THE CYBER TOOLCHAIN

NO. 120

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 25, 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 108 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 31 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.

PROTECT

Find outdated PHP from the outside, without a credentialed scan

vim ? help

ZAP's new passive rule reads server response headers and maps the disclosed PHP version to PHP2Shell and over 60 known CVEs, refreshed by a weekly feed, with built-in suppression so banner noise doesn't bury the real finding. Vigolium adds HTTP/3-aware fuzzing and fingerprint/hosts/dns recon subcommands, so the same pass covers QUIC-fronted hosts that HTTP/1.1-only tooling walks past.

OWASP ZAP · vigolium

IDENTIFY

See which editor extensions your developers actually run

Open VSX is where VS Code forks and AI coding clients pull extensions from, and almost nobody has an inventory of it. Manifold risk-rates roughly 17,000 of them nightly and correlates that against AI agents on the same endpoint; Token Security pairs its Grok Build editor integration with EDR data to surface agent credentials and shadow agents living outside the sanctioned workspace.

Manifold Security · Token Security

RESPOND

Pull evidence from encrypted and volume-managed disk images

Raw images from Linux and Android devices increasingly sit behind LVM2, LUKS or RAID layers that examiners had to unwrap manually before parsing could start. iLEAPP now reads those volumes directly, adds SHA-256 hashing of staged files via `--hash` for chain-of-custody, and ships over 30 new app and system artifacts.

iLEAPP

GOVERN

Give AI agents real identities instead of shared service credentials

Agents today mostly inherit a human's token or a long-lived API key, which makes attribution and revocation impossible. Auth0's CIMD support lets each agent pre-register or dynamically register as its own client with Application Profiles governing authorization centrally; JumpCloud's End User MCP Server moves routine identity chores into AI clients with admin and employee session controls over what the agent may do.

Auth0

◆ AI-Powered Offensive Security

PurpleAILAB Decepticon

Feedback ⓘ v1.6.0 NOTES ↗

RANK

Autonomous Hacking Agent for Red Team

> Decepticon v1.6.0 introduces a versioned update channel, persists per-thread AI and compaction settings, and improves checkpoint diagnostics, SSE recovery, and several smaller CLI, web, and container-level surfaces.

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

01 Versioned update channel via `decepticon update` IMPROVED

80

`decepticon update --channel latest` installs or updates to the latest version, pulling the matching launcher, configuration, and all versioned images in one step.

Upgrade an existing Decepticon installation to v1.6.0, fetching the new launcher, configuration, and all eight versioned images at once.

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

— Exact runnable command given, but no detail on image count or compatibility checks

v1.6.0

02 Persistent per-thread AI controls and compaction threshold

IMPROVED 60

Model, fallback, reasoning effort, and compaction controls now persist per thread across sessions, including a verified 4096-token compaction threshold.

— Names the token threshold but gives no config key or command to set it v1.6.0

THINNER COVERAGE BELOW

03 Checkpoint diagnostics and SSE recovery improvements IMPROVED

40

Improves scoped checkpoint diagnostics, SSE recovery, and transcript navigation across the CLI, web, and server surfaces.

— Names affected surfaces but no mechanism or before/after detail v1.6.0

04 Aligned web and server workspace roots IMPROVED 40

Aligns web and server workspace roots, preventing stale terminal socket handoffs.

— Explains the bug fixed but not how or where to verify it v1.6.0

05 **Persisted Plan and Red drafts across container recreation**

IMPROVED 40

Preserves private Plan and Red drafts and history across container recreation.

— States what persists but no mechanism or storage location v1.6.0

06 **Kali downloads via official CDN** IMPROVED 35

Switches Kali downloads to the official Kali CDN.

— States the source change with no further technical detail v1.6.0

07 **Localized slash commands in CLI** IMPROVED 25

Improves localized slash-command support inside the interactive CLI.

— Very thin description, no commands or languages named v1.6.0

WAS THIS USEFUL? 👍 👎

xalgorix

Feedback ⓘ v4.6.163 NOTES ↗ RANK

Autonomous AI pentesting agents — real-time reconnaissance, vulnerability detection, and exploitation orchestration. Go + TypeScript.

> xalgorix added opt-in native tool calling for llama.cpp, Ollama and vLLM backends, introduced stuck-scan recovery and lean prompt-schema flags, and expanded finding deduplication with new vulnerability classes.

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

01 **Native tool calling for local LLM backends** NEW 85

Adds the `XALGORIX_NATIVE_TOOL_CALLS` environment variable (default off) to enable structured function-calling instead of fragile XML tag parsing, supporting llama.cpp, Ollama, and vLLM backends so models not trained on XML can drive tools reliably during autonomous scans.

Enable native tool calling when using an Ollama backend to allow models not trained on XML to drive tools reliably during autonomous scans.

```
$ XALGORIX_NATIVE_TOOL_CALLS=true xalgorix --web
```

— Names exact env var, default state, and three supported backends. v4.6.163

02 **Stuck-scan recovery for autonomous runs** NEW 85

Adds `XALGORIX_ADVANCE_PAST_STUCK` (default off) to skip un-completable tasks and continue the rest of the scan plan, bounded by `XALGORIX_MAX_STUCK_ADVANCES` (default 3) consecutive no-progress advances before giving up, preventing an unresponsive endpoint from aborting an entire CI run.

Keep a long autonomous scan alive when one task stalls — useful in CI where an unresponsive endpoint would otherwise abort the entire run.

```
$ XALGORIX_ADVANCE_PAST_STUCK=true XALGORIX_MAX_STUCK_ADVANCES=5
xalgorix --web
```

— Two named env vars with defaults and a runnable example. v4.6.160

03 **Lean tool schema for smaller prompts** NEW 73

Adds `XALGORIX_LEAN_TOOL_SCHEMA` environment variable (default off) to strip illustrative example spans from tool and parameter descriptions, shrinking the per-

request prompt while preserving every tool, parameter and core instruction — useful for models with tighter context windows.

Reduce per-request prompt size when using a model with a tighter context window, without losing any tool or parameter coverage.

```
$ XALGORIX_LEAN_TOOL_SCHEMA=true xalgorix --web
```

— Named env var and purpose, but no size/token numbers given. v4.6.160

04 **Finding deduplication classes for recurring issues** IMPROVED

68

Adds a `debug_disclosure` finding class (stack-trace, debug-page and verbose-exception keywords; CWE-209) and a `user_enumeration` finding class (CWE-203/204) so these report once per target instead of re-reporting per path. Also collapses server-wide issue classes — debug/stack-trace disclosure, missing security headers, CORS policy, version/technology disclosure, and clickjacking — to a single finding per target instead of one per path.

— Names all classes and CWEs but offers no flag or command to act on. v4.6.160

WAS THIS USEFUL? 👍 👎

◆ Hacking Hardware

RFID Research Group ChameleonUltra

Feedback ⓘ

v2.3.0 NOTES ↗ RANK

The new generation Chameleon, built on the NRF52840, that emulates, reads, writes and decrypts RFID/NFC cards.

> ChameleonUltra v2.3.0 adds support for new card types including SEOS, MIFARE Ultralight C and Farpointe Pyramid, introduces parity checking during ISO 14443-4 sniffing, and improves LF FSK demodulation robustness and performance.

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

01 **Parity checks in ISO 14443-4 sniffing** NEW — 70

Adds `hf 14a sniff` parity checks to flag malformed frames during ISO 14443-4 sniffing sessions.

— Names exact command and what it checks for v2.3.0

02 **FSK demodulation robustness and performance** IMPROVED — 60

Adds DC-block FSK demodulation for improved weak-coupling robustness with HID and ioProx LF credentials, and removes double-precision `sqrt` from the FSK demodulation hot path to reduce LF processing overhead.

— Explains mechanism and named protocols but no user-facing action v2.3.0

THINNER COVERAGE BELOW

03 **New card type support: SEOS, Ultralight C, Pyramid** NEW — 45

Adds SEOS card emulation support, authentication and emulation for MIFARE Ultralight C, and support for Farpointe Pyramid credentials.

— Names three card types but no mechanism or usage detail given v2.3.0

04 **Low-level ISO 14443-4 API export** NEW — 45

Exports low-level ISO 14443-4 handling for use in client integrations.

— Names the protocol surface but no integration specifics v2.3.0

05 **Configurable long button press threshold** NEW — 30

Adds long button press threshold as a configurable user setting.

— Bare description, no config key or value range given v2.3.0

06 **Improved battery calibration and health reporting** IMPROVED

20

Improves battery calibration and health reporting.

— Generic improvement with no detail on mechanism or metrics v2.3.0

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

ZeroPath

Feedback ⓘ changelog · Oct 25, 2026 CHANGELOG ↗ RANK

ZeroPath is an AI-native SAST platform that finds 2x more real vulnerabilities with 75% fewer false positives.

> ZeroPath's sole release this window deepens its Security Compass/SD Elements integration, adding configurable verification-outcome-to-task-status mapping, new push APIs, an on-demand scan re-push action, and automatic verification notes on every push.

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

01 **On-demand scan re-push and push result stats** NEW 85

Adds a Push latest scan action to re-push a repository's latest finished full scan to SD Elements on demand, for example to apply updated task status mapping rules immediately. Adds a `lastPushStats` object with fields `countermeasuresMarkedDone`, `statusChangesRequested`, `passCount`, `partialCount`, `failCount`, and `notApplicableSkipped` to track Security Compass push results.

Re-push a repo's latest full scan to SD Elements after updating the task status mapping, to immediately reflect the new Pass/Partial/Fail→status rules.

```
$ pushLatestScan
```

— Names the action, its purpose, and all six tracked stat fields.

changelog-20261025-3822bab9

02 **Security Compass API and MCP methods** NEW 78

Adds `securityCompass.listTaskStatuses`, `updateProjectMapping`, and `pushLatestScan` API and MCP methods for the Security Compass integration. `securityCompass.listTaskStatuses` is restricted to users with the admin role, matching the permissions required for the mapping mutation.

List available SD Elements task statuses (admin only) before configuring a new project mapping.

```
$ securityCompass.listTaskStatuses
```

— Names three callable methods and their permission scope. changelog-20261025-3822bab9

03 Configurable SD Elements task status mapping BREAKING

75

Adds per-integration and per-project Task status mapping configuration that controls which SD Elements task status each ZeroPath verification outcome (Pass, Partial, Fail) sets. Integrations created before this change default to Pass→Complete, Partial→Incomplete, Fail→Incomplete, while newly created integrations start with verification notes only and no task status changes by default.

— Explains before/after default behavior but no exact config path.

changelog-20261025-3822bab9

THINNER COVERAGE BELOW

04 Verification notes on SD Elements push NEW

53

Writes one verification note per countermeasure (Pass, Partial, or Fail) with the finding list and links back to ZeroPath on every full scan push to SD Elements.

— Describes automatic behavior with no user-facing command. changelog-20261025-3822bab9

└──▷ BREAKING ON UPGRADE

! Integrations created before this change were given a default task status mapping of Pass→Complete, Partial→Incomplete, Fail→Incomplete; new integrations start with verification notes only and no task status changes by default.

WAS THIS USEFUL? 👍 👎

Mend.io Mend Platform

Feedback ⓘ changelog · Oct 25, 2026

CHANGELOG ↗ RANK

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

> Mend Platform's 26.9.3 release adds aggregated SBOM exports spanning multiple projects and promotes Custom Roles to general availability, alongside a new setting to exclude proprietary components from license workflows.

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

01 Aggregated SBOM export for applications and multi-project scopes

NEW 85

Adds aggregated SBOM export for application and multi-project scopes, consolidating multiple project-level files into one deduplicated document; supports CycloneDX and SPDX standards. Projects in a multi-project scope must belong to the same application, and per-project generation remains the default.

Generate a single deduplicated SBOM for an entire application scope to simplify ingestion into downstream tooling.

📍 In the Mend Platform UI, navigate to your Application, select Export SBOM, choose Aggregated scope, select CycloneDX or SPDX as the standard, and confirm the export.

— Names exact formats, scope rules and UI steps to use it changelog-20261025-8a74e6eb

02 Exclude proprietary libraries from license workflows NEW

75

Adds a new organizational setting, Include Proprietary Libraries in Workflows, to exclude proprietary (in-house) components from automation workflows that check for Empty License conditions.

Prevent proprietary in-house components from triggering false-positive Empty License workflow violations across the organization.

📍 In the Mend Platform UI, navigate to Organization Settings, find the Include Proprietary Libraries in Workflows setting, and toggle it off to exclude proprietary components from license-condition automation.

— Names exact setting and UI path to toggle it changelog-20261025-8a74e6eb

THINNER COVERAGE BELOW

03 **Custom Roles reaches general availability** IMPROVED

45

Promotes Custom Roles from closed beta to general availability, with an enhanced role breakdown and a Create Role wizard for custom role management.

— States GA status and new wizard but no config detail changelog-20261025-8a74e6eb

WAS THIS USEFUL? 👍 👎

SonarSource SonarQube Feedback ⓘ announced · Oct 22, 2026

BLOG ↗

Sonarqube is a enforce consistent quality gates across every pull request and pipeline, from IDE to CI/CD, with automated code review for bugs and security issues.

SonarQube Cloud gains a native `sonar-container-scanner` CLI (open beta) with `--fail-on` threshold gating and a new Container images org sidebar view.

NEW VULNERABILITY CHECK NEW SUPPORTED PLATFORM Introduces `sonar-container-scanner <IMAGE> --organization <ORGANIZATION_KEY> --token <TOKEN>` CLI to scan container images for known OS-layer vulnerabilities and malicious packages, generating an SBOM sent to SonarQube Cloud for analysis.

Adds `--fail-on` flag to `sonar-container-scanner` to set a minimum severity threshold (default: `high`) at which the scanner exits with a failing code, enabling CI pipeline gates on container image risk.

Adds a Container images section in the SonarQube Cloud organization sidebar, listing every analyzed image with its tag, risk count, most severe risk, and last-analyzed timestamp alongside code and dependency risks.

Supports automatic re-checking of previously analyzed images against newly published vulnerabilities without requiring a re-scan, based on organization re-analysis settings.

Available at no additional cost for SonarQube Cloud organizations on Team and Enterprise plans with SonarQube Advanced Security.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

sbomify

Feedback ⓘ v26.10.0 NOTES ↗ RANK

sbomify is a product security artifact hub and a trust center.

> sbomify v26.10.0 ships a new Vulnerabilities page with an AI-agent VEX approval workflow, an Operator role, a Vanta compliance integration, and live WebSocket updates, alongside a large set of breaking API, permission, and self-hosted deployment changes.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 API permission and scope tightening

BREAKING — 95

Several endpoints now enforce stricter permissions and validation: `POST /api/v1/licensing/custom-licenses` is removed; creating an access request now requires authentication and takes no request body, with all access-request routes enforcing token scopes; `GET /workspaces/` now shows members, invitations, and billing data only to admins; the branding `PATCH` endpoint now only accepts `null` for icon and logo fields; `PATCH /advisories/{id}` now rejects unknown fields with 422; and `/releases/{id}/artifacts?mode=available` now requires the `release:manage` scope.

— Every affected endpoint and its new behavior is named v26.10.0

02 Self-hosted deployment tuning env vars

NEW — 90

Adds `REDIS_CACHE_MAX_CONNECTIONS` and `REDIS_CACHE_POOL_TIMEOUT` to tune Redis connection pooling, `LOG_MAX_SIZE` and `LOG_MAX_FILE` for log rotation control, and `KEYCLOAK_PUBLIC_URL` for self-hosted Keycloak configuration.

Tune Redis connection pool limits to support higher concurrency in a self-hosted instance.

```env
REDIS_CACHE_MAX_CONNECTIONS=50
REDIS_CACHE_POOL_TIMEOUT=10
```

— All env vars named with a concrete config example v26.10.0

03 History limit on assessments API

BREAKING — 85

Adds `history_limit` query parameter and `all_runs_total` field to `GET /plugins/assessments/{sbom_id}`, while capping `all_runs` at 50 runs in the default response.

— Exact endpoint, params and cap value are specified v26.10.0

04 **Findings table migration with backfill command** BREAKING

85

The findings migration drops the `vulnerability_scanning_findings` table, which causes old pods in a rolling deploy to fail; run `manage.py backfill_findings --missing-only` to rebuild findings.

— Gives the exact table and a runnable remediation command v26.10.0

05 **Vanta integration for compliance sync** NEW

80

Adds a Vanta integration, enabled via new `VANTA_*` environment variables, to sync frameworks and controls from Vanta with optional Trust Center publishing.

Enable Vanta integration in a self-hosted deployment to sync compliance frameworks and controls.

```env
VANTA_*=<your-vanta-credentials>
```

— Exact env vars and a runnable config example are given v26.10.0

06 **Self-hosted deployment breaking changes** BREAKING

80

Environment variables `TEAMS_*_ROLES` are renamed to `WORKSPACE_*_ROLES` (old names still work as aliases for now); `KEYCLOAK_WEBHOOK_SECRET` is removed; custom `DRAMATIQ_BROKER` overrides must be updated; custom domains now require outbound HTTPS from workers; rotating `SECRET_KEY` requires domains to be re-verified; and the new `ws/user/` WebSocket route must be forwarded in reverse proxy configuration.

— Names every affected env var and infra requirement v26.10.0

07 **VEX approval queue for AI agents** NEW

65

Adds a VEX approval queue where AI agents submit VEX rulings using a new AI-agent token, which a human then approves or rejects on the Review tab.

— Mechanism and UI location (Review tab) are given v26.10.0

THINNER COVERAGE BELOW

08 **Irreversible access token hashing migration** BREAKING

55

Migration `access_tokens.0012` changes token hashing and cannot be rolled back — old releases cannot authenticate API tokens after it runs.

– Migration named but no remediation command given v26.10.0

09 **Pending Owner invitations now grant Admin** BREAKING 55

Pending Owner invitations now grant Admin instead of Owner; affected invitations must be re-sent.

– States the behavior change and the remediation action v26.10.0

10 **Vulnerabilities page with findings dashboard** NEW 50

Adds a Vulnerabilities page showing a deduplicated list of open findings, with release and SBOM-version filters and Fix/Triage prompts for coding agents.

– Names the page and its filters but no exact UI path or API v26.10.0

11 **Operator role** NEW 50

Adds an Operator role that grants read access to the entire workspace and the ability to triage findings, without permission to change anything else.

– Scope of the role is clear but no assignment steps given v26.10.0

12 **Live UI updates over WebSocket** NEW 50

Adds live UI updates over WebSocket via the new `ws/user/` route; self-hosted deployments must forward this route in their reverse proxy configuration.

– Route named but no further mechanism or usage detail v26.10.0

13 **Stripe sync demotes ended paid plans** BREAKING 50

The first Stripe sync moves ended paid plans to the Community plan, making their private products and components public.

– Explains effect but no mitigation steps given v26.10.0

14 **Advisory management improvements** IMPROVED 45

Advisories gain per-product status editing after an advisory is created, bulk delete for advisories, custom ID prefix support for advisories, and Markdown support for advisories.

– Lists four additions without endpoints or UI paths v26.10.0

15 **Invitations API** NEW 25

Adds an Invitations API on both v1 and v2.

– No endpoint paths or fields given v26.10.0

16 **Assessments requirements matrix view** NEW 20

Adds an assessments requirements matrix view.

— Bare name, no further description `v26.10.0`

└──▷ BREAKING ON UPGRADE

- ! Token hashing migration `access_tokens.0012` cannot be rolled back — old releases cannot authenticate API tokens after it runs.
- ! The findings migration drops the `vulnerability_scanning_findings` table; old pods in a rolling deploy will fail. Run `manage.py backfill_findings --missing-only` to rebuild findings.
- ! `POST /api/v1/licensing/custom-licenses` is removed.
- ! Creating an access request now requires authentication and takes no request body; all access-request routes enforce token scopes.
- ! `GET /plugins/assessments/{sbom_id}` caps `all_runs` at 50; use `history_limit` and `all_runs_total` for full history.
- ! `GET /workspaces/` now shows members, invitations, and billing data only to admins.
- ! The branding `PATCH` endpoint now only accepts `null` for icon and logo fields.
- ! `PATCH /advisories/{id}` now rejects unknown fields with 422.
- ! `/releases/{id}/artifacts?mode=available` now requires the `release:manage` scope.
- ! Environment variables `TEAMS_*_ROLES` are renamed to `WORKSPACE_*_ROLES`; old names work as aliases for now.
- ! `KEYCLOAK_WEBHOOK_SECRET` environment variable is removed.
- ! Custom `DRAMATIQ_BROKER` overrides must be updated.
- ! Custom domains now require outbound HTTPS from workers; rotating `SECRET_KEY` requires domains to be re-verified.
- ! The first Stripe sync moves ended paid plans to Community, making their private products and components public.
- ! Pending Owner invitations now grant Admin instead of Owner — re-send affected invitations.
- ! The new `ws/user/` WebSocket route must be forwarded in reverse proxy configuration.

WAS THIS USEFUL?  

cyclonedx-core-java

Feedback ⓘ cyclonedx-core-java-13.4.0

NOTES ↗ RANK

CycloneDX SBOM Model and Utils for Creating and Validating BOMs

> cyclonedx-core-java 13.4.0 moves its full set of deserializers to a streaming architecture for faster, lower-memory SBOM parsing, and drops its transitive dependency on packageurl-java.

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

01 **packageurl-java dependency dropped** BREAKING 60

The `packageurl-java` dependency has been dropped from cyclonedx-core-java; any code relying on transitive access to `packageurl-java` classes via cyclonedx-core-java must now declare that dependency directly.

— Names exact dependency and required migration action cyclonedx-core-java-13.4.0

THINNER COVERAGE BELOW

02 **Streaming deserializers across BOM parsing** IMPROVED 50

Hash and external reference list deserialization now use a unified, streaming path to cut memory overhead, and the remaining simple list, nested component, pedigree, choice, entity, extension, and risk deserializers were converted to the same streaming, low-allocation approach for consistency.

— Explains what changed and lists affected deserializer types, no benchmarks given

cyclonedx-core-java-13.4.0

└──▷ **BREAKING ON UPGRADE**

! The `packageurl-java` dependency has been dropped; any code relying on transitive access to `packageurl-java` classes via `cyclonedx-core-java` must now declare that dependency directly.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

OWASP ZAP

Feedback ⓘ announced · Oct 22, 2026 BLOG ↗ RANK

The ZAP by Checkmarx Core project

> ZAP introduces a passive scan rule that detects PHP2Shell and over 60 known PHP CVEs directly from response headers, with built-in false-positive suppression and an automated weekly vulnerability feed.

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

01 PHP2Shell and PHP CVE detection passive scan rule NEW

85

A new passive scan rule inspects response headers (no extra requests sent) to detect PHP2Shell and 60+ known PHP CVEs, raising a PHP Version with Known Vulnerabilities alert scoped to one alert per site/technology/version, risk-rated by the most severe applicable CVE and listing all affected CVEs with severity and links to a new PHP Known Vulnerabilities help page. False positives (e.g. patched servers) can be suppressed via an Alert Filter matching URL, the PHP Version with Known Vulnerabilities alert name, and the specific version evidence string (e.g. `PHP/8.1.2`), configurable globally, in a context, or in the Automation Framework for CI/CD pipelines. The vulnerability list is kept current via a weekly automated job pulling from official php-src GitHub security advisories, which raises a pull request to update the add-on; the same rule is designed to support additional server-side technologies beyond PHP in future releases.

Suppress a false-positive alert on a host where you know PHP security backports have been applied, without silencing future version changes.

📍 In the ZAP UI, go to Analyse › Alert Filters › Add, set the URL regex to the target site (e.g. `https://www.example.com/.*`), select alert name PHP Version with Known Vulnerabilities, set Evidence to the exact reported version string (e.g. `PHP/8.1.2`), and set New Risk to False Positive.

— Names exact alert, config fields, and UI steps for suppression `launch-20261022-e571b26a`

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ALEAPP

Feedback ⓘ v2026.4.4 NOTES ↗ CODE ↗

Android Logs Events And Protobuf Parser

Adds `--hash` for SHA-256 file staging with container timestamps, WAL-mode SQLite `immutable=1` fallback, and single-file portable Windows releases.

└──▷ GET THIS VERSION

```
$ git clone --branch v2026.4.4 https://github.com/abrignoni/ALEAPP.git
# already have the repo? check out this version:
$ git checkout v2026.4.4
```

Adds `--hash` flag (and a matching GUI checkbox) to record SHA-256 hashes for staged files, capturing container-recorded timestamps from ZIP NTFS/extended timestamp fields, tar headers (GNU, PAX), L01/AD1 entries, FAT/exFAT entries, and ext/APFS/HFS+ times.

Reads Cellebrite CLBX ZIP times from `metadata.msgpack` (adding msgpack as a dependency), and extracts creation time from the fourth value of four-value extended timestamp blocks.

A WAL-mode SQLite database staged without its `-wal` file is now opened read-only with `immutable=1`.

Tar links are now read as the member they point to; corrupt ZIP members are skipped rather than halting the run.

Windows releases ship as a single portable file per architecture (`-windows-x64-portable.zip`, `-windows-arm64-portable.zip`) with no installer required.

- NEW DETECTION Yandex Disk artifact module
- NEW DETECTION KeepSafe artifact module
- NEW DETECTION Google Pay artifact module
- NEW DETECTION Offline DJI TXT flight records artifact module
- NEW DETECTION Shortcut Service artifact module
- NEW DETECTION Chrome Cache - Videos artifact module
- NEW DETECTION Chromium downloads matched to MediaStore
- NEW DETECTION Google app suggestion and key-value caches artifact modules
- NEW DETECTION Discord - Servers artifact module

UPDATED DETECTION Google ODLH gains new Path Points artifact and additional activity fields

UPDATED SUPPORTED PLATFORM Raw image input adds LVM2, LUKS, RAID, FAT12/FAT16 support

└──▷ **BREAKING ON UPGRADE**

- ! Python 3.11 or newer is now required to run ALEAPP from source (previously Python 3.10 was sufficient).
- ! There is no `aleappGUI` executable any more — `aleapp` without arguments now opens the window.
- ! Windows releases no longer ship an installer; the distribution is a single portable file (`aleapp.exe`) that must be placed wherever a calling tool expects it.

WAS THIS USEFUL? 👍 👎

iLEAPP

Feedback ⓘ v2026.4.5 NOTES ↗ CODE ↗ RANK

iOS Logs, Events, And Plist Parser

> iLEAPP v2026.4.5 extends raw-disk image parsing to LVM2, LUKS and RAID volumes, adds SHA-256 hashing of staged files via a new `--hash` flag, and ships over 30 new app and system artifacts, while also moving to Python 3.11+ and retiring the separate GUI executable.

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

01 Raw image support for LVM2, LUKS, RAID, FAT IMPROVED

90

Extends `-t raw` support to LVM2 logical volumes (split, striped and RAID1), LUKS containers opened with a passphrase (via `--image_password_file`, `--image_password_env`, a prompt, or the GUI), RAID member images read together, and FAT12/FAT16 filesystems; also exposes free space, flash-filesystem deleted files, and NTFS directory indexes to artifact modules that request them.

— Names exact flags and volume types but no worked example v2026.4.5

02 SHA-256 staging hashes via --hash NEW

80

Adds `--hash` flag (and a matching GUI checkbox) to record a SHA-256 digest for each staged file, covering zip NTFS and extended timestamp fields, tar headers (GNU and PAX), L01 and AD1 entries, FAT/exFAT entries, and ext/APFS/HFS+ times.

— Concrete flag and coverage list, but no sample output shown v2026.4.5

03 Python 3.15 support and portable Windows builds IMPROVED

75

Supports Python 3.15 (including Windows wheels for Python 3.15 and ARM64) and packages Windows releases as a single portable file per architecture (`-windows-x64-portable.zip` and `-windows-arm64-portable.zip`) with no installer.

— Exact artifact filenames named, straightforward to use v2026.4.5

04 30+ new app and system artifacts NEW

70

Adds new artifacts: logarchive entries from sysdiagnose for unlock, buttons, orientation, gestures, USB host, wake source, AirDrop, calls, weather, family hotspot, charging state, app snapshots, container entries, and data usage; logarchive entries for URL and link opens, camera and microphone capture sessions, TCC authorization cache, location manager state, Screen Time (`ManagedSettings`), `BatteryUI` per-app runtime, and Container Manager State for iOS 27; 32 new Life360 artifacts (including crash-detection diagnostics and recovered history) plus generic Mixpanel and Crashlytics artifacts; and

Apple Watch photo collections, registry history, Mail registry, Weather Cities, Photos Graph Moment Locations, routined visits and vehicle events, Cache Locations (course, altitude and accuracy), KeepSafe photo vault, Reddit post and subreddit history, Anonymous Chat, Emoji Recents/Usage/Preferences, an account-wide encryption key hash artifact, offline DJI TXT flight records, Discord Servers, Battery Center, Alarms sleep alarms, Apple Music search, Media Experience, Media Remote, Photos Recent Searches, Hashtag Images, and supplementary WebKit cache parsers.

— Exhaustive artifact list but no extraction command shown v2026.4.5

---

05 **Cellebrite and GrayKey timestamp parsing** NEW 70

Reads Cellebrite CLBX zip timestamps from `metadata.msgpack` (adding msgpack as a dependency), and reads the creation time from the fourth value of four-value extended timestamp blocks (GrayKey-style `files_full.zip`).

— Names exact source files and mechanism, no usage walkthrough v2026.4.5

THINNER COVERAGE BELOW

06 **GUI executable removed, unified entry point** BREAKING 55

There is no `ileappGUI` executable any more; `ileapp` without arguments opens the window instead.

— Exact replacement command given for the removed binary v2026.4.5

---

07 **Python 3.11 minimum requirement** BREAKING 40

Python 3.11 or newer is now required to run iLEAPP from source; Python 3.10 is no longer supported.

— Clear version requirement, no migration detail beyond that v2026.4.5

---

08 **Logarchive extraction from within sysdiagnose** NEW 25

Reads logarchive from within a sysdiagnose archive.

— One-line description with no mechanism or path given v2026.4.5

└──▷ BREAKING ON UPGRADE

! Python 3.11 or newer is now required to run iLEAPP from source (previously Python 3.10 was supported).

! There is no `ileappGUI` executable any more; `ileapp` without arguments opens the window instead.

WAS THIS USEFUL?

NETWORK SECURITY

◆ Email Security

Material Security

Feedback ⓘ announced · Oct 25, 2026

BLOG ↗ RANK

Material Security is a SaaS platform that protects enterprise email and cloud file collaboration by detecting and remediating data exfiltration, phishing, and account takeover threats.

> Material Security launched a Google Drive Management Agent that automatically investigates sensitive files and remediates over-permissioned sharing based on configurable verdicts.

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

01 **Google Drive Management Agent with verdict-based remediation**

NEW 75

The new Drive Management Agent investigates sensitive Google Drive files across content summary, sensitivity, and collaborator exposure, then assigns one of four verdicts — Cleared, Contained, Exposed, or Overexposed — to determine automated remediation actions. Remediations are configurable per verdict: for example, Exposed files can trigger an owner notification followed by a grace period before access revocation, while Overexposed files can trigger immediate access revocation. Manual search, sort, and filter remain available across the full set of flagged Drive documents for investigations requiring human review.

— Clear mechanism and verdict logic, but no UI path or config key named.

launch-20261025-555e430c

WAS THIS USEFUL?

◆ Network Analysis & Monitoring

ecapture

Feedback ⓘ v2.9.0 NOTES ↗ RANK

Capture SSL/TLS plaintext with eBPF—no MITM proxy or custom CA installation. Supports Linux and Android on x86_64 and arm64.

> eCapture v2.9.0 adds a pcapng-based data source for text mode, letting the openssl/tls, gotls, and gnutls modules reconstruct full HTTP/1.1 and HTTP/2 traffic without switching to pcap mode.

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

01 **pcapng text source for TLS modules** NEW 87

The `openssl` / `tls`, `gotls`, and `gnutls` modules gain a `--text-source pcapng` flag that reconstructs HTTP/1.1 and HTTP/2 request/response pairs with network metadata in text mode. Packet timestamps are converted from TC monotonic time to Unix time before pcapng serialization, and a raw pcapng copy is saved alongside text output only when `--pcapfile` is also specified together with `--text-source pcapng`. `hook` remains the default source; existing scripts relying on prior text-mode behaviour should be validated against this new code path.

Reconstruct full HTTP/1.1 and HTTP/2 request/response pairs with network metadata in text mode — without switching to pcap mode — by feeding pcapng DSB blocks through the decrypt SDK.

```
$ sudo ecapture tls -m text --text-source pcapng
```

Reconstruct HTTP/2 traffic in text mode AND save a raw pcapng copy for later Wireshark analysis in the same run.

```
$ sudo ecapture tls -m text --text-source pcapng --
pcapfile=ecapture.pcapng
```

— Names exact flags and modules with runnable commands and mechanism. v2.9.0

└──▷ BREAKING ON UPGRADE

! The `openssl` / `tls`, `gotls`, and `gnutls` modules now accept `--text-source pcapng` as a data source in text mode; `hook` remains the default, but existing scripts that relied on text-mode behaviour should be validated against the new pcapng code path.

WAS THIS USEFUL?  

◆ Firewalls & Network Defense

Palo Alto Networks

Feedback ⓘ changelog · Oct 21, 2026

BLOG ↗ RANK

Palo Alto Networks is a cybersecurity company offering firewalls, cloud security, and threat-prevention products that protect networks, endpoints, and cloud environments from attacks.

> Palo Alto Networks launched Unit 42 Managed Agentic Endpoint Security, a 24/7 managed service built on Cortex AES, and extended AI Runtime Security to run natively on NVIDIA BlueField DPUs, alongside new deployment, reporting, and identity-governance capabilities for agentic AI.

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

01 AI Runtime Security on NVIDIA BlueField DPUs (GA) NEW

65

AI Runtime Security on NVIDIA BlueField DPUs is now generally available, isolating security processing in dedicated infrastructure silicon so policies are enforced independently of a potentially compromised host, across AI factory and industrial edge deployments.

— Clear mechanism and scope, but no actionable config or command. 2026-10-21-900cd124

02 Managed AES with Standalone Cortex AES deployment option NEW

65

A Managed AES with Standalone Cortex AES deployment option allows organisations to run Cortex AES alongside an existing endpoint security solution to gain AI-native visibility and governance without replacing incumbent tooling.

— Named deployment option gives a concrete starting point. 2026-10-21-900cd124

THINNER COVERAGE BELOW

03 Unit 42 Managed Agentic Endpoint Security service NEW

55

Introduces Unit 42 Managed Agentic Endpoint Security (Managed AES), a 24/7 managed service built on Cortex AES that adds Unit 42 threat hunting, investigation, response, and AI governance expertise to the existing Cortex AES platform.

— Names the service and platform but no setup steps given. 2026-10-21-900cd124

04 Agentic identity governance with NVIDIA Sentry on BlueField DPUs

NEW

50

An integration between agentic identity governance and NVIDIA Sentry on BlueField DPUs moves identity and delegated authority controls to silicon, separate from the host OS, to provide zero-trust protection and tamper-resistant auditing for agents; availability is pending.

— Mechanism described but feature availability is not yet live. 2026-10-21-900cd124

05 **AI Security Posture and Threat Impact Reports** NEW

45

Managed AES adds AI Security Posture Reports and Threat Impact Reports as deliverables, providing periodic policy and control improvement recommendations and threat response summaries.

— Report names given but no detail on access or format. 2026-10-21-900cd124

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

cicd-sensor

Feedback ⓘ releases/v0.0.49 NOTES ↗

Open-source eBPF runtime security sensor for GitHub Actions and GitLab CI/CD.

Agents can now authenticate to the Manager via GitHub Actions OIDC, removing the need to store long-lived Manager tokens.

└──▷ GET THIS VERSION

```
$ git clone --branch releases/v0.0.49 https://github.com/cicd-sensor/cicd-sensor.git
# already have the repo? check out this version:
$ git checkout releases/v0.0.49
```

NEW DETECTION litellm 1.107.0 PyPI compromise exfiltration domain IOC rule

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

File Security Storage template version 2.6.0 for Azure (2026-10-26) BLOG ↗

RANK

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

> Trend Vision One's File Security Storage Azure deployment template now supports deploying into a customer-managed Azure VNet.

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

01 Customer-managed VNet support in File Security Storage NEW

40

The File Security Storage Azure deployment template (version 2.6.0) adds optional customer-managed Azure VNet support, letting the storage deployment run inside a customer-controlled virtual network.

— Names the exact template and feature but no setup steps or parameters.

File Security Storage template version …

NEW SUPPORTED PLATFORM Runtime Malware Scanning for Amazon ECS

> Runtime Malware Scanning for Amazon ECS (2026-10-21) BLOG ↗ STABLE
>
> Terminate mitigation for ECS stops affected tasks and blocks further launches of flagged images, using the same filter rules and log/terminate options as Kubernetes.

WAS THIS USEFUL? 👍 👎

osctrl

Feedback ⓘ v0.6.1 NOTES ↗ RANK

Fast and efficient osquery management

> osctrl v0.6.1 adds per-node osquery config/flag overrides, osquery version management with gradual rollout and per-host pinning, and new osctrld operational monitoring, extension delivery and flag serving, plus a default AWS credential chain for the S3 logger.

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

01 **osquery version management with per-host pinning** NEW

70

Adds osquery version management including a version library, an environment-level version target, per-host version pins, upgrade reporting, and a pin retry rule.

– Mechanism named but no exact UI path or API given v0.6.1

02 **osctrld operational monitoring and health reporting** NEW

70

Adds rate limiting, audit logging, and metrics for all `osctrld` endpoints, plus `osctrld` health reports surfaced via a node-page banner and alertable metrics.

– Names endpoints and banner but no exact metric names or alert config v0.6.1

03 **Per-node osquery config and flag overrides** NEW 60

osctrl now supports per-node overrides for both `config` and `flags`, letting operators apply node-specific osquery configuration and flags independently of the environment default.

– Names config and flags surfaces but no exact UI path or endpoint v0.6.1

THINNER COVERAGE BELOW

04 **Gradual rollout for osquery version targets** NEW 50

Adds gradual rollout for the environment's osquery version target, with cached rollout progress and a node-page stage indicator showing rollout status per host.

– Describes behavior but no config key or command given v0.6.1

05 **S3 logger default AWS credential chain** IMPROVED 45

The S3 logger now uses the default AWS credential chain when no static keys are configured, removing the requirement to supply explicit credentials.

– Clear behavior change but no config key named v0.6.1

06 **osctrld extension delivery and flag serving** NEW 40

`osctrld` can now push osquery extensions down to enrolled nodes and serves the environment's stored `flags` to enrolled nodes; the UI now hides `osctrld` `-only` elements when `osctrld` is not enabled.

— Thin one-line additions without further mechanism detail v0.6.1

WAS THIS USEFUL? 👍 👎

Palo Alto Cortex XDR

Feedback ⓘ changelog · Oct 19, 2026

BLOG ↗ RANK

Palo Alto Networks Cortex XDR is an extended detection and response platform that unifies endpoint, network, and cloud telemetry to detect, investigate, and stop threats.

> Palo Alto Networks launched Cortex XDR One, merging Cortex XDR with Cortex Agentic Endpoint Security into a single sensor and console to extend detection and response to AI agents alongside traditional endpoints.

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

01 **Real-time AI artifact inventory** NEW 59

Delivers real-time inventory of every AI artifact (agents, MCP servers, AI co-pilots, browser tools) the moment it appears on an endpoint, enabling security teams to control which packages, extensions, agents, and AI artifacts are allowed to be installed or used.

— Lists concrete artifact types but no UI/API path given 2026-10-19-7ef992b2

02 **Cortex XDR One unified platform** NEW 53

Cortex XDR One combines Cortex XDR and Cortex Agentic Endpoint Security (AES) under a single sensor and console, giving centralized visibility, management, and investigation across both traditional endpoints and AI agents.

— Names the merged products and console but no setup steps 2026-10-19-7ef992b2

03 **Managed AES with Unit 42** NEW 43

Managed AES with Cortex XDR One provides 24/7 Unit 42-staffed governance, threat hunting, investigation, and response across the full endpoint attack surface including AI agents.

— Describes the managed service scope but no operational detail 2026-10-19-7ef992b2

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

World Monitor

Feedback ⓘ v2.11.4 NOTES ↗

Real-time global intelligence dashboard. AI-powered news aggregation, geopolitical monitoring, and infrastructure tracking in a unified situational awareness interface

Pins Iran–GCC geopolitical tension pairs with live scoring over a rolling 168-hour window.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.11.4
https://github.com/koala73/worldmonitor.git
# already have the repo? check out this version:
$ git checkout v2.11.4
```

NEW INTEGRATION FAA NOTAM ingestion via seed-notam-nms

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ announced · Oct 22, 2026 BLOG ↗ RANK

Tines lets teams put AI to work across their most sensitive systems, data, and processes, with the security and governance to do it safely.

> Tines' September 2026 release ships a faster v2 Records API, a new MCP server for AI clients, expanded AI model choice, and several case-management and admin usability upgrades.

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

01 v2 Records API for cases and child records IMPROVED

60

The v2 Records API enables single-call retrieval of related cases and child record data, with record searches running twice as fast as before.



— Names the API and performance gain but no endpoint syntax launch-20261022-5eb9a1cc

THINNER COVERAGE BELOW

02 Expanded AI model selection NEW

50

Adds AI model selection, supporting OpenAI GPT-6 Astra, Sol, and Luna, and Anthropic `Claude Opus 5.5`.

— Names exact models but not how selection is configured launch-20261022-5eb9a1cc

03 Tines MCP server for AI clients NEW

45

Adds the Tines Model Context Protocol (MCP) server, letting compatible AI clients search and summarize records and manage cases.



– Describes capability but no setup or connection details launch-20261022-5eb9a1cc

04 **Case and Agent Management via API** NEW 45

Adds case field updates and comment-change notifications, and agent skill management—including renaming a skill without reconnecting AI Agent actions that reference it—through the Tines API.

– Names the API surface but no endpoint or payload detail Explains behaviour change but no endpoint syntax given

launch-20261022-5eb9a1cc

05 **Markdown formatting formulas for story data** NEW 35

Adds eight new markdown formulas for formatting story data into headings, lists, and other structures without writing markdown syntax manually.

– Mentions count and purpose but no formula names launch-20261022-5eb9a1cc

06 **Bulk file download from case blocks** NEW 35

Adds the ability to download all files from a case block together for faster evidence handoff.

– Clear user action but no UI path or limit given launch-20261022-5eb9a1cc

07 **Case queue column controls and priority labels** IMPROVED 25

Adds column controls and priority labels to case queues to keep important details visible.

– Brief UI description, no navigation specifics launch-20261022-5eb9a1cc

08 **Admin controls for audit retention and AI usage** IMPROVED 20

Gives admins controls over how long audit data is retained and a clearer view of AI usage.

– Generic description, no config keys or thresholds named launch-20261022-5eb9a1cc

WAS THIS USEFUL?  

◆ Detection Engineering

Tenzir

Feedback ⓘ announced · Oct 22, 2026 BLOG ↗ RANK

Tenzir is a data pipeline engine for security teams that collects, transforms, enriches, and routes telemetry.

> Tenzir shipped a new isMalicious threat-intelligence integration built on open STIX 2.1/TAXII 2.1 standards, adding Library packages for in-stream indicator enrichment, local lifecycle tracking of indicators, and on-demand reputation lookups.

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

01 isMalicious TAXII enrichment via Tenzir Library NEW

78

Adds `ismalicious`, `stix`, and `tenzir` packages to the Tenzir Library, enabling TAXII collection pulls from isMalicious and local indicator lookup without per-event API calls. Fetches TAXII collections (starting with `malicious-domains` and `malicious-ips`) and maps each STIX Indicator to OCSF OSINT Inventory Info for in-stream DNS and network event enrichment, keeping `risk` and `confidence` in separate OCSF `osint` fields per match to preserve the distinction between severity and source certainty; unmatched events are marked unknown, not safe.

— Names packages and OCSF mapping but no exact command or endpoint shown

launch-20261022-78b305a2

THINNER COVERAGE BELOW

02 Local indicator lifecycle tracking for isMalicious feeds NEW

59

Tracks indicator lifecycle locally: enforces `valid_until` expiry and removes revoked indicators after incremental polls using `added_after` between scheduled walks.

— Explains mechanism with named fields but no runnable config shown

launch-20261022-78b305a2

03 On-demand isMalicious API lookups NEW

54

Supports optional on-demand API lookups of single indicators against the isMalicious API during investigations, returning the full report as an event; these lookups count against the monthly reputation-check quota.

— Describes behaviour and quota impact but no concrete call syntax

launch-20261022-78b305a2

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Ory

Feedback ⓘ changelog · Oct 25, 2026 CHANGELOG ↗ RANK

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

> Ory shipped configurable one-time code formatting, external-ID event matching, admin API attribution in event logs, and a redesigned Console layout rolled out to all workspaces.

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

01 Redesigned Ory Console layout and navigation IMPROVED

50

Ory introduces a new layout and navigation for the Ory Console, rolled out automatically to all Ory Network workspaces and projects with no migration or opt-in required.

Navigate the redesigned Ory Console — no migration or opt-in required, it is live for every workspace.

> 📍 1. Go to `https://console.ory.sh` 2. Log in to your Ory Network account 3. Select any workspace or project — the new layout and navigation are active automatically

— Clear navigation path given, but no detail on what changed in layout.

changelog-20261025-b670a14d

02 External-ID matching in event stream NEW

48

The Ory Network event stream now supports external-ID matching for identity, login, registration, recovery, settings, verification, and session events.

— Names all seven event types but no mechanism or config shown.

changelog-20261025-b670a14d

03 Admin API attribution in event logs IMPROVED

38

Events emitted by Ory Identities, Ory OAuth2, and Ory Permissions now show who made an admin API request.

— Names three products affected but no field or log format detail.

changelog-20261025-b670a14d

04 Configurable one-time code length and charset NEW

32

Ory Kratos Identities via Ory Network now support configurable length and character set for one-time codes (OTP).

— No config key or exact values given, just prose description.

changelog-20261025-b670a14d

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ announced · Oct 23, 2026 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 launched an End User MCP Server that lets employees run self-service identity tasks — profile updates, SSO app lookup/launch, access requests, approvals, and password/MFA status checks — from AI clients like ChatGPT, Claude, and Cursor, with admin and employee-level session controls.

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

01 End User MCP Server for AI-client self-service NEW

85

JumpCloud added an End User MCP Server built on the Model Context Protocol, letting employees perform User Portal self-service (profile updates, SSO app lookup and launch, access requests, approval tasks, password/MFA status review) from compatible AI clients such as ChatGPT, Claude, and Cursor. Admins enable it org-wide from `Settings > JumpCloud AI` in the Admin Portal, configuring session duration and automatic revocation, and can view and bulk-revoke active sessions there (connected client, client ID, redirect URI, employee, creation time, last-used time, expiry); employees can view and individually revoke their own sessions from `Security > JumpCloud MCP` in the User Portal. Connections authenticate via OAuth only (no API-key auth), each session scoped to the connecting employee's existing JumpCloud permissions, and sensitive credential actions — password reset/change, MFA enrollment/removal, SSH key management, device enrollment — are explicitly excluded from MCP scope and remain User Portal-only.

Enable the End User MCP Server for your org so employees can start connecting AI clients to their User Portal accounts.

> 📍 In the Admin Portal, go to Settings › JumpCloud AI and enable the End User MCP Server. Set the session duration to control how long connections remain active before automatic revocation.

Audit and revoke suspicious or stale end-user AI client connections across the organization.

> 📍 In the Admin Portal, go to Settings › JumpCloud AI. Review the list of active user sessions (client name, client ID, redirect URI, employee, last-used date, expiry). Select one or more sessions and choose Revoke to terminate them immediately.

Let an employee disconnect an AI client they no longer use without involving IT.

> 📍 In the User Portal, go to Security › JumpCloud MCP. Locate the session for the client to disconnect and select Revoke.

— Exact admin/user UI paths and scope limits given, but no API/flag names.

launch-20261023-80ebbfb8

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ announced · Oct 22, 2026 BLOG ↗ RANK

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

> Auth0 introduced CIMD-based identity management for AI agents, including manual client pre-registration, an Early Access dynamic registration flow, and Application Profiles for centrally managing security and authorization settings across CIMD clients.

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

01 **Application Profiles for CIMD client management** NEW

68

Adds Application Profiles (Early Access) for centrally managing security and authorization settings — including refresh token policies, Demonstrating Proof of Possession (DPoP), and Organizations — across dynamically registered CIMD clients.

— Names concrete settings covered but no configuration steps given.

launch-20261022-1cd42ba9

02 **CIMD manual pre-registration for clients** NEW 63

Adds CIMD (client ID metadata URL) support with manual pre-registration: an administrator submits a client metadata URL to Auth0, which retrieves and validates the metadata before creating the client after admin approval.

— Names mechanism and approval flow but no exact admin UI/API steps.

launch-20261022-1cd42ba9

03 **CIMD dynamic client registration (Early Access)** NEW

60

Adds CIMD dynamic registration (Early Access), enabling just-in-time client discovery and registration during an OAuth authorization flow using a client ID metadata URL, without requiring admin pre-registration of every client.

— Describes mechanism but lacks concrete setup steps; Early Access.

launch-20261022-1cd42ba9

NEW INTEGRATION Terraform Auth0 Modules — b2b-saas-starter

> announced · Oct 22, 2026 BLOG ↗
>
> New `b2b-saas-starter` Terraform module provisions a full B2B identity stack — RBAC, MFA, attack protection, enterprise SSO, and custom domains — from one module block.

> Single `module` block provisions app client, API, RBAC, MFA, attack protection, and branding on Auth0.
>
> `roles` and `permissions` inputs declare role-permission mappings as plain data with pre-apply validation.
>
> `organizations` input creates per-customer Auth0 Organizations with SAML (`samlp`) or OIDC enterprise SSO connections.
>
> `custom_domain` and `enable_custom_domain_verification` inputs with `dns_verification_record`, `dns_verification_command`, and `custom_domain_next_step` outputs guide multi-step custom-domain setup.
>
> `manage_tenant`, `manage_attack_protection`, and `manage_mfa` flags let multiple module instances share a tenant without ownership conflicts.

WAS THIS USEFUL? 👍 👎

◆ Non-Human & Agent Identity

Token Security

Feedback ⓘ announced · Oct 22, 2026 BLOG ↗

RANK

Token Security accelerates secure adoption of Agentic AI by managing & governing AI agents & non-human identities.

> Token Security launched a Grok Build integration that lets developers query and remediate agent credential and MCP risk directly inside the editor, paired with EDR data to catch shadow agents and endpoint credentials outside the sanctioned workspace.

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

01 Grok Build Marketplace plugin for in-editor findings NEW

75

A Token Security plugin in the Grok Build Marketplace lets developers query Token findings — secrets, MCP server permissions, and credential ownership — without leaving the editor. It integrates with the Grok Build Admin API to enumerate which users are enabled in a workspace and which MCP servers and tools each team has configured, and the Grok Build agent itself can call these plugin tools during remediation workflows to query secret and credential-ownership metadata directly from the editor.

Ask Token Security which secrets are expiring soon and who owns them, directly from inside Grok Build, so the right developer can rotate credentials before they become a risk.

> 📍 Which secrets expire in the next 30 days, and who owns them?

— Names plugin, Admin API and agent tool-calling with a working example query

launch-20261022-dad6e37a

THINNER COVERAGE BELOW

02 EDR pairing surfaces endpoint credential risk in Grok Build

NEW 55

Pairs EDR integrations (e.g. CrowdStrike) with Grok Build workspace data to surface MCP config files and hardcoded credentials found on developer endpoints, including agents running outside the sanctioned Grok Build workspace.

— Names mechanism and a concrete EDR example but no command or config path

launch-20261022-dad6e37a

03 Per-finding attribution for Grok Build risks NEW 45

Each risky Grok Build config item — hardcoded PATs, over-permissioned MCP servers, or unreviewed remote endpoints — is routed to the developer or identity owner responsible for it.

— Describes routing behaviour and named risk types but no interface detail

launch-20261022-dad6e37a

WAS THIS USEFUL?

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Compyl

Feedback ⓘ Oct 23rd, 2026 NOTES ↗ RANK

Compyl is the AI-powered GRC platform that maps one control library to SOC 2, ISO 27001, HIPAA, PCI DSS and 70+ frameworks, so you stay audit-ready year-round.

> Compyl added custom fields support to the Clients module, allowing administrators to define and manage custom fields on client records.

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

01 **Custom fields for Clients module** NEW 40

Administrators can now define and manage custom fields on client records within the Clients module.

. . . .

— States the capability but no field types, limits, or config details. Oct 23rd, 2026

WAS THIS USEFUL?

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.36.0 (2026-10-24) BLOG ↗ RANK

Runlayer is an MCP gateway that enforces identity-mapped permissions and detects threats across the agents and MCP servers employees use.

> Runlayer 1.36.0 lets skill admins detach published skills from their source repo for direct editing, while also sharpening the Go agent's task follow-through and retraining ToolGuard's removal recommendations.

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

01 **Skill detachment from source repository** NEW 80

Skill owners and skill admins can detach a skill published with `runlayer skills push` from its source repository so it can be edited directly in Runlayer, preserving files, version history, and sharing. Editors can also detach a skill once its owner is deactivated, preventing orphaned skills when a publisher leaves. After detachment, later pushes from the source repository no longer update the skill, automatic updates are turned off, and pushing a skill with the same name fails until one is renamed.

– Names the push command and full detach/rename behavior 1.36.0 (2026-10-24)

THINNER COVERAGE BELOW

02 **Go agent task follow-through improvements** IMPROVED 50

The Go agent now checks its own work before marking tasks complete, waits on slow results such as CI or deploys, reports the real outcome, and supports timed reminders (e.g. 'ping me in 30 min').

– Describes behavior changes but no config or command 1.36.0 (2026-10-24)

03 **ToolGuard tool-removal recommendation retraining** IMPROVED 45

Retrains the ToolGuard tool-removal model so recommendations are scoped to the tools a deployment actually sends and are proportionate to each agent's tool set, rather than occasionally suggesting removing every tool.

– Explains the fix but gives no actionable surface 1.36.0 (2026-10-24)

└──▷ **BREAKING ON UPGRADE**

! Pushing a skill with the same name as an already-detached skill fails until one of them is renamed.

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Agentic Infrastructure

microsandbox Feedback ⓘ v0.7.11 NOTES ↗ RANK

easy, fast, programmable and local-first microVM runtime and library

> microsandbox v0.7.11 adds support for authenticated HTTP CONNECT proxies in its SDK, letting sandboxes reach external services through a proxy requiring credentials.

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

01 **HTTP CONNECT proxy authentication in SDK** NEW 35

The SDK now supports HTTP CONNECT proxy authentication, enabling sandboxes to reach external services through authenticated proxies.

— Names the capability but no config keys, flags, or usage shown v0.7.11

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI Runtime Defense & Guardrails

Manifold Security

Feedback ⓘ announced · Oct 22, 2026

BLOG ↗ RANK

Manifold Security is an AI detection and response platform that monitors what autonomous AI agents do at runtime on enterprise endpoints and blocks risky actions.

> Manifold Security launched Open VSX editor-extension scanning, inventorying and risk-rating roughly 17,000 extensions nightly and tying that risk directly to AI agents running on the same endpoint.

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

01 Open VSX editor extension scanning and inventory NEW

83

Manifold adds Open VSX editor-extension scanning to the public Manifest intelligence site, covering ~17,000 extensions refreshed nightly and searchable without a Manifold account under Editor Extensions. Any extension detected on a managed machine by the endpoint collector (version `0.997.2` or newer) appears in the Manifold inventory alongside MCP servers, skills, and plugins, with a Manifest tab showing the rating for the installed version and flagging when the installed version (e.g. 9.0.0) differs from the scanned version. Extensions rated critical, high, or medium automatically raise a supply-chain risk on every AI agent running on the same machine. Tracked editor surfaces are VS Code, VS Code Insiders, Cursor, Windsurf, Kiro, and Antigravity.

Check a Cursor or Windsurf extension for supply-chain risk before anyone on your team installs it.

📍 1. Go to the Manifest site 2. Select Editor Extensions 3. Search for the extension by name to view its rating, findings, and whether the publisher owns the claimed code repository

Triage a flagged extension on a developer machine — verify the endpoint collector meets the minimum version, then view the extension's Manifest rating inside the inventory.

📍 1. Confirm the endpoint collector is at version 0.997.2 or newer on the developer machine 2. In the Manifold console, open Inventory 3. Locate the editor extension (VS Code, Cursor, Windsurf, Kiro, or Antigravity) 4. Select the extension and open its Manifest tab to see the rating, findings, and scanned-version notice

— Names collector version, editor list, and exact rating/risk mechanism.

launch-20261022-fe7b4ea8

WAS THIS USEFUL?  

OTHER / UNCATEGORIZED

◆ UNCATEGORIZED

vigolium

Feedback ⓘ v0.5.5 NOTES ↗ RANK

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

> Vigolium v0.5.5 adds HTTP/3-aware fuzzing, three new reconnaissance subcommands (fingerprint, hosts, dns), a self-describing schema/introspection layer, and broad TSV output and field-filtering support across core commands.

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

01 **HTTP/3 fuzzing with protocol fallback** NEW 90

Adds `--http-version`, `--http3`, and `--http3-only` flags to `fuzz`, multiplexing over HTTP/3 → HTTP/2 → HTTP/1.1 with the negotiated protocol reported as `pace.protocol`.

Fuzz an endpoint over HTTP/3 first, falling back through HTTP/2 and HTTP/1.1, to test behaviour across protocol versions.

```
$ vigolium fuzz --http3 -u https://target.example.com/FUZZ -w
builtin:default --fields pace.protocol,status,length
```

— Named flags, fallback mechanism, and a runnable example v0.5.5

02 **Fingerprint, hosts, and dns subcommands** NEW 90

Adds `vigolium fingerprint` with `--depth light|standard|full` to store per-origin host facts, `vigolium hosts` with `--sans` and `--soft-404` flags to read stored host facts, and `vigolium dns` to resolve and store DNS records, with `dig` and `host` shims routing to the new dns subcommand.

Collect deep per-origin host facts — TLS certs, banners, favicons — for a set of targets before a pen-test engagement.

```
$ vigolium fingerprint --depth full -u https://target.example.com
```

Dump all stored hosts with SAN data and soft-404 signatures in TSV for easy import into a spreadsheet or SIEM.

```
$ vigolium hosts --sans --soft-404 --format tsv
```

— Three named subcommands with flags and two runnable examples v0.5.5

03 **Schema and self-documentation subcommands** NEW 70

Adds `vigolium schema storage [table]` to describe tables, columns, and row counts, a `--usage` flag available on any command, `vigolium schema flags <command> -j` for machine-readable flag docs, and JSON schemas for `fuzz`, `replay`, `ingest`, `traffic body/headers`, `hosts`, `dns`, `fingerprint`, and `schema storage/flags`.

— Exact subcommand and flag syntax but no worked example v0.5.5

04 **TSV output and shared field-control flags** NEW 70

Adds `--format tsv` output to `traffic`, `finding`, `endpoints`, `params`, `hosts`, and `dns` commands, and adds `--fields`, `--compact`, and `--max-output-bytes` flags to `fuzz`, `replay`, `endpoints`, `params`, `log ls`, `hosts`, and `dns`.

— Enumerates every command and flag affected v0.5.5

05 **Traffic field system additions** IMPROVED 60

Adds `traffic --fields matches_not_found` and `--exclude-not-found` flags, plus new `uuid_short` and `response_headers` fields to the shared `--fields` system, with nearest-name suggestions offered on typos.

— Named fields and flags but no usage shown v0.5.5

THINNER COVERAGE BELOW

06 **Favicon hashing and protocol pinning in discovery** IMPROVED 55

Adds favicon mmh3 hash output (Shodan/FOFA compatible) alongside existing MD5 and SHA-256 hashes, and pins each origin's discovered protocol in `discovery.engine.protocol` with a dedicated `brute_force_concurrency` pool.

— Named fields and config key but no usage example v0.5.5

07 **UUID prefix matching for record flags** IMPROVED 55

Record UUID flags now accept a unique 8+ character prefix, returning `ambiguous_id` on collision.

— Concrete behaviour and error name, no example v0.5.5

08 **Ingest listing and WAF block-page capture** NEW 50

Adds `ingest -j` to list stored items and failure reasons, plus `--keep-blocked` to store WAF block pages.

— Named flags with minimal explanation of behaviour v0.5.5

09 **PHP vulnerable-version passive detection** NEW 45

Adds a `php2shell-version` passive module flagging vulnerable PHP 8.2–8.5 banners.

— Named module and version range, no usage shown v0.5.5

WAS THIS USEFUL? 👍 👎

Arcade

Feedback ⓘ announced · Oct 22, 2026 BLOG ↗ RANK

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

> Arcade launched the Arcade Plugin, giving agents in popular coding and chat clients access to over 1,000 tools through a single shared gateway, along with built-in tool recommendation and authorization handling.

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

01 **Arcade Plugin for major agent clients** NEW 75

The Arcade Plugin can be installed in Claude Code, Claude Desktop, Cursor, VS Code, GitHub Copilot CLI, and Codex, connecting agents to a single shared Arcade gateway without requiring any MCP gateway configuration.

Install the Arcade Plugin in a supported agent client to start making tool calls without configuring an MCP gateway.

> 📍 In your agent client (Claude Code, Claude Desktop, Cursor, VS Code, GitHub Copilot CLI, or Codex): 1. Install the Arcade Plugin from the client's plugin/extension marketplace. 2. Sign in with your Arcade account. 3. Ask your agent something like 'summarize my unread email' or 'post the release notes to #eng' — the plugin routes the call through the shared Arcade gateway automatically.

— Names clients and install flow but not config internals launch-20261022-d21b34e1

THINNER COVERAGE BELOW

02 **Tool Recommendation catalog** NEW 55

Tool Recommendation keeps a catalog of over 1,000 agent-optimized tools out of the model's context window, enabling a single Arcade gateway to serve all users dynamically instead of loading the full tool list into context.

— Explains mechanism but no direct usage step given launch-20261022-d21b34e1

03 **Skills and lifecycle hooks for tool authorization** NEW 50

Skills and lifecycle hooks shipped inside the plugin teach agents when to invoke Arcade tools and how to authorize services like Gmail or Slack directly from within the agent client.

— Names services but no concrete config or command shown launch-20261022-d21b34e1

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 25, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue