THE CYBER TOOLCHAIN

NO. 061

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED SEPTEMBER 1, 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 98 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 23 tools

▾ // THE BRIEF

HIDE SAME ISSUE, SAME PROMPT, TWO WRITERS: Opus 5 | GPT-5.5

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. Cycode added agent-orchestrated scanning aimed at authorization and logic flaws no rule engine can

vim ? help

express; Aembit now inspects and controls MCP tool traffic from the same plane that issues agent credentials; and Auth0's Cross-App Access pulls third-party API tokens through IdP trust instead of per-integration consent screens.

DETECT

Catch broken-authorization and logic bugs that pattern-matching never sees

Cycode's Agentic Code Scanning adds a reasoning pass over code for flaws — missing tenant checks, broken object-level auth — that have no signature to match, plus attack-chain synthesis to show what an individual finding actually reaches. Checkmarx's Fusion engine takes the multi-model route to the same problem. Both target the class of bug that today gets found by a human reading the diff.

Cycode · Checkmarx

PROTECT

Put AI agent traffic under the same access policy as everything else

Aembit's CrowdStrike AIDR integration means MCP tool calls from an agent are inspected and policy-checked at the same control plane that issues that agent's identity and credentials, rather than being an unmonitored side channel. nono v0.75.0 pushes the other end: glob-based network filtering, stricter PATH-overlap protection, and shaping of ambient credential tokens so a sandboxed agent cannot reach past its profile.

Aembit · nono

PROTECT

Stop handing every integration a user consent screen and a long-lived token

Auth0's Cross-App Access Requesting App lets an app obtain a third-party API token on the strength of IdP trust — the app-to-app fan-out that normally produces dozens of separately consented, rarely revoked grants now runs through one identity provider that can see and cut them. Clerk's configurable reverification window narrows how long a session stays privileged enough to perform a sensitive action after the last auth.

Auth0 · Clerk

Does **Opus 5** read better? 👍 👎

◆ Exploitation & C2

NetSPI OCInferno

Feedback ⓘ changes since 2026-08-03

CODE ↗

A pentesting tool for enumeration/download/graphical analysis of OCI content. Includes an OpenGraph generator for Bloodhound-style analysis.

OCInferno 0.7.0 adds RPST extraction modules for OKE Kubernetes clusters and OCI Data Science pipelines.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-03
https://github.com/NetSPI/OCInferno.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-03
```

└──▷ TRY IT

Extract an RPST from an OKE cluster node without needing the OCI CLI or workload-identity federation.

```
$ run exploit_kubernetes_clusters_rpst --cluster-id
ocid1.cluster.oc1.phx..
```

Autonomously inject a `CUSTOM_SCRIPT` step into a new Data Science pipeline and retrieve the RPST via OCI Logging.

```
$ run exploit_datascience_pipelines_rpst --project-id
ocid1.datascienceproject.oc1.. --log-group-id ocid1.loggroup.oc1.. --
cleanup
```

New `exploit_kubernetes_clusters_rpst` module extracts OCI resource principal RPST from OKE node instance principal by spinning up a temporary pod that accesses IMDS at `169.254.169.254`; supports both Enhanced and Basic OKE clusters, accepts `--cluster-id` or `--kubeconfig`, and generates a Kubernetes Bearer token without requiring the OCI CLI.

New `exploit_datascience_pipelines_rpst` module extracts RPST from OCI Data Science Pipeline runs via a `CUSTOM_SCRIPT` step that reads `OCI_RESOURCE_PRINCIPAL_RPST`, `OCI_RESOURCE_PRINCIPAL_PRIVATE_PEM`,

`OCI_RESOURCE_PRINCIPAL_VERSION`, and `OCI_RESOURCE_PRINCIPAL_REGION`; operates in autonomous mode (`--project-id`) or targets an existing pipeline (`--pipeline-id`), with optional exfil to a listener via `--output-url` or retrieval from OCI Logging via `--log-group-id`, and optional teardown with `--cleanup`.

Adds 6 new or rewritten RPST modules under the DevOps exploit surface, live-validated in PHX.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

Acunetix 360

Feedback ⓘ 1 RELEASE · 2026-09-01 CHANGELOG ↗

RANK

Invicti Web + API is an end-to-end web security scanner that offers a 360 view of an organization's security. Allowing you to take control of the security of all you web applications, web services, and APIs to ensure long-term protection.

> Acunetix 360 added an opt-in HAR file capture option for authentication verification, giving teams a way to diagnose credential or session issues during scans.

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

01 HAR file capture for authentication verification NEW

70

A new opt-in Generate HAR files option can be toggled on under authentication verification settings to capture HAR files for inspecting auth flows and diagnosing credential or session issues; it is disabled by default to avoid unnecessary HAR file generation.

Enable HAR file capture during authentication verification to inspect auth flows and diagnose credential or session issues.

> 📍 In the scan or authentication settings, locate Generate HAR files under authentication verification and toggle it on.

— Names exact setting and default, gives navigation steps changelog-20260901-2d29b713

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Cycode

Feedback ⓘ 2 RELEASES · 2026-09-01 BLOG ↗ RANK

Cycode's Agentic Development Security Platform unites security and development teams with actionable, code-to-runtime context to identify, prioritize, and fix the software risk that matters.

> Cycode launched Agentic Code Scanning, a fourth scanning dimension that uses agent-orchestrated reasoning to catch authorization and logic-level flaws that rule engines and pattern-matching cannot express, paired with attack-chain synthesis, automated remediation workflows, and a self-improving rule-generation flywheel.

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

01 **Agentic Workflows for automated remediation** NEW 70

Consumes exploitability-qualified findings and automates remediation: teams define triggering events, agent action sequences, and confidence thresholds once, and the workflow fires automatically, routing findings to the owning engineer and returning fixes as reviewable pull requests with a full audit trail.

— Describes a definable workflow with clear inputs and PR output

Introducing Agentic Code Scanning: The …

02 **Benchmark caught CVEs other scanners missed** NEW 65

Detected authorization CVEs (gitea `CVE-2025-68941` and next.js `CVE-2025-29927`) that Semgrep OSS, CodeQL, and Cycode's own rule engine all missed, across a ten-repository, six-language benchmark run on an affordable open-weights reasoning model.

— Highly specific results but not an actionable feature itself

Introducing Agentic Code Scanning: The …

03 **Per-scan mode, model, and scope controls** NEW 60

Introduces per-scan configuration of `mode`, `model`, and `scope` as three independent dials, letting teams target expensive reasoning only at code classes rules cannot describe while keeping deterministic scanning on every commit; every scan records what ran, when, on which code, and with which model for auditability.

— Names the three dials but not exact flag syntax

Introducing Agentic Code Scanning: The … Introducing Agentic Code Scanning: The …

04 **Attack-chain synthesis across files** NEW 60

Links related findings across files into a single multi-step exploit path, carrying data flow from untrusted source to dangerous sink and tagging the resulting risk category, e.g.

Code Injection — Critical.

— Concrete mechanism and example tag, no user-facing command

Introducing Agentic Code Scanning: The …

THINNER COVERAGE BELOW

05 **Agentic Code Scanning as a fourth dimension** NEW 55

Adds an orchestration layer that plans from a threat model, ranks code by vulnerability class, and disproves each finding before surfacing it, running above the three existing deterministic/AI scanning tiers to catch logic-level flaws like broken access control that have no matchable pattern.

— Explains mechanism and scope but no exact config surface given

Introducing Agentic Code Scanning: The … Introducing Agentic Code Scanning: The …

06 **Self-improving flywheel promotes findings to SAST rules** NEW 50

Agentic Code Scanning findings that generalize are automatically converted into new deterministic SAST rules (Dimension 1/2), so a vulnerability class caught once by the model re-runs at near-zero cost indefinitely, compounding coverage as repo count grows.

— Mechanism explained but no configurable surface named

Introducing Agentic Code Scanning: The … Introducing Agentic Code Scanning: The …

07 **Context Intelligence Graph grounding for attack chains**

IMPROVED 40

Integrates with Cycode's Context Intelligence Graph to ground attack chains in real call relationships, ownership, and reachability rather than speculative model inference.

— Names the graph integration but stays high-level

Introducing Agentic Code Scanning: The …

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗ RANK

Veracode is a cloud-based application security platform that performs static, dynamic, and software composition analysis to identify vulnerabilities.

> Veracode's CLI (v2.38–v2.52) adds a toolchain discovery command, OAuth2-based login, SBOM export via SCA, and a wide range of packaging flags, exit codes, and language/framework support, alongside scan policy evaluation and reporting improvements — with breaking changes to the SCA download URL default and static scan verbosity.

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

01 **Packaging control flags, environment variables, and exit codes**

NEW 95

`veracode package` gained finer control over what gets packaged: `--exclude` prevents specified Java packages or JAR files from being scanned, `--skip-packagers` excludes specific unsupported artifacts (e.g. unsupported SQL) without affecting other languages, `--java-source` and `VERACODE_PACKAGE_JAVA_SOURCE=true` enable Java source code packaging, `VERACODE_PACKAGE_GOLANG_GENERATE` enables client-side CGO code generation for Go, and `SRCLR_MAKE_WINDOWS` enables preprocessing C/C++ source with Make on Windows using a native gcc or clang compiler. The command also added `--strict`, which returns exit code 4 on build failures, and new exit codes 5, 6, and 7 to improve CI/CD pipeline error handling.

— Every flag, env var and exit code named verbatim CLI updates (2026-08-31)

02 **Toolchain discovery via `veracode package discover`** NEW

80

The new `veracode package discover` command detects toolchains required to produce scannable artifacts, with detection later extended to Java projects using Gradle and JavaScript projects using npm or Yarn; the command also gained exit codes to support automation.

— Names command, extensions and exit codes but no example output CLI updates (2026-08-31)

03 **Expanded language, framework and build-system support in packaging**

NEW 80

`veracode package` added support for `requirements-dev.txt` in Python packaging to include development dependencies, the Poetry package manager for Python autopackaging, a generic packager for Apex, Go v1.24, .NET 9, MSBuild .NET framework

packaging, preprocessor header files used to build PCH, and an option to process all Makefiles in a directory tree. Preprocessed source code scanning is now the default method for C/C++ projects on Windows and is now also available on Linux.

— Lists every added language/tool but thin per-item mechanism CLI updates (2026-08-31)

04 **Scan reporting and data-source improvements** IMPROVED

70

`veracode scan` now categorizes secret-type findings as Critical, prioritizes vulnerability results from the NVD database when available, allows specifying a custom location for downloading vulnerability databases, adds Dockerfile heredoc syntax parsing support, introduces schema enhancements for container and IaC scan results, and creates a separate asset for each repository in ADO workflows for easier individual scan result tracking. `veracode static scan` now displays the Static Engine version in the scan summary.

— Many named improvements, each described in one line only CLI updates (2026-08-31)

05 **SBOM generation via `veracode sca`** NEW

70

The `veracode sca` command now supports SBOM generation via SCA agent-based scans, with results publishable in CycloneDX 1.6 and SPDX 2.3 formats.

— Names command and exact output formats CLI updates (2026-08-31)

06 **Custom policy application and per-finding evaluation in scans**

NEW

65

`veracode scan` allows applying a custom policy defined in the Veracode Platform and enables viewing policy evaluation for individual findings alongside the overall policy result.

— Two-part capability named clearly, no config example CLI updates (2026-08-31)

07 **Default SCA download URL changed** BREAKING

65

The `SRCCLR_DOWNLOAD_URL` environment variable now defaults to `https://sca-downloads.veracode.com`, replacing the previous default of `https://download.sourceclear.com`, which is now deprecated; all instances in your environment should be updated.

— Exact old and new URLs given, clear migration step CLI updates (2026-08-31)

THINNER COVERAGE BELOW

08 **Target UUID support in `veracode dynamic scan`** IMPROVED

55

`veracode dynamic scan` now supports target UUIDs, in addition to webhook UUIDs, for triggering scans.

– Clear named capability but no usage example CLI updates (2026-08-31)

09 **OAuth2 onboarding via `veracode auth login`** NEW 50

Introduces the `veracode auth login` command for onboarding new users via OAuth2.

– Names command but no further detail on flow CLI updates (2026-08-31)

10 **Automatic output artifact generation in `veracode policy`**

IMPROVED 45

`veracode policy` now automatically determines and generates the appropriate output artifacts based on the policy created.

– Names command but mechanism stated only briefly CLI updates (2026-08-31)

11 **CLI version requirement for `veracode scan` data completeness** 40

BREAKING

Using CLI versions prior to v2.48.0 with the `veracode scan` command results in incomplete data displayed in the Veracode Platform.

– States the threshold version but no further remediation detail CLI updates (2026-08-31)

12 **`--verbose` no longer default in static scan** BREAKING 40

The `veracode static scan` command no longer includes the `--verbose` flag by default.

– Named flag change but no workaround or reasoning given CLI updates (2026-08-31)

└──▷ **BREAKING ON UPGRADE**

! The previous `SRCCLR_DOWNLOAD_URL` default of `https://download.sourceclear.com` is now deprecated; all instances in your environment should be updated to `https://sca-downloads.veracode.com` (v2.39.0).

! Using old CLI versions (prior to v2.48.0) with the `veracode scan` command will result in incomplete data displayed in the Veracode Platform (v2.48.0).

! The `veracode static scan` command no longer includes the `--verbose` flag by default (v2.38.2).

WAS THIS USEFUL? 👍 👎

cve-lite-cli

Feedback ⓘ 1 RELEASE · 2026-09-01 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 v1.32.0 adds EPSS score enrichment and an npm registry advisory fallback to widen vulnerability coverage, plus new controls for incomplete detection data and license scanning.

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

01 **Incomplete-data policy control** NEW 90

New `--incomplete-policy warn|error` flag (default: `warn`) controls exit behavior when detection data is incomplete; `error` exits with code 3. The GitHub Action also gains a matching `incomplete-policy` input for CI pipelines.

Treat incomplete detection data as a hard error (exit 3) so CI pipelines never silently pass on partial results.

```
$ cve-lite /path/to/project --incomplete-policy error
```

— Exact flag, values, exit code, and CI example given. v1.32.0

02 **EPSS score enrichment** NEW 70

Findings are enriched with FIRST.org EPSS scores after the OSV pass: the highest score appears in the terminal top-3 block, a verbose table column, and the HTML report detail panel, while the full score array is included in JSON output. Enrichment is skipped in offline mode and for GHSA-only findings.

— Behavior well described but no flag to trigger it. v1.32.0

03 **npm registry advisory fallback** NEW 60

Adds the npm registry advisory API as a secondary source for packages with zero OSV matches, closing coverage gaps for CVEs without an npm ecosystem mapping in OSV. The fallback is skipped in offline mode and is non-fatal on error.

— Automatic mechanism explained but no user-facing control. v1.32.0

THINNER COVERAGE BELOW

04 **License detection in scans** NEW 50

Adds license detection via the `--check-licenses` flag, letting scans flag copyleft or unknown licenses alongside CVE findings in the same run, e.g. `cve-lite`

`/path/to/project --check-licenses --fail-on high`.

Audit a project for CVEs and copyleft/unknown licenses in one pass, failing CI on high-severity CVEs.



— Only known via a usage example, not described in changelog text. v1.32.0

WAS THIS USEFUL? 👍 👎

Checkmarx

Feedback ⓘ 1 RELEASE · 2026-08-17 BLOG ↗ RANK

Checkmarx agentic application security software platform Combines Hybrid scanning, AI-powered agents, and unified risk intelligence across every attack surface

> Checkmarx unveiled a wave of AI-driven AppSec capabilities at Black Hat USA 2026, headlined by the multi-model Fusion scanning engine, an autonomous Developer Assist guardrail for coding agents, AI Inventory/AI-BOM asset discovery, AI-assisted triage and remediation, and cross-surface risk correlation via One Risk Orchestration.

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

01 One Risk Orchestration cross-surface risk correlation NEW

76

One Risk Orchestration correlates signals from SAST, SCA, infrastructure-as-code security, API security, secrets detection, containers, repository health, DAST, and third-party tools, enriched with exploitability, reachability, exposure, ownership, and business criticality for enforcement across pull requests, builds, and releases.

— Lists many integrated surfaces but no config or API named

The Next Model Won't Fix This: Three Ap…

02 Remediation Assist automated fix suggestions and PRs NEW

75

Remediation Assist provides review-ready guidance and suggested code changes for prioritized findings and, for supported GitHub Code Repository Integration projects, can open a separate pull request containing the proposed fix for developer review — nothing is merged automatically.

— Names GitHub PR integration as a concrete action path

The Next Model Won't Fix This: Three Ap…

03 AI Inventory and AI-BOM for AI asset discovery NEW

73

AI Inventory and `AI-BOM` use deterministic analysis to identify models, agents, MCP servers, AI libraries, and SDKs across repositories, with findings traceable to specific files and lines of code, and versioned AI-BOM documentation generated per release.

— Names scope and output but no command or endpoint given

The Next Model Won't Fix This: Three Ap…

04 Checkmarx Fusion multi-model AI scanning engine NEW

70

Checkmarx Fusion (early access) is a Multi-Model AI Engine that runs curated frontier models over the hybrid SAST engine's already-validated findings to surface novel and AI-generated-code vulnerabilities with no known rule, CVE, or signature. It is available as an optional, usage-based layer where teams choose higher- or lower-cost models per project.

— Mechanism and pricing model described but no config surface named

The Next Model Won't Fix This: Three Ap…

05 **Developer Assist autonomous coding-agent guardrail** NEW 69

Developer Assist is a new autonomous mode that uses hooks and MCP to intercept coding-agent actions such as file writes and package installations before they take effect, returning remediation guidance directly to the agent from an independent AppSec layer.

— Interception mechanism named, no exact setup steps given

The Next Model Won't Fix This: Three Ap…

06 **Triage Assist AI-based finding triage** NEW 62

Triage Assist applies AI reasoning to eligible SAST and SCA findings, grounded in reachability, exploitability, code context, risk context, policy, and mitigating controls, to distinguish actionable findings from acceptable risks and suspected false positives.

— Clear criteria listed but no UI or API entry point named

The Next Model Won't Fix This: Three Ap…

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ

2 RELEASES · 2026-08-24 → 2026-08-27 BLOG ↗ RANK

Finite State Platform is a software supply chain security tool that analyzes firmware and software artifacts for vulnerabilities and compliance issues.

> Finite State Platform added a License Policy rule for unlicensed components, expanded firmware unpacking to ESP32 and Airoha Bluetooth-audio images, upgraded SBOM exports to CycloneDX 1.6 with author metadata, and made reachability results persistent across binary scans.

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

01 Empty-license policy flags on Bill of Materials NEW

65

A new License Policy setting lets teams flag components with no license, assigning a severity and a label that then appears on the Bill of Materials for any component lacking a license.

Flag unlicensed components in a License Policy so they surface with a severity on the Bill of Materials.

> 📍 In the Finite State Platform, go to License Policy, add a rule for components with no license, set the desired severity and label, then save — the label will appear on the Bill of Materials for any component lacking a license.

— UI steps given but no API/config surface named Changelog 2026-08-27 - Empty-license po…

THINNER COVERAGE BELOW

02 ESP32 firmware unpacking NEW

55

Finite State now unpacks Espressif ESP32 firmware images, including full flash dumps, so ESP32 devices are analyzed as firmware rather than unidentified blobs.

— Describes scope and mechanism, no invocation steps

Changelog 2026-08-24 - ESP32 and Blueto…

03 Airoha Bluetooth-audio firmware unpacking NEW

55

Finite State unpacks Airoha Bluetooth-audio firmware and identifies vendor SDK and DSP components, enabling complete SBOMs from those images.

— Names identified component types but no usage steps

Changelog 2026-08-24 - ESP32 and Blueto…

04 **CycloneDX 1.6 SBOM export with author metadata** IMPROVED

50

Binary scan SBOMs can now be exported as CycloneDX 1.6 with author metadata, producing auditor-ready documents directly from binary analysis.

— Names format and metadata but no export command shown

Changelog 2026-08-24 - ESP32 and Blueto…

05 **Persistent reachability results before binary assessment**

IMPROVED

50

Reachability results are now saved with the scan and applied before binary assessment, so all later analysis uses a consistent verdict.

— Explains behavior change but no concrete config or command

Changelog 2026-08-24 - ESP32 and Blueto…

WAS THIS USEFUL?  

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Kubernetes cri-tools

Feedback ⓘ 1 RELEASE · 2026-09-01

NOTES ↗ RANK

CLI and validation tools for Kubelet Container Runtime Interface (CRI) .

> cri-tools v1.37.0 adds connection retry with backoff and registry credential support to crictl, improves OpenTelemetry tracing, and expands critest conformance coverage for image volumes and NRI plugin scenarios.

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

01 CRICTL_CREDS/CRICTL_AUTH support on create and run NEW

81

Adds `CRICTL_CREDS` and `CRICTL_AUTH` environment variables and a `-u` alias for `--username` to the `crictl create` and `crictl run` commands, aligning their authentication handling with `crictl pull`.

Pull and run a container from a private registry using credentials via environment variables, matching the behaviour already available in `crictl pull`.

```
$ CRICTL_CREDS='myuser:mypassword' crictl run container-config.json pod-config.json
```

— Named env vars, flag alias, and a runnable example command. v1.37.0

02 Connection retry with exponential backoff in crictl NEW

72

Adds a `--max-retries` flag to `crictl` to control the number of connection retry attempts, using exponential backoff, with a default of 3 retries.

— Named flag and default given, but no config file or edge-case detail. v1.37.0

THINNER COVERAGE BELOW

03 New critest conformance tests for image volumes and NRI NEW

55

New `critest` conformance tests cover image volumes, NRI callbacks during container and pod sandbox lifecycle events, and NRI plugin failure scenarios in `CreateContainer` and `RunPodSandbox`.

— Names test targets but no command or run instructions given. v1.37.0

04 **OpenTelemetry root span per crictl invocation** IMPROVED

45

OpenTelemetry tracing now emits a root span per `crictl` execution that groups all gRPC calls made to the runtime during that invocation.

— Explains mechanism but no config key or flag to act on. v1.37.0

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Threat Intelligence

Yeti

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗ RANK

Your Everyday Threat Intelligence

> Yeti 2.10.0 reshapes the semantic search API response to be dramatically smaller, but breaks existing clients that parse the old field layout.

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

01 **Slimmed semantic search response shape** BREAKING 91

The `POST /api/v2/search/semantic` endpoint now returns each hit as `{score, matched, object_summary}` instead of a full object with `semantic_score` and `matched_on` merged in. The new `matched` field returns only the specific matching document (e.g. for a DFIQ question matched via an approach, just that approach with its description and steps, not all approaches). The new `object_summary` field contains `id`, `root_type`, `type`, `name`, `description`, and `tags` — the `id` can be used to fetch the full object separately, and `dfiq_yaml` is no longer included in responses. This cuts response size from 19,845 to 2,836 bytes (~86% reduction) on a 5-result query with identical results and ordering. Existing clients parsing the old shape will break.

Query semantic search and access the new slimmed response — use `object_summary.id` to fetch the full object only when needed.

```
$ curl -s -X POST https://<yeti-host>/api/v2/search/semantic \
  -H 'Content-Type: application/json' \
  -H 'Authorization: Bearer <token>' \
  -d '{"query": "lateral movement via scheduled tasks", "count":
5}' | \
  jq '.hits[] | {score: .score, name: .object_summary.name, type:
.object_summary.type, id: .object_summary.id}'
```

— Names exact endpoint, fields, and before/after byte sizes 2.10.0

└──▷ BREAKING ON UPGRADE

! The `POST /api/v2/search/semantic` response shape changed: hits are now `{score, matched, object_summary}` — the old `semantic_score` and `matched_on` fields written directly into the object are gone, and `dfiq_yaml` is no longer returned. Any client parsing the old shape will break.

WAS THIS USEFUL?  

◆ Endpoint Security (EDR/XDR)

Sophos Central

Feedback ⓘ 1 RELEASE · 2026-08-10 NOTES ↗

RANK

Sophos Central is a cloud-based security management platform that provides centralized endpoint protection, threat detection, and device management.

> Sophos Central's latest release focuses on AI-related visibility and identity risk, adding a Generative AI DNS filtering category, an AI Defense early access program, and ML-driven Identity Risk Scores in ITDR alongside expanded ITDR regional coverage and a refined partner bulk-license workflow.

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

01 **ITDR Identity Risk Score** NEW 60

Sophos ITDR now surfaces a Risk Score per identity, combining breach likelihood with potential impact, powered by a patent-pending ML model that updates as findings and accounts change.

— Explains mechanism well but no config or UI path given August 10, 2026

THINNER COVERAGE BELOW

02 **Generative AI category in DNS Protection filtering** NEW 55

A new Generative AI category in DNS Protection filtering policies gives administrators visibility into and control over AI service usage across their networks.

— Names the exact filtering category but not usage limits August 10, 2026

03 **Bulk license add-on selection in partner portal** IMPROVED 50

The partner portal bulk license workflow now lets administrators explicitly select add-ons and products when applying changes to multiple monthly-usage customers simultaneously.

— Clear workflow described but no exact UI steps August 10, 2026

04 **Sophos AI Defense early access program** NEW 48

Sophos AI Defense (Early Access Program) lets EDR, XDR, and MDR customers on Sophos Fusion monitor, control, and secure AI usage across their environment.

— Describes scope but EAP details and mechanism thin August 10, 2026

05 **ITDR regional expansion** IMPROVED 30

Sophos ITDR expands regional availability to Canada, Australia, Japan, India, and Brazil.

— Just a list of new regions, no other detail August 10, 2026

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ 1 RELEASE · 2026-09-01 CHANGELOG ↗ RANK

runZero is an asset discovery and inventory platform that identifies and maps all devices and services across networks without agents.

> runZero's 5.1.260826.0 release centers on a major expansion of its custom integration scripting platform, alongside MCP OAuth support, Cyclone JSON export for cryptographic bills of materials, and an upgraded Censys Platform v3 integration.

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

01 **Custom Integration scripting platform expanded** NEW

85

Custom Integration scripts now support `WMI` (WinRM, SMB, DCERPC) for deep Windows host collection, custom application protocols over `TCP`, `UDP`, and Unix streams, and shared common settings for protocols like `HTTP` and `TLS`. Scripts can now define detailed credential fields and validation rules, track a maturity mode (`alpha`, `beta`, `stable`) for filtering, and be run in console instances in addition to Explorers and the CLI. Scripts can be authored and tested in-product using an AI model of choice, submitted for inclusion as native runZero sources, and are now available to all administrators including those with narrow access scopes; the community custom integration script library now covers more than 100 products and is accessible directly in the console.

– Names protocols, credential and maturity mechanics; console path is starting point only changelog-20260901-ff9c41ad

02 **MCP service OAuth and shared AI tools** NEW

65

The MCP service now supports OAuth, including Export Token and Organization Key as OAuth Client Secrets, in addition to API Clients. The MCP service and in-product AI Threads now share the same tools.

– Names OAuth secret types but no setup steps given changelog-20260901-ff9c41ad

THINNER COVERAGE BELOW

03 **CBOM Cyclone JSON export** NEW

55

The Cryptographic Bill of Materials (CBOM) is now exportable as a Cyclone JSON format and detailed report.

– Names format but not export path or file structure changelog-20260901-ff9c41ad

04 **Censys integration v3 API and retry handling** IMPROVED

50

The Censys integration now supports the Censys Platform v3 API, and the Censys Platform integration now retries API calls on connection errors and other retryable HTTP codes.

— Names version and retry behavior, no config detail changelog-20260901-ff9c41ad

05 **AI Threads and Artifacts enhancements** IMPROVED 50

The AI Threads feature now supports charting and diagrams via Mermaid, and the AI Artifacts feature now supports full-page viewing. AI-generated content (threads, artifacts, exported reports, downloaded files, and custom integration scripts) is now labelled as AI-generated, following EU guidance.

— Names Mermaid and labeling policy but no usage steps changelog-20260901-ff9c41ad

06 **Annotator role in SSO settings** NEW 45

The Annotator role has been added to the list of default roles selectable in SSO settings.

— Names role and setting location, minimal further detail changelog-20260901-ff9c41ad

07 **Dashboards usage column in query library** NEW 45

Dashboard query library now shows how many personal dashboards each saved query is used on, and whether it is used on other dashboards, in a Dashboards column.

— Names exact column but no navigation path given changelog-20260901-ff9c41ad

08 **Promiscuous ARP probe toggle on scan page** NEW 45

A configuration option has been added to the scan page to toggle the promiscuous ARP probe.

— Names the toggle and page but no further mechanism changelog-20260901-ff9c41ad

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Didit

Feedback ⓘ 3 RELEASES · SEEN 2026-09-01 CHANGELOG ↗ RANK

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

> Didit's biggest change this window is a breaking overhaul of session-deletion responses alongside new opt-in biometric-template retention and a management API; it also shipped programmatic transaction-monitoring rule management with MCP tools and AI-assisted rule migration, event-driven monitoring, standalone risk-scoring APIs for email/phone/IP, a general-purpose Document AI endpoint, and a large expansion of Copilot and identity-document coverage.

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

01 Biometric template retention and management

BREAKING

100

Applications can opt into biometric-template retention via `face_retention_policy` and `face_retention_days` fields on `PATCH /v3/webhook/`, or via the Retain biometric template toggle and required retention-period field under App Settings → Data (also visible at Lists → Biometric templates), keeping one image-free face template anchored to a user after session deletion so duplicate detection, Face Search, and biometric authentication keep working. A new Biometric Templates API — `GET /v3/biometric-templates/`, `GET /v3/biometric-templates/count/`, `GET` / `DELETE /v3/biometric-templates/{template_uuid}/`, and `POST /v3/biometric-templates/delete/` — lists, counts, inspects, and purges retained templates; `DELETE /v3/session/{sessionId}/delete/` also accepts `retain_face_embeddings` to override the policy per session, and `deletion_instruction` / `instruction_id` to satisfy privacy-erasure requests, while `POST /v3/users/delete/` now purges a user's templates before deleting them and returns `503` if a template can't be purged. As a breaking change, `DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` now return `200 OK` with a JSON body (`face_retention_outcome`, `biometric_template_uuid`, per-session `results[]`) instead of `204 No Content`.

Enable biometric-template retention on your app so duplicate detection and Face Search keep working after session data is purged.

```
$ curl -X PATCH https://api.didit.me/v3/webhook/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"face_retention_policy": "retain", "face_retention_days":
365}'
```

Enable biometric-template retention on a webhook so Face Search and duplicate detection keep working after you purge session data.


```
$ curl -X PATCH https://api.didit.me/v3/webhook/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"face_retention_policy": "retain", "face_retention_days":
365}'
```


Delete a session while overriding the app-level retention policy to force immediate purge of that person's biometric template.

```
$ curl -X DELETE
https://api.didit.me/v3/session/{sessionId}/delete/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"retain_face_embeddings": false}'
```

Satisfy a privacy-erasure request by purging every retained biometric template for a person, regardless of application retention policy.

