user, and per-target customisation of HTTP headers sent by tools during executions; proxy configuration individually configurable for HTTP, HTTPS, FTP, and all protocols; a path field on target ports to limit execution scope; and redirect-following when resolving URLs so tools execute against the final destination.

— Names all four configuration surfaces added to executions `2.0.0-1`

06 **Finding triage and fixed-status tracking** NEW 60

Practitioners can mark each finding as true positive, false positive, or won't fix with a supporting comment. Rekono also adds automatic and manual `fixed` status: a finding is marked fixed automatically when subsequent executions stop reporting it.

— Explains triage states and the automatic-fixed mechanism `2.0.0-1`

07 **Multi-format report generation with template customisation**

NEW 60

Adds report generation in PDF, JSON, or XML filtered by the user's criteria, with support for customising the HTML template used to produce PDF reports.

— Names formats and the customisable template mechanism `2.0.0-1`

08 **Message queue rename and broker replacement** BREAKING 60

Renames the `tasks-queue`, `executions-queue`, and `findings-queue` queues to `tasks`, `executions`, and `findings`, and replaces `Redis` with `Valkey` as the message broker.

— Names old and new queue names and the broker swap `2.0.0-1`

THINNER COVERAGE BELOW

09 **Official multi-arch Docker images** NEW — 55

Publishes official Docker images on Docker Hub for both `linux/amd64` and `linux/arm64`.

— Named registry and architectures, directly pullable 2.0.0-1

10 **SMTP configuration moved to database** BREAKING — 55

Moves SMTP configuration into the database, removing it from the configuration file and from `RKN_EMAIL_*` environment variables.

— Names removed env vars and new storage location 2.0.0-1

11 **Telegram bot inline scan commands replace dedicated commands** — 50

BREAKING

Removes the `/newauth`, `/newtechnology`, and `/newvulnerability` Telegram bot commands; these values are now requested inline when a scan is created.

— Names removed commands and replacement flow, no exact new syntax 2.0.0-1

12 **Pentesting notes with rich-text editor** NEW — 45

Adds pentesting notes linkable to any Rekono entity, with a rich-text editor supporting code blocks, tables, task lists, and images.

— Names editor capabilities but no entry point given 2.0.0-1

13 **Automatic finding creation from user-provided data** NEW — 45

Adds automatic creation of findings from user-provided data (targets, target ports, input technologies) to maintain consistent finding relationships.

— Names input types but not the resulting workflow 2.0.0-1

14 **Per-finding detail pages** NEW — 40

Adds per-finding detail pages showing finding metadata, exposure window, and related findings.

— Clear UI destination but limited mechanism 2.0.0-1

15 **Alerting system for matching findings** NEW — 35

Notifies each alert's subscribers when a discovered finding matches the alert's criteria.

— Describes the trigger but no configuration surface 2.0.0-1

16 **Execution transparency: original reports and commands** NEW

35

Adds download of the original reports generated by each tool and visualisation of the exact commands executed by each execution.

— States what is now visible/downloadable, no path given `2.0.0-1`

17 **New default wordlists** NEW 35

Adds new default wordlists covering CMS, APIs, GraphQL, and specific services.

— Names wordlist categories, no counts or file paths `2.0.0-1`

18 **Administrative queue monitoring** NEW 30

Administrators can now inspect queue status and statistics.

— No detail on what statistics or where to view them `2.0.0-1`

19 **Expanded metrics pages** NEW 25

Adds expanded metrics pages with additional KPIs and statistics.

— No specifics on which KPIs were added `2.0.0-1`

20 **Light and dark UI themes** NEW 20

Adds light and dark UI themes.

— Bare mention with no toggle location given `2.0.0-1`

└──▷ **BREAKING ON UPGRADE**

! The `/api/system/` endpoint is replaced by `/api/settings/` and per-integration endpoints — any client calling `/api/system/` will break.

! The `address` field on host objects is renamed to `ip` — any code or integration referencing `address` will break.

! The `scheduled_in` field is removed from scans; use `scheduled_at` instead.

! Finding fields `first_seen`, `last_seen`, `detected_by`, and `is_active` are removed from the API.

! The `enable` and `disable` operations on findings are removed.

! The Telegram bot commands `/newauth`, `/newtechnology`, and `/newvulnerability` are removed.

! SMTP configuration is removed from the configuration file and from `RKN_EMAIL_*` environment variables; it must now be set via the database/UI.

! The queues `tasks-queue`, `executions-queue`, and `findings-queue` are renamed to `tasks`, `executions`, and `findings` — any external tooling targeting the old queue names will break.

! Redis is replaced by Valkey — existing Redis-only deployments must migrate to Valkey.

! Required Python version is raised to 3.11.

! Required Dirsearch version is raised to 0.5.0.

! TLS certificates are now stored in a Docker volume instead of the `docker/nginx/tls` directory.

! Rekono Desktop (Electron apps and Debian package) is removed; the only supported deployment methods are Docker and from source.

! The `rootpath` sub-path deployment option is removed.

WAS THIS USEFUL? 👍 👎

## adscan

Feedback ⓘ 1 RELEASE · 2026-08-23 NOTES ↗

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

> adscan v11.2.0 adds a dedicated DNS server flag so scans in segmented networks resolve names correctly instead of silently falling back from Kerberos to NTLM.

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

01 **Separate DNS server for segmented networks** NEW 88

The `--dns-server` flag is added to `adscan ci`, `adscan execute`, and `adscan doctor` to direct all name resolution to a DNS host separate from the domain controller when the AD-zone DNS server isn't the DC. It covers domain and PDC-hostname discovery, SMB and ADCS CA name resolution, Kerberos SPN lookup, cross-forest trust resolution, and the network-preflight DNS check. `--dc-ip` remains the authentication and enumeration target, and omitting `--dns-server` preserves existing behaviour.

Run a CI scan against a segmented network where the AD DNS server is a separate host from the domain controller, preventing silent fallback from Kerberos to NTLM.

```
$ adscan ci auth --type audit --interface eth0 --domain
corp.local --dc-ip 10.0.0.1 --dns-server 10.0.0.5 -u alice -p
'S3cr3t!'
```

— Names exact flag, three subcommands, and behaviour scope precisely v11.2.0

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

## Nmap

Feedback ⓘ changes since 2026-08-18 CODE ↗

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

Nmap's rsync-list-modules NSE script gains structured XML output

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-18
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-18
```

WAS THIS USEFUL? 👍 👎

# ProjectDiscovery subfinder

Feedback ⓘ

1 RELEASE · 2026-08-22 NOTES ↗

Fast passive subdomain enumeration tool.

> Subfinder v2.16.0 adds a new passive source and makes result-count limits reliable across sources that paginate.

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

## 01 Max-results enforcement across paginating sources IMPROVED

72

The `-mr` / `-max-results` flag is now enforced across all paginating sources, so result caps are reliably honored when a source pages through results instead of being ignored past the first page.

Cap enumeration to 100 results per source across paginating APIs to keep recon runs predictable in time and volume.

```
$ subfinder -d example.com -mr 100 -all
```

— Names exact flag and behavior change, mechanism clear v2.16.0

## 02 New scanmalware passive source NEW

60

Adds `scanmalware` as a new passive subdomain discovery source, usable via `-s scanmalware`.

Use the new scanmalware source alongside crtsh to cross-reference malware-reported subdomains during passive recon.

```
$ subfinder -d example.com -s scanmalware,crtsh -v
```

— Names flag and source but no detail on data provided v2.16.0

WAS THIS USEFUL? 👍 👎

## APPLICATION SECURITY

◆ SAST / SCA / Code Security

# cve-lite-cli

Feedback ⓘ 1 RELEASE · 2026-08-23 NOTES ↗

Fast, developer-friendly JS/TS dependency vulnerability scanner with local lockfile scanning, OSV matching, direct vs transitive visibility, --fix, JSON output, and practical remediation guidance.

> cve-lite-cli v1.30.0 extends `--fix` to auto-apply safe transitive dependency fixes without touching parent versions, adds a Root column tracing transitive findings back to their direct dependency across all report formats, and speeds up pnpm monorepo scans.

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

## 01 Auto-apply safe transitive fixes with --fix IMPROVED

85

`--fix` now auto-applies safe within-range transitive refreshes when the fix is available inside the parent's declared range, without requiring a parent-level upgrade — landing the smallest safe change automatically.

Auto-fix transitive vulnerabilities in-range without upgrading parents — useful in CI to land the smallest safe change.

```
$ cve-lite /path/to/project --fix
```

— Exact flag, mechanism, and runnable CI command given v1.30.0

THINNER COVERAGE BELOW

## 02 Root column for transitive findings NEW

55

A new Root column shows the top-level direct dependency behind each transitive finding, across terminal, JSON, and HTML report outputs, including multi-folder scans.

— Names output formats but no exact command or flag v1.30.0

## 03 Faster pnpm lockfile parsing IMPROVED

50

pnpm lockfile parsing now runs once per scan instead of up to three times, delivering approximately 14% faster warm scans on large monorepos.

— Concrete before/after and number but nothing to configure v1.30.0

WAS THIS USEFUL? 👍 👎

## Waratek

Feedback ⓘ 1 RELEASE · 2026-08-18 BLOG ↗

Waratek is a Java application security platform that provides runtime protection against vulnerabilities and attacks without code modification.

> Waratek shipped Reflection Protection, a new zero-configuration RASP rule that blocks unsafe Java reflection and dynamic class-loading attacks, including zero-days, at the JVM layer.

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

01 **Reflection Protection rule for Waratek RASP** NEW 72

A new Reflection Protection rule for Waratek RASP for Java blocks exploitation of unsafe Java reflection and dynamic class-loading vulnerabilities — including zero-days — by hooking into the Reflection API and tracking untrusted data at the JVM layer (covering patterns like Class.forName() and .newInstance() abuse). It requires no source code changes, redeployment, or application downtime, and can be enabled directly from the Waratek Portal's rule wizard with no configuration, taking effect immediately across all protected applications.

Enable Reflection Protection across all protected Java applications without any configuration or downtime.

> 📍 In the Waratek Portal, open the rule wizard, locate the Reflection Protection rule, and enable it. Protection takes effect immediately at the JVM layer across all connected applications.

— Names mechanism, attack patterns and exact enablement path via Portal wizard.

Waratek Launches Reflection Protection,…

WAS THIS USEFUL? 👍 👎

# Crash Override Chalk

Feedback ⓘ 1 RELEASE · 2026-08-11

NOTES ↗

Chalk allows you to follow code from development, through builds and into production.

> Chalk v1.2.0 adds a DNS-based telemetry sink for restricted environments, sink fallback chains and auto-disabling, richer sink metadata headers, per-chalk reporting, new timing/error runtime keys (with a breaking template change), a configurable process nice level, and collapsible GitLab CI log output.

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

## 01 DNS telemetry sink for restricted environments NEW

90

New `dns` sink type encodes chalk key values into DNS lookup hostnames via a `domain_template` with `{KEY}` placeholders, providing a telemetry channel for environments where HTTPS egress is restricted; supports `dns_server`, `require_chalk_mark`, and per-key `normalize.<KEY>.callback` fields in `sink_config`.

Beacon chalk metadata over DNS in air-gapped or HTTPS-restricted environments, truncating the commit ID to fit within the 63-character DNS label limit.

con4m

```
func trunc16(v: string) {
  return slice(v, 0, 16)
}

sink_config my_dns_beacon {
  sink:               "dns"
  domain_template:    "{_COMMIT_ID}.
{METADATA_ID}.beacons.example.com"
  dns_server:         "10.0.0.1:5353"
  require_chalk_mark: true
  normalize._COMMIT_ID.callback: func trunc16
}
subscribe("report", "my_dns_beacon")
```

— Full config example and named fields provided v1.2.0

## 02 Sink fallback chains via fallback_sink_config NEW

88

New `fallback_sink_config` field on all `sink_config` sections names a backup sink to use when the primary is unavailable or disabled; chalk walks the full fallback chain on each delivery failure, with each sink tracking its own `disable_after_errors` counter independently.

Route reports to a regional failover API when the primary is down, so chalk delivery continues without manual intervention.

```con4m
sink_config primary_api {
  sink:                "post"
  uri:                 "https://api.us-east-
1.example.com/report"
  disable_after_errors: 3
  fallback_sink_config: "backup_api"
}

sink_config backup_api {
  sink:                "post"
  uri:                 "https://api.eu-west-
1.example.com/report"
  disable_after_errors: 3
}

subscribe("report", "primary_api")
```

— Config example shows primary/backup sink setup v1.2.0

---

03 **Metadata headers and header forwarding on sinks** NEW

80

`post` and `presign` sinks now send chalk metadata headers with every request: `X-Chalk-Version`, `X-Chalk-Operation`, `X-Chalk-Action-Id`, `X-Content-Length`, `X-Chalk-Digest-Sha256`, and `X-Chalk-Attempt` (1-based retry counter). The `presign` sink also supports `x-forward-headers`, copying header names listed in that response header from the redirect response into the upload PUT, enabling metadata injection (e.g. `X-Amz-Meta-*`) that cannot be encoded in the presigned URL.

— All header names listed, no runnable example v1.2.0

---

04 **Post-mark error keys replace FAILED_KEYS** BREAKING

80

Two new runtime artifact keys capture post-mark errors: `_ERR_INFO` (superset of `ERR_INFO`, full error log through end of operation) and `_FAILED_KEYS` (superset of `FAILED_KEYS`, includes post-build step failures); all report templates except `report_all` now subscribe to these instead of the chalk-time originals. `report_default`, `report_large`, and `insertion_default` templates have had

`FAILED_KEYS` removed and `_FAILED_KEYS` added, so consumers reading `FAILED_KEYS` from these templates must switch to `_FAILED_KEYS`.

— Migration note given, exact keys named v1.2.0

05 **Auto-disabling sinks on repeated errors** NEW 78

`post`, `presign`, and `s3` sinks now automatically disable themselves on repeated errors via a configurable `disable_after_errors: <int>` threshold (default 3); 4xx responses (except 429) disable immediately, while 5xx, 429, and network errors increment a consecutive-failure counter.

— Named config key and status-code behavior but no example v1.2.0

06 **Configurable nice level for reporting process** NEW 65

New `exec.nice` config field sets the nice level (range `0` - `19`, default `5`) for the chalk reporting process after fork, reverted before postexec and heartbeat so each phase uses its own independent value.

— Named config field with range, no usage example v1.2.0

07 **Runtime timing and heartbeat keys** NEW 65

Four new runtime host keys for timing and heartbeat tracking: `_MONOTIME` (monotonic clock at operation start, in ms), `_SINCE_TIMESTAMP` (wall-clock ms since process start), `_SINCE_MONOTIME` (monotonic ms since process start), and `_HEARTBEAT_COUNT` (1-based counter, present only in heartbeat reports).

— Keys named with meaning, no usage example v1.2.0

08 **Collapsible GitLab CI log sections** NEW 65

GitLab CI log output is now automatically wrapped in `section_start` / `section_end` markers with `collapsed=true` when `GITLAB_CI` is set, making chalk reports collapsible in the GitLab job log UI.

— Exact markers and trigger env var named v1.2.0

THINNER COVERAGE BELOW

09 **Per-chalk reporting mode** NEW 55

`outconf` sections now support `per_chalk_reports: true` and `custom_report` sections support `per_chalk: true` to emit one report per chalk mark instead of bundling all marks into a single report.

— Config keys named, no example given v1.2.0

└──▷ **BREAKING ON UPGRADE**

! `report_default`, `report_large`, and `insertion_default` templates have had `FAILED_KEYS` removed and `_FAILED_KEYS` added; consumers reading `FAILED_KEYS` from these report templates must switch to `_FAILED_KEYS`.

---

WAS THIS USEFUL? 👍 👎

# StackHawk

Feedback

1 RELEASE · 2026-08-23

CHANGELOG ↗

StackHawk is a dynamic application security testing platform that automatically scans web applications for vulnerabilities during development and CI/CD pipelines.

> StackHawk 6.4.0 introduces flexible multi-profile scan modes, a new auth-bypass detection rule, and an overhaul of Team ownership management for applications, alongside a fix that makes `hawk validate config` correctly fail CI builds on invalid configuration.

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

01 **Team ownership management for applications** NEW 83

The Applications table now shows Team ownership for each application and supports inline reassignment, with apps lacking a Team labeled `Unassigned`; the same field appears in `Application Settings > General` with `Unassigned` listed first. Members can now view Teams pages for Teams they belong to, and View Only users can review every Team in the organization; Team management controls (Create a Team, Add Members, Add Applications, rename, Delete) now render disabled rather than hidden for users lacking permission, and users who cannot manage Teams but belong to at least one Team must assign a newly created application to one of their Teams during the Create Application flow. Separately, `PUT /api/v1/org/{orgId}/team/{teamId}/application` was fixed so team assignments actually persist — it previously returned success without saving.

Re-confirm application-to-Team assignments via the API after the fix to `PUT /api/v1/org/{orgId}/team/{teamId}/application` — previously these calls returned success without saving.

```
$ curl -X PUT
https://api.stackhawk.com/api/v1/org/{orgId}/team/{teamId}/applic
ation   -H 'Authorization: Bearer <token>'   -H 'Content-Type:
application/json'   -d '{"applicationId": "<appId>"}'
```

— Names endpoint, UI locations and permission behavior in detail

changelog-20260823-f86f3f67

02 **`hawk validate config` exit code fix** BREAKING 70

`hawk validate config` previously exited `0` even on hard validation errors and missing configuration files; it now exits non-zero on those conditions, so CI pipelines chaining commands like `hawk validate config stackhawk.yml && ./deploy.sh` will correctly fail when the configuration is invalid.

— Exact command and before/after exit behavior given changelog-20260823-f86f3f67

THINNER COVERAGE BELOW

03 **New auth-bypass scan rule** NEW 15

StackHawk 6.4.0 adds a new scan rule for detecting authentication bypass vulnerabilities.

— Only mentioned in summary, no mechanism or ID given changelog-20260823-f86f3f67

└──▷ **BREAKING ON UPGRADE**

! `hawk validate config` previously exited `0` even on hard validation errors and missing configuration files; it now exits non-zero on those conditions. CI pipelines relying on the always-zero exit behavior (e.g., `hawk validate config stackhawk.yml && ./deploy.sh`) will now correctly fail when the configuration is invalid.

WAS THIS USEFUL? 👍 👎

◆ **Web Application Security**

## caido

Feedback ⓘ 1 RELEASE · 2026-08-22 NOTES ↗

Caido releases, wiki and roadmap

> Caido v0.58.2 adds live progress logging to the project upgrade window, giving visibility into long-running migrations.

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

01 **Progress logs in project upgrade window** IMPROVED 35

The project upgrade window now displays live progress logs during migration, giving visibility into long-running upgrades.

— Names the UI surface but no mechanism or example given. v0.58.2

WAS THIS USEFUL? 👍 👎

# Salt Security

Feedback

ⓘ 1 RELEASE · 2026-08-04 BLOG ↗

Salt Security is an API security platform that detects and protects against API threats, vulnerabilities, and misuse in real-time.

> Salt Security launched its first AWS WAF Partner Managed Ruleset purpose-built for API and AI agent protection, adding MCP traffic awareness, deeper API attack detection, and request enrichment for downstream analytics.

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

## 01 AWS WAF Managed Ruleset for API and AI agent protection NEW

85

Salt Managed Rules for AWS WAF is the first AWS WAF Partner Managed Ruleset built specifically to protect APIs and AI agents. It is subscribable from AWS Marketplace and attachable to existing web ACLs with no new infrastructure, and is available across all commercial AWS Regions and globally via Amazon CloudFront.

Subscribe to and attach Salt Managed Rules to an existing AWS WAF web ACL to immediately gain API and AI agent threat protection.

📍 1. Open the AWS Console and navigate to AWS Marketplace. 2. Search for `Salt Managed Rules for AWS WAF` and subscribe. 3. Go to AWS WAF › Web ACLs and select your target ACL. 4. Choose `Add managed rule groups`, locate `Salt Managed Rules`, and attach it to the web ACL. 5. Save the ACL configuration — no new agents or architecture changes required.

— Includes exact AWS console steps to subscribe and attach the ruleset.

Salt Debuts First AWS WAF Managed Rules…

THINNER COVERAGE BELOW

## 02 MCP endpoint awareness in AWS WAF NEW

50

Adds industry-first MCP endpoint awareness that identifies and labels MCP traffic, blocks unauthenticated MCP access, and provides observability into MCP interactions inside AWS WAF environments.

— Describes behavior but gives no config key or flag to enable it.

Salt Debuts First AWS WAF Managed Rules…

## 03 API attack vector detection beyond standard WAF rules NEW

50

Detects API-layer attack vectors missed by standard WAF rules, including credential brute force, excessive GraphQL queries, SSRF, prototype pollution, and JWT-based

anomalies.

— Names specific attack types but gives no configuration or activation steps.

Salt Debuts First AWS WAF Managed Rules…

04 **Downstream analytics enrichment via labeled request attributes**

NEW

45

Enriches downstream analytics by labeling critical request attributes including auth headers, user identifiers, and GraphQL queries, flowing security signals into existing detection stacks.

— Explains enrichment purpose but no endpoint or field schema given.

Salt Debuts First AWS WAF Managed Rules…

WAS THIS USEFUL? 👍 👎

## REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

## CAPEv2 Feedback ⓘ changes since 2026-08-21 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 gains an API call trace diff tool, psycopg3 PostgreSQL support with auto-upgrade, and iterative process tree cycle detection.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-21
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-21
```

Adds automatic in-memory upgrade of `postgresql://` connection strings to `postgresql+psycopg://` in `lib/cuckoo/core/database.py` when psycopg v3 is installed, preventing startup `ImportError` without any config change.

Adds `psycopg` (v3) support as the PostgreSQL database driver, enabling modern async capability and significant performance gains over the previous psycopg2 driver.

Introduces an API call trace diff comparison tool in the web compare view (`web/compare/views.py`) with basic sequence alignment, replacing the earlier PR #3093 approach.

Routes demultiplexed archive/container submissions by the platform detected per extracted file rather than the outer container, deferring `machine='all'` expansion until the extracted platform is known.

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

# Mitiga

Feedback ⓘ 3 RELEASES · 2026-07-24 → 2026-08-19 BLOG ↗

Mitiga is a cloud security platform that detects and responds to threats in AWS, Azure, and GCP environments.

> Mitiga launched Helios AIDR, a new AI-driven detection and response platform anchored by AI Insights for automated triage, and rolled out the HeliosAI Agent with guided investigation, natural-language log search, AI-generated event summaries, and API access, alongside a new Investigation Workbench for unified multi-cloud and SaaS incident review.

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

01 **AI Summary for Workbench Events** NEW 75

Adds AI Summary for Workbench Events, which reads a selected group of forensic events and generates a narrative covering what happened, in what order, and why it matters, with one-click access to the underlying raw log entries and detectors.

Use AI Summary for Workbench Events to turn dozens of scattered forensic events from a cloud incident into a single readable narrative.

> 📍 In the Mitiga Helios console, open Workbench, select the forensic events you want to summarize, then click AI Summary to generate a narrative of what happened, in what order, and why it matters — with a link to the underlying raw log entries.

— Exact UI steps to generate a summary are given HeliosAI Agent: Finally, a Tier 2/3 Ana…

02 **HeliosAI Investigation guided cross-cloud incident review**

NEW 70

Adds HeliosAI Investigation to the HeliosAI Agent, providing guided incident investigation across AWS, GCP, Azure, Microsoft 365, Okta, and SaaS applications — surfacing leads, connecting identity activity, and showing the forensic evidence and confidence reasoning behind every verdict.

— Names six supported platforms but no navigation or command

HeliosAI Agent: Finally, a Tier 2/3 Ana…

03 **Investigation Workbench in IR2 Platform** NEW 70

Adds Investigation Workbench to the IR2 Platform, providing a single-pane view of activity chains across multi-cloud and SaaS environments without requiring deep cloud investigation expertise. It gives automated timelines of cloud and SaaS user activity to

help SOC teams assess incident scope — such as tracking post-phishing actions like file downloads or email access — and lets internal IR and SOC teams determine incident materiality and choose between in-house containment or external expert escalation, supporting SEC 4-day disclosure requirements.

— Describes workflow and use case but no config or command

Introducing Investigation Workbench (20…

04 **Helios AIDR platform with AI Insights triage** NEW 60

Introduces Helios AIDR, an AI-driven cloud detection and response platform performing real-time triage across multi-cloud and SaaS environments. AI Insights, its first capability, automatically executes contextual triage tests against every potential threat lead and significant cloud activity, dynamically adjusts risk scores to separate routine activity from genuine concerns without manual analyst intervention, and delivers pre-assembled evidence packages with automatically crafted attacker timelines spanning multiple cloud environments.

— Detailed mechanism but no UI path or command shown

Helios AIDR: Why Cloud Security Needs I…

THINNER COVERAGE BELOW

05 **Natural-language log search in HeliosAI Agent** NEW 50

Adds natural-language log search to the HeliosAI Agent, letting analysts query the Cloud Security Data Lake in plain English with no query syntax or schema knowledge required.

— Names the Cloud Security Data Lake but gives no query example

HeliosAI Agent: Finally, a Tier 2/3 Ana…

06 **AI integrated across the detection lifecycle** NEW 45

Helios AIDR applies AI across the full detection lifecycle: research intelligence for formulating Indicators of Attack (IoAs), automated detection rule validation and tuning before deployment, and real-time runtime analysis in production.

— Names lifecycle stages but no concrete usage detail

Helios AIDR: Why Cloud Security Needs I…

07 **HeliosAI Agent API access** NEW 30

Exposes the HeliosAI Agent via API for teams building their own agentic SOC workflows.

— No endpoint, method, or path named HeliosAI Agent: Finally, a Tier 2/3 Ana…

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

# Realm Security

Feedback ⓘ 2 RELEASES · 2026-07-28 → 2026-08-04

BLOG ↗

> Realm Security added a direct Databricks Zerobus Ingest pipeline for structured, enriched security telemetry, and debuted Detection Integrity alongside Data Haven Search to let teams cut SIEM log volume without losing threat coverage.

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

01 **Databricks Zerobus Ingest integration** NEW 79

Realm streams security telemetry directly into Unity Catalog Delta tables via Databricks' serverless push-based Zerobus Ingest API, with no customer-managed staging bucket or message bus required. It sends both raw logs and parsed, OCSF-normalized versions simultaneously so teams can route raw events to Bronze and structured data to Silver within a medallion architecture, applying real-time enrichment (GeoIP, threat intel, and customer-supplied CMDB or HRIS context) before data lands in the lakehouse. The feed is configured entirely through the Realm UI once the Databricks target table, Zerobus Ingest endpoint, and access credentials are set up, requiring no additional pipeline to build or maintain.

— Names integration and data flow but only a UI setup path, no API call shown

Realm + Databricks Zerobus Ingest: Clea…

02 **Data Haven Search for historical events** NEW 79

Search inside Realm Data Haven lets analysts query their full immutable security history directly against OCSF-normalized observables and resupply exact events to any destination, with no restore step, schema setup, or data-lake infrastructure required. It launches with unmetered queries, with pricing based on ingestion volume rather than how much analysts query or retrieve.

During an active investigation, search Data Haven to pull historical events that were never sent to the SIEM and resupply them to your analysis destination.

> 📍 In the Realm Platform, go to Data Haven › Search, query against OCSF-normalized observables for the relevant time window, then select Resupply to forward the matching raw events to your chosen destination.

— Includes concrete navigation steps and pricing model detail

Realm Security Debuts Detection Integri…

03 **Detection Integrity for SIEM rule coverage** NEW 78

Detection Integrity reads every SIEM detection, maps each rule to its dependent log sources and fields, and auto-generates pipeline protection rules to preserve coverage as log volume is reduced, supporting Sigma, KQL, and SPL rule formats. It produces a shareable report showing what was reduced, what was protected, projected savings, and MITRE ATT&CK coverage before and after each reduction, and flags a finding instead of cutting it when a reduction cannot be applied safely.

— Clear mechanism and report contents, but no config key or command given

Realm Security Debuts Detection Integri…

---

WAS THIS USEFUL? 👍 👎

◆ Security Data Fabric & Federated Search

# Vega

Feedback ⓘ 1 RELEASE · 2026-08-04 NOTES ↗

Vega is a decentralized protocol for derivatives trading that enables peer-to-peer trading of futures and options on a blockchain.

> Vega launched Detection Skills, a new open standard for attaching agentic triage, investigation, and tuning logic to detections, along with a public skill library, sandbox, and a matching workflow inside the Vega Cyber Defense Platform.

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

## 01 Detection Skills Library, sandbox, and GitHub repo NEW

83

The Detection Skills Library at `detectionskills.io` ships 50+ skills from Vega Research and partners — the same skills that run inside the Vega platform. A sandbox at `detectionskills.io` lets users build a detection, attach a skill, preview the agentic verdict, and export spec-compliant YAML. A GitHub repository hosts the library so the community can contribute Detection Skills back to the shared set.

Try the Detection Skills sandbox to attach a skill to a detection, preview the agentic verdict, and export a spec-compliant YAML skill file.

> 📍 In the detectionskills.io sandbox: (1) Build or paste in a detection. (2) Select a skill from the Detection Skills Library (e.g. a triage skill). (3) Preview the verdict output. (4) Export the result as spec-compliant YAML for use in your own platform.

— Names sandbox, library site, YAML export, and repo with steps snapshot-20260823

## 02 Detection Skills workflow in Vega Cyber Defense Platform NEW

67

A new three-step Detection Skills workflow inside the Vega Cyber Defense Platform lets teams compose detections in a detection notebook, attach skills from the Skill Library to run triage, investigation, and optimization, and approve or reject tuning proposals that write approved verdicts back into the rule. Triage skills run before an alert surfaces rather than after escalation, preventing noise before it reaches an analyst.

— Describes workflow steps and before/after behavior, no exact UI path snapshot-20260823

## 03 Detection Skills open standard NEW

60

Detection Skills is a new open standard, built on Anthropic's Agent Skills format, for writing triage, investigation, and optimization/tuning reasoning as versioned skills that

attach directly to detections. Skills execute as a continuous agentic loop rather than a one-off check.

– Standard explained conceptually but no runnable surface given `snapshot-20260823`

THINNER COVERAGE BELOW

04 **Detection Skills for agentic AI threats** NEW 35

New Detection Skills add coverage for agentic AI threat scenarios, specifically targeting dangerous model activity as a detection use case.

– Names the use case but gives no mechanism or example `snapshot-20260823`

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ **Detection Engineering**

## Tenzir

Feedback ⓘ **1 RELEASE · 2026-08-23** **BLOG ↗**

Tenzir is a data pipeline platform that ingests, transforms, and analyzes security telemetry and logs for threat detection.

> Tenzir Test 2.1.0 introduces format-aware sorting that keeps multiline record blocks intact during test output comparison, stabilizing nondeterministic event ordering.

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

01 **Format-aware test output sorting** NEW 52

Tenzir Test now applies pre-compare transforms consistently across the shell, Python fixture, and diff runners before baseline comparison, and keeps multiline record blocks intact when sorting test output, which stabilizes nondeterministic event ordering.

– Describes mechanism and scope but no config key or flag given.

`Tenzir Test v2.1.0: Format-aware test o…`

WAS THIS USEFUL? 👍 👎

## IDENTITY, ACCESS & DATA PROTECTION

◆ Public Key Infrastructure (PKI)

# DigiCert CertCentral

Feedback ⓘ 1 RELEASE · 2026-08-23

CHANGELOG ↗

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

> CertCentral shipped three distinct capabilities in this window: a reusable DNS TXT record method for bulk domain revalidation, multi-year validity options for X9 PKI TLS certificates, and in-place ICA rotation on existing ACME Directory URLs.

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

## 01 Multi-year validity for X9 PKI TLS certificates IMPROVED

83

X9 PKI for TLS certificate validity options are extended from a 397-day maximum to selectable 1-, 2-, or 3-year terms, with custom validity capped at 1,095 days. Available on the Dashboard under `Request a Certificate > Request X9 PKI for TLS Certificate > Certificate Settings > Validity period`.

Request a 3-year X9 PKI for TLS certificate for a long-lived internal system to reduce renewal overhead.

> 📍 In CertCentral, go to Dashboard › Request a Certificate › Request X9 PKI for TLS Certificate, then under Certificate Settings › Validity period, select the 3-year option and complete the order.

— Exact before/after limits and precise navigation path changelog-20260823-0e8cddb0

## 02 In-place ICA rotation on ACME Directory URLs NEW

75

Lets you update the issuing intermediate CA (ICA) certificate on an existing ACME Directory URL without creating a new URL, via the ACME Directory URL Restrictions panel using the `Intermediate chains [Intermediate CA] > [Root CA]` menu.

Switch an existing ACME Directory URL to a different intermediate CA chain without recreating it — useful when rotating ICA certificates.

> 📍 In CertCentral, open the ACME Directory URL settings, then navigate to ACME Directory URL Restrictions › Intermediate chains [Intermediate CA] › [Root CA] menu, select the desired certificate chain, and save changes.

— Names exact menu path and behavior change changelog-20260823-0e8cddb0

THINNER COVERAGE BELOW

### 03 Persistent DNS TXT record for bulk revalidation NEW

53

CertCentral now supports the Persistent DNS TXT record DCV method in bulk domain revalidation — a one-time DNS TXT record containing a DigiCert-provided persistent URI that is reused on every subsequent revalidation without creating new records or values.

— Explains mechanism but no UI/config steps given changelog-20260823-0e8cddb0

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

# vopono

Feedback ⓘ 2 RELEASES · 2026-08-22 → 2026-08-23 NOTES ↗

Run applications through VPN tunnels with temporary network namespaces

> vopono added a health-check subcommand and made outputs scriptable with JSON, alongside AirVPN WireGuard sync support and IPv6 fixes.

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

01 **vopono check connectivity subcommand** NEW 65

New `vopono check` subcommand runs daemon connectivity and health checks covering IPv4, IPv6, and DNS lookups.

— Names exact subcommand and checks covered 0.10.22

02 **AirVPN WireGuard config sync** NEW 60

Adds AirVPN WireGuard config syncing support via API key through `vopono sync`.

— Names provider, mechanism and command 0.10.21

THINNER COVERAGE BELOW

03 **JSON output flag for commands** NEW 55

Adds a `--json` flag to most commands for programmatic/scriptable output.

— Names exact flag but scope is 'most commands' 0.10.21

04 **IPv6 handling improvements** IMPROVED 25

Improves IPv6 handling in network namespaces, carried over from an unpublished 0.10.20 release.

— No mechanism or scope beyond general claim 0.10.21

WAS THIS USEFUL? 👍 👎

# Cato Networks

Feedback

2 RELEASES · 2026-08-10 → 2026-08-13

BLOG ↗

Cato Networks is a cloud-native security platform providing unified networking and security services through a global SD-WAN infrastructure.

> Cato Networks shipped Agentic Threat Prevention, a new inline and agentic-decision engine that detects and blocks multi-stage AI-driven attacks in real time, and extended Cato AI Security to Claude via Anthropic Inference Hooks with MCP/skill visibility, full-lifecycle policy enforcement, and scoped rollout controls.

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

## 01 Agentic Threat Prevention for multi-stage attacks NEW

78

New Agentic Threat Prevention capability combines inline conditional enforcement — which dynamically tags hosts based on recent activity and denies specific behaviors mapped to MITRE ATT&CK tactics, for example blocking executable downloads from cloud services on hosts that have contacted known C2 servers — with an agentic decision-making layer that continuously aggregates telemetry, reasons about host patterns, and automatically applies prevention policies in real time without manual intervention, for example denying lateral movement vectors on hosts exhibiting both internal discovery and suspicious external communications. Findings, supporting evidence, activated controls, and mitigation status are surfaced inside the Cato Management Application.

Review an active agentic attack finding — including evidence, controls activated, and mitigation status — as it unfolds in the Cato Management Application.

> 📍 In the Cato Management Application, navigate to Agentic Threat Prevention to view the attack timeline, supporting evidence, activated controls, and current mitigation status for the affected host.

— Rich mechanism and MITRE mapping, but no policy config surface named

Defeating the Agentic Attacker: Agentic…

## 02 Claude support via Anthropic Inference Hooks NEW

75

Extends Cato AI Security coverage to Claude Chat, Claude Code, and Claude Cowork via Anthropic Inference Hooks, requiring no endpoint agent or network inline access. Adds a `scout-debug` command to verify hook registration inside supported agents for in-workflow deployment troubleshooting, and supports pre-emptive hook configuration so

hook policy can be written before a supported agent is detected, ensuring the first session on a new device is not outside enforcement.

— Names a real command and hook mechanism but no full syntax

Inside Cato’s Latest AI Security Innova…

THINNER COVERAGE BELOW

03 **Scoped agent policy rollout via Cato Scout** NEW 52

Supports scoped rollout by applying controls to specific users or groups, with Cato Scout client identity tying policies and audit records to the correct endpoint user or device.

— Explains identity-based scoping mechanism but no exact config path

Inside Cato’s Latest AI Security Innova…

04 **MCP and skill visibility for agentic Claude** NEW 45

Adds MCP and tool visibility so teams can enumerate the MCP servers and tools available to agents and assess their combined risk surface, and adds skill visibility and posture assessment across the environment to help identify skills that require review before use in business workflows.

— Describes visibility scope but no UI path or query surface given

Inside Cato’s Latest AI Security Innova…

05 **Full-lifecycle policy enforcement for Claude interactions**

NEW 43

Enables policy enforcement across user prompt, AI response, tool input, and tool output for agentic Claude interactions.

— Names the four enforcement stages but no config or command shown

Inside Cato’s Latest AI Security Innova…

WAS THIS USEFUL?  

◆ Identity & Access Management

# Aembit

Feedback ⓘ 2 RELEASES · 2026-08-05 → 2026-08-17 BLOG ↗

Aembit is a zero-trust identity security platform that manages and authenticates machine identities across cloud infrastructure.

> Aembit extended its Workload Identity Federation credential provider ecosystem to cover both Anthropic's Claude API and OpenAI's ChatGPT API, letting workloads authenticate with short-lived federated tokens instead of static API keys.

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

01 **Claude Workload Identity Federation credential provider** NEW

98

Adds a Claude Workload Identity Federation Credential Type in the Aembit Credential Provider UI with required fields `federation_rule_id`, `organization_id`, and `service_account_id`, and optional fields `audience`, `scope` (defaults to `workspace:developer`), `workspace_id`, and `lifetime` (defaults to `3600` seconds, range `300` – `43200`). A matching `claude_wif` block is added to the Aembit Terraform Provider with the same required/optional arguments. The `EdgeController` transparently performs the federation token exchange against Claude's `POST /v1/oauth/token` OAuth endpoint and injects the resulting short-lived bearer token into outbound requests with no application code changes, replacing static `sk-ant-…` keys. Aembit also auto-generates a dedicated OIDC Issuer URL (serving `/.well-known/openid-configuration` and JWKS) for Claude Console to use when registering Aembit as a federation issuer, and this extends Aembit's existing WIF pattern (previously AWS STS, Azure Entra, GCP) so any OIDC-compliant identity source (AWS IAM, Google Cloud, Microsoft Entra ID, GitHub Actions, Kubernetes service accounts, SPIFFE, Okta) can authenticate to Claude under the same Access Policy model.

Provision a Claude WIF Credential Provider in Terraform so workloads get short-lived Claude tokens instead of static `sk-ant-…` keys.

```hcl
resource "aembit_credential_provider" "claude_wif" {
  name = "claude-wif-prod"
  claude_wif = {
    federation_rule_id = "fdrl_xxxxxxxxxxxxxxx"
    organization_id    = "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxx"
    service_account_id = "svac_xxxxxxxxxxxxxxx"
    workspace_id       = "wrkspc_xxxxxxxxxxxxxxx"
    scope              = "workspace:developer"
    lifetime           = 3600
  }
}
```

Verify that a newly configured Claude WIF Credential Provider can successfully exchange tokens with Claude's OAuth endpoint before attaching it to a policy.

📍 # In the Aembit Admin Portal: # 1. Go to Credential Providers # 2. Select the Claude WIF Credential Provider you created # 3. Click `Verify` # A green `Verified Successfully` banner confirms Aembit can reach # POST `/v1` /oauth/token and exchange tokens on behalf of workloads.

— Names exact fields, Terraform block, and OAuth endpoint with runnable example.

Aembit Adds Workload Identity Federatio…

THINNER COVERAGE BELOW

02 **OpenAI Workload Identity Federation credential provider** NEW

58

Aembit Cloud v1.33 adds an OpenAI WIF credential provider enabling secretless, keyless access to the OpenAI ChatGPT API, using short-lived scoped access tokens instead of static `sk-proj-...` API keys. The EdgeController intercepts outbound OpenAI API requests at runtime, exchanges an Aembit-issued federation token at OpenAI's token endpoint, and injects the resulting short-lived access token, requiring no application code changes, SDK modifications, or credential management logic. This extends Aembit's WIF credential provider ecosystem, which already covers AWS STS, Azure Entra, GCP, and Anthropic Claude, under a single secretless access policy model.

— Describes mechanism but lacks named config fields or exact endpoint.

OpenAI Workload Identity Federation: Ae…

WAS THIS USEFUL? 👍 👎

# Token Security

Feedback

ⓘ 2 RELEASES · 2026-08-06 BLOG ↗

Token Security is a blockchain security analysis platform that audits smart contracts and tokens for vulnerabilities, scams, and security risks.

> Token Security expanded its identity control plane with a new integration into Anthropic's Claude Compliance API for governing AI agent identities, and shipped an MCP Server Discovery capability that hunts down hidden MCP servers and embedded secrets across enterprise endpoints.

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

## 01 MCP Server Discovery for hidden AI servers and secrets NEW

83

New capability automatically retrieves MCP configuration files (e.g. `claude_desktop_config.json`, `.mcp.json` configs) from endpoints via CrowdStrike Real Time Response (RTR) without user disruption, then continuously scans and inventories MCP servers and AI agent processes across enterprise environments by integrating with CrowdStrike and other endpoint ecosystems without requiring additional agents. It analyzes CrowdStrike logs for IDE and Claude Code instances to pinpoint which MCP servers they communicate with in near real time, scans collected config files to identify hardcoded secrets (API keys, plaintext tokens, credentials), correlates them to identities, and redacts sensitive values in the platform UI. Results are shown in a Risk Graph displaying each MCP server's connections, access scope, and authentication classification.

— Names exact config file paths and CrowdStrike RTR mechanism, no CLI/API given.

Token Security's MCP Server Discovery C…

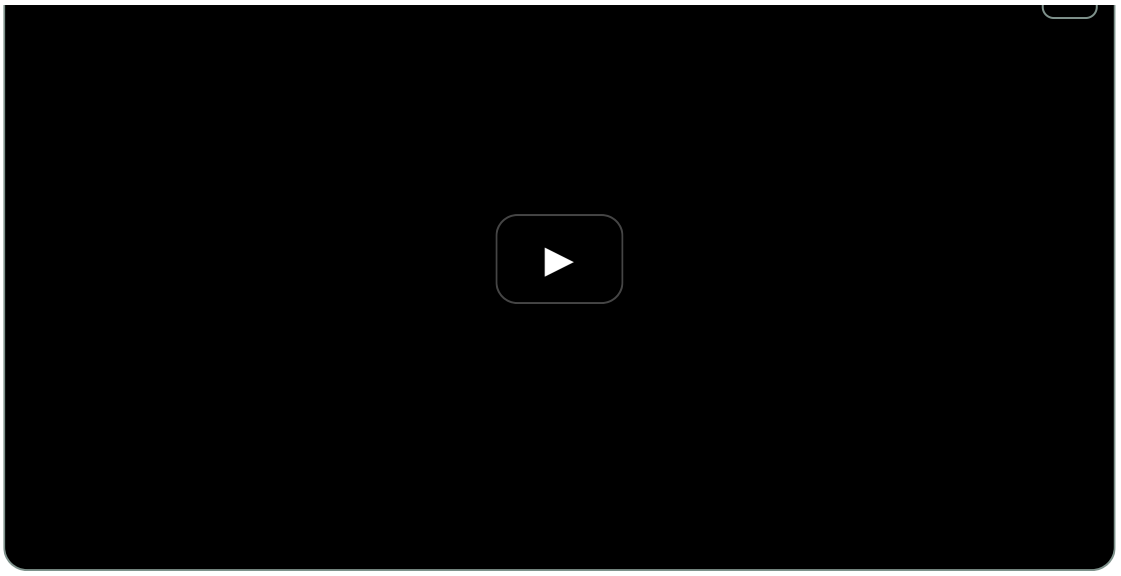
## 02 Claude Compliance API integration for Anthropic identities

NEW

72

Integrates with Claude's Compliance API to pull Anthropic identity and activity signals — API keys, workspaces, projects, users, service accounts, model activity, tool usage, administrative events, and key lifecycle events — into Token's identity control plane. Enables discovery and ownership mapping of Anthropic API keys across the full access chain (creation in Claude, storage in repository secrets, invocation by CI/CD pipelines, execution inside cloud workloads), surfaces risks like stale keys, keys without owners, abandoned workspaces, over-privileged access, and unrotated credentials, and correlates Claude activity with identity/access context from cloud, SaaS, source control, CI/CD, and MCP servers into a unified identity graph for blast-radius analysis. Claude-related identities now feed into Token's posture scoring and remediation workflows alongside other non-human and AI agent identities.





— Rich mechanism and scope but no config keys or commands to act on directly.

Token Security Brings Claude Compliance…

---

WAS THIS USEFUL? 👍 👎

# Yubico

Feedback

1 RELEASE · 2026-08-03

BLOG ↗

Yubico manufactures hardware security keys that provide two-factor authentication and passwordless login across websites and applications.

> Yubico partnered with Devolutions to let admins program vaulted credentials directly onto YubiKey hardware slots, enabling clipboard-free password injection in air-gapped environments.

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

01 **'Copy to YubiKey' credential injection in Devolutions RDM**

NEW

78

A new Copy to YubiKey action in Devolutions Remote Desktop Manager (RDM) programs a vaulted credential directly into a YubiKey hardware slot ( Slot 1 or Slot 2 ) as a static password, with the YubiKey acting as a USB HID keyboard to type it in. This allows credential injection into air-gapped server consoles and BIOS prompts where clipboard redirection is disabled, without software or network access, while credentials stay vaulted in RDM and slots are only written deliberately to preserve audit trails.

— Names mechanism, slots, and use case; lacks UI navigation steps.

The ‘Air-Gap Conundrum’: When Password …

WAS THIS USEFUL? 👍 👎

# Nudge Security

Feedback

1 RELEASE · 2026-08-23

CHANGELOG ↗

Nudge Security is a SaaS security platform that discovers, monitors, and manages cloud applications and data exposure risks.

> Nudge Security extended its browser-based AI agent discovery research preview to surface AI agents built on Datadog and HubSpot.

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

01 **AI agent discovery for Datadog and HubSpot** NEW 55

Nudge Security's browser extension can now discover AI agents built on Datadog and HubSpot as part of an AI agent discovery research preview. Access to the research preview can be requested from the AI Agents tab under Identities.

Enable browser-based AI agent discovery to surface Datadog and HubSpot agents your team has built.

> 📍 In the Nudge Security browser extension, go to Settings and turn on browser-based agent discovery. Then navigate to Identities › AI Agents to review discovered agents or request research preview access.

— Names UI path and providers but no mechanism detail. changelog-20260823-5a226f98

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

# Zama fhEVM

Feedback ⓘ 1 RELEASE · 2026-08-14 NOTES ↗

FHEVM, a full-stack framework for integrating Fully Homomorphic Encryption (FHE) with blockchain applications

> fhEVM v0.14.0 ships a unified EIP-712 decryption model, a full KMS epoch/context lifecycle with new required deployment tasks, LayerZero-based cross-chain confidential bridging, and new Solidity FHE operations, alongside ciphertext migration/attestation and coprocessor reliability work.

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

## 01 Unified EIP-712 decryption via signDecryptionPermit

BREAKING

95

Ships unified EIP-712 user decryption via signDecryptionPermit(), merging self and delegated decryption into one request model with per-handle ownership, `durationSeconds`, protocol-versioned permits, and ERC-1271 smart-account signature support. This is now the required SDK decryption permit path: `durationDays` is replaced by `durationSeconds`, permit objects carry a protocol version, several low-level SDK exports were removed in favor of high-level encrypt/decrypt/sign actions, and user-decryption authorization moves to the unified v2 flow, with legacy on-chain request shapes remaining only for the deprecation window.

Request a unified EIP-712 decryption permit from the SDK with the new per-handle, duration-in-seconds model.

```javascript
const permit = await signDecryptionPermit({ handles: [handle], durationSeconds: 3600, signer });
```

— Runnable example plus exact field rename and deprecation path v0.14.0

02 **New KMS deployment and rollout Hardhat tasks** BREAKING

85

Adds `deployProtocolConfigFromMigration`, `deployKMSGenerationFromMigration`, `upgradeKMSGeneration`, and `verify` Hardhat tasks as the new required deployment and rollout path, replacing the removed KMSGeneration Gateway migration path; deployment and rollout flows must use these tasks along with the KMS/wave1 migration verification scripts.

— Named tasks given but no exact CLI invocation v0.14.0

03 **FHE.mulDiv widened multiply-divide operation** NEW 83

Adds FHE.mulDiv(...) for widened multiply-divide operations on encrypted integers, with HCU accounting including the scalar-type check for `fheMulDiv`.

Call the new widened multiply-divide operation on encrypted integers inside a confidential Solidity contract.

```solidity
euint64 result = FHE.mulDiv(encryptedA, encryptedB,
encryptedDivisor);
```

— Runnable example and named HCU check v0.14.0

04 **KMS epoch and context lifecycle management** NEW 80

Adds `destroyKmsEpoch` to the KMS connector for epoch destruction, alongside single-in-flight enforcement, signature-checked epoch activation confirmations, and getters for context-switch monitoring. Adds per-context KMS threshold setters with caps on threshold and signer set for hardened decryption consensus, and new reinitializer versions for FHEVMExecutor, HCULimit, and `ProtocolConfig` host/gateway contracts, adding protocol storage for KMS epochs, context anchors, decryption invalidation, and confidential bridge configuration.

— Named contracts and function but no usage example v0.14.0

05 **ConfidentialBridge for cross-chain encrypted handles** NEW

78

Adds the `ConfidentialBridge` (LayerZero-based) for encrypted handle bridging across chains, including bridge events, Solidity helper APIs, deploy/prepareUpgrade tasks, LZ-delegate auto-resync to the ACL owner, and post-upgrade endpoint verification.

— Named components but no usage example given v0.14.0

06 **FHE.toExternal for all encrypted types** NEW 75

Adds FHE.toExternal(...) for all encrypted types, enabling conversion of any encrypted value to an external representation from Solidity.

Convert any encrypted value to an external representation for off-chain consumption or bridging.

```solidity
externalEuint64 ext = FHE.toExternal(encryptedValue);
```

— Runnable example included v0.14.0

07 **Ciphertext attestation and shadow verification in KMS connector**

NEW 65

Adds ciphertext attestation plumbing and shadow verification in the KMS connector, including a shared ciphertext-attestation crate, embedded ciphertext format, coprocessor registry refresh, and S3 material verification.

— Named components, no direct command v0.14.0

08 **SNS/S3 ciphertext migration support** NEW 62

Adds SNS/S3 ciphertext migration support with dedicated concurrency and DB pool controls, stalled-error backoff, and a wave1 migration verification script.

— Names controls but no exact config keys v0.14.0

09 **SDK automatic protocol/WASM version resolution** IMPROVED 62

Adds automatic protocol/WASM version resolution and bundled TFHE to the SDK, with cleartext mode, multi-version host-contracts support, and a narrower public API surface.

— Describes changes but no example call v0.14.0

10 **Coprocessor reliability and observability improvements**

IMPROVED 60

Adds full block payloads with adaptive Redis drain and stalled dependency-chain parking for improved migration and live-operation reliability, replay of quorum-accepted input proofs and drift-recovery hardening with documented quorum requirements, and latency/gap metrics for observability of block-processing performance.

— Mechanism described but no exact metric or config names v0.14.0

11 **Configurable replay start block for host-listener** NEW 60

Adds `--seed-start-block` flag to the host-listener poller for controlling the coprocessor block replay start point.

— Exact flag named but minimal mechanism v0.14.0

---

12 **Canonical protocol-config chain ID for coprocessor** BREAKING

60

The coprocessor's Ethereum chain-ID flag was renamed to the canonical protocol-config chain ID, and `host_chains` seeding is now declarative — coprocessor configuration must be updated accordingly.

— Names config change but not the exact flag names v0.14.0

└──▷ **BREAKING ON UPGRADE**

! The signDecryptionPermit() API is now the required SDK decryption permit path; `durationDays` is replaced by `durationSeconds`, permit objects carry a protocol version, and several low-level SDK exports were removed in favor of high-level encrypt/decrypt/sign actions.

! User-decryption authorization moves to the unified v2 flow; legacy on-chain request shapes remain only for the deprecation window.

! The KMSGeneration Gateway migration path was removed; deployment and rollout flows must use `deployProtocolConfigFromMigration`, `deployKMSGenerationFromMigration`, `upgradeKMSGeneration`, `verify` tasks, and the KMS/wave1 migration verification scripts.

! The coprocessor's Ethereum chain-ID flag was renamed to the canonical protocol-config chain ID; `host_chains` seeding is now declarative — coprocessor configuration must be updated accordingly.

---

WAS THIS USEFUL? 👍 👎

## MIND

Feedback

1 RELEASE · 2026-07-29

BLOG ↗

> MIND introduced a suite of AI DLP agents alongside a Model Context Protocol interface, letting teams direct policy creation, incident remediation, and override review through plain-language agent control.

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

### 01 Model Context Protocol interface for agent control NEW

60

A new Model Context Protocol (MCP) interface lets teams direct MIND agents in plain language from any MCP-compatible AI client, removing the need to switch consoles or learn query syntax.

— Names the MCP protocol but no exact client setup or command shown

MIND AI DLP agents: Your team. One miss…

THINNER COVERAGE BELOW

### 02 Policy Producer Agent for autonomous DLP policy creation NEW

45

The Policy Producer Agent autonomously creates and continuously refines DLP policies based on observed environment activity and direct instructions, aiming to eliminate manual policy engineering.

— Explains mechanism but no config, policy syntax or UI path given

MIND AI DLP agents: Your team. One miss…

### 03 Rapid Response Agent for automated remediation NEW

45

The Rapid Response Agent executes approved remediation actions immediately and escalates to human approval when required, removing queue delays in incident handling.

— Describes behavior but no trigger conditions or setup steps named

MIND AI DLP agents: Your team. One miss…

### 04 Reason Reviewer Agent for override justification review NEW

45

The Reason Reviewer Agent evaluates policy override justifications against organizational guidance in real time, replacing rubber-stamp exception reviews.

— Explains purpose but no policy source or workflow details given

MIND AI DLP agents: Your team. One miss…

WAS THIS USEFUL? 👍 👎

## GOVERN

◆ AI/LLM Security

# Aurascape

Feedback ⓘ 1 RELEASE · 2026-08-20 BLOG ↗

> Aurascape shipped the Zero-Bypass MCP Gateway, a new architecture pairing an MCP Gateway with an AI Proxy to block AI agent tool calls that skip policy enforcement, along with discovery, access control, authorization, auditing and threat-coverage capabilities built around it.

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

01 **Zero-Bypass MCP Gateway and AI Proxy enforcement** NEW 58

Introduces the Zero-Bypass MCP Gateway, which governs MCP tool access and marks every approved call so the AI Proxy can block any tool invocation that bypassed the Gateway before the model acts.

— Explains mechanism (marking approved calls, blocking bypass) but no config surface named.

Zero-Bypass MCP Gateway Data Sheet (202…

02 **OAuth 2.1 authorization for MCP connections** NEW 48

Implements OAuth 2.1 authorization with tokens scoped per user per server for MCP connections.

— Names auth standard and token scoping but no setup steps.

Zero-Bypass MCP Gateway Data Sheet (202…

03 **Tool-level access control and policy actions** NEW 46

Supports tool-level sanctioning bound to users and groups for fine-grained access control over which agents can call which MCP tools, with policy actions ranging from coaching to hold-for-approval to blocking for MCP tool calls.

— Names policy action tiers but no config keys or UI path given.

Zero-Bypass MCP Gateway Data Sheet (202…

04 **OWASP MCP Top 10 threat coverage** NEW 42

Detects and maps coverage against the OWASP MCP Top 10, including tool-poisoning attacks and credential-chaining scenarios across multiple systems.

— Names standard and two attack types but no detection mechanism.

Zero-Bypass MCP Gateway Data Sheet (202…

05 **Automatic MCP server and tool discovery** NEW 38

Adds automatic MCP server and tool discovery with automatic enrollment, eliminating manual registration of AI agent tool surfaces.

— States what it does but no mechanism or config detail.

Zero-Bypass MCP Gateway Data Sheet (202…

06 **Audit-ready records of MCP conversations** NEW 28

Generates audit-ready records of every observed MCP conversation for compliance and investigation.

— Bare description of audit capability, no format or access detail.

Zero-Bypass MCP Gateway Data Sheet (202…

WAS THIS USEFUL?  

## INFRASTRUCTURE

◆ Internet Services

# Cloudflare

Feedback ⓘ 1 RELEASE · 2026-08-23 CHANGELOG ↗

Cloudflare is a web performance and security platform providing DDoS protection, CDN, DNS, and developer tools.

> Cloudflare shipped CASB auto-remediation for SaaS file-sharing findings, richer API error context, granular OAuth scope control, and higher Browser Rendering limits, alongside dashboard and local-dev improvements for Workers.

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

01 **CASB auto-remediation policies** NEW 80

Adds CASB remediation policies in Cloudflare One to automatically remediate file-sharing findings in Microsoft 365 and Google Workspace — or send webhooks to Slack, ServiceNow, or any destination — the moment CASB detects a finding, with no manual triage.

Create a CASB remediation policy that auto-revokes external sharing on Google Workspace file-sharing findings and sends a webhook to Slack.

📍 In Cloudflare One, go to Cloud & SaaS findings › Policies, select Create a policy, enter a Policy name, set Vendor to Google Workspace and the relevant Integration and Finding type, then under 'Define what to do with findings that match your trigger' enable both Run Remediation and Send webhooks, then toggle Enable policy and select Create policy.

— Names product area, vendors, and destinations with a full setup example.

changelog-20260823-6cf4c75f

02 **Documentation link in API 403 errors** IMPROVED 80

Cloudflare API `403 Forbidden` responses now include a `documentation_url` field pointing directly to the denied endpoint's API documentation, surfacing required roles for faster self-service troubleshooting and enabling agents to drive approval workflows from error context.

Inspect a Cloudflare API 403 response to immediately identify which docs page lists the required roles for the denied endpoint.

```
$ curl -s -o response.json -w "%{http_code}" -X GET \
  "https://api.cloudflare.com/client/v4/workers/scripts" \
  -H "Authorization: Bearer <token>" && \
  cat response.json | jq '.errors[0].documentation_url'
```

— Names exact field and includes a runnable curl+jq example. changelog-20260823-6cf4c75f

03 **Optional OAuth scope controls** NEW 75

OAuth clients can now classify configured scopes as `required` or `optional` in the Cloudflare dashboard; users see Read Only and Full Access templates and a scope search on the consent screen, and can decline optional scopes to enforce least-privilege access. `wrangler login` now presents an Edit optional permissions step during authorization, letting users choose which optional OAuth scopes to grant Wrangler or the Cloudflare API MCP server instead of approving every requested scope.

— Covers both dashboard and wrangler surfaces with UI screenshots as evidence.

changelog-20260823-6cf4c75f

04 **Deployments tab for Durable Object namespaces** NEW 70

Durable Object namespaces gain a Deployments tab in the Cloudflare dashboard showing live Worker versions, traffic split, requests, error rate, and wall time per version — including actual observed traffic share alongside the configured percentage.

— Describes exact metrics shown with a screenshot but only a UI path to reach it.

changelog-20260823-6cf4c75f

05 **Soft-navigation tracking in Web Analytics** BREAKING 70

Web Analytics extends `navigationType` values to `navigate`, `soft-navigation` (Chrome Soft Navigation API), and `routing-apis` (Navigation API / History API fallback), adding native LCP measurement for SPAs on supported browsers. This may change reported pageview volume depending on front-end architecture and visitor traffic patterns.

— Names exact new field values and flags a reporting-volume side effect.

changelog-20260823-6cf4c75f

06 **Higher Browser Rendering limits** IMPROVED 65

Browser Rendering (Workers Paid plan) default limits raised: concurrent browsers 120 → 200, new browser instances/second 1 → 3, Quick Actions requests/second 10 → 30.

— Precise before/after numbers but no usage steps given. changelog-20260823-6cf4c75f

07 **FUSE Docker privileges in Miniflare** IMPROVED 65

Miniflare now automatically grants FUSE (Filesystem in Userspace) Docker privileges for local Containers development on macOS, WSL, and Linux rootless Docker when

`/dev/fuse` is available, enabling use cases such as mounting R2 buckets with FUSE.

— Names platforms and device path but no command example. changelog-20260823-6cf4c75f

└──▷ **BREAKING ON UPGRADE**

! Web Analytics may report a different volume of pageviews after the soft navigation measurement update — the extent depends on front-end architecture and visitor traffic patterns.

---

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 23, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue