THE CYBER TOOLCHAIN

NO. 095

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 5, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 141 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 32 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. Anchore turned CRA compliance into an artifact teams can actually ship, binding SBOMs, VEX and KEV-gated y to a release and exporting it for ENISA. pySigma 2.0 swapped in

vim ? help

RE2 so an untrusted Sigma rule can no longer hang your detection pipeline. Netcap cut its vulnerability index from roughly 4.1 GB to 299 MB.

GOVERN

Produce the CRA evidence pack for a release instead of assembling it by hand

Anchore's product/version model in `anchorectl` attaches images, SBOMs, VEX and KEV-gated policy to a named release and exports VDR and SBOM documents for ENISA SRP submission, with alerts when compliance status flips — so the evidence is generated continuously rather than reconstructed at audit time. Scrut closes the matching people-side gap with a structured offboarding workflow synced to HRIS/IDP, so departed-employee access revocation leaves a record instead of a Slack thread.

Anchore · Scrut Automation

DETECT

Run community Sigma rules you did not write without risking the pipeline

pySigma 2.0 defaults to the RE2 regex engine, which has no backtracking, so a crafted or merely sloppy rule regex can no longer stall a backend compile — and `SigmaPolicy` puts the trust-model knobs in one place instead of scattered flags. Velociraptor adds Azure Monitor streaming and a VQL language server, so hunt queries get authored with completion and errors caught before they hit endpoints.

pysigma · Velociraptor

PROTECT

Gate a pipeline on severity and keep the scanner's index small enough to ship

Strix's `--fail-on` turns agentic testing into a build gate rather than a report someone reads later, and `budget_policy=pause` stops a runaway agent burning spend mid-run. Netcap's move from bleve to SQLite+FTS5 drops the vulnerability index from roughly 4.1 GB to 299 MB, which is the difference between an index you can bake into a container and one you fetch at runtime; Sysdig now scores non-OS packages against the GitHub Advisory Database across 11 language ecosystems, so expect existing severity ratings to move.

Strix · Netcap

◆ AI-Powered Offensive Security

Strix

Feedback ⓘ v1.7.0 NOTES ↗ RANK

Open-source AI penetration testing tool to find and fix your app's vulnerabilities.

> Strix v1.7.0 adds CI-friendly severity gating with `--fail-on`, a `budget_policy=pause` cost safeguard, an interactive `--resume` run picker, and tightens finding evidence requirements alongside several CLI, container and TUI refinements.

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

01 **CI exit-code gating via --fail-on** NEW 85

Adds the `--fail-on` flag to the CLI to gate headless scan exit code on finding severity, letting CI pipelines fail builds when results meet or exceed a configurable severity threshold, e.g. `strix --target ./app-directory --fail-on high`.

Fail a CI job automatically when Strix finds any high- or critical-severity vulnerability during a headless scan.

```
$ strix --target ./app-directory --fail-on high
```

— Exact flag plus runnable example and clear threshold mechanism. v1.7.0

02 **STRIX_API_TYPE forces API endpoint** NEW 75

Adds the `STRIX_API_TYPE` environment variable to force use of the Responses API vs chat completions endpoint, overriding Strix's model-based auto-detection.

Force Strix to use the Responses API instead of chat completions when your provider or model requires it.

```
$ export STRIX_API_TYPE=responses
export STRIX_LLM="openai/gpt-4o"
export LLM_API_KEY="your-api-key"
strix --target ./app-directory
```

— Named env var with a runnable example, but no detail on detection logic. v1.7.0

03 **Interactive --resume run picker** NEW 70

Adds an interactive run picker shown when `--resume` is supplied with no run name, letting operators choose from prior runs directly in the terminal, e.g. `strix --resume`.

Resume a previous scan interactively when you do not remember the exact run name.

```
$ strix --resume
```

— Flag plus runnable example showing exact usage. v1.7.0

04 **budget_policy=pause cost safeguard** NEW 65

Adds `budget_policy=pause` so every agent is parked at the cost limit until the operator explicitly resumes it, preventing runaway spend.

— Named config value and behaviour, but no example of setting it. v1.7.0

05 **Stronger evidence requirements for findings** IMPROVED 65

Requires `http_exchange_ids` for proxy-validated findings, linking HTTP exchange evidence directly to reported vulnerabilities, and prompts agents to include local Git blame in the technical details of findings.

— Names the required field but not the validation workflow. v1.7.0

THINNER COVERAGE BELOW

06 **Read-only local sources as :ro bind mounts** IMPROVED 50

Makes `read_only` local sources become `:ro` bind mounts in the container runtime.

— Names the config field and resulting mount mode. v1.7.0

07 **Per-attempt provider request logging** NEW 45

Adds a structured per-attempt provider request log that records provider request IDs, giving traceability for individual LLM calls.

— No log location, format or access method given. v1.7.0

08 **Interactive scope definition in chat** NEW 35

Lets the verified user define and change scope interactively in chat.

— States capability without the chat command or UI path. v1.7.0

09 **Lazy MCP connection initialization** IMPROVED 30

Adds lazy initialization of MCP connections so connections are not established until actually needed.

— Bare description of the optimization with no further mechanism. v1.7.0

10 **Agents can delete filed vulnerability reports** NEW

30

Adds the ability for agents to delete a vulnerability report they previously filed.

— States the capability but not how deletion is triggered. v1.7.0

11 **Animated `wait_for_agents` TUI indicator** IMPROVED 20

Animates the `wait_for_agents` indicator in the TUI.

— Cosmetic change with no further detail. v1.7.0

WAS THIS USEFUL? 👍 👎

◆ **Hacking Hardware**

JustCallMeKoko LLC ESP32 Marauder v8

Feedback ⓘ

price change · Oct 5, 2026 STORE ↗ STABLE

The official handheld ESP32-C5 board for the ESP32 Marauder firmware, with a touchscreen, dual-band Wi-Fi 6, Bluetooth 5 and GPS for Wi-Fi and BLE auditing and wardriving.

> Marauder v8 `marauder-v8`
>
> **$145** ~~$100~~ ▲ $45 · 45%
>
> list price at justcallmekokollc.com · Oct 5, 2026

▾ **1 more release in this issue** · 2026-10-04 → 2026-10-05

> back in stock · Oct 4, 2026 STORE ↗ STABLE
>
> Marauder v8 (marauder-v8) back in stock at justcallmekokollc.com on 4 Oct 2026.

WAS THIS USEFUL? 👍 👎

Flipper Devices Flipper Zero

Feedback ⓘ

price change · Oct 5, 2026 STORE ↗ STABLE

A portable multi-tool for pentesters and geeks; this is its official firmware.

> Flipper Zero `flipper-zero`
>
> **$239** ~~$169~~ ▲ $70 · 41%
>
> list price at flipper.net · Oct 5, 2026

▾ **1 more release in this issue** · 2026-10-04 → 2026-10-05

> back in stock · Oct 4, 2026 STORE ↗ STABLE
>
> Flipper Zero (flipper-zero) back in stock at flipper.net on 4 Oct 2026.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

Invicti

Feedback ⓘ changelog · Oct 5, 2026 CHANGELOG ↗

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

Invicti adds a credit ledger, HTML and archive report downloads, agentic pentest platform management, scope enforcement, automatic false positive marking, and secret masking in recorders.

Adds AI Credit Usage section under the Pentests nav item to review every credit addition and deduction alongside current available balance, with pending credits reserved at assessment start.

Adds real-time credit usage display in the Summary tab for each assessment, reflecting consumed credits as the run progresses.

Enables HTML format download for assessment reports, ready to open in any browser or share with stakeholders.

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

Configures how long internal agents retain scan logs and where they store scan data.

+10 more

NEW DETECTION Ruby deserialization, JWT authentication bypass, and path traversal blocklist bypass detection

└─▷ BREAKING ON UPGRADE

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

! Credit values are now whole numbers instead of decimal values, which may affect any integrations or displays that expected decimal credit balances.

WAS THIS USEFUL? 👍 👎

HUMAN Security

Feedback ⓘ announced · Oct 5, 2026 PRESS ↗

HUMAN Security is a cybersecurity company protecting digital platforms against fraud, abuse, and malicious bot traffic.

HUMAN Ad Integrity Suite for Platforms lets DSPs, SSPs, and exchanges embed structured IVT, viewability, and brand-safety signals — with explainable reason codes and custom brand safety agents — into their own reporting and supply-quality workflows.

NEW DETECTION NEW INTEGRATION Adds HUMAN Ad Integrity Suite for Platforms, a new offering for DSPs, SSPs, and programmatic exchanges that unifies IVT, viewability, and brand-safety verification in a single architecture.

Provides explainable signals and reason codes so platform teams can see why an impression or placement received a particular classification and surface that rationale to buyers and sellers.

Delivers supply-path and inventory intelligence with granular analysis across publishers, URLs, traffic sources, and supply paths to support supply-path optimization and distinguish trusted inventory.

Introduces custom brand safety agents that move beyond keyword lists, supporting flexible, brand-specific safety categories without requiring platforms to build their own detection engines.

WAS THIS USEFUL? 👍 👎

ProjectDiscovery Katana

Feedback ⓘ v1.8.0 NOTES ↗

CODE ↗ RANK

A next-generation crawling and spidering framework.

> Katana v1.8.0 adds headless auth replay driven by recorded Chrome DevTools flows, plus new hybrid-engine extensibility points for correlating browser context and hooking into page lifecycle events.

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

01 Recorded-flow auth replay for headless crawls NEW

85

The new `-rf` / `--recorded-flow` flag replays a Chrome DevTools Recorder JSON export (or an explicit steps file) before a headless crawl starts, letting katana authenticate automatically instead of requiring manual session setup. Usage: `katana -u https://app.example.com -headless -rf ./login-flow.json`.

Replay a recorded login flow before crawling an authenticated app — record the login once in Chrome DevTools and pass the exported JSON so katana handles auth automatically.

```
$ katana -u https://app.example.com -headless -rf ./login-flow.json
```

— Named flag, file format, and runnable command given v1.8.0

THINNER COVERAGE BELOW

02 Hybrid engine extensibility: context ID and page hooks NEW

48

The hybrid crawler now exposes `BrowserContextID` so callers can correlate crawl traffic to a specific browser context, and adds optional per-page `BeforeNavigate` and `AfterLoad` hooks so custom logic can run at each page lifecycle event during a crawl.

— Names fields and hooks but no usage example v1.8.0

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Cloudsmith

Feedback ⓘ changelog · Oct 5, 2026 CHANGELOG ↗

RANK

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

> Cloudsmith added package traceability to client logs and extended CLI support to a new package format.

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

01 **Connected repositories in client log exports** NEW 55

Client log exports now include a `connected_repositories` field, listing the connected repositories a package was served through so downloads can be traced back to their source.

— Names the exact field but no usage example changelog-20261005-783fbc32

02 **Nix package support via CLI** NEW 50

The Cloudsmith CLI (v1.21.0–v1.26.0) now supports pushing Nix packages.

— Names format and CLI version range, no command shown changelog-20261005-783fbc32

WAS THIS USEFUL? 👍 👎

cve-lite-cli

Feedback ⓘ v1.38.0 NOTES ↗ RANK

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 v1.38.0 adds per-finding context signals for smarter triage, a new override-hygiene rule (with a fix making a previously dead rule actually run), and several reporting-output clarity improvements.

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

01 **Override-hygiene rule OA011** NEW 75

New override-hygiene rule `OA011` detects an override floor that exceeds every declaring package's range, joining the `OverrideRuleId` union alongside OA001–OA010.

Catch overrides pinned above every parent's declared range (OA011) alongside the full CVE scan, and emit SARIF for GitHub Code Scanning.

```
$ cve-lite /path/to/project --check-overrides --sarif
```

— Named rule id plus a runnable example command v1.38.0

02 **Contextual prioritization signals per finding** NEW 70

Adds `reachability`, `exposure`, and `usage` contextual prioritization signals to each verbose finding, surfaced separately from the EPSS `prioritySignal`. `reachability` reports `path_depth_2` and above or falls back to `direct_dependency`; `exposure` flags findings whose path runs through a request-handling package; `usage` is omitted entirely on a default scan rather than printed as `n/a`.

— Clear mechanism and field names, but no example run shown v1.38.0

03 **Reporting output disclosure improvements** IMPROVED 60

The `--all` findings table drops the `IDs` column, though advisory IDs remain available in `--verbose`, JSON, SARIF, CycloneDX, and the HTML report. The usage scan's 5000-file cap is now disclosed in the terminal, in `--json`, and in both HTML reporters, scoped per folder in multi-folder mode; a truncated scan records `unknown` for the contextual usage signal rather than claiming a package is unused. Compact output now discloses withheld findings and command groups with counts and instructions, e.g.

`Showing 3 of 53 findings. Run --verbose --all to see them all.`, instead of silently truncating.

— Names every surface changed but no single runnable example v1.38.0

THINNER COVERAGE BELOW

04 OA010 override rule now actually reachable

BREAKING

55

OA010 was previously unreachable from the scan path and no-oped silently, meaning a prior `clean` `--check-overrides` result may not have been a genuine pass. Teams gating CI on `--check-overrides` are advised to re-run it after upgrading.

— Clear remediation step but limited mechanism detail v1.38.0

└──▷ **BREAKING ON UPGRADE**

! If you gate CI on `--check-overrides`, re-run it after upgrading: OA010 was previously unreachable from the scan path and no-oped silently, so a prior `clean` result may not have been a genuine pass.

WAS THIS USEFUL? 👍 👎

Mend.io Mend Platform

Feedback ⓘ

Version 26.9.2 (04-October-2026) NOTES ↗ RANK

Traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend.io's window centered on finer-grained Mend AI triage, a new scan-independent email notification workflow with dedicated zero-day alerts, a doubled Mend Container scan size limit, plus incremental CLI SAST and Renovate manager improvements.

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

01 CLI scan output and exclusion controls in Mend SAST IMPROVED

90

The `mend code` CLI (v26.9.1) gains a `--print-to-console` flag to tail SAST scan output live in CI logs without reading a separate log file. Default exclusions now automatically ignore `/_MACOSX/..*` metadata files across all SAST scan engines, preventing macOS system files from interfering with the TreeSitter parser on Windows.

Tail SAST scan output live in CI logs without separately reading a log file — useful in pipelines where log file access is inconvenient.

```
$ mend code --print-to-console
```

— Names exact flag and exclusion pattern with runnable command. product docs

02 Send Email action in Automation Workflows NEW

75

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

Receive immediate notification when a newly published CVE affects your organization, without waiting for the next scheduled scan.

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

Set up a scan-independent email alert so your team is notified immediately when a newly published CVE affects your inventory, without waiting for the next scan.

📍 In the Mend Platform, go to Administration › Automation Workflows, create or edit a workflow, add a new action of type `Send Email`, select the target users or groups, configure the delivery frequency, and set the trigger to `New Findings` (e.g. newly published CVEs).

— Exact UI navigation steps given, no API/config key. Version 26.9.2 (04-October-2026)

03 **Image scan size limit increased to 4 GB in Mend Container**

IMPROVED 70

Raises the image scan size limit for Mend Container from 2 GB to 4 GB, enabling pulling, extraction, and scanning of larger container images across all registry integrations including ECR, with a clear error message for images exceeding 4 GB.

— Names exact before/after limit and registry scope, no command needed. product docs

THINNER COVERAGE BELOW

04 **Rich version constraints for Gradle manager in Renovate** NEW 45

Supports rich version constraints in the Groovy/Kotlin DSL for the `gradle` manager.

— Names manager and DSL but gives no example syntax. product docs

05 **Azure PR auto-complete re-application in Renovate** IMPROVED 45

Re-applies auto-complete settings to existing PRs on the `azure` platform.

— Names the platform and behaviour, no further mechanism. product docs

06 **Finding-level Suppress and Review triage in Mend AI** IMPROVED 35

Extends the Suppress and Review triage actions to the individual finding level, allowing security teams to triage specific findings within an entity independently rather than only at the whole-entity level.

— Describes behaviour change but no UI path or API named. product docs

07 **Dedicated zero-day vulnerability alerts** NEW 25

Adds dedicated zero-day alerts to provide immediate awareness when an organization is exposed to emerging threats.

— No mechanism or trigger details given beyond the name. Version 26.9.2 (04-October-2026)

08 **Improved vulnerability fix PR raising in Renovate** IMPROVED 10

Renovate improves how vulnerability fix PRs are raised, per the release summary, though no further mechanism or scope is described.

— Only a vague summary mention, no mechanism given. product docs

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

Velocidex Velociraptor

Feedback ⓘ v0.77.3 NOTES ↗

CODE ↗ RANK

Velociraptor is an endpoint visibility and collection tool that gathers host-based state information using the Velociraptor Query Language (VQL).

> Velociraptor v0.77.3 adds Azure Monitor log streaming, a VQL language server, ARM64 Windows builds, and new inventory and logging configuration options, while tightening backup output and dropping x86 Windows support.

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

01 **Azure Monitor upload integration** NEW 78

Adds the azure_monitor_upload() VQL plugin and `Azure.Monitor.Upload` server-event artifact to stream completed-flow results to an Azure Log Analytics workspace via the Logs Ingestion API, supporting service principal, `AZURE_*` environment variables, managed identity, and the Azure SDK default credential chain. A new Azure Monitor Creds secret type stores Log Analytics credentials and destination in the secrets service.

– Names plugin, artifact and secret type but no sample config shown v0.77.3

02 **VQL Language Server Protocol support** NEW 65

Adds a VQL Language Server Protocol (LSP) server, invoked via `velociraptor lsp`, providing diagnostics, syntax highlighting, hover, completions, and document symbols for VQL artifact development.

– Exact command given but no editor integration steps v0.77.3

03 **Inventory download transforms** NEW 60

Adds a `download_transforms` option to the inventory service, currently supporting `gunzip` to decompress gzip-compressed files before storing them in the inventory.

– Names config option and transform but no usage example v0.77.3

04 **Backup output restricted to client/flow fields** BREAKING

60

`Server.Utils.BackupDirectory`, `Server.Utils.BackupGCS`, and `Server.Utils.BackupS3` now emit only `ClientId`, `FlowId`, and fields derived from them rather than all columns — downstream consumers reading any other column will break and must be updated.

— States before/after but no migration path given v0.77.3

THINNER COVERAGE BELOW

05 **ARM64 Windows build target** BREAKING

50

Adds an ARM64 Windows build target, with MSI and offline collectors updated to support ARM64; x86 support has been removed from these collectors.

— Describes change but no migration guidance for x86 users v0.77.3

└──▷ **BREAKING ON UPGRADE**

! `Server.Utils.BackupDirectory`, `Server.Utils.BackupGCS`, and `Server.Utils.BackupS3` now emit only `ClientId`, `FlowId`, and derived fields — downstream consumers reading any other column will break and must be updated.

! The ARM64 Windows MSI and offline collectors no longer include x86 support.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

Netcap

Feedback ⓘ v0.10.0 NOTES ↗ CODE ↗ RANK

A framework for secure and scalable network traffic analysis - https://netcap.io

> Netcap v0.10.0 replaces its bleve-based vulnerability indexes with a single SQLite+FTS5 database, cutting disk usage from roughly 4.1 GB to 299 MB and requiring an upgrade step to fetch the new format.

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

01 SQLite+FTS5 vulnerability database replaces bleve indexes

95

BREAKING

Vulnerability and exploit lookups move from the `nvd.bleve` and `exploit-db.bleve` indexes to a single `netcap.sqlite` (SQLite + FTS5) database, shrinking disk usage from ~4.1 GB to 299 MB. After upgrading, users must run `net util -download-dbs` to fetch the new layout 2 database from `https://dbs.netcap.io/dbs/v2/` with sha256 verification, since the old `https://dbs.netcap.io/dbs/latest` endpoint now serves a frozen bleve revision and is no longer rebuilt. `net util -index` no longer calls `IndexData`; it now calls the new `BuildVulnDB` path and rebuilds `netcap.sqlite` from NVD and exploit-db feeds regardless of the index name passed (e.g. `-index all`). Query results are now capped at 10 hits ordered by BM25 score, replacing bleve's score-threshold filtering.

After upgrading to v0.10, download and verify the new SQLite vulnerability database before running any vuln lookups.

```
$ net util -download-dbs
```

Rebuild `netcap.sqlite` locally from downloaded NVD and exploit-db feeds, for air-gapped or custom database workflows.

```
$ net util -index all
```

— Names every command, file, URL, and the exact disk/result-count changes. v0.10.0

└──▷ **BREAKING ON UPGRADE**

! Vulnerability database format changed: `nvd.bleve` and `exploit-db.bleve` are no longer read; you must run `net util -download-dbs` after upgrading to fetch the new `netcap.sqlite` (layout 2) from `https://dbs.netcap.io/dbs/v2/`.

! The `https://dbs.netcap.io/dbs/latest` endpoint now serves a frozen bleve revision and is no longer rebuilt; netcap v0.10+ must use the v2 URL.

! `net util -index` no longer calls `IndexData`; it now calls `BuildVulnDB` and rebuilds `netcap.sqlite` — passing a specific index name has no effect; any value (e.g. `-index all`) rebuilds both NVD and exploit-db.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ Preferences and dashboards over the API:

NOTES ↗ RANK

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

> Plerion added API access to manage user preferences and dashboards, letting teams programmatically control home dashboards and sidebar navigation across users, groups, and tenants.

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

01 **Preferences and dashboards API** NEW 46

New API key-authenticated endpoints let callers set a user's or user group's home dashboard and sidebar navigation, set the tenant default home dashboard, and list custom dashboards.

– Describes capabilities but no exact endpoint paths or methods given

Preferences and dashboards over the API:

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Mend.io Mend Container

Feedback ⓘ

Version 26.9.2 (04-October-2026) NOTES ↗ RANK

Traditional application security misses AI risk. Standalone AI security tools miss your code.

> Mend Container added configurable email notifications and zero-day threat alerts to its workflow engine, and raised the container image scan size limit from 2 GB to 4 GB.

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

01 **Image scan size limit raised to 4 GB** IMPROVED 70

Raises the container image scan size limit from 2 GB to 4 GB across all registry integrations, including ECR, enabling successful pulling, extraction, and scanning of larger images. Surfaces a clear error message when a scanned image exceeds the new 4 GB limit, instead of failing silently.

— Names exact before/after limit, scope, and error behaviour

Version 26.9.2 (04-October-2026)

THINNER COVERAGE BELOW

02 **Send Email workflow action** NEW 55

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

— Describes behaviour and configuration but no UI path or API name product docs

03 **Zero-day vulnerability alerts** NEW 20

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

— Bare description with no mechanism or scope detail product docs

WAS THIS USEFUL? 👍 👎

Anchore

Feedback ⓘ

What EU CRA’s 24-hour deadline means for your SBOM (2026-09-29) BLOG ↗ RANK

Anchore's software supply chain solutions automate vulnerability scanning, strengthen container security, and support compliance with NIST, FedRAMP & more.

> Anchore Enterprise ships an EU CRA quickstart built around a new product/version model in anchorectl: teams can now attach images, SBOMs, VEX data and KEV-gated policies to a release, get notified the instant compliance status shifts, and export VDR and SBOM documents for ENISA SRP submission.

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

01 CRA policy attachment and continuous evaluation NEW

93

Adds `anchorectl app update --policy-id` and `anchorectl app version policy status get --fail-based-on-results` to attach and evaluate a CRA-focused policy (blocking on any KEV-listed finding) at the product level in CI. Anchore also continuously re-evaluates every stored SBOM against updated vulnerability data and policy automatically, flagging a newly non-compliant release the moment underlying data changes rather than waiting for the next scheduled scan.

Attach a live container image and a supplier-delivered CycloneDX SBOM as assets to a release version, then enforce a KEV-blocking CRA policy in CI so the build fails if any exploited vulnerability is present.

```
$ anchorectl app version asset add container-image-remote
registry.example.com/iot-gateway/web-console:2.4.0 \
  --app iot-gateway --version 2.4.0 --asset web-console --wait

anchorectl app version asset add sbom telemetry-module-
1.8.0.cdx.json \
  --app iot-gateway --version 2.4.0 --asset telemetry-module -t
device --wait

anchorectl app update iot-gateway --policy-id eu-cra-baseline
anchorectl app version policy status get 2.4.0 --app iot-gateway
--fail-based-on-results
```

— Names flags, CI behavior, and before/after evaluation timing

What EU CRA’s 24-hour deadline means fo…

02 **VDR export for ENISA SRP submission** NEW 88

Adds `anchorectl app version export vdr` subcommand (with `-f` for output file) to generate a Vulnerability Disclosure Report bundling product identity, component inventory, and VEX analysis into a single CycloneDX document suitable for ENISA SRP submission.

When a KEV match triggers a `policy_status_changed` alert, record a VEX assessment and export a Vulnerability Disclosure Report ready for ENISA SRP submission — all before the 24-hour CRA clock expires.

```
$ anchorectl app version vex add 2.4.0 --app iot-gateway \
  --vuln-id CVE-2022-22965 \
  --pkg-name spring-beans \
  --pkg-type java-archive \
  --pkg-version 5.3.17 \
  --status affected \
  --action-statement "Security update 2.4.1 ships the supplier's
fixed component."

anchorectl app version export vdr 2.4.0 --app iot-gateway -f iot-
gateway-2.4.0-vdr.cdx.json
```

— Exact command, flag, and output format specified

What EU CRA's 24-hour deadline means fo…

03 **Engineering VEX annotations on releases** NEW 85

Adds `anchorectl app version vex add` subcommand with `--vuln-id`, `--pkg-name`, `--pkg-type`, `--pkg-version`, `--status`, and `--action-statement` flags to record VEX annotations against a specific release version.

When a KEV match triggers a `policy_status_changed` alert, record a VEX assessment and export a Vulnerability Disclosure Report ready for ENISA SRP submission — all before the 24-hour CRA clock expires.

```
$ anchorectl app version vex add 2.4.0 --app iot-gateway \
  --vuln-id CVE-2022-22965 \
  --pkg-name spring-beans \
  --pkg-type java-archive \
  --pkg-version 5.3.17 \
  --status affected \
  --action-statement "Security update 2.4.1 ships the supplier's
fixed component."

anchorectl app version export vdr 2.4.0 --app iot-gateway -f iot-
gateway-2.4.0-vdr.cdx.json
```

— Every flag named alongside a runnable example

What EU CRA's 24-hour deadline means fo…

04 **Asset attachment to product versions** NEW 83

Adds `anchorectl app version asset add container-image-remote` and `anchorectl app version asset add sbom` subcommands to attach heterogeneous assets — direct container images or supplier-delivered SBOMs — to a versioned product release.

Attach a live container image and a supplier-delivered CycloneDX SBOM as assets to a release version, then enforce a KEV-blocking CRA policy in CI so the build fails if any exploited vulnerability is present.

```
$ anchorectl app version asset add container-image-remote
registry.example.com/iot-gateway/web-console:2.4.0 \
  --app iot-gateway --version 2.4.0 --asset web-console --wait

anchorectl app version asset add sbom telemetry-module-
1.8.0.cdx.json \
  --app iot-gateway --version 2.4.0 --asset telemetry-module -t
device --wait

anchorectl app update iot-gateway --policy-id eu-cra-baseline
anchorectl app version policy status get 2.4.0 --app iot-gateway
--fail-based-on-results
```

— Two named subcommands demonstrated in a working example

What EU CRA's 24-hour deadline means fo…

05 **Product and release version modeling via anchorectl** NEW 80

Adds `anchorectl app add` and `anchorectl app version add` subcommands to model a product as an App and each release as a Version, enabling product-level (not scan-level) compliance tracking.

Model a product and released version in Anchore Enterprise so every subsequent scan, policy evaluation, and VDR export is scoped to that release rather than an ad-hoc image scan.

```
$ anchorectl app add iot-gateway \
  --contact-name "Product Security Officer" \
  --contact-email psirt@example.com \
  --description "Industrial IoT gateway - product with digital
elements"

anchorectl app version add 2.4.0 --app iot-gateway --status
released --release-date 2026-06-15
```

— Named commands with a runnable example, but modeling mechanism is brief

What EU CRA's 24-hour deadline means fo…

THINNER COVERAGE BELOW

06 **Policy status change notifications** NEW 55

Emits `policy_status_changed` notifications at the Application Version level whenever a release's compliance outcome shifts, with subscription endpoints including Webhooks, Slack, Teams, Jira, GitHub, and email.

— Event name and channels listed but no config surface shown

What EU CRA's 24-hour deadline means fo…

07 **Unified SBOM export in CycloneDX and SPDX** NEW 40

Exports a unified SBOM per release version in CycloneDX and SPDX formats once all assets are attached.

— Formats named but no command or trigger given What EU CRA's 24-hour deadline means fo…

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ Secure SaaS - October 05, 2026

BLOG ↗ RANK

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

> Sysdig Secure introduced a Technical Preview Supply Chain Findings page for license and end-of-life tracking, surfaced CVE lookup and AI-based threat evaluation workflows, and switched non-OS package vulnerability scoring to the GitHub Advisory Database across 11 language ecosystems — a change that can shift existing severity ratings.

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

01 GitHub Advisory Database CVSS sourcing for non-OS packages

78

BREAKING

Sysdig Secure now sources severity and CVSS scores for non-OS package vulnerabilities from the GitHub Advisory Database, falling back to NVD when GitHub has no score, covering Java, JavaScript, Python, Go, Rust, PHP, Ruby, .NET, Erlang, Dart, and Swift packages. After upgrade, severity ratings on existing non-OS package findings may change (commonly from High to Medium), meaning vulnerability policies that trigger on High or Critical may pass or fail on different packages following the next scan.

— Names data sources, 11 ecosystems and concrete before/after impact

Secure SaaS - October 05, 2026

02 Supply Chain Findings page (Technical Preview) NEW

62

A new Supply Chain Findings page lists license and end-of-life findings across pipeline, registry, and runtime resources. Findings can be filtered by severity, stage, finding type, license category, or license name, and grouped by package, image, or runtime resource.

— Names filters and grouping options but no exact navigation path

changelog-20261005-e0b53142

└──▷ **BREAKING ON UPGRADE**

! Severity ratings on existing non-OS package findings may change (commonly from High to Medium) after upgrade; vulnerability policies that trigger on High or Critical may now pass or fail on different packages following the next scan.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

CrowdStrike

Feedback ⓘ changelog · Oct 5, 2026 API ↗ RANK

CrowdStrike provides cloud-based endpoint security and threat intelligence services protecting organizations from cyberattacks and malware.

> CrowdStrike expanded its developer API surface with new endpoint groups for prevention policy management and audit event querying.

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

01 **Audit API group for querying audit events** NEW 50

Adds an Audit API group with 86 operations for querying audit events.

— Gives operation count but no specific endpoints or usage. 2026-10-05-22053f66

WAS THIS USEFUL? 👍 👎

Huntress

Feedback

SIEM Log Sources API is now available

NOTES ↗

RANK

Give feedback to the Huntress team so we can make more informed product decisions. Powered by Canny.

> Huntress shipped a new SIEM Log Sources API for programmatic coverage audits and added manual-only Microsoft 365 hardening controls that are excluded from automatic Managed Deployments.

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

01 **SIEM Log Sources API** NEW 70

A new SIEM Log Sources API endpoint returns each source's name, type, organization, agent, status, last activity, and 30-day data volume — the same data shown on the Source Management page — letting partners audit coverage and detect silent sources without opening the portal.

— Describes returned fields and use case but no exact method/path.

SIEM Log Sources API is now available

WAS THIS USEFUL?

◆ Threat Intelligence

Flashpoint

Feedback ⓘ changelog · Oct 5, 2026 CHANGELOG ↗

RANK

Flashpoint empowers the world's leading organizations with the best data for threat intelligence to protect people, places, and assets.

> Flashpoint shipped a cluster of major additions this window: an External Attack Surface Management module in Ignite, a full Managed Attribution redesign with a built-in Anonymous Browser, and a new MCP Server, alongside Natural Language Search in Echosec, an ATM Receipts Dashboard for fraud analysis, AI-summarized alerting, and a streamlined Ignite-to-Managed Attribution analysis workflow.

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

01 On-demand Exit Location change in Managed Attribution

IMPROVED 70

Managed Attribution virtual machines support on-demand Exit Location changes without interrupting workflow: open the active VM session controls, select Exit Location (or Change Exit Location), choose a new country or region, and confirm — the session continues uninterrupted.

Change your virtual machine's Exit Location in Managed Attribution mid-session to maintain anonymity without restarting your workflow.

> 📍 In Flashpoint Managed Attribution, open the virtual machine session controls and select Change Exit Location, choose the desired country or region, and confirm — the session continues without interruption.

Switch your Managed Attribution Exit Location mid-session to mask attribution without dropping your current workflow.

> 📍 In Flashpoint Managed Attribution, open your active virtual machine session, select Exit Location, choose the desired new exit country or region, and confirm — the session continues without interruption.

— Exact UI steps given for mid-session change changelog-20261005-45de6c42

02 Ignite-to-Managed Attribution analysis workflow NEW

70

A streamlined workflow lets analysts transition directly from discovering a threat in Flashpoint Ignite to safely analyzing it within Managed Attribution: locate a threat item, then select `Analyze in Managed Attribution` to open it inside the built-in Anonymous Browser with a chosen Exit Location.

Pivot from a threat discovered in Ignite directly into Managed Attribution's Anonymous Browser for safe, attribution-controlled analysis.

> 📍 In Flashpoint Ignite, locate a threat item of interest, then select `Analyze in Managed Attribution` to open it inside the built-in Anonymous Browser with a chosen Exit Location.

— Exact UI steps including button name given `changelog-20261005-45de6c42`

---

03 **ATM Receipts Dashboard in Fraud Intelligence** NEW 65

The ATM Receipts Dashboard in Flashpoint Fraud Intelligence analyzes compromised ATM receipt images, exposing card and account data, deposits, and fraud trends; it is reached via Ignite > Fraud Intelligence > ATM Receipts Dashboard, with filters by date range or card type.

Navigate to the ATM Receipts Dashboard in Flashpoint Fraud Intelligence to analyze compromised ATM receipt images and surface exposed card data and fraud trends at scale.

> 📍 In Flashpoint Ignite, go to Fraud Intelligence › ATM Receipts Dashboard, then filter by date range or card type to review compromised receipt images and detected account data.

Detect exposed card and account data from compromised ATM receipt images surfaced in illicit sources.

> 📍 In Flashpoint Fraud Intelligence, navigate to the ATM Receipts Dashboard to review flagged receipt images, filter by date or fraud trend, and identify exposed card numbers, account data, and deposit patterns at scale.

— Exact navigation path given, mechanism described `changelog-20261005-45de6c42`

---

04 **Natural Language Search and AI Results Refinement in Echosec**

NEW 60

Echosec's `Locate` view adds Natural Language Search, letting analysts type plain-language geospatial queries instead of boolean strings, paired with AI Results Refinement to narrow the returned posts/signals for geospatial OSINT discovery.

Use Natural Language Search in Echosec to find geospatial intelligence without crafting boolean queries.

📍 In Echosec, open the Locate view, type a plain-language query such as 'protests near central station last 48 hours', then use AI Results Refinement to narrow the returned posts.

Use Natural Language Search in Echosec's Locate view to find geospatial signals without crafting boolean queries.

📍 In Echosec, open the Locate view, select Natural Language Search, type a plain-English query such as 'protests near government buildings in Bogota', then use AI Results Refinement to filter the returned signals.

Use Natural Language Search in Echosec's Locate view to find geospatial intelligence without crafting keyword queries manually.

📍 In Echosec, go to the Locate view, then enter a natural language query such as 'protests near the central train station last week' into the search bar and review AI Results Refinement suggestions to narrow the results.

— UI path and query example given, no API detail `changelog-20261005-45de6c42`

THINNER COVERAGE BELOW

05 **Optimize AI-powered keyword suggestions** NEW 55

Optimize is an AI-powered feature that automatically suggests and inserts location-based keywords — business names, street names, landmarks — into real-time location-based searches.

— Named feature with concrete keyword types, no UI path `changelog-20261005-45de6c42`

06 **AI Summarization for Alerting in Ignite** NEW 45

Flashpoint Ignite adds AI Summarization for Alerting, delivering structured, reference-backed summaries for Community alerts via in-platform view or daily email briefings.

— Describes output channels but no config steps `changelog-20261005-45de6c42`

07 **External Attack Surface Management module in Ignite** NEW 40

Flashpoint launches External Attack Surface Management (EASM) as an add-on module in Ignite, providing a continuous attacker's-eye view of the perimeter.

— Describes module purpose but no setup or config detail `changelog-20261005-45de6c42`

08 **Managed Attribution redesign with Anonymous Browser** NEW 40

Flashpoint ships a full product redesign of Managed Attribution with a built-in Anonymous Browser capability for safely analyzing threats without leaving the platform.

— Named feature but no usage steps for Anonymous Browser itself

changelog-20261005-45de6c42

09 **Filter Sharing in Echosec** NEW 25

Echosec introduces Filter Sharing to synchronize teams and reduce noise in large datasets.

— One-line description, no mechanism or UI path changelog-20261005-45de6c42

10 **Flashpoint MCP Server** NEW 15

Flashpoint introduces an MCP Server.

— Name only, no mechanism or usage described changelog-20261005-45de6c42

11 **Intelligence Requirements in Ignite** NEW 15

Flashpoint Ignite introduces Intelligence Requirements.

— Name only, no further description provided product docs

WAS THIS USEFUL? 👍 👎

◆ SIEM

Google SecOps

Feedback ⓘ October 05, 2026 (2026-10-05) BLOG ↗

RANK

Meet your business challenges head on with AI and cloud computing services from Google, including security, data management, and hybrid & multi-cloud.

> Google SecOps added Terraform provider and Google Cloud client library support, letting teams manage and automate core Chronicle resources as code.

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

01 **Terraform and client library support for Chronicle API resources**

NEW 73

Google SecOps now offers a Terraform provider and Google Cloud client library for managing and automating Google SecOps resources programmatically across the Chronicle API, covering Feeds, Parsers, Rules, Findings Refinements, Retrohunts, Dashboards, Data Tables, Watchlists, Custom Lists, Environments, Environment Groups, Data Export, SOAR Domains, and SOAR Networks.

— Names all 14 resource types but gives no commands or exact config keys.

October 05, 2026 (2026-10-05)

WAS THIS USEFUL? 👍 👎

Elastic Security Labs

Feedback ⓘ

Behind the tags: How Elastic SIEM grades 1,781 detection rules on noise, speed, and threa

BLOG ↗ RANK

Power insights and outcomes with The Elasticsearch Platform. See into your data and find answers that matter with enterprise solutions designed to help you accelerate time to insight.

> Elastic Security 9.5 ships the AlertZero suite, adding an AI-driven alert analysis workflow with auto-closure of false positives and deeper integration of Attack Discovery into existing SOC workflows.

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

01 **Security alert analysis workflow with auto-closure** NEW

65

A new Security alert analysis workflow in Elastic Security 9.5 performs a first-pass classification of alerts as true positive, false positive, or inconclusive, attaching a confidence score and a rationale note; it applies to detection rules selected by a SIEM engineer or SOC manager. High-confidence false positives can be auto-closed within this workflow, reducing the volume of alerts that reach human review before they hit the analyst queue.

— Names confidence score and roles but no config key or API. product docs

THINNER COVERAGE BELOW

02 **Attack Discovery integration and workflow embedding** NEW

50

Security alert analysis results now feed Attack Discovery to give it a cleaner, noise-reduced starting point for deeper attack investigation. Attack Discovery can also be embedded inside existing Elastic Workflows for enrichment, escalation, and case creation, letting AI-driven investigation plug into established SOC playbooks without a separate implementation.

— Describes integration and embedding but no concrete surface named. product docs

UPDATED DETECTION Monthly tag refresh for prebuilt detection rules — Noise, Performance, Threat, and Profile

UPDATED DETECTION Standardised prebuilt-rule tag taxonomy — Domain, Platform, Data Source, Service, Vuln, Resources

NEW DETECTION Planned MITRE ATLAS tags for prebuilt rules

Behind the tags: How Elastic SIEM grades 1,781 detection rules on noise, speed, and th

BLOG ↗ LATEST STABLE

Elastic SIEM introduces monthly-recalculated Noise, Performance, Threat, and Profile tags on prebuilt rules, scored from real fleet telemetry via an automated pipeline with human review.

An automated monthly pipeline scores every prebuilt rule across `Noise`, `Performance`, `Threat`, and `Profile` dimensions from a 30-day telemetry window and opens a draft PR on `elastic/detection-rules` for human review before merging.

`Profile: Recommended` and `Profile: Aggressive` synthesise severity, noise, performance, and threat coverage into a single enablement-posture label, with hard gates preventing `Noise: High` rules from ever being `Recommended`.

LLMs are used only as a fallback for `Threat` tag classification where heuristics find no match, capped at five tags per rule and filtered against a human-maintained catalog; manually curated tags (e.g. `Threat: Cobalt Strike`) are never overwritten by the pipeline.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

pysigma

Feedback ⓘ v2.0.0 NOTES ↗ RANK

Python library to parse and convert Sigma rules into queries (and whatever else you could imagine)

> pySigma 2.0 is a major security-hardening release, centralizing all trust-model settings into a new `SigmaPolicy` dataclass and switching the default regex engine to RE2 to eliminate ReDoS from untrusted Sigma rules.

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

01 SigmaPolicy for centralized security controls

BREAKING

95

Introduces the `SigmaPolicy` dataclass (importable from `sigma.policy`) to declare trust-model settings — `allow_template_vars`, `allow_external_sources`, `vars_allowed_paths`, `regex_engine` — once, with those settings propagating automatically through every pipeline call. Adds predefined `SafePolicy` (RE2 engine, all security features disabled) and `TrustedPolicy` (Python engine, fully trusted sources) profiles, passed via a new `policy` parameter on ProcessingPipeline.from_yaml() and ProcessingPipeline.from_dict(), replacing the prior individual-flag parameters. Legacy `PYSIGMA_ALLOW_VARS_EXECUTION` and `PYSIGMA_ALLOW_EXTERNAL_SOURCES` environment variables remain as a fallback for security settings.

Harden a public-facing application that accepts user-uploaded Sigma rules — use `SafePolicy` to block ReDoS and disable all external-source features in one line.

```python
from sigma.policy import SafePolicy
from sigma.processing.pipeline import ProcessingPipeline

pipeline = ProcessingPipeline.from_yaml(yaml_string,
policy=SafePolicy())
```

Enable template variables and external sources for a private CI/CD pipeline where you fully control the rule sources, using `SigmaPolicy` for explicit opt-in.

```python
from sigma.policy import SigmaPolicy
from sigma.policy.regex_engine import PythonRegexEngine
from sigma.processing.pipeline import ProcessingPipeline

policy = SigmaPolicy(
    allow_template_vars=True,
    allow_external_sources=True,
    vars_allowed_paths=("/trusted/rules",),
    regex_engine=PythonRegexEngine()
)
pipeline = ProcessingPipeline.from_yaml(yaml_string,
policy=policy)
```

— Names exact dataclass, fields, profiles, parameter, with runnable code v2.0.0

02 **RE2 regex engine as new default** BREAKING 75

Switches the default regex engine from Python's `re` to `RE2RegexEngine` (Google RE2) to eliminate ReDoS from untrusted Sigma rules; `PythonRegexEngine` remains available for trusted rule sources by setting `regex_engine` via `SigmaPolicy`. Rules using regex features unsupported by RE2 will break unless `PythonRegexEngine` is explicitly selected.

— Names exact engine classes and migration path, no runnable example v2.0.0

THINNER COVERAGE BELOW

03 **Sandboxed template processing** IMPROVED 55

Sandboxes template processing so that only safe built-in types (`str`, `int`, `dict`, `list`, etc.) can have methods called — arbitrary object methods are blocked.

— Names allowed types but no way to configure it v2.0.0

04 **Security key stripping in pipeline YAML** BREAKING 50

Strips security-related YAML and dictionary keys before pipeline instantiation, blocking runtime opt-in of security features through untrusted pipeline definitions; unknown security-related keys in top-level pipeline YAML now raise an error rather than being silently ignored, so pipelines with inline security settings will fail to load.

— Describes behavior change but no specific key names given v2.0.0

05 **Correlation and pipeline loading additions** NEW 30

Supports fractional thresholds in correlation condition definitions, enables loading nested query post-processing items from YAML, and validates backend compatibility

when loading YAML pipelines.

— Three one-line additions with no concrete mechanism given `v2.0.0`

└──▷ **BREAKING ON UPGRADE**

! `allow_template_vars`, `allow_external_sources`, `vars_allowed_paths`, and `regex_engine` are no longer accepted as individual parameters on ProcessingPipeline.from_yaml() and from_dict(); they must now be passed via a `SigmaPolicy` instance using the new `policy` parameter.

! Security-related keys in top-level pipeline YAML are now rejected (unknown keys raise an error rather than being silently ignored), so any pipeline YAML that contained inline security settings will fail to load.

! The default regex engine changes from Python's `re` to `RE2RegexEngine`; rules using regex features unsupported by RE2 will break unless `PythonRegexEngine` is explicitly selected via `SigmaPolicy`.

! Features previously enabled by default (template vars, external sources) now default to disabled; code that relied on those defaults must explicitly opt in through `SigmaPolicy`.

WAS THIS USEFUL?  

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Cloudmersive

Feedback ⓘ

Cloudmersive Private Cloud DLP AI API v2-10-0 Released (2026-10-04) BLOG ↗

RANK

Cloudmersive is a cloud API platform providing developer tools for document processing, data validation, conversion, and security tasks.

> Cloudmersive's Private Cloud DLP AI API now handles audio redaction batch jobs whose output files exceed 2 GB.

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

01 **Large output support for audio redaction batch jobs** IMPROVED

40

The Private Cloud DLP AI API now supports audio redaction batch job responses where the output file exceeds 2 GB in size.

— Names the exact limit but no endpoint or usage detail

Cloudmersive Private Cloud DLP AI API v…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

iDenfy

Feedback ⓘ changelog · Oct 5, 2026 CHANGELOG ↗ RANK

Explore iDenfy documentation for identity verification (KYC), business verification (KYB), AML screening, fraud prevention, and more.

> iDenfy's Sprint 136 release introduces an MCP integration that lets AI assistants like Claude query account data, lets Shopify merchants buy age estimation directly in-store, and brings the verification flow into European Accessibility Act compliance, alongside smaller additions such as company-verification risk scoring, a new dashboard upsell menu, Spanish KYB support, and Brazil CPF date-of-birth capture.

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

01 **MCP integration for AI assistants** NEW 75

Exposes iDenfy account data over MCP (Model Context Protocol) so AI assistants such as Claude and Claude Code can connect to an account, with a consent screen letting the user select which data areas (e.g. verifications, company risk) the assistant may view or edit.

Let Claude query your iDenfy account data — such as verification sessions or risk results — without leaving the AI assistant interface.

> 📍 In Claude or Claude Code, connect a new MCP integration, supply your iDenfy credentials, and use the consent screen to select which data areas (e.g. verifications, company risk) the assistant may view or edit.

— Names protocol and assistants and gives connection steps, but no API endpoint

changelog-20261005-7e21bf07

THINNER COVERAGE BELOW

02 **Shopify age-estimation signup and billing** NEW 50

Enables Shopify merchants to sign up for and pay for age estimation directly within Shopify.

— Names the Shopify integration but no setup steps changelog-20261005-7e21bf07

03 **iFrame fallback to browser tab on camera failure** IMPROVED

45

Allows iFrame integrations to fall back to opening a new browser tab when the device camera fails during verification.

– Explains trigger and behaviour but no config flag changelog-20261005-7e21bf07

04 **'Explore products' menu in dashboard** NEW

45

Adds a new Explore products menu in the dashboard that shows products the account does not yet have.

– Names exact UI location, a clear starting point changelog-20261005-7e21bf07

05 **Risk assessment results in company verification API** NEW

40

The company verification API now returns risk assessment results as part of its response.

– No endpoint path, response fields, or example given changelog-20261005-7e21bf07

06 **Dedicated accounting email for Pay-As-You-Go partners** NEW

40

Adds a dedicated accounting email setting for Pay-As-You-Go partners so they can receive invoices and payment receipts.

– Describes purpose but not where the setting is configured changelog-20261005-7e21bf07

07 **Date-of-birth collection in Brazil CPF checks** NEW

30

Adds date-of-birth collection to Brazil CPF checks.

– Names specific check but no field or API detail changelog-20261005-7e21bf07

08 **Spanish language support in business verification** NEW

25

Adds Spanish language support to the business verification (KYB) flow.

– Bare addition with no flow or config detail changelog-20261005-7e21bf07

09 **European Accessibility Act compliance for verification flow**

IMPROVED

25

Brings the verification flow into compliance with European Accessibility Act requirements.

– States compliance claim without specifics of what changed changelog-20261005-7e21bf07

WAS THIS USEFUL?  

Keycloak

Feedback ⓘ changelog · Oct 5, 2026 CHANGELOG ↗ RANK

Keycloak - the open source identity and access management solution. Add single-sign-on and authentication to applications and secure services with minimum effort.

> Keycloak's latest changelog adds wallet-based OID4VP authentication, a database-backed stateless multi-cluster v2 that deprecates the Infinispan-based v1, a preview Client Admin API v2 for declarative OIDC/SAML client management, and a new client policy executor restricting the `may_act` delegation claim, alongside deprecation of four client certificate admin endpoints.

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

01 **OID4VP wallet-based authentication** NEW 80

The OID4VP verifier lets users authenticate by presenting verifiable credentials from their wallets instead of a password, enabled with `bin/kc.sh start --features=oid4vc-vci,oid4vc-vp`.

Enable the OID4VP verifier so users can authenticate by presenting verifiable credentials from their wallets instead of a password.

```
$ bin/kc.sh start --features=oid4vc-vci,oid4vc-vp
```

— Runnable command given, underlying mechanism lightly described

changelog-20261005-91d2f674

THINNER COVERAGE BELOW

02 **Client policy executor for may_act delegation claim** NEW 45

A new client policy executor restricts the `may_act` claim in access tokens, giving per-client control over which clients can participate in delegation.

— Describes purpose but no config key or flag given changelog-20261005-91d2f674

└──▷ **BREAKING ON UPGRADE**

! Deprecated endpoint POST `/admin` /realms/{realm} `/clients` /{client-uuid} `/certificates` /{attr} `/download` (still works)

! Deprecated endpoint POST `/admin` /realms/{realm} `/clients` /{client-uuid} `/certificates` /{attr} `/generate` (still works)

! Deprecated endpoint POST `/admin` /realms/{realm} `/clients` /{client-uuid} `/certificates` /{attr} `/generate-and-download` (still works)

! Deprecated endpoint POST `/admin` /realms/{realm} `/clients` /{client-uuid} `/certificates` /{attr} `/upload` (still works)

! Multi-cluster v1 (the `infinispan` feature) is deprecated and will be removed in a future major release; migrate to multi-cluster v2 using `--features=stateless`.

WAS THIS USEFUL? 👍 👎

Didit

Feedback

changelog · Oct 5, 2026

CHANGELOG ↗

RANK

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

> Didit shipped tiered KYB shareholder/UBO data endpoints, bulk AML monitoring and sandbox-testing APIs, new MCP tools for agent-driven session and workflow management, a Web SDK 0.4.0 with transaction submission and wallet-ownership checks, and live digital ID wallet verification across several European providers.

└──▷ **WHAT SHIPPED · 28 FEATURES** most completely described first ? what's the number?

01 **Bulk AML monitoring toggle API** NEW 88

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

Bulk-enable ongoing AML monitoring for up to 1,000 users in a single API call, with a dry run first to estimate scope.

```
$ curl -X POST https://api.didit.me/v3/users/aml-monitoring/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"user_ids": ["u_1", "u_2", "u_3"], "monitoring": true,
"dry_run": true}'
```

— Named endpoints, limits, and a runnable example given. changelog-20261005-84557804

02 **KYB shareholder and UBO data tiers** NEW 83

Didit adds `POST /v3/kyb/shareholders/` and `POST /v3/kyb/ubo/` endpoints to fetch shareholder and ultimate-beneficial-owner data at Lite, Shareholders, or UBOs tiers, with the response naming the tier requested and the tier delivered; the same three-tier KYB registry data purchasing is now selectable per country in the KYB Registry Check workflow step, with price shown before selection.

— Endpoints and tier mechanism named, no schema detail. changelog-20261005-84557804

03 **Transaction monitoring in Web SDK 0.4.0** NEW 83

Web SDK 0.4.0 adds a submitTransaction() method to submit transactions for monitoring from the browser using a scoped transaction token, with automatic device intelligence collection, plus wallet-ownership checks that run the check and report completion through a callback when a transaction requires it.

— Named method and callback mechanism, no code sample. changelog-20261005-84557804

04 **Didit MCP tools for session and workflow management** NEW

80

Didit MCP adds `didit_session_explain_decision`, `didit_session_webhooks`, `didit_session_webhook_resend`, `didit_workflow_build_graph`, and `didit_workflow_get_kyb_registry_catalog` for AI-agent-driven session debugging and workflow building; the webhook tools let an agent list and resend a session's webhooks.

— All five tool names given, usage via agent not shown. changelog-20261005-84557804

05 **Session PDF report endpoint** NEW

75

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

Fetch a branded PDF audit report of a user's 20 most recent completed sessions for compliance evidence.

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

— Exact endpoint and runnable example provided. changelog-20261005-84557804

06 **Sandbox AML hit simulation** NEW

75

The `sandbox_scenario` parameter on `POST /v3/aml/` simulates an AML hit in sandbox without requiring an end-user session.

Simulate an AML hit in sandbox to test your alert-handling workflow without a real end-user session.

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

— Named parameter, endpoint, and runnable example. changelog-20261005-84557804

07 **Webhook signature verification v2** BREAKING

65

Didit adds an X-Signature-V2 webhook verification header; the previous X-Signature-Simple header is now deprecated and should be migrated away from.

– Named headers given, no migration steps detailed. changelog-20261005-84557804

08 **Live digital ID wallet verification** NEW 63

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

– Named providers and countries, no integration steps. changelog-20261005-84557804

09 **Document liveness scores in decision API** NEW 62

Every decision API result now includes `document_liveness` scores covering portrait replacement, printed copy, and screen replay detection, even when no warning fired.

– Named field and sub-scores, no response sample. changelog-20261005-84557804

10 **Skip liveness after register lookup** NEW 62

A `skip_liveness_and_face_match` per-country switch, on by default, skips the Liveness and Face Match steps after a successful government-register lookup that returns a portrait.

– Named switch and trigger condition, no config path. changelog-20261005-84557804

11 **Workflow path explanation field** NEW 62

The decision API (with `include=events`) now exposes a `workflow_terminal_node_id` field showing why a session took its particular path through the workflow.

– Named field and query param, no sample response. changelog-20261005-84557804

THINNER COVERAGE BELOW

12 **Expanded document type and back-side capture coverage** NEW 58

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

– Named documents and countries, no API/flag given. changelog-20261005-84557804

13 **Branded SMS sender customization** NEW 58

A new SMS tab under Customization adds an SMS sender name field with a live preview to brand verification-code messages.

– Named UI location and field, clear navigation. changelog-20261005-84557804

14 **Detected codes field for documents** NEW 57

A new `detected_codes` field on Proof of Address and Document AI results surfaces QR codes and barcodes found on the document.

— Named field and affected result types. changelog-20261005-84557804

15 **Covered critical field warning** NEW 56

A new warning is raised when a document number or other critical field is obscured, configurable on both the workflow step and the standalone ID Verification API.

— Behaviour and two config surfaces named, not a field name. changelog-20261005-84557804

16 **AML screening of document-extracted entities** NEW 53

AML screening can now run against any entity named in a document processed by an earlier Document AI step, with a per-step choice of which entity to screen.

— Mechanism described, no step configuration name. changelog-20261005-84557804

17 **Workflow graph contradiction validation** NEW 53

Workflow graph validation now flags contradicting branches with a warning badge on the canvas, with the same warning surfaced through the graph-validation API.

— Mechanism and API surfaced, no endpoint path given. changelog-20261005-84557804

18 **Face Search toggle for Liveness/Age steps** NEW 50

A `face_search_enabled` switch on the Rules tab lets workflows turn Face Search off per workflow for Liveness and Age Estimation steps.

— Named switch and location, thin description. changelog-20261005-84557804

19 **Branch routing on document subtype** NEW 48

Workflows can now route sessions based on the exact document subtype, such as national ID card versus residence permit.

— Example subtype given, no config surface named. changelog-20261005-84557804

20 **Copilot slash-command skill picker** NEW 48

Didit Copilot adds a slash-command skill picker — typing `/` in the composer lists templates for the current page.

— Exact trigger key named, clear starting action. changelog-20261005-84557804

21 **Per-country non-document source selection** NEW 45

Where multiple non-document lookup sources exist for a country, the source can now be selected per country, configurable in the console tab and via the API.

— Two config surfaces named, no specifics on sources. changelog-20261005-84557804

22 **iGaming rule presets for transaction monitoring** NEW

45

The Transaction Monitoring rules library adds iGaming rule presets such as bonus abuse and chip dumping.

— Named presets, no config or rule syntax given. changelog-20261005-84557804

23 **On-demand wallet screening in console** NEW 45

Console users can run on-demand wallet screening as a single check, now listed as Wallet Verification & Screening.

— Named console item, thin behavioural detail. changelog-20261005-84557804

24 **Back-side-missing warning for ID verification** NEW

45

A new `DOCUMENT_BACK_SIDE_NOT_AVAILABLE` informational warning appears in ID Verification results when only the front side of a document could be collected.

— Named warning code, no handling guidance. changelog-20261005-84557804

25 **Current-value velocity filters** NEW 42

Velocity conditions in Transaction Monitoring can now filter on the current transaction's value.

— Brief description, no field or rule name given. changelog-20261005-84557804

26 **Copilot change cards with undo** NEW 42

Copilot shows change cards after workflow edits, with an Undo action to revert the last change.

— Named UI action, little mechanism beyond that. changelog-20261005-84557804

27 **Web SDK modal language additions** NEW 35

The Web SDK modal now supports Bosnian, Kazakh, Kyrgyz, and Albanian.

— Named languages only, no activation detail. changelog-20261005-84557804

28 **Ongoing AML monitoring evidence display** NEW 25

Each AML result now displays ongoing-monitoring evidence to prove that monitoring ran.

— Bare description, no UI path or field named. changelog-20261005-84557804

└──▷ **BREAKING ON UPGRADE**

! X-Signature-Simple is deprecated — migrate webhook signature verification to X-Signature-V2.

WAS THIS USEFUL?  

SailPoint

Feedback ⓘ

New Capability: HashiCorp Vault (On-Premise) Connector for NHI Discovery (2026-10-05)

BLOG ↗ RANK

SailPoint is an identity governance and access management platform that automates user provisioning, access reviews, and compliance across enterprise systems.

> SailPoint's Delimited File Connector now supports IAM role-based authentication to AWS S3, replacing long-lived access keys with assumable roles.

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

01 **IAM role authentication for S3 file connector** NEW INTEGRATION

NEW 90

The Delimited File Connector adds an Authentication Type dropdown (IAM User or IAM Role) when `File Transport = AWS S3` , letting practitioners use IAM role auth instead of static access keys. New IAM Role configuration fields include Role Name (e.g. `sailpoint/S3ReaderRole`), Account ID (for cross-account bucket access), External ID (stored encrypted when the role's trust policy requires `sts:ExternalId`), Role Session Name , and Role Session Duration (default `3600` seconds). Credentials refresh automatically across long aggregation runs, and Test Connection was extended to validate IAM role access and return specific success/failure messages for trust-policy, permission, and missing instance-profile errors.

Configure an AWS EC2-hosted VA Delimited File source to authenticate via a dedicated cross-account IAM role, avoiding any stored access keys.

> 📍 In Identity Security Cloud, open your Delimited File source › Connection Settings. Set File Transport to `AWS S3` , then set Authentication Type to IAM Role . Fill in: S3 Bucket: my-identity-bucket S3 Region: us-east-1 Role Name: sailpoint/S3ReaderRole Account ID: 123456789012 External ID: <your-sts-external-id> Role Session Name: sailpoint-aggregation Role Session Duration: 3600 Click Test Connection to confirm trust-policy and permissions before scheduling aggregation.

— Names exact config fields, defaults, and UI path with a worked example.

Enhancement: Delimited File Connector -…

NEW INTEGRATION HashiCorp Vault (On-Premise) NHI Source Connector

> New Capability: HashiCorp Vault (On-Premise) Connector for NHI Discovery (2026-10-05)

BLOG ↗ LATEST STABLE

New HashiCorp Vault (On-Premise) ISC Source connector brings Vault token and userpass NHI metadata into Identity Security Cloud via Virtual Appliance, without reading secret payloads or altering the existing Credential Provider.

Adds a new HashiCorp Vault (On-Premise) ISC Source connector that discovers and aggregates Vault token metadata (issue/expiration time, policies, token type) and userpass account metadata (username, policies, token TTL/period settings) as non-human identities — without reading secret payloads.

Connects to on-premises HashiCorp Vault Community edition through the Virtual Appliance using the Vault host URL and an API token, keeping the existing HashiCorp Vault Credential Provider unchanged.

Surfaces NHI metadata as `Dataset` resources in Identity Security Cloud: Vault tokens modeled as credentials, and Vault userpass accounts modeled as machine accounts.

Supports both Express and standard source setup, with a `Test Connection` step to validate connectivity before aggregation.

`Enhancement: Delimited File Connector - SMB File Transport with Kerberos Authenticatio`

BLOG ↗ STABLE

Delimited File Connector gains SMB2/SMB3 + Kerberos file transport, removing the need to stage Windows share files onto SFTP or S3 before aggregation.

Adds `SMB` as a `File Transport` option on the VA-based Delimited File Connector, enabling direct aggregation from Windows File Server and Samba shares over SMB2/SMB3 without staging files to SFTP or S3.

Supports Kerberos authentication for SMB transport, configured via `Host` (FQDN), `Port` (default `445`), `Share Name`, `Kerberos Realm`, `Username`, `Password`, and optional `Kerberos Configuration Directory` (folder on the VA containing `krb5.conf`) fields in source configuration.

Returns explicit success or failure messages for authentication, path, permission, and network errors when using `Test Connection` to validate Kerberos login, share access, and file read before scheduling aggregation.

WAS THIS USEFUL? 👍 👎

Aembit

Feedback ⓘ

Aembit Launches Identity and Access Security Controls for Meta Muse (2026-10-05)

BLOG ↗ RANK

Aembit is IAM for agentic AI and workloads. Enable secretless access and real-time policy enforcement across any environment to protect data and accelerate AI adoption.

> Aembit added Meta Muse as a supported AI agent workload, brokering short-lived credentials through its MCP Identity Gateway and MCP Authorization Server, and folded Muse into its unified access-policy model alongside Copilot Studio, Claude, ChatGPT and Gemini.

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

01 **Meta Muse support via MCP Identity Gateway** NEW 90

Aembit adds Meta Muse as a supported client workload in its IAM for Agentic AI Platform, routing Muse's enterprise resource requests through the MCP Identity Gateway and MCP Authorization Server for policy-evaluated, short-lived credential issuance. Muse authenticates using Dynamic Client Registration against the Aembit authorization server, avoiding manual app creation, client ID, or secret setup, and connects to the gateway by requesting the gateway URL as an OAuth scope, matching the gateway's protected-resource metadata advertisement.

Verify the Muse-to-gateway connection and confirm tool availability by calling tools/list after the OAuth flow completes.

📍 In the Muse chat window, type: 'connect my Aembit MCP gateway'. After Dynamic Client Registration and the OAuth flow complete, Muse will invoke tools/list against https://<tenant>.`mcpgateway.aembit.io`/mcp and return the available tools (e.g. 49 tools for GitHub).

— Names gateway, authz server and DCR mechanism with verification steps

Aembit Launches Identity and Access Sec…

02 **Muse skill file for persistent MCP tool calls** NEW 85

Aembit generates a reusable Meta Muse skill file (e.g. `~/workspace/skills/aembit-mcp-gateway/bin/mcp.py`) using `CREDENTIAL_NAME`, `MCP_URL`, and `ALLOWED_HOSTS` constants, enabling persistent `tools/list` and `tools/call` invocations through the gateway.

Configure the Muse skill file so the agent can list and invoke tools through the Aembit MCP Identity Gateway using short-lived credentials.

```python
CREDENTIAL_NAME = "custom.aembit-mcp-gateway"
MCP_URL = "https://<tenant>.mcpgateway.aembit.io/mcp"
ALLOWED_HOSTS = ["<tenant>.mcpgateway.aembit.io", "
<tenant>.id.useast2.aembit.io"]
```

— Exact file path, constants and runnable config snippet given

Aembit Launches Identity and Access Sec…

THINNER COVERAGE BELOW

03 **Unified access policy extended to Meta Muse** IMPROVED

55

Aembit extends centralized access policy management to cover Meta Muse alongside Microsoft Copilot Studio, Claude, ChatGPT, Gemini, and custom LLM-based agents under a single policy model.

— Names covered agents but no policy mechanism or config details

Aembit Launches Identity and Access Sec…

04 **MCP Authorization Tracing for Muse setup** NEW

35

Aembit provides MCP Authorization Tracing to diagnose authorization flow issues during Meta Muse-to-gateway setup.

— Named but no detail on how tracing is accessed or read

Aembit Launches Identity and Access Sec…

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 1Password for Mac 8.12.40 (2026-10-05)

BLOG ↗ RANK

All the latest releases from the 1Password team

> 1Password for Mac 8.12.40 adds the ability to save one-time passwords via macOS Autofill, streamlines the New Item screen, and improves GitHub Actions setup guidance and export safety warnings.

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

01 **macOS Autofill one-time password saving** NEW 75

Users can now save one-time passwords generated via macOS Autofill directly into 1Password by navigating to `System Settings > General > AutoFill & Passwords > Set Up Codes In` and selecting 1Password, making OTPs available across devices instead of only in the system keychain.

Store macOS-generated one-time passwords directly in 1Password instead of the system keychain, so OTPs are available across all your devices.

> 📍 Open System Settings › General › AutoFill & Passwords › Set Up Codes In, then select 1Password.

— Exact UI path and mechanism given, plus runnable steps example

1Password for Mac 8.12.40 (2026-10-05)

THINNER COVERAGE BELOW

02 **GitHub Actions setup instructions in Developer Environments** NEW 45

Adds a button in `Developer > Environments > GitHub Actions` destination to view setup instructions directly in the app.

— Names exact UI location but no further mechanism 1Password for Mac 8.12.40 (2026-10-05)

03 **New Item screen discoverability improvements** IMPROVED 45

The New Item screen now surfaces the API Credential item type near the top for faster access, makes additional item types easier to discover generally, and updates the search field placeholder text to Search for a website or category.

— Names concrete UI changes but limited mechanism detail

1Password for Mac 8.12.40 (2026-10-05)

04 **Unencrypted export safety warning** IMPROVED 30

Adds a clearer warning about storing unencrypted files safely, with a link to learn more, shown on the export confirmation screen.

— Brief description of a UI warning with no further specifics

1Password for Mac 8.12.40 (2026-10-05)

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ai-safe2-framework

Feedback ⓘ 2026-10-5_CLI_1.0.0 NOTES ↗

RANK

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

> Ships a new local-first CLI (AI SAFE² CLI 1.0.0) with ten commands spanning agent assessment, environment inventory, skill gating, evidence correlation, and remediation planning, alongside a debut NEXUS Sovereign Payment Gateway SDK that enforces fail-closed, deterministic controls over AI-agent-initiated payments.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

01 **Sovereign Payment Gateway package for agentic payments** NEW

95

The `nexus-a2a-sdk==0.5.0` Python package (import `nexus_sdk`) delivers a monotonic, fail-closed transaction lifecycle that prevents AI agents from self-granting authority, accessing unrestricted payment keys, or overriding deterministic controls. It composes: the Key Guardian, which isolates signing authority from the agent and signs only approved canonical transactions; the Policy Authority, for independent deterministic evaluation of payment policy; the Runtime Verifier, which confirms the executing workload matches an approved, sufficiently fresh measurement before value is allowed to move; the Human Intent Authority, binding approval to actual transaction constraints and trusted rendering rather than agent-generated descriptions; the Settlement Truth Authority, which obtains payment outcomes from authoritative external sources instead of trusting the agent; the Gateway State Vault, for durable storage of authority, exposure, replay, revocation, reservation, and settlement state; the Policy Parity Guard, which detects divergence between policy implementations before activation; the Sovereign Payment Coordinator, composing all controls into an ordered, fail-closed execution path; Governed Recovery, to handle ambiguous settlement without unsafe retries or premature restoration of spending capacity; Sandbox Rail Readiness enforcement, to prevent unverified payment integrations from being promoted into active use; and Deployment Proof and Readiness, producing a signed, independently verifiable activation decision.

— Package, import, and every architectural component named. nexus-v0.5.0

02 **Evidence subcommands for identity, claims, and correlation**

NEW

85

`safe2 evidence system` records the complete agent system identity (model, harness, tools, skills, memory, environment, policy, evaluator, authority) behind an assessment. `safe2 evidence claims` compares explicit agent claims against receipts, disclosures, conflicts, and missing evidence. `safe2 evidence truth` correlates attributed evidence across multiple harnesses while preserving coverage gaps. `safe2 evidence diagnose` ranks evidence-backed investigation paths without asserting a verified root cause.

— Four named commands each with explicit mechanism. 2026-10-5_CLI_1.0.0

03 **Sealed assessment bundles via `safe2 assess`** NEW 80

`safe2 assess` creates a sealed project assessment bundle containing canonical JSON, an evidence manifest, and a human-readable Decision Card. The command also gains `--scan-content` and `--inspect-config` flags for explicit opt-in to content and configuration inspection.

— Names command, bundle contents, and both opt-in flags. 2026-10-5_CLI_1.0.0

04 **Multi-harness environment inventory via `safe2 doctor`** NEW 60

`safe2 doctor` inventories multi-harness environments — including WSL distributions and explicit SSH targets — and exposes missing or partial coverage.

— Command and scope named, no flags detailed. 2026-10-5_CLI_1.0.0

05 **Installation verification and bounded project init** NEW 60

`safe2 self-check --strict` performs installation self-verification, and `safe2 init . --profile local` performs bounded project initialization.

— Two runnable commands with exact flags given. 2026-10-5_CLI_1.0.0

THINNER COVERAGE BELOW

06 **AISM remediation planning via `safe2 aism plan`** NEW 55

`safe2 aism plan` converts confirmed AISM findings into a reviewable remediation plan without granting execution authority.

— Command named with scope limit, no example shown. 2026-10-5_CLI_1.0.0

07 **Challenge Lab evidence via `safe2 challenge`** NEW 55

`safe2 challenge` runs or imports bounded Challenge Lab evidence while preserving provider verdicts, provenance, and disagreements.

— Command named, behaviour described in one line. 2026-10-5_CLI_1.0.0

08 **Agent-consumable documentation index under `/NEXUS`** NEW

55

Adds an `llms.txt` retrieval index and a versioned machine-readable documentation manifest under the canonical `/NEXUS` directory for agent-consumable documentation discovery.

— Named file and directory, no retrieval example given. nexus-v0.5.0

09 **Named payment threat model** NEW

55

Counters prompt-injection-driven purchases, replay and duplicate-payment attacks, authority escalation by child agents, spending split across agents to evade limits, protocol and checkout downgrades, and revocation races as named threat categories in the included payment threat model.

— Threats enumerated, but no mitigation mechanism tied in. nexus-v0.5.0

10 **Skill screening via `safe2 gate skill`** NEW

50

`safe2 gate skill` statically screens a skill before installation or use without executing it.

— Single command, mechanism stated briefly. 2026-10-5_CLI_1.0.0

11 **Evidence intake from Codex, OpenTelemetry, SkillSpector, and NEXUS**

NEW

50

Includes Codex JSONL and OpenTelemetry reference translators, plus attributed NVIDIA SkillSpector and NEXUS evidence intake.

— Four integrations named, no configuration detail. 2026-10-5_CLI_1.0.0

12 **Offline acceptance bundles via `safe2 acceptance`** NEW

45

`safe2 acceptance` creates and replays an offline first-user acceptance bundle.

— Command named, minimal mechanism given. 2026-10-5_CLI_1.0.0

13 **Provider-neutral adapter and evidence contracts** NEW

40

Delivers provider-neutral adapter and evidence contracts that bind attributed evidence across agents and scanners without conflating source confidence levels.

— Mechanism described abstractly, no named surface. 2026-10-5_CLI_1.0.0

14 **Canonical result-state vocabulary for CLI outputs** NEW

40

Introduces four result states — Evidence-consistent, Contradicted, Unverifiable, and Limited — as the canonical vocabulary for all command outputs.

— Named states, but no command tie-in shown. 2026-10-5_CLI_1.0.0

WAS THIS USEFUL? 👍 👎

Docker Sandboxes

Feedback ⓘ v0.47.0 NOTES ↗ RANK

Docker Sandboxes isolates applications in lightweight containers, providing secure, reproducible environments for development and deployment.

> Docker Sandboxes v0.47.0 adds automatic sandbox cleanup, a long-running kit capability, richer cloud TTL reporting, stronger kit pull safety limits, and refresh-token support for MCP authorization.

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

01 Auto-cleanup with `sbx run --rm` NEW

85

`sbx run --rm` automatically removes a local or cloud sandbox after its agent session ends. It cannot be combined with `--detached`, and cloud sandboxes interrupted by a dropped connection are instead kept, with the CLI printing the removal command.

Run an agent session in a cloud sandbox and have it automatically cleaned up when the session ends — useful in CI pipelines where you want no leftover sandboxes.

```
$ sbx run --rm --cloud my-agent-sandbox
```

— Flag, interaction rules, and fallback behaviour all specified v0.47.0

02 Long-running kit capability for local sandboxes NEW

85

The `com.docker.sandbox/long-running@1` kit capability declaration keeps local sandboxes running after all sessions disconnect. Cloud sandboxes and `sbx kit add` cannot provide this capability: required entries declaring it are rejected and optional entries are skipped.

Declare the long-running capability in a kit so local sandboxes it provisions stay alive after agents disconnect, letting the next session reconnect without a cold start.

yaml

```
capabilities:
  - com.docker.sandbox/long-running@1
```

— Exact capability string and rejection/skip rules given v0.47.0

03 Refresh tokens for MCP authorization IMPROVED

85

MCP authorization now requests `offline_access` when the server advertises it, using saved defaults or resource-required scopes so servers can issue refresh tokens.

Running `sbx mcp auth` again obtains a grant with a refresh token for existing credentials.

Re-authenticate an existing MCP server connection to obtain a refresh token now that `offline_access` is requested automatically.

```
$ sbx mcp auth
```

— Command, scope mechanism, and refresh behaviour all named v0.47.0

04 **Cloud TTL status fields in JSON output** IMPROVED 70

`sbx --cloud ttl` JSON output now includes `stopped` and `ttl_paused` fields; stopped sandboxes are reported as stopped rather than expired, and the TTL restarts on resume.

— Named fields and behaviour change described, no command example v0.47.0

05 **Kit signing compatibility with registries blocking deletion** IMPROVED 60

`sbx kit sign` and `sbx kit push --sign` now succeed on registries that refuse manifest deletion, including GitHub Container Registry and Docker Hub.

— Names commands and affected registries, no mechanism detail v0.47.0

06 **Cloud policy enforcement on sandbox move or creation** BREAKING 60

Creating or moving a sandbox to the cloud now applies cloud policy and the kit's network rules without copying locally added network rules; moving a sandbox with local HTTP method or path restrictions warns that those restrictions will not apply in the cloud.

— Explains before/after behaviour and warning, no exact command shown v0.47.0

THINNER COVERAGE BELOW

07 **Port collision handling in `sbx ports --publish`** IMPROVED 55

Host-port collisions from `sbx ports --publish` now identify the occupied binding and offer a retry with an automatically allocated port.

— Names the flag and new recovery behaviour, no exact UX shown v0.47.0

08 **Validated upstream proxy in settings** IMPROVED 45

`sbx settings set` now rejects invalid upstream proxy values before saving them.

— Names the command but gives no validation detail v0.47.0

09 **Limits on kit pull blob and archive sizes** NEW 45

Kit pulls now enforce limits on registry-declared blob sizes, decompressed content, and archive entry counts.

— Describes the safeguard but no concrete limit values v0.47.0

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ changelog · Oct 5, 2026 CHANGELOG ↗ RANK

The free and flexible app for your private thoughts.

> Obsidian's latest update adds a command to open and navigate files that live outside the vault, brings templates and iOS theming to Quick Capture, and ships smaller refinements to Bases kanban columns, cross-view search matching, Quick Look previews, RTL workspace flipping, and font-size controls.

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

01 **Full-path search matching across views** IMPROVED 70

Search in the File Explorer, Bookmarks, Outline, and Sync views now matches the full path of each item, and excluding a term (e.g. `-Archive`) hides the matching folder and everything inside it.

— Names affected views and exact exclude syntax. changelog-20261005-69756476

THINNER COVERAGE BELOW

02 **Hide column option in Bases kanban view** NEW 45

Adds Hide column to the right-click menu of a Bases kanban column.

— Exact UI action named, but no further mechanism described. changelog-20261005-69756476

WAS THIS USEFUL? 👍 👎

◆ Internet Services

Cloudflare

Feedback ⓘ changelog · Oct 5, 2026 CHANGELOG ↗

RANK

Connect, protect, and build everywhere.

Web Search API beta, strict service token auth, Pi Durable agent harness, 30-day analytics for all plans, KV namespace jurisdiction GA, D1 US jurisdiction, and hash_in_range() availability.

NEW INTEGRATION Adds Web Search API (beta) at `POST /{account_id}/ai/websearch/` supporting `provider` (`ceramic`, `exa`, `linkup`), `limit`, and `options.gateway.id` fields; routes through AI Gateway with Zero Data Retention and no markup over provider list price.

Adds `strict_service_token_auth` field to `PATCH https://api.cloudflare.com/client/v4/accounts/{account_id}/access/organizations` to enforce consistent service-token-only auth — blocking redirects to login, ignoring Allow policies and `CF_Authorization` cookies, and logging failures in Access authentication logs.

Adds `--jurisdiction=us` flag to `wrangler d1 create` for creating D1 databases that run and persist data within the United States.

Adds `--jurisdiction` flag (values: `eu`, `us`, `fedramp`) to `wrangler kv namespace create` and `cf kv namespaces create`, and a `jurisdiction` field to `POST /storage/kv/namespaces` REST API, for Workers KV namespace data localization — now generally available.

Adds `PiHarness` class to the Cloudflare Agents SDK (`agents/harness/pi`) for building durable long-running agents using the Pi Durable harness, with lifecycle persistence across restarts and crashes in Durable Objects.

+5 more

└──▷ BREAKING ON UPGRADE

! Zero Trust organizations created on or after October 5, 2026 have `strict_service_token_auth` enabled by default and cannot turn it off — service-token requests will never redirect to login and Allow policies will be ignored for those requests.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 5, 2026

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

▷ Subscribe Feedback on this issue