```
$ curl -X DELETE
https://api.didit.me/v3/session/{sessionId}/delete/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"deletion_instruction": "privacy_erasure",
"instruction_id": "<your-instruction-id>"}'
```

— Full endpoint list, config fields, and exact response-shape change given

changelog-20260901-2e1362ed · snapshot-20260901 · changelog-20260901-a9a1815e

02 **Standalone risk APIs for email, phone, and IP** NEW

90

New `POST /v3/email/risk/`, `POST /v3/phone/risk/`, and `POST /v3/ip/risk/` endpoints return standalone risk signals for an email address, phone number, or IP address without starting a full verification session, each priced per call.

Check whether an email address is high-risk before allowing account creation, without starting a full verification session.

```
$ curl -X POST https://api.didit.me/v3/email/risk/ \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"email": "user@example.com"}'
```

— Three exact endpoints named with a runnable sample call `snapshot-20260901`

---

03 **Transaction-monitoring rule management API and MCP tools** NEW

85

New `/v3/transactions/rules/` Management API endpoints let you create, update, delete, and backtest transaction-monitoring rules programmatically, with matching `library` and `install` endpoints to browse, install, and remove preset rules from the rule library — mirroring what's available in the Business Console. Nine `didit_transaction_rule_*` MCP tools expose the same list, create, update, delete, backtest, and library workflows to any MCP-connected AI agent, and every manual transaction status change is now attributed to the console user who made it.

— Endpoints and MCP tool names given, no example call shown `changelog-20260901-2e1362ed` `snapshot-20260901` `changelog-20260901-a9a1815e`

---

04 **Document AI extraction endpoint** NEW

75

New `POST /v3/document-ai/` endpoint extracts structured data from any uploaded document, not just identity documents, and accepts uploads up to 50 MB.

Extract structured data from an uploaded business document without building a verification session.

```
$ curl -X POST https://api.didit.me/v3/document-ai/ \
  -H 'Authorization: Bearer <token>' \
  -F 'file=@invoice.pdf'
```

— Exact endpoint, size limit, and runnable curl example given `snapshot-20260901`

---

05 **Identity document coverage expansion** NEW

75

New identity documents: Moldova ID card (2025), Pakistan ID card (2025), Uruguay ID card (2026), Greece asylum seeker card (2025), and a third Belgium driving license

variant. Malaysia MyKad is now selectable as its own subtype rather than a generic ID card, and extraction is upgraded for Belgium, Germany, India, Israel, Madagascar, New Zealand, Norway, Sierra Leone, South Africa, and the United States. Firearm licenses are now a selectable document category, covering Victoria (Australia) and Illinois, Massachusetts, and two New Jersey permits (United States). Digital driving licenses are added for Argentina (Mi Argentina and Jujuy) and Australia (New South Wales including provisional, South Australia including learner, and Victoria including heavy vehicle). The ePassport trust store is expanded to 136 issuing countries with new signing-certificate generations for 32 of them, workflows can now route on the specific NFC skip reason when a chip is not read, and the barcode on North American driving licenses now reports its signed fields with each signature validated against its certificate's validity window.

— Every added document, country, and category enumerated `changelog-20260901-2e1362ed` `snapshot-20260901`

06 **Didit Copilot in Business Console** IMPROVED 75

Didit Copilot is now a resizable floating panel available on every Business Console page, with fullscreen layout on mobile and voice plus typed input. Copilot can build a complete workflow from a natural-language description — adding and configuring steps, branches, and final-status rules — and validate it before publishing; build a questionnaire from an attached CSV or Excel file, including very long answer lists; configure Database Validation by conversation for any supported country; and answer questions with inline citations that open the underlying documentation without leaving the panel.

— Rich behavior described, but no config key or API to script it

`changelog-20260901-2e1362ed` `snapshot-20260901`

07 **AI-assisted transaction-rule migration** NEW 70

The Didit assistant, or any MCP-connected agent, can accept a rule set exported as CSV or Excel from another provider, map it onto Didit's rule schema, backtest it against your own transaction history, and create the rules in test mode.

— Clear mechanism described but no command or endpoint named `changelog-20260901-2e1362ed` `snapshot-20260901`

08 **Structured provider-outage error codes** IMPROVED 65

Registry-search outages now return HTTP `502` with `code: 'kyb_registry_provider_unavailable'`, and phone-provider outages now return HTTP `502` with `code: 'phone_provider_unavailable'`, so integrations can detect and react to provider downtime programmatically.

— Exact status codes and error codes named for both providers `snapshot-20260901`

09 **ISO/IEC 27566-1 age-assurance workflows** NEW 60

Age-assurance workflows can now support ISO/IEC 27566-1: a workflow that declares the standard cannot return identity attributes alongside the age decision. The Age Estimation borderline band is now exposed in the workflow graph, enabling routing of uncertain results to a document check.

— Standard and behavior named; workflow-graph path not detailed

changelog-20260901-2e1362ed snapshot-20260901

10 **Database Validation gains BVN source and flexible inputs** NEW

60

Nigeria Bank Verification Number (BVN) is now a Database Validation source, gated on a selfie face match with names as optional inputs. Database Validation can also run without a document capture step, with each field drawing from its own input source, such as a selfie plus a national ID number.

— Concrete gating and input mechanism described, no config key named

changelog-20260901-2e1362ed snapshot-20260901

11 **Support section in Business Console** NEW 60

A new Support section in the Business Console gives organization-scoped tickets with live status, a standard tickets table, category/severity/session-ID filing, attachment support, per-ticket permalinks, and escalation to a human agent in the same thread.

— Concrete ticket fields named; only a UI path, not an API changelog-20260901-2e1362ed snapshot-20260901

THINNER COVERAGE BELOW

12 **Event-driven transaction monitoring** NEW 55

A new Event Ingestion endpoint accepts non-monetary event types — login, gameplay, bonus, payment, dispute, security, and responsible gambling. New Rules on Events support velocity, ratio, ordered-sequence, and cross-subject correlation rule types for event-driven monitoring logic, and SDK Event Submission lets client devices capture events using a scoped monitoring credential.

— Event types and rule types named, but no path or sample call product docs

13 **Pixel-level document occlusion detection** NEW 50

A partially obscured document is now detected at pixel level and returns a new `DOCUMENT_OCCLUSION_DETECTED` result, translated into every supported language.

— Named result code given, no route or config to act on changelog-20260901-2e1362ed snapshot-20260901

14 **Per-feature engine selector in workflow editor** NEW

50

A per-feature engine selector in the workflow editor (both graph and simple views) lets you choose which engine runs each individual check, and the choice persists across both editor types.

— Describes UI behavior and persistence, no config key named `snapshot-20260901`

15 **Sessions listable ordered by update time** IMPROVED 45

Sessions can now be listed and ordered by `updated_at`, surfacing recently imported or modified sessions.

— Names the sort field but no endpoint or query param shown `changelog-20260901-2e1362ed` `snapshot-20260901`

16 **Idempotency keys on the AML API** IMPROVED 40

Idempotency keys are now supported on the AML API, so a retried screening request is never charged or screened twice.

— Names the AML API but no header or field name given `snapshot-20260901`

17 **Risk-factor breakdown for AML and wallet screening** IMPROVED 40

AML and Wallet Screening results now show the factors behind a risk score, and wallet screening gains a shareable report.

— States what changed but no field names or report format `changelog-20260901-2e1362ed`

18 **Registry coverage analytics in Reports** NEW 40

Registry coverage analytics are now available in Reports, showing registry search coverage and success by country and workflow version.

— Names the report dimensions but not the navigation path `changelog-20260901-2e1362ed` `snapshot-20260901`

19 **Marketplace section in Business Console** NEW 35

A new Marketplace section in the Business Console lets you browse and enable partner modules.

— Bare description of a new UI section, no further detail `snapshot-20260901`

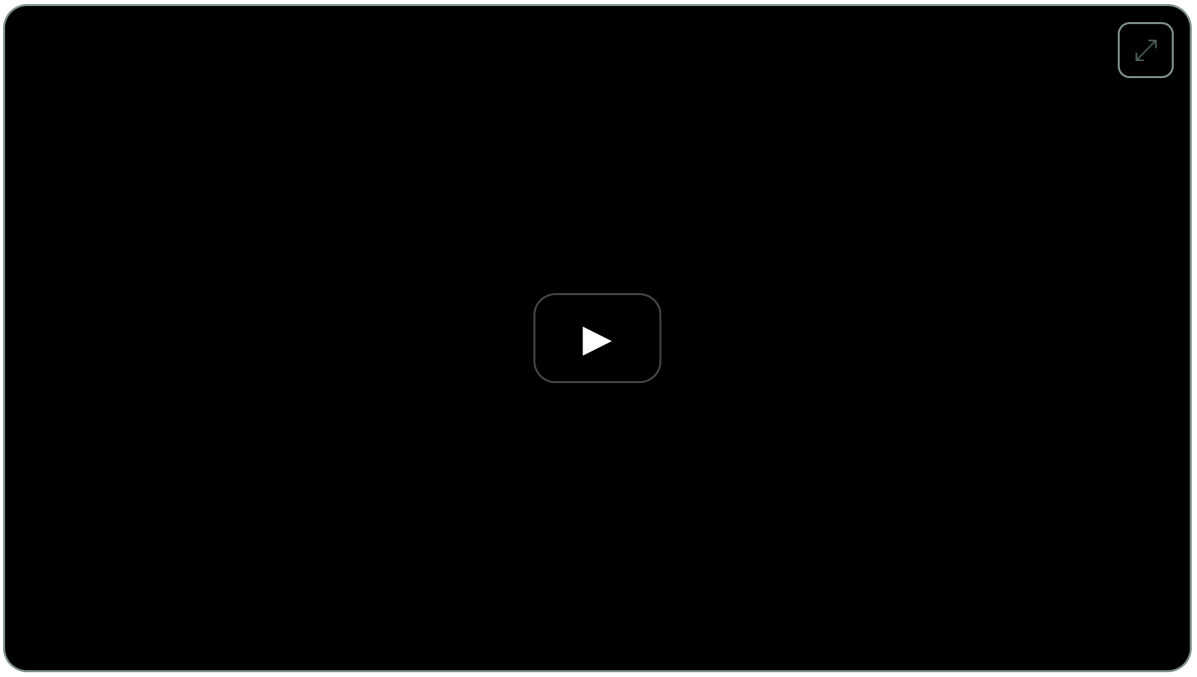
20 **French overseas departments in company registry search** NEW 35

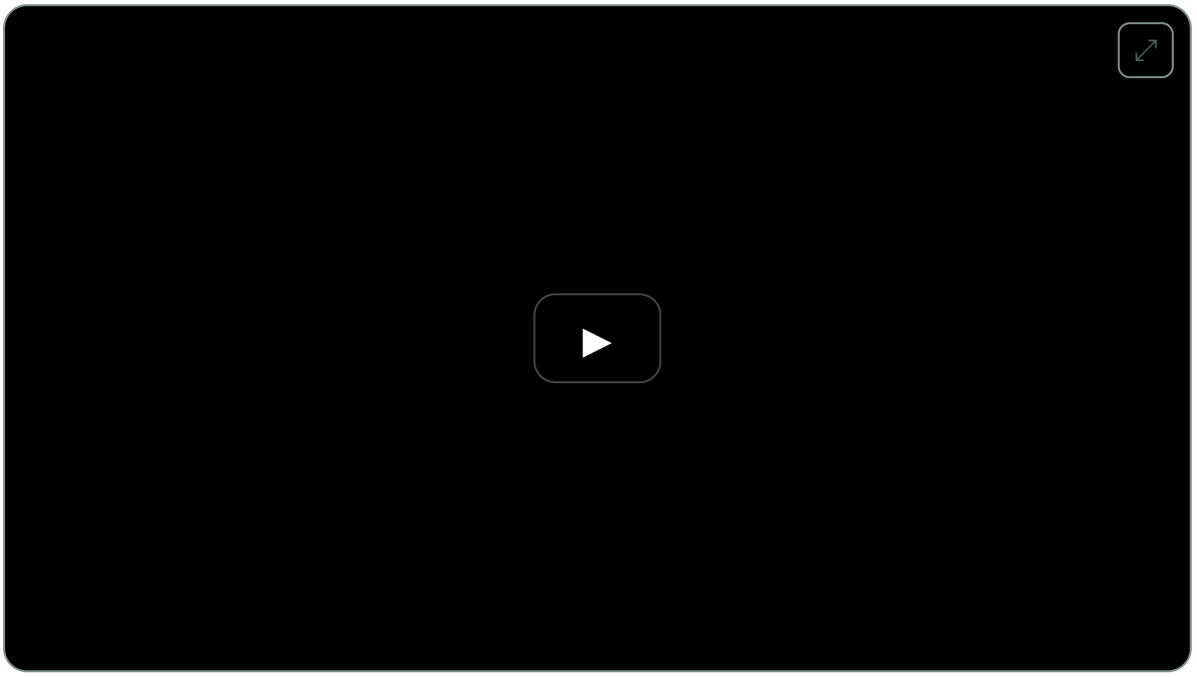
French overseas departments are now searchable in the French company registry.

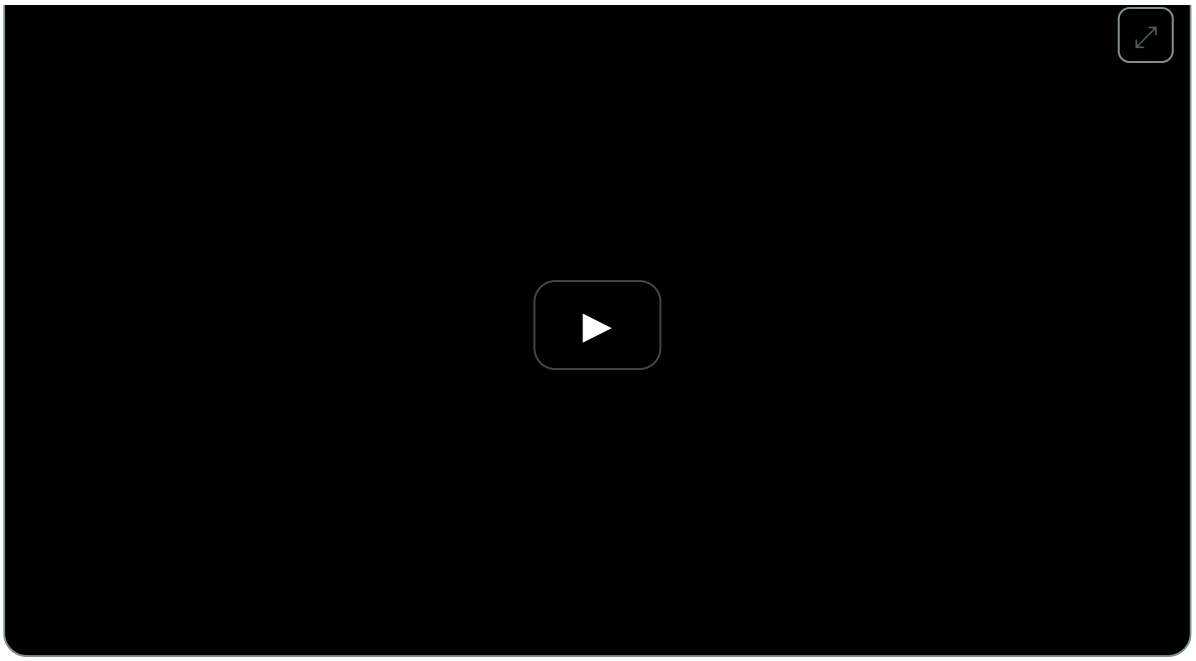
— One-line scope addition with no further mechanism given `snapshot-20260901`

└──▷ ALSO FROM THESE RELEASES









└──▷ **BREAKING ON UPGRADE**

! `DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` now return `200 OK` with a JSON body containing `face_retention_outcome`, `biometric_template_uuid`, and per-session `results[]` for batch requests, instead of `204 No Content`. Clients that assert on `204` must be updated.

! Session deletion behavior has changed: applications that relied on full session data persisting after deletion must adopt the new biometric-template retention policy, as session data is now purged while only an image-free face template may be retained.

! `DELETE /v3/session/{sessionId}/delete/` and `POST /v3/sessions/delete/` now return `200 OK` with a JSON body containing `face_retention_outcome`, `biometric_template_uuid`, and per-session `results[]` (for batch), instead of `204 No Content` — clients that assert on `204` must be updated.

! `POST /v3/users/delete/` now purges the user's retained biometric templates before deleting the user, and returns `503` (retryable, nothing deleted) if a template cannot be purged.

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗ RANK

JumpCloud makes it simple to manage Windows, Apple, Linux, and Android devices and to deliver secure access with SSO, MFA, Zero Trust, RADIUS, and more.

> JumpCloud's agent update expands browser support for Device User Trust enrollment and adds on-demand update capability for required applications.

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

01 **On-demand updates for required applications** IMPROVED

38

Required software applications configured with immediate delivery can now be updated on demand through local agent triggers on macOS and Windows, improving responsiveness for administrator and self-service update flows.

– Describes mechanism and platforms but no exact trigger command.

JumpCloud Agent Release Notes for Augus…

02 **Firefox Developer Edition support for Go native messaging host**

NEW

35

The JumpCloud Go native messaging host on macOS now supports Firefox Developer Edition, extending browser coverage for Device User Trust enrollment.

– Names browser and platform but no config or command shown.

JumpCloud Agent Release Notes for Augus…

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗ RANK

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

> SailPoint improved Workflows testing by validating trigger filters before workflows go live, helping avoid misconfigured triggers firing unintended events.

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

01 **Trigger filter testing in Workflows** IMPROVED 63

SailPoint Workflows now warns in the workflow editor and on the Test Workflow experience when a filterable trigger has no filter configured. During test runs, the configured trigger filter is now applied against the test input and shown as matching or not — previously test runs ignored the trigger filter entirely.

Validate a JSONPath trigger filter against a real test payload before enabling a workflow, to confirm only intended events will fire it.

📍 In Identity Security Cloud, open your Workflow › select a filterable event trigger › add your JSONPath filter expression › go to Test Workflow, provide your test input JSON, and run the test. The result now shows whether your trigger filter matched the input, letting you confirm scope before enabling.

— Names UI flow and behaviour change but no API/config keys.

Enhancement: Easier Trigger Filter Test…

WAS THIS USEFUL? 👍 👎

Auth0

Feedback ⓘ 1 RELEASE · 2026-08-31 BLOG ↗ RANK

Secure users, AI agents, and more with Auth0, an easy-to-implement, scalable, and adaptable authentication and authorization platform.

> Auth0 introduces Cross-App Access Requesting App in Early Access, letting applications silently fetch third-party API tokens through IdP trust instead of per-integration user consent.

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

01 Cross-App Access Requesting App (Early Access) NEW

81

A new Dashboard toggle on enterprise connections enables the Cross-App Access Requesting App flow, letting apps silently fetch third-party API tokens for authenticated users without requiring per-integration Connect your account clicks. It is built on Auth0 Token Vault using the ID-JAG (Identity Assertion Authorization Grant) protocol, currently supported on OIDC and Okta Workforce connections. In multi-org tenants, token resolution is org-aware: requests are scoped by `org_id` and rejected if the enterprise connection or target resource app is not enabled for that organization, or if the org context is ambiguous across multiple identities.

.
.
.
.

— Names protocol, toggle location, and org scoping rules but early access UI only

Cross-App Access Requesting App now in …

WAS THIS USEFUL? 👍 👎

Castle

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗ RANK

A complete suite of SDKs and APIs to monitor users, block bots, and stop in-app fraud.

> Castle added an SDKs page in its dashboard giving teams a centralized view of SDK integration versions and update status.

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

01 **SDKs page for integration version tracking** NEW 50

A new SDKs page provides a centralized view of all SDK integrations, showing which are outdated so teams can quickly identify where to upgrade. Navigate to the SDKs page in the Castle dashboard to see the version status of each integrated SDK and identify any that need updating to ensure full detection model coverage.

Check which of your SDK integrations are behind the latest release to ensure full detection model coverage.

> 📍 In the Castle dashboard, navigate to the SDKs page to see the version status of each integrated SDK and identify any that need updating.

— Names UI location and purpose but no mechanism details SDKs page

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ 1 RELEASE · 2026-09-01 CHANGELOG ↗ RANK

The easiest way to add authentication and user management to your application. Purpose-built for React, Next.js, Remix, and “The Modern Web”.

> Clerk shipped a configurable reverification window for sensitive actions, a new Admin Logs audit trail for Dashboard/API changes, and biometric sign-in support across its mobile SDKs.

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

01 Admin Logs audit trail NEW

90

A new Admin Logs section in the Clerk Dashboard provides a reverse-chronological audit trail of configuration changes made via the Dashboard, Backend API, or Platform API — tracking events such as creating OAuth applications, updating instance settings, rotating secrets, and managing Roles and Permissions. Logs support filtering by Event type (with trailing wildcard support, e.g. `oauth_application.*`), Instance, Application, Actor, IP address, and Time range; available on Business and Enterprise plans, recording from August 24, 2026 onward.

Filter Admin Logs to see all OAuth application events across your workspace using a trailing wildcard event type.

> 📍 In the Clerk Dashboard, open Admin Logs, select the Event type filter, and enter `oauth_application.*` to see all OAuth application events.

— Details filters, plans, and coverage start date precisely changelog-20260901-634d332a

02 Configurable reverification window NEW

85

A new Reverification window setting on the Sessions page in the Clerk Dashboard lets admins configure how long a successful reverification remains valid, from 1 to 10 minutes (default: 10 minutes). This controls how frequently users must re-authenticate before sensitive actions such as changing a password, adding/removing an email address, revoking a session, or deleting an account.

Shorten the reverification window to 2 minutes so users must re-authenticate frequently before sensitive actions like password changes.

> 📍 In the Clerk Dashboard, go to Sessions › Session lifetime › Reverification window and set the value to 2 (minutes).

— Names exact setting, path, range, and default with steps changelog-20260901-634d332a

03 **Biometric credential enrollment in mobile SDKs** NEW

80

Clerk mobile SDKs now support biometric credential enrollment and sign-in (Face ID, Touch ID, Android biometrics) via `useBiometricCredentials` (React Native), `Clerk.shared.biometricCredentials.signInWithBiometrics` and `Clerk.shared.biometricCredentials.biometricCredentialId` (Swift), and `Clerk.biometricCredentials.signInWithBiometrics` and `Clerk.biometricCredentials.biometricCredentialId` (Android), with the private key remaining device-bound.

— Names exact SDK APIs per platform but no usage steps given `changelog-20260901-634d332a`

WAS THIS USEFUL?  

Aembit

Feedback ⓘ 1 RELEASE · 2026-09-01 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 integrated CrowdStrike AIDR into its access policy engine, letting security teams inspect and control MCP tool traffic for AI agents directly from the same control plane that issues identities and credentials.

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

01 **CrowdStrike AIDR content inspection for MCP access policies**

NEW 81

Aembit access policies can now attach a CrowdStrike AIDR content security configuration to inspect MCP traffic — tool listings, tool inputs, and tool outputs — from the same control plane that handles identity and credential issuance. Inspection is opt-in per policy: only policies explicitly configured with a CrowdStrike AIDR content security configuration trigger checks, and each policy can be set to fail-open or fail-closed if AIDR is unreachable, with every failure recorded as a workload event. Based on AIDR's decision, Aembit allows, blocks, or forwards a transformed version of the MCP content (e.g., sensitive data redacted or defanged), logging the outcome as a workload event tied to the agent identity, target MCP server, and access policy. The integration is configurable through the Aembit UI, API, or Terraform, and requires a CrowdStrike AIDR for Agents subscription.

Attach a CrowdStrike AIDR content security configuration to a sensitive MCP access policy via the Aembit UI so that tool listings, inputs, and outputs are inspected before reaching your AI agent.

> 📍 In the Aembit console, open the target Access Policy › Content Security › attach a CrowdStrike AIDR content security configuration › set fail behavior to `fail-closed` › Save.

— Full mechanism, named config surfaces, and UI steps given.

Aembit + CrowdStrike: More Judgment Beh…

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Knocknoc

Feedback ⓘ 1 RELEASE · SEEN 2026-09-01 NOTES ↗ RANK

Knocknoc is a security tool that detects and prevents unauthorized access attempts through anomaly detection and threat intelligence.

> Knocknoc's v26.08 release adds new enforcement backends for nftables and Fastly, alongside improved Palo Alto visibility and admin diagnostics tooling.

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

01 **New enforcement backends: nftables and Fastly** NEW

48

Knocknoc adds native `nftables` integration to dynamically control network access on a Linux host based on authenticated users, and adds Fastly ACL enforcement as a new supported enforcement target to control which source IPs an edge service allows through.

— Names two integrations and their mechanism but no config steps v26.08

02 **Improved Palo Alto visibility for administrators** IMPROVED

18

Knocknoc improves Palo Alto visibility for administrators.

— One-line mention with no mechanism or scope given v26.08

03 **Improved admin diagnostics tooling** IMPROVED

16

Knocknoc improves operational diagnostics tooling for administrators.

— Bare mention with no specifics on what changed v26.08

WAS THIS USEFUL? 👍 👎

Cato Networks

Feedback ⓘ 1 RELEASE · 2026-09-01 BLOG ↗ RANK

Cato SASE Cloud converges SD-WAN, security, global backbone, and remote access into a global cloud-native service. WE ARE SASE™

> Cato rolled out a new managed SASE offering for SMBs, giving partners a pooled licensing model and a self-service portal to onboard and manage SMB customers faster.

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

01 **MSASE Business Center partner portal** NEW 60

The new Cato MSASE Business Center adds a consolidated Business Overview page showing business accounts, orders, license usage, invoices, and customer lifecycle activity across a partner's portfolio. Authorized partners can self-provision eligible services and activate new customer accounts directly in the portal without waiting for manual approval workflows.



— Names the portal and page but no exact navigation steps.

Cato Managed SASE for SMB: Flexible Cap…

THINNER COVERAGE BELOW

02 **SMB FlexPool license pooling for partners** NEW 55

The new `Cato SMB FlexPool` gives MSPs and service providers a partner-level pool of Cato licenses, enabling unlimited portability to allocate and reallocate capacity across eligible SMB customers without needing per-customer license purchases.

— Describes mechanism but no limits or numbers given.

Cato Managed SASE for SMB: Flexible Cap…

WAS THIS USEFUL? 👍 👎

DNSFilter

Feedback ⓘ 1 RELEASE · 2026-09-01 BLOG ↗ RANK

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

> DNSFilter's September 1 web application and API update makes Customizable Dashboards the default landing view, adds live status indicators to deployment widgets, and fixes TLD categorization in lists and logs.

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

01 **Customizable Dashboards default with merged filters drawer**

IMPROVED 62

Customizable Dashboards is now the default landing view for all users, replacing the retired legacy Overview page; existing bookmarks to the old page redirect automatically. The new Dashboard Filters drawer merges the previous Quick Filters and Advanced Filters into a single experience, grouped by Source (Organizations, Sites, Roaming Clients/Relays, Users) and Traffic (Result).

— Explains before/after and filter grouping but no exact UI path

09/01 September 1 – Web Application and…

THINNER COVERAGE BELOW

02 **Live status indicators on deployment widgets** NEW 52

Deployment widgets — Roaming Clients Protected, Sites Protected, Relays, Users, and Collections — now show a pulsating green dot indicating live data, which is unaffected by the dashboard's time range filter.

— Names widgets and mechanism but no config or command

09/01 September 1 – Web Application and…

03 **Top Level Domain categorization for bare TLDs** IMPROVED 52

Bare TLD entries (e.g. `.ru`, `.xyz`) in Allow/Block lists and the Query Log now display as Top Level Domain instead of Uncategorized, distinguishing intentional TLD-level rules from genuinely uncategorized domains.

— Concrete example values but only a display fix, no action needed

09/01 September 1 – Web Application and…

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Zama fhEVM

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗ RANK

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

> fheVM v0.14.1 adds tasks to mirror KMS context and epoch state onto replica chains and introduces Polygon mainnet chain support in the SDK and Solidity library.

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

01 **KMS canonical context mirroring for host contracts** NEW

45

New tasks mirror the canonical KMS context and epoch onto replica chains in host contracts, keeping replicas synchronized with the canonical KMS state.

— Names the mechanism but not the actual task command v0.14.1

02 **Polygon mainnet chain support** NEW

35

Adds Polygon mainnet chain configuration to both the SDK and the Solidity library.

— States what was added but no config details v0.14.1

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NoLabs nono

Feedback ⓘ 1 RELEASE · 2026-09-01 NOTES ↗ CODE ↗

RANK

The nono runtime isolates AI agents with policy-controlled filesystem, network, and credential access.

> nono v0.75.0 extends its policy-controlled sandbox with profile composition on the proxy command, stricter PATH-overlap protection, glob-based network filtering, and a way to shape ambient credential tokens for prefix-sniffing clients.

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

01 `format` field for ambient credential phantom tokens NEW

85

`CommandCredentialConfig` gains a `format` field for ambient credentials — a literal template such as `sk-ant-oat01-{}` that shapes the visible phantom token so prefix-sniffing clients recognise it as valid while nono brokers the real secret behind it.

Shape a phantom credential token so a prefix-sniffing client (e.g. an Anthropic SDK checking for `sk-ant-`) still recognises it as valid while nono brokers the real secret.

```json
{
  "command_policies": {
    "credentials": {
      "claude-key": {
        "type": "ambient",
        "env_var": "ANTHROPIC_API_KEY",
        "format": "sk-ant-oat01-{}"
      }
    }
  }
}
```

— Named config field with concrete template value and worked JSON example. v0.75.0

02 `--strict-broker-path` flag hardens sandbox startup NEW

83

New sandbox arg `--strict-broker-path` causes nono to refuse to start, instead of merely warning, when a filesystem grant overlaps a directory on `PATH`, blocking planted-binary attacks on host processes outside the sandbox.

Harden a sandbox so it refuses to start if any filesystem grant overlaps a PATH directory — catch planted-binary risk in CI before the agent runs.

```
$ nono run --profile nolabs-ai/opencode --strict-broker-path --
opencode
```

— Named flag, clear before/after behaviour, and exact CLI usage. v0.75.0

03 **Profile composition via `--extends` in `nono proxy`** NEW

75

`nono proxy` gains the `--extends <PROFILE>` flag, giving it the same repeatable profile-composition semantics already available on `nono run`; it requires `--profile` to be set.

Compose an extra-domains profile on top of an existing proxy profile at runtime — useful when a specific session needs additional allow-listed hosts without permanently editing the base profile.

```
$ nono proxy -p default --extends extra-domains
```

— Named flag with runnable example, but composition mechanics not detailed further.

v0.75.0

THINNER COVERAGE BELOW

04 **PATH sanitization for host-side brokers** IMPROVED 50

Host-side credential and URL brokers now sanitize `PATH`, preventing bare-name binary resolution from grants that overlap `PATH` directories.

— Describes mechanism but offers no config surface or example to act on. v0.75.0

05 **Glob patterns for network allow/deny lists** IMPROVED

40

Profile and proxy network configuration now supports glob pattern matching in env var and hostname allow/deny lists.

— Names the surfaces affected but gives no syntax or example. v0.75.0

06 **Initial tool sandbox examples for `command_policies`** NEW

30

Ships an initial set of tool sandbox examples demonstrating `command_policies` configurations.

— Bare mention of examples shipped, no specifics given. v0.75.0

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · September 1, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue