THE CYBER TOOLCHAIN

NO. 117

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 23, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 400 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 72 tools

▾ // THE BRIEF HIDE

What stands out across today's releases, grouped by what it lets you do. Every tool named links to its entry below.

Three unrelated things stand out. ScubaGear v2.0.0 swaps the retired Defender baseline for a Security Suite one and adds an Entra ID policy blocks risky AI agents. ntoseye can now set breakpoints and step

vim ? help

inside live Hyper-V guests. Auth0 lets admins stop email scanners from silently burning recovery links by unblocking brute-force lockouts by hand.

GOVERN

Re-baseline M365 against the Security Suite and prove what drifted since last run

ScubaGear's retired Defender baseline left a gap in CISA SCuBA assessments; the Security Suite baseline closes it, `Invoke-SCuBADiff` turns two assessment runs into a delta so you can show an auditor what changed rather than re-reading two full reports, and a new Entra ID policy checks whether risky AI agents are blocked. Auth0's Application Profiles centralize the settings that govern how agents and MCP clients register and authenticate, instead of per-app drift.

ScubaGear · Auth0

DETECT

Debug a live Hyper-V guest from the host instead of reproducing the bug outside the VM

ntoseye adds breakpoints and single-stepping inside Hyper-V and Sandbox guests plus partition/VTL inspection, EPT walking and hypercall/VM-exit breakpoints — analysis of a sample or rootkit that only misbehaves under virtualization no longer requires getting an agent into the guest. Adalanche's MCP server lets an assistant query AD ACLs directly, and new REACH/TRAIL query modes plus LAPS, GPO and SDDL edges widen what an attack path can be traced through.

ntoseye · adalanche

PROTECT

Run an AI coding agent without giving it your host

coi wraps agent execution in Incus containers with nftables network isolation and supply-chain guardrails, so an agent that fetches a malicious package or shells out cannot reach the developer's workstation or the rest of the network. Cycode's Custom Maestro Agents push the other direction — user-defined agents with confidence gating and human-approval controls triage violations, with token usage tracked so the cost of automated triage is visible rather than a surprise.

coi · Cycode

◆ Reconnaissance & OSINT

taranis-ai

Feedback ⓘ 1.5.2 NOTES ↗ RANK

Taranis AI is an advanced Open-Source Intelligence (OSINT) tool, leveraging Artificial Intelligence to revolutionize information gathering and situational analysis.

> Taranis AI 1.5.x overhauls its LLM backbone with a centralized endpoint manager and OpenRouter batch processing, replaces binary cybersecurity classification with eight-category topic classification, and adds fuzzy news deduplication alongside a string of breaking infrastructure changes (PostgreSQL 18, Redis persistence, secure cookies/HSTS).

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 **Central LLM endpoint management** BREAKING 90

A new Admin Settings → LLM Endpoints UI centrally manages LLM providers, models, credentials and timeouts, with per-feature or per-bot assignment and private CPU inference support. This replaces the standalone `llm-bot` service: the `BOT_ENDPOINT`, `BOT_API_KEY`, `SUMMARY_ENDPOINT` and `TITLE_ENDPOINT` parameters are removed on startup, and `LLM_BASE_URL`, `LLM_API_KEY`, `LLM_MODEL` and `LLM_TIMEOUT` environment variables no longer configure NER, sentiment, cybersecurity classification, clustering or summarization/title bots — admins must recreate providers in LLM Endpoints and remove retired bot containers or Kubernetes workloads before resuming automatic jobs.

Navigate to the central LLM endpoint manager to register an OpenRouter provider and assign it to background bots.

> 📍 In the Taranis AI admin panel, go to Admin Settings → LLM Endpoints. Click Add Provider, enter your OpenRouter base URL, API key, and desired model, then set a timeout. Under each bot (e.g. clustering, summarization), select the new provider from the assignment dropdown. Enable batch processing for background bots where compatible.

— Exact UI path, removed env vars and migration steps named 1.5.1

02 **Fuzzy news deduplication** NEW 90

Opt-in source-scoped fuzzy deduplication for collected news: changed articles at an existing URL update the original news item with revision history, and similar articles from the same source can join an existing story while each article is retained separately. Configurable under Collection Similarity Matching with a default score threshold of 85 and a 30-day lookback.

Tune fuzzy news deduplication thresholds to prevent boilerplate-heavy sources from over-grouping stories.

> 📍 In the Taranis AI admin panel, go to Collection Similarity Matching. Set the score threshold above the default of 85 (e.g. 92) and reduce the lookback from the default 30 days to 7 days for sources with repeated boilerplate.

— Names config area, default threshold and lookback with tuning example `1.5.1`

---

03 **PostgreSQL 18 upgrade for bundled deployments** BREAKING

85

Bundled Compose deployments are upgraded to PostgreSQL 18 with a changed data mount path; existing PostgreSQL 14–17 clusters cannot start directly with the new image, so operators must run `./database/upgrade-database.sh` per the database upgrade procedure before starting.

— Gives exact upgrade command and affected version range `1.5.1`

---

04 **Eight-category topic classification replaces binary cybersecurity flag**

80

BREAKING

Taranis AI replaces binary cybersecurity classification with eight-category LLM-bot topic classification, which now requires a dedicated Decisions endpoint instead of a Responses or Chat Completions endpoint. `CLASSIFICATION_THRESHOLD` no longer controls topic classification and has no effect, and existing cybersecurity classification attributes remain readable but are no longer updated.

— Names endpoint type and config key but no exact migration steps `1.5.2`

---

05 **OpenRouter batch processing for background bots** NEW

80

Background LLM bots — NER, sentiment, classification, clustering and summarization — can now use OpenRouter batch processing via the shared LLM endpoint; realtime remains the default and Chat still requires a realtime endpoint.

Navigate to the central LLM endpoint manager to register an OpenRouter provider and assign it to background bots.

> 📍 In the Taranis AI admin panel, go to Admin Settings → LLM Endpoints. Click Add Provider, enter your OpenRouter base URL, API key, and desired model, then set a timeout. Under each bot (e.g. clustering, summarization), select the new provider from the assignment dropdown. Enable batch processing for background bots where compatible.

— Names affected bots and provider with a usage walkthrough 1.5.1

06 **Top news sources on Assess dashboard** NEW 70

The Assess dashboard card now shows a top news sources distribution, with a configurable per-profile display limit (default 3; set to 0 to hide).

Control how many top news sources appear on the Assess dashboard card for a given user profile.

> 📍 In the Taranis AI user profile settings, locate the top source distribution limit field. Set it to 5 to show the five leading sources on the Assess card, or 0 to hide the distribution entirely.

— Names exact default and zero-to-hide behaviour with example 1.5.1

07 **Enforced secure cookies and HSTS** BREAKING 70

Secure browser cookies and HSTS (one-year duration) are now enforced outside `DEBUG=true` mode; `JWT_COOKIE_SECURE=false` no longer disables this protection, so TLS and HTTP-to-HTTPS redirects must be configured before upgrading since browser HSTS persists after an application rollback.

— Names env vars and duration but no exact migration command 1.5.1

08 **Timeout handling for publishers and bot jobs** BREAKING 70

Email, FTP and SFTP publishers now enforce a 30-second network inactivity timeout. Separately, an explicit bot `REQUESTS_TIMEOUT` now limits both individual requests and the entire bot job, so installations processing many stories may need a larger value.

— Names affected publishers and a config key without exact values 1.5.1

THINNER COVERAGE BELOW

09 **Redis persistence disabled by default** BREAKING 50

Bundled Redis deployments now have persistence disabled; external Redis instances need equivalent configuration, since pending jobs, retries, unfinished bot chains, and pending LLM batch tracking are lost on Redis restart.

— Explains consequence but no exact config to restore persistence 1.5.1

10 **Resumable bot pipelines** NEW — 40

Bot DAGs now run in one resumable worker pipeline, so interrupted pipelines can pick up where they left off instead of restarting from scratch.

— Explains mechanism but no concrete surface to act on 1.5.2

11 **Live story tasks in My Tasks and Scheduler** NEW — 40

Story tasks now appear live in the My Tasks and Scheduler views.

— Names two views but describes little mechanism 1.5.2

12 **TLP default inheritance changed** BREAKING — 35

Sources without an explicit TLP level now inherit the global default instead of using a source-level default.

— States behaviour change with no configuration detail 1.5.1

13 **Combined tag generation in story editor** NEW — 30

The story editor adds combined tag generation.

— Bare one-line addition with a named UI surface 1.5.2

14 **Method selection for new news items** NEW — 30

A method selection step is now offered before creating a news item.

— Thin description of a UI workflow step 1.5.2

15 **Collapsible admin sidebar submenus** IMPROVED — 30

Admin sidebar destinations are grouped into collapsible submenus for easier navigation.

— Names the UI area but no deeper mechanism 1.5.2

16 **Top source icons on story cards** NEW — 25

Story cards now show icons for their top sources.

— Cosmetic addition with minimal detail 1.5.2

17 **Local timezone chat timestamps** IMPROVED — 20

Chat timestamps now display in the user's local timezone instead of a fixed timezone.

— Minor UI fix with little detail 1.5.2

└──▷ **BREAKING ON UPGRADE**

! Binary cybersecurity classification is replaced by eight-category topic classification: a dedicated Decisions endpoint is now required instead of a Responses or Chat Completions endpoint.

- ! `CLASSIFICATION_THRESHOLD` no longer controls topic classification and has no effect.
- ! Existing cybersecurity classification attributes remain readable but are no longer updated.
- ! Bundled Compose deployments now use PostgreSQL 18 with a changed data mount; existing PostgreSQL 14–17 clusters cannot start directly with the new image — run `./database/upgrade-database.sh` per the database upgrade procedure before starting.
- ! The standalone `llm-bot` service is replaced: `BOT_ENDPOINT`, `BOT_API_KEY`, `SUMMARY_ENDPOINT`, and `TITLE_ENDPOINT` parameters are removed on startup; `LLM_BASE_URL`, `LLM_API_KEY`, `LLM_MODEL`, and `LLM_TIMEOUT` environment variables no longer configure NER, sentiment, cybersecurity classification, clustering, or summarization/title bots — recreate providers in `LLM Endpoints` and remove retired bot containers or Kubernetes workloads before resuming automatic jobs.
- ! `JWT_COOKIE_SECURE=false` no longer disables Secure cookies or HSTS; only `DEBUG=true` bypasses these protections — configure TLS and HTTP-to-HTTPS redirects before upgrading, as browser HSTS persists after an application rollback.
- ! Bundled Redis deployments now have persistence disabled; external Redis instances need equivalent configuration — pending jobs, retries, unfinished bot chains, and pending LLM batch tracking are lost on Redis restart.
- ! An explicit bot `REQUESTS_TIMEOUT` now limits both individual requests and the entire bot job — installations processing many stories may need a larger value.
- ! Sources without an explicit TLP level now inherit the global default instead of using a source-level default.

WAS THIS USEFUL? 👍 👎

---

nerva

Feedback ⓘ v1.77.1 NOTES ↗

Fast service fingerprinting CLI for 170+ protocols (TCP/UDP/SCTP) - built by Praetorian

└──▷ **GET THIS VERSION**

```
$ git clone --branch v1.77.1 https://github.com/praetorian-inc/nerva.git
# already have the repo? check out this version:
$ git checkout v1.77.1
```

NEW DETECTION PHP fingerprinter with version/SAPI detection and php2shell findings

---

WAS THIS USEFUL? 👍 👎

favirecon

Feedback ⓘ v1.1.0 NOTES ↗

Use favicons to improve your target recon phase. Quickly detect technologies, WAF, exposed panels, known services.

Favicon detection via `<link>` tags improved with better HTML fallback that resolves icons correctly in more cases.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v1.1.0
https://github.com/edoardottt/favirecon.git
# already have the repo? check out this version:
$ git checkout v1.1.0
```

NEW DETECTION Fingerprint database expansion

WAS THIS USEFUL? 👍 👎

adalanche

Feedback ⓘ r1014 NOTES ↗ RANK

Attack Graph Visualizer and Explorer (Active Directory) ...Who's *really* Domain Admin?

> Adalanche r987-r996 adds an MCP server for querying and inspecting ACLs, overhauls AQL with new REACH/TRAIL/UNION query modes (replacing the removed SIMPLE mode), and expands attack-path analysis with Windows LAPS, GPO, SDDL and session-edge detection, alongside WASM-based graph layouts and smaller UI and query-accuracy fixes.

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

01 AQL REACH, TRAIL and UNION query modes replace SIMPLE

NEW DETECTION BREAKING 80

Adds AQL `REACH` mode (returns every edge on any valid route between nodes), `TRAIL` mode (avoids edges already present in the result, replacing the removed `SIMPLE` mode), `UNION` parsing support for combining multiple query result sets, and traversal priorities `SHORTEST`, `LONGEST`, `LIKELY`, and `UNLIKELY` for steering path results. REACH `CHEAPEST` and `SHORTEST` path analysis now keeps fewer, more relevant routes between start and end nodes; REACH flow counting now excludes routes a deny ACE refuses to the account using them, and route analysis omits steps a deny ACE refuses to the account acting at that point. AQL `SIMPLE` mode is removed and queries using it must be migrated to `TRAIL`.

— Migration path noted but no example query syntax r995 r987

02 (new) ACL, LAPS, GPO, and session-edge attack-path analysis

NEW DETECTION UPDATED SUPPORTED PLATFORM NEW 75

Adds Windows LAPS analysis (requiring read and control access, with DSRM support), tags stale GPO links, delegation targets, and published templates as findings, adds task SDDL parsing for scheduled task attack-path analysis, adds session-type edges for `Batch`, `Service`, `Negotiate`, `SessionNetworkKerberos`, and `SessionNetworkNTLMV2` logon paths, adds Service Control Manager permission analysis for privilege escalation on local machines, adds GPO targeting for site-linked GPOs with machine-own policy results taking precedence, and adds ACL analysis for partial denies, owner rights, forest heuristics, and AdminSDHolder.

— Rich mechanism detail but no command or UI path given Lists analysis targets but no mechanism or usage detail

r987

03 MCP server with ACL inspection and documentation NEW

70

Exposes an MCP server at `/mcp` on the web service, exposing routes and causes without exposing secrets. Adds `MCP get_acl` to display a security descriptor in plain, human-readable terms, and MCP now serves the built-in documentation.

– Named endpoint and command but no usage walkthrough given r995 r987

04 **WASM-based graph layout engine** NEW 65

Adds a WASM layout engine and dynamic graph layout UI for in-browser graph computation, new deterministic, overlap-free Layered and Stress graph layout algorithms, and a graph edge width preference in the UI.

– Named algorithms and UI preference but no exact navigation r987

THINNER COVERAGE BELOW

05 **Bounded profiling and memory diagnostics** NEW 35

Adds bounded profiling and memory diagnostics, accessible with Go 1.27.

– No flag or command given to access diagnostics r987

06 **Query result counting and limit reporting** IMPROVED 35

Query results now report when a node limit cut them short, and query totals now correctly count the nodes that merged or folded nodes represent.

– Behaviour fix described without a config surface r995 r987

07 **Independent node dragging in graph UI** IMPROVED 35

Dragging a node in the graph UI now moves that node independently, rather than panning the entire view.

– Clear UI behaviour change with no further mechanism r996

08 **Edge cause provenance tracking** NEW 30

Every edge records its cause, enabling full edge provenance tracking.

– States capability without usage or query detail r987

09 **OpenGraph loader for custom graph data** NEW 25

Adds an OpenGraph loader for importing custom graph data.

– Bare mention with no format or command details r987

UPDATED DETECTION WriteAltSecIdent targets computer accounts

NEW DETECTION NTLM and LDAP client hardening attributes

NEW DETECTION Stored-credential locations from assessment

NEW DETECTION Observed sessions and process tokens from assessment

NEW DETECTION DCOM and PowerShell-remoting rights from assessment

NEW DETECTION Shared-service DLL hijack from assessment

NEW DETECTION Machine hardening facts from assessment

NEW DETECTION Trusted-domain Authenticated Users membership modeling

└──▷ BREAKING ON UPGRADE

! AQL `SIMPLE` mode is removed; queries using `SIMPLE` must be migrated to `TRAIL`.

r1014 NOTES ↗ LATEST STABLE

AQL defaults to REACH (CHEAPEST) to prevent OOM; edge provenance shown as tooltip; MCP clients survive restarts; details cards show attribute pictures.

└──▷ GET THIS VERSION

```
$ git clone --branch r1014
https://github.com/lkarlslund/Adalanche.git
# already have the repo? check out this version:
$ git checkout r1014
```

AQL defaults to REACH (CHEAPEST) mode instead of full path enumeration, preventing out-of-memory crashes on large graphs.

MCP clients remain connected and functional across adalanche server restarts.

Details cards now render pictures stored in object attributes.

Route window shows edge provenance as a tooltip instead of an inline list.

Machine import now extracts data from the assessment file rather than retaining the raw assessment blob.

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

Datadog pathrunner

Feedback ⓘ v1.0.0 NOTES ↗ RANK

> Pathrunner hit v1.0.0 with a full Metasploit-style REPL/CLI, 80+ AWS privilege-escalation modules, runtime-selectable payloads, PMapper integration, and session tooling, then followed up in v0.2.3 with a profile-loading alias, purple-team CloudTrail logging, and expanded IAM-role backdoor support.

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

01 REPL and CLI command interface NEW

90

Pathrunner ships an interactive Metasploit-style REPL (`pathrunner`) and a scriptable CLI that share identical command handlers via an adapter pattern, supporting `identity add` , `identity list` , `identity switch` , `search` , `info` , `use` , `show options` , `show payloads` , `set` , and `exploit` . `identity add` also accepts a `--from-profile` alias to load AWS identities by profile name for scripted or multi-account workflows.

Pivot through a Lambda privilege-escalation path end-to-end: add your AWS profile, pick a module, set a target role, choose a credential-exfiltration payload, and fire the exploit.

```
$ pathrunner
pathrunner> identity add --profile pentest-profile
pathrunner> search lambda
pathrunner> use lambda-001
pathrunner> show options
pathrunner> set ROLE_ARN
arn:aws:iam::123456789012:role/TargetRole
pathrunner> show payloads
pathrunner> set PAYLOAD exfil/response
pathrunner> exploit
```

Add an AWS identity directly by named profile without typing the full flag, useful in scripted or multi-account workflows.

```
$ pathrunner> identity add --from-profile my-aws-profile
```

— Full command set and alias named, directly runnable. v1.0.0 v0.2.3

02 Runtime-selectable exploit payloads IMPROVED

90

Provides dozens of interchangeable payloads chosen at runtime via `set PAYLOAD`, including credential exfiltration (`exfil/response`), HTTPS exfiltration, backdoor role/user/policy creation, and reverse shells. The `backdoor/attach-policy` payload was extended to target IAM roles in addition to IAM users.

Pivot through a Lambda privilege-escalation path end-to-end: add your AWS profile, pick a module, set a target role, choose a credential-exfiltration payload, and fire the exploit.

```
$ pathrunner
pathrunner> identity add --profile pentest-profile
pathrunner> search lambda
pathrunner> use lambda-001
pathrunner> show options
pathrunner> set ROLE_ARN
arn:aws:iam::123456789012:role/TargetRole
pathrunner> show payloads
pathrunner> set PAYLOAD exfil/response
pathrunner> exploit
```

Backdoor an IAM role (not just a user) by attaching a policy — now supported as a target type for the attach-policy payload.

```
$ pathrunner> use <module>
pathrunner> set PAYLOAD backdoor/attach-policy
pathrunner> set ROLE_ARN
arn:aws:iam::123456789012:role/TargetRole
pathrunner> exploit
```

— Named payload types and exact flag with before/after scope. v1.0.0 v0.2.3

---

03 **80+ AWS privilege-escalation exploit modules** NEW 65

Ships 80+ exploit modules across 20+ AWS services — including IAM, EC2, Lambda, STS, ECS, Glue, CloudFormation, SSM, and Bedrock — sourced from the pathfinding.cloud catalog, selectable via `search` and `use` in the REPL.

Pivot through a Lambda privilege-escalation path end-to-end: add your AWS profile, pick a module, set a target role, choose a credential-exfiltration payload, and fire the exploit.

```
$ pathrunner
pathrunner> identity add --profile pentest-profile
pathrunner> search lambda
pathrunner> use lambda-001
pathrunner> show options
pathrunner> set ROLE_ARN
arn:aws:iam::123456789012:role/TargetRole
pathrunner> show payloads
pathrunner> set PAYLOAD exfil/response
pathrunner> exploit
```

— Names service count and source but no per-module mechanism. v1.0.0

04 **Automatic credential capture into new identities** NEW

65

Auto-extracts captured credentials from successful exploits and registers them as new identities immediately accessible via `identity list` and `identity switch`, enabling chained escalation across sessions.

After a successful exploit captures credentials, switch to the new identity and immediately map onward escalation paths with PMapper.

```
$ pathrunner> identity list
pathrunner> identity switch lambda_AB12
pathrunner> pmapper analyze
```

— Mechanism explained but no limits or format given. v1.0.0

05 **External tool integrations** NEW

65

Integrates PMapper graph analysis via `pmapper analyze` to enumerate reachable escalation paths from any captured identity, and supports importing CloudFox output as a resource store for use within pathrunner sessions.

After a successful exploit captures credentials, switch to the new identity and immediately map onward escalation paths with PMapper.

```
$ pathrunner> identity list
pathrunner> identity switch lambda_AB12
pathrunner> pmapper analyze
```

— Exact command given but limited mechanism detail. Names integration but no command or flag given.

v1.0.0

THINNER COVERAGE BELOW

06 **Per-session workspace isolation** NEW 45

Provides workspace isolation so each session maintains independent identities, command history, and tracked resources.

— Describes scope but no config or command surface named. v1.0.0

07 **Purple team CloudTrail event logging** NEW 45

Adds purple team CloudTrail event logging across all exploit modules and the discovery layer, recording attacker actions for defensive detection engineering.

— Describes scope across modules but no config or flag named. v0.2.3

08 **Attacker infrastructure helpers for callback exploits** NEW 35

Includes attacker-infrastructure helpers for trigger-and-callback exploits that require out-of-band listeners.

— Brief description, no commands or config named. v1.0.0

09 **CloudTrail-aware cleanup reports** NEW 35

Generates CloudTrail-aware cleanup reports of tracked resources created during exploitation sessions.

— States output but no command or format specifics. v1.0.0

WAS THIS USEFUL?  

◆ AI-Powered Offensive Security

Dreadnode

Feedback ⓘ Agentic attack trace parity NOTES ↗ RANK

Documentation for the Dreadnode AI security platform — self-hosting guides, platform concepts, SDK reference, and CLI documentation.

> Dreadnode added a tree-search-based SLIP jailbreak attack to AI red-teaming, enriched the findings parquet export with audit and compliance columns, introduced session pinning in Chat and Agents, and made hosted inference SDK streams more resilient with automatic retries.

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

01 **SLIP jailbreak attack tool** NEW 75

Adds `SLIP` self-jailbreaking attack tool to AI red-teaming, enabling tree-search-based prompt jailbreaks against a target model with configurable depth, branching, and judge models. Launch by selecting `SLIP` from the attack tool list, setting the target model, and configuring depth, branching, and judge models.

Run a SLIP tree-search jailbreak campaign against a target model to measure attack success rate with controlled branching.

> 📍 In the AI red-teaming UI, select your project, choose `SLIP` from the attack tool list, set your target model, configure depth, branching, and judge models, then launch the campaign.

— Names mechanism and config parameters plus launch steps Agentic attack trace parity

02 **Session pinning in Chat and Agents** NEW 70

Pin sessions from the `...` menu in the Chat sidebar or from `Agents › Sessions`; pinned sessions appear at the top of the list and are visible only to the user who pinned them.

— Exact menu and navigation path given Agentic attack trace parity

03 **Resilient hosted inference SDK streaming** IMPROVED 60

Hosted inference SDK streams dropped mid-connection are now retried up to 3 times with exponential backoff (1s, 2s, 4s) to reduce failures on unstable network paths.

— Concrete retry numbers but no user-facing control Agentic attack trace parity

THINNER COVERAGE BELOW

04 **Enriched findings parquet export** IMPROVED 57

The findings parquet export gains new columns for audit trail, ASR, compliance tags, and adversarial-ML metrics, plus project-level rollups, and now reflects the latest operator-

reviewed severity and status.

— Names columns added but no export command shown `Agentic attack trace parity`

WAS THIS USEFUL? 👍 👎

apex

Feedback ⓘ v2.6.1 NOTES ↗ RANK

AI-powered offensive security testing using autonomous agents, directly in your terminal.

> Apex v2.6.1 adds custom inference endpoint and Hoonify provider support, a `--resume` flag for session restore, and a new `targeted-pentest` command for scoped, objective-driven testing, alongside smaller additions to evidence capture, observability, credential tracking, sandboxed execution, and Herdr-based session restore.

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

01 Custom inference endpoints and Hoonify provider support NEW

90

Apex now supports custom OpenAI-compatible inference endpoints via `customProviders` configuration or the `APEX_CUSTOM_PROVIDERS` environment variable, selectable in headless commands with `--model-provider <id> --model <model>`. A new built-in Hoonify inference provider can be connected via `/providers` in the TUI or by setting `HOONIFY_API_KEY`, and headless commands accept `--model-provider hoonify --model <catalog-model-id>`.

— Names config key, env var, TUI path and exact CLI flags for both providers. v2.6.1

02 Session resume via `--resume` flag NEW

85

The `--resume <session-id>` CLI flag (with optional `--model <model-id>`) reopens a saved session, restoring operator mode and approval settings after a server restart.

Restore a previous Apex session after a server restart, keeping the same model and conversation state.

```
$ pensar --resume <session-id> --model claude-opus-5
```

— Exact flag with runnable resume command example. v2.6.1

03 Scoped endpoint pentests via `targeted-pentest` command NEW

83

A new `targeted-pentest` CLI command with `--target` and `--objective` flags launches a scoped, objective-driven pentest against a single endpoint instead of a full-surface sweep.

Run a scoped pentest against a single endpoint with a specific objective instead of a full-surface sweep.

```
$ pensar targeted-pentest --target https://example.com/api/login
--objective 'Test authentication bypass'
```

— Exact command and flags with runnable example. v2.6.1

04 **Agent state reporting and session restore via Herdr** NEW

65

Apex now reports `working`, `blocked`, and `idle` agent states to Herdr 0.9.2+, enabling session restore without extra configuration.

— Names states and Herdr version but no usage command given. v2.6.1

THINNER COVERAGE BELOW

05 **OpenTelemetry span fidelity for LLM calls** IMPROVED

50

Deepens OpenTelemetry span coverage and improves span fidelity for LLM calls through four additions: trace compaction and inference continuity, bounded compaction research evidence capture, root IO and consistent identity tracking, and retained bounded model-call failure metadata.

— Names four tracing improvements but no configurable surface. v2.6.1

06 **Isolated per-command shell executor** NEW

40

Adds an isolated per-command shell executor to sandbox command execution within agent runs.

— Describes sandboxing mechanism but no configurable surface named. v2.6.1

07 **Export of validated ATIF trajectory bundles** NEW

35

Adds the ability to export structured agent trajectory data as validated ATIF trajectory bundles.

— Names ATIF format but no endpoint or flag given. v2.6.1

08 **Pentest findings enriched with credential IDs and per-run rollout evidence**

NEW

30

Pentest output is enriched with credential IDs recorded on documented findings, and now captures native per-run rollout evidence within pentest sessions.

— No flag or mechanism beyond a brief description. Brief description without mechanism or usage surface.

v2.6.1

09 **Scoped hosted storage backends for agent memory** NEW

30

Supports scoped hosted storage backends for persistent, scoped memory across engagements.

— Describes capability without naming backend types or config keys. v2.6.1

WAS THIS USEFUL? 👍 👎

◆ **Exploitation & C2**

ssh-mitm

Feedback ⓘ 6.0.1 NOTES ↗ RANK

SSH-MITM - ssh audits made simple

> ssh-mitm 6.0.1 adds Flatpak sandbox warnings for unreachable paths and flags, and replaces four external dependencies with in-house implementations.

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

01 **Flatpak sandbox access warnings** NEW

65

Adds a startup warning in Flatpak builds when running without `home` / `host` filesystem access, alerting that `--scp-replace-file`, `--sftp-replace-file`, `--ssh-mirrorshell-key`, and custom host key paths outside the app's private data directory are unreachable. Also adds a permission-denied host key read hint under Flatpak, alongside the existing Snap `ssh-keys` hint.

— Names exact flags and conditions but it's just a warning message 6.0.1

02 **In-house replacements for external dependencies** IMPROVED

65

Replaces `python-json-logger` with an in-house JSON log formatter, replaces `colored` with an in-house ANSI styling module (`sshmitm.colors`) that now respects `NO_COLOR` / `FORCE_COLOR` environment variables, replaces `psrpcore` with an in-house PSRP wire-format parser (`sshmitm.plugins.powershell.psrp_wire`), and replaces `pytz` with the standard library's `datetime.UTC`.

— Names all replaced modules and new internal paths but no user-facing action 6.0.1

WAS THIS USEFUL? 👍 👎

◆ Wireless & RF

airgorah

Feedback ⓘ v0.9.0 NOTES ↗ RANK

A WiFi security auditing software

> Airgorah v0.9.0 introduces a full PMKID attack workflow and removes its dependency on the aircrack-ng suite and mergecap.

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

01 **PMKID attack workflow** NEW 45

Adds PMKID detection, display, solicitation, and cracking support as a new attack workflow, giving users an alternative to handshake-based attacks.

– Names the workflow stages but no UI path or command given. v0.9.0

02 **Removal of aircrack-ng and mergecap dependencies** BREAKING 45

Removes the aircrack-ng suite and mergecap as external dependencies; any workflow or wrapper scripts that assumed these dependencies are present may need to be updated.

– Clear before/after but no migration specifics. v0.9.0

03 **Status bar channel indicator** NEW 30

Shows the channel the interface is currently listening on in the status bar.

– Bare description of a small UI addition. v0.9.0

└──▷ **BREAKING ON UPGRADE**

! The aircrack-ng suite and mergecap are no longer required or used; any workflow or wrapper scripts that assumed these dependencies are present may need to be updated.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Aikido Security

Feedback ⓘ

#58 - Cloud Security Audits, External Network Pentesting and Aikido is now FedRAMP Modera

NOTES ↗ RANK

Aikido Security is a developer-focused application security platform that scans code, dependencies, containers, and cloud infrastructure to detect and help fix vulnerabilities.

> Aikido adds AI-driven Cloud Security Audits and External Network Pentesting, achieves FedRAMP Moderate authorization for government customers, and rounds out the release with a Google SecOps SIEM integration, six new API endpoints, Zen Firewall agents for five languages, Terraform/OpenTofu provider updates, and Code Quality scan improvements.

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

01 **Google SecOps SIEM integration** NEW 70

New Google SecOps integration lets Aikido optionally forward Activity logs to a SIEM for centralized alerting and correlation; enable via Settings > Integrations > SIEM by turning on Transmit Activity logs and selecting Google SecOps as the destination.

Forward Aikido Activity logs to Google SecOps for centralized SIEM alerting and correlation.

📍 Steps: In Aikido, go to Settings › Integrations › SIEM, enable Transmit Activity logs, and select Google SecOps as the destination.

— Exact UI steps given for setup. #58 - Cloud Security Audits, External N…

02 **New API endpoints for CWE, invites, and assessments** NEW 70

Aikido adds `API - List CWE descriptions` to retrieve CWE metadata, `API - List pending invites` to query outstanding user invitations, `API - List AI code analysis assessments` to fetch AI code analysis results, `API - List pentest assessments` to retrieve pentest findings, `API - Disable user` to deactivate users, and `API - Remove attach surface subdomain` to remove subdomains from the attack surface.

— Named endpoints but no HTTP method or path given.

#58 - Cloud Security Audits, External N…

03 **Cloud Security Audit for multi-step attack paths** NEW 60

New Cloud Security Audit feature uses AI agents that reason across cloud and IaC configurations to surface multi-step attack paths, linking each finding back to the causative template.

— Explains mechanism but no UI path or config given.

#58 - Cloud Security Audits, External N…

THINNER COVERAGE BELOW

04 **External Network Pentesting for internet-exposed assets** NEW

55

New External Network Pentest capability scans internet-exposed assets — including forgotten subdomains and open admin panels — producing evidence-confirmed findings.

— Describes scope and output but no setup steps. #58 - Cloud Security Audits, External N…

05 **Code Quality scan limits removed and auto-PR creation**

IMPROVED

55

Removes the maximum-rule restriction for Code Quality scans and enables auto-creation of PRs for issues found in full repository scans.

— Describes before/after but no config or flag named.

#58 - Cloud Security Audits, External N…

06 **Terraform provider gains container resources, OpenTofu listing**

NEW

50

Adds Terraform Container management resources and publishes the Aikido provider to the OpenTofu registry.

— Named resource type but no further configuration detail.

#58 - Cloud Security Audits, External N…

07 **FedRAMP Moderate authorization for Aikido for Government** NEW

45

Aikido for Government achieves FedRAMP Moderate authorization and is listed on the FedRAMP Marketplace, allowing federal agencies to adopt Aikido without a full security review.

— States authorization status but no action path for readers.

#58 - Cloud Security Audits, External N…

08 **Zen Firewall agents for Dotnet, Python, PHP, Ruby, Node** NEW

40

Adds Zen Firewall agent releases for Dotnet, Python, PHP, Ruby, and Node.

— Lists supported languages but no setup or mechanism detail.

#58 - Cloud Security Audits, External N…

09 **Hugging Face added to Aikido Intel malware ecosystems** NEW

30

Adds Hugging Face to supported malware ecosystems in Aikido Intel.

— Bare addition with no mechanism or scope detail.

#58 - Cloud Security Audits, External N…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

skylos

Feedback ⓘ v5.0.0 NOTES ↗

Open source local-first PR scanner that finds dead code, security bugs, secrets, quality regressions, and AI-code mistakes before merge. For first timers refer to https://duriantaco.github.io/skylos/repo-map/

Adds `skylos agent recap`, `skylos cicd agent-review` with GitHub OIDC and bounded source bundles, and a live README agent-activity badge.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v5.0.0 https://github.com/duriantaco/skylos.git
# already have the repo? check out this version:
$ git checkout v5.0.0
```

Adds `skylos agent recap` subcommand and hooks agent activity into `skylos doctor`, plus a live README badge.

Adds `skylos cicd agent-review` for approved Cloud PR reviews, with GitHub OIDC support, bounded source bundles, and a generated trusted workflow.

NEW DETECTION Security detections: Dependabot execution/cooldown, decoded-secret exposure, workflow obfuscation, trusted-publishing advice

NEW DETECTION Quality idiom-advice rules for repeated imports, boolean branches, f-string conversions, and re-raise-only handlers

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

GitHub Copilot in Visual Studio — September update (2026-10-23) BLOG ↗ RANK

GitHub is a web-based platform for version control and collaborative software development using Git repositories.

> GitHub's biggest window features were in Copilot: Visual Studio gained PR exploration via a Git agent, a Bring Your Own Model preview, and one-click NuGet vulnerability fixing, while the usage metrics API now names custom skills/agents/MCP servers and organizations got a global default policy for GA Copilot features. Pull request review also picked up persistent whitespace preferences and a full-width split view, the Copilot impact dashboard added opt-in industry benchmarking, and the meta API listed new Actions blob storage hosts that self-hosted allowlists must add.

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

01 Industry comparisons in Copilot impact dashboard NEW

93

An opt-in Industry comparison section in the Copilot impact dashboard benchmarks agentic adoption and leverage against a peer median and top decile over a 28-day period, filterable by company size (1–999, 1,000–4,999, or 5,000+ employees) and one of ten industries, with automatic widening to all industries when a selected industry's peer group is too small. A Copilot metrics industry comparison policy — set from the dashboard or centrally from enterprise/organization Copilot policies — controls both data contribution and visibility with a single toggle, and comparisons appear only in the dashboard UI, excluded from the usage metrics API, NDJSON exports, and downloads.

Opt your enterprise into industry comparisons so leaders can benchmark agentic Copilot adoption against peers of the same size and sector.

> 📍 In the Copilot impact dashboard, go to the Industry comparison section and enable the Copilot metrics industry comparison policy. Alternatively, navigate to your enterprise or organization Copilot policies and enable the same policy there. Then select your company size (e.g., 1,000–4,999) and one of the ten available industries to set your peer group.

— Exact bands, policy name, and opt-in steps all given.

Industry comparisons in the Copilot imp…

02 Named custom skills, agents, and MCP servers in usage metrics API

IMPROVED

88

The `totals_by_skill`, `totals_by_custom_agent`, `totals_by_mcp`, `totals_by_slash_cmd`, and `totals_by_plugin` arrays in enterprise and organization usage reports now show each customer-defined item's configured name instead of grouping it under `other` or `custom`, and each array now returns up to 25 named entries per record (up from 5), with remainder collapsed into a trailing `other` entry. Named entries appear in per-user 1-day and 28-day reports, aggregated 1-day reports, and `day_totals` inside aggregated 28-day reports, and require Copilot CLI 1.0.97+ (interactive and prompt mode), the GitHub Copilot app, ACP integrations like JetBrains, or Copilot SDK v1.0.21+.

— Exact field names, limits, and client version requirements given.

Custom skills, agents, and MCP servers …

---

03 **Global default policy for GA Copilot features** BREAKING

83

A global default policy, reachable via AI Controls → Copilot → Default policy for new features, sets unconfigured eligible Copilot settings on Business and Enterprise plans to Enabled, Disabled, or Let organizations decide, applying to the Features & clients and MCP servers in Copilot configuration surfaces. The rollout is a one-time, irreversible change for features left Unconfigured; later changes to the default policy will not revert features already set.

Review and set the global default before the October 26 rollout completes, to ensure unconfigured Copilot features land in the state your organization requires.

📍 In the enterprise or organization console, go to AI Controls → Copilot → Default policy for new features and choose Enabled, Disabled, or Let organizations decide before the rollout takes effect for your account.

— Exact UI path given but no CLI/API equivalent. Global default for Copilot features gen…

---

04 **Persistent whitespace preference in PR review** NEW

83

A new Settings → Whitespace changes preference offers Always show, Always hide, or Hide for this diff only, applied per user across all pull requests. When whitespace is hidden, a Hide whitespace filter badge appears beside All commits with an X to temporarily reveal whitespace without changing the saved preference.

Set a personal whitespace preference that persists across all PRs so you never have to re-hide noisy formatting changes.

📍 In a pull request, go to the Files changed tab → open Settings → Whitespace changes → select Always hide. Your preference is saved and applied to all future PRs automatically.

Temporarily reveal whitespace on a specific diff without losing your saved `Always hide` preference.

> 📍 When reviewing a PR with whitespace hidden, locate the `Hide whitespace` filter badge beside `All commits` and click its X to show whitespace for the current diff only.

— Exact menu path, option names, and badge behavior given.

`Pull requests: Saved whitespace prefere…`

---

05 **One-click NuGet vulnerability fixing via Copilot** NEW

82

A fixer action on NuGet Audit warnings in Visual Studio's Error List opens Copilot Chat in agent mode and uses the NuGet MCP server to remediate vulnerable packages. It requires signing in to GitHub Copilot and enabling the NuGet MCP server from the Copilot Chat tools menu.

Trigger an AI-assisted fix for a vulnerable NuGet package the moment Visual Studio flags it in the Error List.

> 📍 1. In the Error List, locate a NuGet Audit warning for a vulnerable package. 2. Click the fixer action on the warning. 3. Copilot Chat opens in agent mode with the NuGet MCP server and proposes a remediation — review and apply.

— Exact UI location and step-by-step remediation flow given.

`GitHub Copilot in Visual Studio — Septe…`

---

06 **New Actions blob storage hosts in meta API** BREAKING

80

The GitHub meta API now lists `productionresultssa30.blob.core.windows.net` through `productionresultssa49.blob.core.windows.net` and `productionresultssa100.blob.core.windows.net` through `productionresultssa199.blob.core.windows.net` as new GitHub Actions blob storage hosts. Organizations using per-host network allowlists must add these ranges or job summaries and logs will be interrupted.

— Exact host ranges and required allowlist action named.

`Meta API endpoints updated for Actions …`

---

07 **PR exploration via Git agent in Visual Studio Copilot** NEW

75

Copilot Chat's Git agent can attach an open pull request as context directly from the Git Repository window, returning a summary with navigable links to its comments and files. This requires connecting a GitHub or Azure DevOps MCP server.

Surface a summary of an open pull request — with links to its comments and files — without leaving Visual Studio.

> 📍 1. Open GitHub Copilot Chat. 2. Select the Git agent. 3. In the Git Repository window, right-click the pull request and choose `Attach as context`. 4. Ask: 'Summarize this pull request and list the files changed.'

— Clear mechanism and named MCP requirement, with exact steps.

GitHub Copilot in Visual Studio — Septe…

---

08 Bring Your Own Model preview in Visual Studio Copilot NEW

75

Bring Your Own Model (BYOM) is now in preview and enabled by default across Community, Professional, and Enterprise editions, supporting Microsoft Foundry, OpenAI, Anthropic, and Ollama providers, with or without a GitHub sign-in. BYOM now requires the new Agent (Preview) harness built on the GitHub Copilot SDK — the previous Ask and Agent modes no longer support it, and previously configured Ollama models must be re-added after updating.

— Names providers and editions but no setup steps given.

GitHub Copilot in Visual Studio — Septe…

THINNER COVERAGE BELOW

09 Full-width split view for new and deleted files IMPROVED

30

In split diff view, new and deleted files now span both columns (full width), matching unified view behavior to reduce line wrapping.

— One-line layout change with no further mechanism given.

Pull requests: Saved whitespace prefere…

NEW DETECTION Public monitoring for secret scanning in GHEC DR

└──▷ BREAKING ON UPGRADE

! BYOM now requires the new Agent (Preview) harness built on the GitHub Copilot SDK; the previous Ask and Agent modes no longer support BYOM, and previously configured Ollama models must be re-added after updating.

! Eligible generally available Copilot features and capabilities left `Unconfigured` will be automatically set to `Enabled`, `Disabled`, or `Let organizations decide` according to the selected global default — this is a one-time irreversible change; subsequent changes to the default policy will not revert features already set by the rollout.

! Organizations using per-host network allowlists for GitHub Actions blob storage must add `productionresultssa30` through `productionresultssa49` and `productionresultssa100`

through `productionresultssa199` (all under `.blob.core.windows.net`) or job summaries and logs will be interrupted.

> Public monitoring now available with data residency (2026-10-22) BLOG ↗
>
> STABLE
>
> Public monitoring for GHEC DR uses verified domain matching to attribute exposed secrets and deliver encrypted alerts within the enterprise's designated geography.
>
> Enables public monitoring for GitHub Enterprise Cloud with data residency (GHEC DR), extending secret scanning to public commits across github.com with attribution via verified domain matching.
>
> Delivers secret scanning alerts encrypted directly to the enterprise dashboard, keeping data within the designated geography through tenant isolation.
>
> Enterprise admins and enterprise security managers can enable public monitoring under enterprise security settings, included at no additional cost with an active GitHub Advanced Security or GitHub Secret Protection license.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Fix updates (2026-10-22) BLOG ↗ RANK

Veracode is a cloud-based application security platform that scans code, binaries, and running applications to find and help remediate software vulnerabilities.

> Veracode expanded its Fix for SCA capabilities with a new CLI command for dependency remediation and asynchronous processing for GitHub-based fix requests.

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

01 **veracode fix sca CLI command** NEW 85

New `veracode fix sca` CLI command (and accompanying SCM integration) analyzes vulnerable dependencies from SCA findings, generates remediation patches, and identifies safe upgrade paths.

Trigger SCA-based dependency fix suggestions from the CLI against a project with vulnerable dependencies.

```
$ veracode fix sca
```

— Names runnable CLI command and its exact mechanism Fix updates (2026-10-22)

02 **Async GitHub Workflow Integration for Fix for SCA** IMPROVED 65

Veracode Fix for SCA now processes fix requests from the GitHub Workflow Integration asynchronously, freeing the GitHub Actions runner immediately and delivering fixes as a new pull request with a results comment on the original PR.

— Describes workflow mechanism but no exact config/flag given Fix updates (2026-10-22)

WAS THIS USEFUL? 👍 👎

zizmor Feedback ⓘ v1.31.0 NOTES ↗ RANK

Static analysis for GitHub Actions (and more)

> zizmor v1.31.0 sharpens several of its GitHub Actions audits — cache-poisoning, dangerous-triggers, typosquat-uses, self-hosted-runner, and unpinned-images — for broader and more precise detection, and improves its PowerShell parsing.

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

01 Cache-mode awareness in cache-poisoning audit IMPROVED

72

The `cache-poisoning` audit now recognises the `cache-mode` setting: it skips findings when a job's effective `cache-mode` is `none`, and flags `cache-mode: write` and `cache-mode: write-only` usage when combined with dangerous triggers that grant an attacker privileged code execution. The audit also gains awareness of the `pnpm/setup` action.

— Names exact config key and values but no runnable command v1.31.0

THINNER COVERAGE BELOW

02 issue_comment flagged as dangerous trigger IMPROVED

48

The `dangerous-triggers` audit now treats the `issue_comment` trigger as fundamentally dangerous and produces more precise finding/diagnostic spans.

— Names the trigger but no concrete usage guidance v1.31.0

03 Pre-commit config inputs in typosquat-uses audit IMPROVED

40

The `typosquat-uses` audit now supports pre-commit config inputs.

— Clear scope addition but no detail on mechanism v1.31.0

04 More precise unpinned-images audit for matrix refs IMPROVED

35

The `unpinned-images` audit produces more precise findings for workflows containing indirect matrix references.

— Names the scenario but not the mechanism v1.31.0

05 More capable PowerShell parsing grammar IMPROVED

35

zizmor now uses a more capable tree-sitter grammar for PowerShell, capturing more non-trivial PowerShell expressions.

— Describes benefit but not which expressions newly covered v1.31.0

06 **Awareness of image_version workflow trigger** NEW 27

zizmor gains awareness of the `image_version` workflow trigger.

— Bare mention of a new recognised trigger v1.31.0

07 **More precise self-hosted-runner audit** IMPROVED 23

The `self-hosted-runner` audit is now significantly more precise in its findings.

— No specifics on what changed or how v1.31.0

WAS THIS USEFUL? 👍 👎

apkhunt

Feedback ⓘ v2.0.0 NOTES ↗ RANK

APKHunt is a comprehensive static code analysis tool for Android apps that is based on the OWASP MASVS framework. Although APKHunt is intended primarily for mobile app developers and security testers, it can be used by anyone to identify and address potential security vulnerabilities in their code.

> APKHunt v2.0.0 is a ground-up rewrite that replaces the Go CLI with a local Docker web workbench for Android APK static analysis, adding background scanning, an OWASP MASVS coverage map, entry-point and threat-model views, an ADB runtime-proof workbench, API route evidence collection, and CycloneDX SBOM generation with OSV enrichment.

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

01 Docker web workbench replaces Go CLI

BREAKING

95

The `go run apkhunt.go -p ...` and `-m ...` CLI commands are gone from `main` and replaced by a Docker Compose web workbench served at `http://localhost:5005`, launched via `export SECRET_KEY="$(openssl rand -hex 32)" && docker compose up --build`. The Docker image bundles JADX, Apktool, Syft, TruffleHog, curl, and ADB, so no separate host installation is required; existing CLI text reports are not imported into the new web app.

Stand up the full APKHunt web workbench locally for the first time — generates a fresh application secret and builds the Docker image.

```
$ git clone https://github.com/Cyber-Buddy/APKHunt.git
cd APKHunt
export SECRET_KEY="$(openssl rand -hex 32)"
docker compose up --build
```

— Names exact commands, port, env var and bundled tools v2.0.0

02 OWASP MASVS control coverage page NEW

75

A new control page maps all 24 current OWASP MASVS controls across eight groups, including which checks are missing, exposed at `/masvs-coverage`.

— Named route gives exact navigation target v2.0.0

03 CycloneDX SBOM generation with OSV enrichment NEW

70

Generates a CycloneDX SBOM via Syft and enriches versioned package identities with per-package OSV vulnerability data, preserving unavailable or unknown query states explicitly.

— Names SBOM format and tooling but no invocation shown v2.0.0

04 **API route evidence workbench** NEW 67

Inventories HTTPS routes from decompiled source without contacting hosts, and supports opt-in scoped anonymous `GET` requests against one authorized hostname for evidence collection.

— Names HTTP method and scope restriction concretely v2.0.0

05 **Coverage ledger for static analysis signals** NEW 65

A coverage ledger records APK member extraction states, recovered-file rule evaluations, skipped files, and decompiler attempts across 38 active static signals plus one inventory-only rule.

— Specific signal counts but no way to invoke or view it v2.0.0

06 **Deterministic Android threat model** NEW 62

A deterministic (no AI used) threat model generates Android trust boundaries, STRIDE review questions, counter-hypotheses, and source references, with a downloadable JSON output.

— Describes method and output format but no command shown v2.0.0

THINNER COVERAGE BELOW

07 **Background APK scanning with persistent status dock** NEW 57

Uploaded APKs (up to 200 MB by default) are scanned in the background with a status dock that persists across page navigation; the web UI queues one scan at a time.

— Gives size limit and queue behaviour but no exact UI path v2.0.0

08 **Entry point graphs from manifest components** NEW 55

Entry point graphs connect manifest-declared Android components to handler-class, permission-check, WebView, API, and sensitive-action evidence recovered from JADX-decompiled source.

— Explains evidence types but no access path given v2.0.0

09 **ADB runtime proof workbench** NEW 50

An ADB runtime proof workbench checks the installed base APK hash and records launch results for explicitly exported activities.

– States what it checks but no interface detail v2.0.0

└──▷ **BREAKING ON UPGRADE**

! The `go run apkhunt.go -p ...` and `-m ...` CLI commands are no longer present on `main`; existing CLI text reports are not imported into the web app.

WAS THIS USEFUL? 👍 👎

◆ **Secrets Detection & Management**

infisical

Feedback ⓘ v0.167.1 NOTES ↗ RANK

Infisical is the open-source platform for securing secrets, certificates, access, and AI Agents.

> Infisical introduced Agent Vault, letting AI agents operate with needed access and context without ever holding a credential, and extended PKI certificate discovery to Linux Servers.

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

01 **Agent Vault for AI agents** NEW 25

Agent Vault enables AI agents to run with the access, context, and capabilities they need without ever holding a credential.

.
.

– Describes purpose but no mechanism or config surface named product docs

02 **Linux Server certificate discovery in PKI** NEW 25

Adds Linux Server certificate discovery to the PKI certificate management surface.

– Names the surface but no detail on setup or scope v0.167.1

WAS THIS USEFUL? 👍 👎

Truffle Security trufflehog

Feedback ⓘ v3.99.4

NOTES ↗ CODE ↗

TruffleHog finds, verifies, and analyzes leaked credentials across Git, chat, wikis, cloud storage, and more

Adds per-source `debug` config field for targeted logging and a `Result.VerificationSkipped` output field to distinguish unverified from unverifiable results.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v3.99.4
https://github.com/trufflesecurity/trufflehog.git
# already have the repo? check out this version:
$ git checkout v3.99.4
```

Adds a `debug` field to `LocalSource` in source config files, enabling debug-level logging for a single source without raising the global log level — old configs default to off and are unaffected.

Adds `Result.VerificationSkipped` field to result output, distinguishing results where no verifier ran (due to `--no-verification`, a per-detector override, or a source verify setting) from genuine unverified verdicts.

NEW DETECTION Lambda Cloud API key detector

WAS THIS USEFUL? 👍 👎

kingfisher

Feedback ⓘ v2.14.0 NOTES ↗ RANK

Detect secrets + live validation. Map blast radius.

> Kingfisher's biggest addition this window is a native wizard GUI for configuring and reviewing scans, alongside Azure Blob Storage scanning, OAuth2 authentication for HTTP validators, Amazon Bedrock discovery in AWS blast-radius mapping, and stabilized 1.0.0 embedding APIs, plus a long tail of CLI flag, Git-scanning, and Jira/Confluence refinements.

└──▷ **WHAT SHIPPED · 24 FEATURES** most completely described first ? what's the number?

01 **Scan limit and timeout control flags** IMPROVED 91

Adds `--no-limits` flag to remove all scan budget and timeout caps in a single switch, later refined to separate `--no-limits` (lifts content caps, preserves timeouts) from `--no-timeouts` (disables scan operation deadlines). Library calls timeout(Duration::ZERO) and max_response_bytes(0) now disable their respective limits instead of returning a build error.

Remove scan operation deadlines when scanning very large repositories where timeouts fire before completion, while still keeping content caps lifted.

```
$ kingfisher scan /path/to/large-repo --no-limits --no-timeouts --format json
```

— Runnable example shows both flags and their distinct effects. v2.9.0 v2.14.0

02 **GitHub/GitLab org-wide scan parallelization and `--user-file`**

NEW INTEGRATION IMPROVED 86

Adds `--user-file` for GitHub and GitLab scans to supply a list of users or groups from a file. GitHub user/org and GitLab user/group scans now overlap discovery, cloning, and scanning for faster org-wide sweeps. Clone limits for these org scans now select the first unique repositories discovered per provider in API order; explicit Git URLs and added wiki URLs are not counted toward the limit.

Feed a list of GitHub usernames to scan all their repos in one invocation, with discovery, cloning, and scanning running in parallel.

```
$ KF_GITHUB_TOKEN="..." kingfisher scan github --user-file
users.txt --format json
```

— Exact flag plus runnable example and clone-limit mechanism. v2.9.0

03 **OAuth2 authentication for custom HTTP validators** NEW

85

Adds OAuth2 password-grant authentication for custom HTTP validators, with environment-backed credentials, token reuse/refresh, failure cooldowns, and fixed-origin delivery. Implemented via new `kingfisher-rules` 2.0.0 and `kingfisher-scanner` 2.0.0 packages with an exhaustive `HttpValidation.auth` field and opt-in OAuth2 service authentication; existing YAML and TOML rules load without changes. Rust code constructing `HttpValidation` struct literals must add `auth: None` or a populated Some(Box::new(HttpValidationAuth { oauth2: config })) when upgrading.

— Gives exact struct field and migration code, no CLI example. v2.14.0

04 **`KF_GIT_BINARY` and `KF_GIT_SSL_BACKEND` environment variables**

NEW 81

Adds `KF_GIT_BINARY` and `KF_GIT_SSL_BACKEND` environment variables for controlling Git binary selection and SSL backend on Windows.

Override the Git binary and SSL backend on Windows when the system Git causes timeout or certificate issues.

```
$ KF_GIT_BINARY="C:\Program Files\Git\bin\git.exe"
KF_GIT_SSL_BACKEND=schannel kingfisher scan /path/to/repo --
format json
```

— Exact env vars with runnable Windows example. v2.10.0

05 **`--since-commit` default scan scope changes** BREAKING

80

`--since-commit` now scans all refs by default instead of only the previously scoped branch; use `--branch HEAD` to retain the prior scope. This was further changed so `--since-commit` scans full commit ranges by default instead of net-diff; existing workflows relying on net-diff behaviour must add `--git-history none`.

— Exact flags and migration steps across two releases. v2.8.0 v2.9.0

06 **Deduplication now opt-in via `enable_dedup`** BREAKING

78

Changes `ScannerConfig::enable_dedup` default from `true` to `false`; deduplication is now opt-in, keying on both blob ID and source path to preserve identical content found at different paths.

— Exact config field, default change, and keying mechanism. v2.8.0

07 **Commit metadata redaction flag** NEW

78

Adds `--commit-metadata=false` flag to preserve repository links, commit IDs, and file paths while hiding author, committer, and date details.

Audit a full Git repository history while stripping author, committer, and date metadata from findings — useful when sharing reports externally.

```
$ kingfisher scan /path/to/repo --commit-metadata=false --format
json
```

— Exact flag with runnable example command. v2.14.0

08 **New Liquid filters: rsa_sha256, luhn_valid, iban_valid** NEW

75

Adds `rsa_sha256`, `luhn_valid`, and `iban_valid` Liquid filters: `rsa_sha256` enables RS256 JWT validation with PKCS#1/PKCS#8 PEM keys, while `luhn_valid` and `iban_valid` provide offline checksum validation for rule validation.

— Names exact filters and use case, no example rule snippet. v2.14.0

09 **`--since-hours` recent history scan flag** NEW

73

Adds `--since-hours N` to scan only recent Git history within the last N hours.

Scan only the last 24 hours of Git history across a repo to catch secrets introduced in recent commits.

```
$ kingfisher scan https://github.com/my-org/my-repo.git --since-hours 24 --format json
```

— Exact flag with runnable example. v2.9.0

10 **Native wizard GUI for scanning and report review** NEW

70

Adds the optional native `wizard` (`gui`) for configuring scans, tracking progress, and inspecting JSON/JSONL or SARIF reports with filters, exports, and Git provenance. Includes independent report tabs and native Kingfisher desktop icons for macOS, Linux, and Windows.

— Names the GUI and report formats, no launch command given. v2.8.0

11 **`--extraction-depth 0` disables archive depth limit** BREAKING

70

`--extraction-depth 0` now disables the archive depth limit instead of its previous behaviour; use `--no-extract-archives` to disable extraction entirely.

— Exact flag and alternative flag for full disable. v2.9.0

12 **`rules check` validates negative examples** IMPROVED

70

Improves `rules check` to verify `negative_examples` and apply scan-time capture selection, entropy, safe-list, pattern requirements, checksums, and rule filters to both positive and negative examples.

— Names subcommand and checks applied, no example invocation. v2.14.0

13 **Amazon Bedrock discovery in AWS blast-radius mapping**

UPDATED DETECTION NEW

65

Adds bounded Amazon Bedrock discovery to AWS blast-radius mapping with regional foundation-model and inference-profile visibility. JSON and both HTML viewers retain region/model context without invoking or enabling models; the standalone graph uses model IDs as labels with full ARNs shown in the details.

— Describes viewer behavior but no CLI flag or command given. v2.13.0

14 **Jira API stops following redirects** BREAKING

65

Jira API requests no longer follow redirects, including same-origin redirects. Reverse-proxy installations must pass the final canonical `--url` , including any context path such as `/jira` .

– Clear migration instruction naming the exact flag. v2.14.0

15 **Shared path/extension filtering across all storage scans**

IMPROVED 63

Adds shared path/extension selection across files, Git sources, and archive members. S3, GCS, Azure Blob, and Hugging Face storage scans now filter before downloading and support resumable listings and incremental caching.

– Explains filtering mechanism but no flag or config name given. v2.14.0

16 **Stabilized 1.0.0 embedding APIs** NEW 63

Stabilizes 1.0.0 embedding APIs for `kingfisher-core`, `kingfisher-rules`, `kingfisher-scanner`, and `kingfisher-bin`, with runnable examples and version-aware crates.io publishing.

– Names exact crates and versioning but no code sample given. v2.8.0

17 **Vectorscan crate switch; Windows builds move to MinGW**

60

BREAKING

Switches to the published `kingfisher-vectorscan` crate with checksum-verified prebuilt archives, removing the local Vectorscan CMake build on supported targets. Windows builds now use GNU/LLVM MinGW; MSVC is unsupported.

– Names the crate and build toolchain change, no migration steps. v2.8.0

18 **Git scan performance improvements** IMPROVED 60

Speeds up large Git history scans with candidate-first metadata, subtree pruning, and exact ancestor summaries; replaces Git subprocesses with `gix` for faster repository audit metadata; and improves `--branch` scan performance by avoiding per-commit Git index setup, with branch-only remote scans now fetching just the selected branch and keeping branch-specific clones separate from full-repository clones.

– Describes mechanism but no flag or benchmark given. Names the library swap but no benchmark or flag. Names the flag and mechanism, no benchmark numbers.

v2.14.0 v2.10.0 v2.8.0

THINNER COVERAGE BELOW

19 **Azure Blob Storage scanning with multiple auth methods** NEW

55

Adds Azure Blob Storage scanning with SAS, account-key, Entra-token, and anonymous authentication.

– Names auth methods but no CLI flag or config example given. v2.14.0

20 **Minimum supported Rust version raised to 1.99.0**

BREAKING

45

Minimum supported Rust version is now 1.99.0; builds on older Rust toolchains will fail.

– Exact version number, no migration steps beyond upgrading. v2.10.0

---

21 **Filter evaluation errors now fail the scan** BREAKING

45

Betterleaks filter evaluation errors now fail the scan call instead of silently keeping the finding.

– Describes behavior change, no example or config name. v2.8.0

---

22 **Python and Rust SDK embedding examples** IMPROVED

40

Adds Python rule-detail inspection and a Rust embedding example covering regexes, filters, dependencies, validation, and revocation definitions. Expands Python SDK examples and documentation for directory scans, filtering, validation, deduplication, and cancellation.

– Lists topics covered but no code path or API name. v2.10.0

---

23 **Jira access checks and Jira/Confluence rate-limit handling**

IMPROVED

38

Adds Jira access checks and project exclusions, and expands Jira/Confluence rate-limit handling, pagination, and API context paths.

– Names the areas touched but no config key or flag given. v2.14.0

---

24 **Git author name and email in finding metadata** NEW

35

Adds Git author name and email to finding metadata alongside existing committer information.

– Describes field addition without example output. v2.8.0

NEW DETECTION SMTP credential validation

NEW INTEGRATION Local Docker image scanning via Bollard (Docker Engine API)

UPDATED DETECTION Betterleaks catalog updated to v2.0.0-rc.1

└──▷ BREAKING ON UPGRADE

! Jira API requests no longer follow redirects, including same-origin redirects. Reverse-proxy installations must pass the final canonical `--url`, including any context path such as `/jira`.

! `--no-limits` now lifts content caps but preserves timeouts. To also disable scan operation deadlines, add `--no-timeouts` explicitly.

- ! Rust code constructing `HttpValidation` struct literals must add `auth: None` (or a populated Some(Box::new(HttpValidationAuth { oauth2: config })) when upgrading to `kingfisher-rules` 2.0.0; Rust scanner consumers must use rules 2.x types.
- ! Minimum supported Rust version is now 1.99.0; builds on older Rust toolchains will fail.
- ! `--extraction-depth 0` now disables the archive depth limit instead of its previous behaviour; use `--no-extract-archives` to disable extraction entirely.
- ! `--since-commit` now scans all refs by default instead of only the previously scoped branch; use `--branch HEAD` to retain the prior scope.
- ! Clone limits for GitHub/GitLab org scans now select the first unique repositories discovered per provider in API order; explicit Git URLs and added wiki URLs are not counted toward the limit.
- ! Library: timeout(Duration::ZERO) and max_response_bytes(0) no longer return a build error — they now disable their respective limits.
- ! `--since-commit` now scans full commit ranges by default instead of net-diff; existing workflows relying on net-diff behaviour must add `--git-history none`.
- ! `ScannerConfig::enable_dedup` now defaults to `false` (previously `true`); library consumers relying on automatic deduplication must explicitly opt in.
- ! Betterleaks filter evaluation errors now fail the scan call instead of silently keeping the finding.
- ! Windows builds now use GNU/LLVM MinGW; MSVC is unsupported.

v2.12.0 NOTES ↗ STABLE

Adds offline Python revocation preparation, batch revocation, result-based Python revocation, and opt-in AWS responder credentials with explicit deactivate/delete actions.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.12.0
https://github.com/mongodb/kingfisher.git
# already have the repo? check out this version:
$ git checkout v2.12.0
```

Adds opt-in AWS responder credentials via a named profile or the default credential chain, with explicit `deactivate` or `delete` actions — deletion is never an automatic fallback.

Adds offline Python revocation preparation with exact rule IDs/names or deterministic provider families, explicit candidate validation, and advisory-only typo suggestions.

Adds source-text/file revocation preparation, structured missing-input diagnostics, and explicit batch revocation with complete-input deduplication and retained

interruption outcomes.

Adds resolved credential inputs and result-based Python revocation, automatically reusing validated AWS pairs and other supporting values.

Releases `kingfisher-scanner` 1.4.0 with structured AWS revocation options and outcome API.

+3 more

v2.11.0 NOTES ↗ STABLE

Adds composable Python scan inputs, Git-scope enumeration with `GitScope`/git_inputs(), `DetectionScanner`, expand_content(), SDK caching via Rules(cache_dir=...)/Rules(cache=False), and cancellable validation deadlines.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.11.0
https://github.com/mongodb/kingfisher.git
# already have the repo? check out this version:
$ git checkout v2.11.0
```

Enables Python SDK caching by default with Rules(cache_dir=...) and Rules(cache=False) controls.

Adds composable Python filesystem and Git-history enumeration through `ScanInput`, filesystem(), git_history(), and Scanner.scan_inputs().

Adds native git_inputs() with `GitScope` history baselines, inclusive branch roots, time windows, snapshots, net diffs, staged index content and unreachable object coverage; `GitInput` retains author/committer identities, times, parents and messages.

Adds opt-in Python `DetectionScanner` and Rust detection options for shared CLI confirmation/component-window semantics, credential-URI fallback, secret containment suppression, bounded two-level Base64 decoding, inline-ignore and HTML/CSS context filtering.

Adds composable expand_content() for shared SQLite SQL and Python bytecode string extraction from files, historical blobs and archive members, with source provenance and per-input budgets.

+5 more

v2.7.0 NOTES ↗ STABLE

Candidate verification caps searches at 16 combinations in 30 seconds, keeps unresolved pairings inconclusive, disables HTTP redirects, and excludes unselected

secrets from reports.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.7.0
https://github.com/mongodb/kingfisher.git
# already have the repo? check out this version:
$ git checkout v2.7.0
```

v2.6.0 NOTES ↗ STABLE

Adds `--disk-offload` to cap finding memory across large multi-repo scans; reduces peak memory via digest-only deduplication and removal of the Bloom prefilter.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.6.0
https://github.com/mongodb/kingfisher.git
# already have the repo? check out this version:
$ git checkout v2.6.0
```

Adds `--disk-offload` flag to offload accumulated finding memory to disk across repositories, falling back to memory with a warning when temporary storage fills up.

Reduces peak scan memory by deduplicating findings on exact cryptographic digests alone, removing the Bloom prefilter and its dependency, and slimming Git object index and raw matcher bookkeeping.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

Cycode

Feedback ⓘ

Introducing Custom Maestro Agents in Agentic Workflows (2026-10-22) 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 introduced Custom Maestro Agents — user-defined AI agents with configurable instructions and outcomes that plug into Agentic Workflows to automate violation triage — along with confidence gating, human-approval and stability controls, and token-usage tracking for these agents.

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

01 **Custom Maestro Agents in AI agents library** NEW 87

Users can define custom Maestro agents in the AI agents library by specifying analysis instructions, evidence to weigh, and up to three named outcomes (plus a reserved Failure outcome) — covering use cases like upgrade impact analysis, credential validity, and code ownership mapping. Agents can be tested against a real violation before publishing, returning the verdict, reasoning, and token consumption.

Route a widespread CVE to one coordinated Jira ticket while sending isolated findings to the owning team — built entirely in the Agentic Workflows UI using a custom agent.

📍 1. Go to AI Agents Library › New Agent. 2. Name it (e.g. Widespread or Isolated CVE), describe its purpose, and in the analysis definition tell it to count open violations sharing the same CVE across the org against a threshold you specify. 3. Define two outcomes: Widespread (CVE count exceeds threshold) and Isolated (below threshold). Maestro adds Failure automatically. 4. Test the agent against a real violation; review the verdict, reasoning, and token consumption. 5. Publish the agent, then open the target Agentic Workflow › Add Step › select your new agent. 6. Under the Widespread branch: open one coordinated Jira ticket + post to Slack. Under Isolated : route to the owning team's board.

— Names library UI path, outcome limits, and test-before-publish flow.

Introducing Custom Maestro Agents in Ag…

02 **Custom Maestro agents as Agentic Workflow steps** NEW 79

Custom Maestro agents can be added as steps in Agentic Workflows, where each outcome branch drives a downstream action: open a Jira ticket, message a Slack channel, set severity or status, add a label, open a fix pull request, or hand off to another agent.

Route a widespread CVE to one coordinated Jira ticket while sending isolated findings to the owning team — built entirely in the Agentic Workflows UI using a custom agent.

📍 1. Go to AI Agents Library › New Agent. 2. Name it (e.g. `Widespread or Isolated CVE`), describe its purpose, and in the analysis definition tell it to count open violations sharing the same CVE across the org against a threshold you specify. 3. Define two outcomes: `Widespread` (CVE count exceeds threshold) and `Isolated` (below threshold). Maestro adds `Failure` automatically. 4. Test the agent against a real violation; review the verdict, reasoning, and token consumption. 5. Publish the agent, then open the target Agentic Workflow › Add Step › select your new agent. 6. Under the `Widespread` branch: open one coordinated Jira ticket + post to Slack. Under `Isolated`: route to the owning team's board.

— Lists every action type and shows the workflow step UI path.

Introducing Custom Maestro Agents in Ag…

03 **Human approval and stable-logic safety controls for agents**

NEW 62

A `Require human approval` toggle can be set on any outcome branch's first action, surfacing the agent's recommendation in Required Approvals before anything executes. `Stable logic` locks an agent's outcomes while a workflow uses it, preventing branch-breaking changes as the agent's instructions evolve.

Prevent accidental action on low-confidence verdicts and require a human sign-off before silencing a real leaked credential.

📍 1. In the Agentic Workflow outcome branch for your `Live Secret vs Test Credential` agent, set the confidence threshold above 70% (e.g. 90%) on the `Real Credential` branch. 2. Toggle `Require human approval` on the first action (e.g. `Post to incident channel`). 3. When the agent fires, the recommendation appears in Required Approvals — no action occurs until a reviewer approves.

— Names both controls and where approvals surface, but no deeper mechanism.

Introducing Custom Maestro Agents in Ag…

04 **Confidence thresholds per outcome branch** NEW 60

Each outcome branch in an Agentic Workflow can be given a configurable confidence threshold, defaulting to 70%, so the branch only executes when the agent's confidence meets the configured minimum.

Prevent accidental action on low-confidence verdicts and require a human sign-off before silencing a real leaked credential.

> 📍 1. In the Agentic Workflow outcome branch for your `Live Secret vs Test Credential` agent, set the confidence threshold above 70% (e.g. 90%) on the `Real Credential` branch. 2. Toggle `Require human approval` on the first action (e.g. `Post to incident channel`). 3. When the agent fires, the recommendation appears in Required Approvals — no action occurs until a reviewer approves.

— Gives the default value but limited mechanism beyond that.

Introducing Custom Maestro Agents in Ag…

THINNER COVERAGE BELOW

05 **Token usage tracking for custom agents** NEW 41

The AI usage dashboard now shows per-workflow token attribution, with a ready-made template that alerts when usage reaches 80% of quota.

— Names the dashboard and 80% threshold but no setup steps given.

Introducing Custom Maestro Agents in Ag…

WAS THIS USEFUL? 👍 👎

Cloudsmith

Feedback ⓘ

CLI v1.30.0: Error reporting, TLS-intercepting proxy support, and a username and passwor

BLOG ↗ RANK

Cloudsmith is a fully-managed artifact management, software supply chain security, and software distribution platform.

> Cloudsmith CLI v1.30.0 adds anonymous error reporting and TLS-intercepting proxy support, moves credential helper installation to a more durable location, and deprecates username/password login ahead of a November 2026 cutoff.

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

01 **Username/password login deprecated** BREAKING 75

Username and password login via `cloudsmith login` and `cloudsmith token` is deprecated and stops working on November 20, 2026.

– Exact commands and cutoff date given, but no migration path named

CLI v1.30.0: Error reporting, TLS-inter…

THINNER COVERAGE BELOW

02 **TLS-intercepting proxy support** NEW 50

The CLI now supports TLS-intercepting proxies without disabling certificate verification, enabling CLI use in enterprise environments with SSL inspection.

– Describes mechanism but no flag or config to enable it

CLI v1.30.0: Error reporting, TLS-inter…

03 **Anonymous error reporting** NEW 40

The CLI sends anonymous reports of unexpected errors to Cloudsmith by default, allowing crash detection without manual support tickets.

– Explains behavior but no opt-out flag or setting named

CLI v1.30.0: Error reporting, TLS-inter…

04 **Persistent credential helper installation** IMPROVED 40

Credential helpers are now installed in a location that survives package-manager upgrades, preventing accidental removal.

– States the change but not the install path itself

CLI v1.30.0: Error reporting, TLS-inter…

└──▷ **BREAKING ON UPGRADE**

! Username and password login via `cloudsmith login` and `cloudsmith token` is deprecated and stops working on November 20, 2026.

WAS THIS USEFUL?  

plumber

Feedback ⓘ v0.7.11 NOTES ↗ RANK

Plumber detects CI/CD security issues in your GitHub workflows and gives you a score

> Plumber's biggest change this window is making contextual scoring the default output everywhere (with a flag and env var to opt back to the old formula), alongside a new capability recognition engine, remote repository scanning without cloning, and deeper GitHub Actions and GitLab CI pipeline analysis.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 Contextual score becomes the default output BREAKING

90

Score letters, `--min-score` and `--min-points` gates, and badges now read the contextual score by default, backed by a labelled corpus and a migration report. Run with `--score-profile v3` or set `PLUMBER_ANALYZE_SCORE_PROFILE=v3` to keep the previous formula.

Preserve the pre-v0.6.0 scoring formula in CI so existing score gates and badges are not disrupted during migration.

```
$ plumber analyze --score-profile v3 github.com/owner/repo
```

Set the legacy score profile via environment variable for all plumber invocations in a CI runner without changing command flags.

```
$ export PLUMBER_ANALYZE_SCORE_PROFILE=v3
plumber analyze
```

— Exact flag and env var given with runnable migration commands. v0.6.0

02 Capability recognition engine and platform API NEW

85

Adds a capability recognizer with five evidence rules that detects capabilities from pipeline context or `.plumber.yaml` declarations, a `plumber capabilities defaults` command to print the built-in capability list ready to copy into `.plumber.yaml`, and a `--pbom` flag so a push carries recognized capabilities in the PBOM output. Platform mode now exposes the full built-in capability list rather than a restricted subset.

Bootstrap a capability declaration in `.plumber.yaml` by printing the full built-in list as a starting point to copy and trim.

```
$ plumber capabilities defaults
```

— Names the command and flag with a runnable bootstrap example. v0.7.5 v0.7.8

03 **Positional repository target for `plumber analyze`** NEW

75

`plumber analyze` now accepts a repository as a positional argument, e.g. `github.com/owner/repo` or `gitlab.com/group/project`, directly on the command line without a local clone.

Audit a public GitHub repo for CI/CD security gaps without cloning it locally — useful for supply-chain reviews or scripted audits.

```
$ plumber analyze github.com/owner/repo
```

Scan a self-hosted GitLab project remotely in one step, with a token in the environment for authenticated access.

```
$ GITLAB_TOKEN=glpat_xxxx plumber analyze
gitlab.com/group/project
```

— Two runnable commands including authenticated remote scan example. v0.5.17

04 **GitLab CI authorized component and function sources**

NEW INTEGRATION NEW

65

Adds a control requiring GitLab CI components and `run:` step functions to come from authorized sources, collects `run:` step functions and the registry host in the IR, and includes component and function sources as inputs to the contextual score. Plumber's own GitLab component is trusted by default in the built-in configuration.

— Mechanism and scoring impact explained but no exact config key shown. v0.6.1

THINNER COVERAGE BELOW

05 **Dedicated per-job environment and GITLAB_TOKEN scoping**

IMPROVED

55

Runs each job in a dedicated environment, isolating execution for improved security posture, and scopes `GITLAB_TOKEN` to the `plumber` environment in GitLab CI/CD variables to limit token exposure.

– Names the token and environment but no config steps to apply it. v0.5.21

06 **GitHub Actions trigger and permission analysis**

UPDATED DETECTION NEW 50

Adds GitHub Actions analysis for who can trigger a job, environment protections, inert jobs, and the real default token permissions.

– Lists analysis targets but no flag or command to invoke it. v0.7.1

07 **Situation analysis facts engine** NEW 50

The `situation` command gains job facts for attack paths, entry subjects, push trigger filters, and cache inheritance. Per-job situation facts and their engine entry point are evaluated and attached to the control analysis output.

– Names the facts added but no exact invocation shown. v0.5.19 v0.6.0

08 **Pipeline model (IR) field enrichment** IMPROVED 50

The internal pipeline model gains `visibility`, `runs-on`, `needs`, `caches`, `artifacts`, and `environment` fields for GitHub Actions and GitLab CI jobs, `if` conditions for GitHub Actions jobs, actions, and run scripts, and `artifacts.reports` key tracking so each job carries the keys of its reports and the include that defined it.

– Many named fields but purely internal, nothing to action directly. v0.5.18 v0.5.13 v0.7.5

09 **Claim-state tracking and three-valued score evaluator**

IMPROVED 45

Adds claim-state tracking to every path claim in the report, using the Radar corpus as a regression gate, and introduces a three-valued evaluator for job conditions and claim states in the scoring engine.

– Mechanism described but no reader-facing command or config. v0.7.1

10 **GitLab CI `only`/`except` keyword support** NEW 35

Adds support for GitLab CI `only` and `except` keywords on jobs.

– Bare feature name with no further mechanism given. v0.5.19

NEW DETECTION Invalid GitHub workflow file detection

UPDATED DETECTION Forbidden tag control — image exemption

└──▷ **BREAKING ON UPGRADE**

! Score letters, `--min-score` and `--min-points` gates, and badges now read the contextual score by default. Run with `--score-profile v3` or set `PLUMBER_ANALYZE_SCORE_PROFILE=v3` to keep the previous formula.

v0.7.11 NOTES ↗ STABLE

└──▷ GET THIS VERSION

```
$ git clone --branch v0.7.11
https://github.com/getplumber/plumber.git
# already have the repo? check out this version:
$ git checkout v0.7.11
```

v0.7.9 NOTES ↗ STABLE

└──▷ GET THIS VERSION

```
$ git clone --branch v0.7.9
https://github.com/getplumber/plumber.git
# already have the repo? check out this version:
$ git checkout v0.7.9
```

v0.6.4 NOTES ↗ STABLE

└──▷ GET THIS VERSION

```
$ git clone --branch v0.6.4
https://github.com/getplumber/plumber.git
# already have the repo? check out this version:
$ git checkout v0.6.4
```

v0.5.12 NOTES ↗ STABLE

└──▷ GET THIS VERSION

```
$ git clone --branch v0.5.12
https://github.com/getplumber/plumber.git
# already have the repo? check out this version:
$ git checkout v0.5.12
```

WAS THIS USEFUL?  

sbom-operator

Feedback ⓘ 0.47.0 NOTES ↗ RANK

Catalogue all images of a Kubernetes cluster to multiple targets with Syft

> sbom-operator 0.47.0 adds support for downloading SBOMs from OCI cosign attestations.

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

01 **SBOM download from cosign attestations** NEW 45

sbom-operator can now download SBOMs from OCI cosign attestations, including any additional SBOMs embedded within the attestation.

— Describes what changed but no config or command shown. 0.47.0

WAS THIS USEFUL? 👍 👎

cyclonedx-core-java

Feedback ⓘ cyclonedx-core-java-13.3.0

NOTES ↗ RANK

CycloneDX SBOM Model and Utils for Creating and Validating BOMs

> cyclonedx-core-java 13.3.0 extends the Dependency model for CycloneDX 1.6 spec compliance and updates license data to the latest SPDX list.

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

01 **Dependency `provides` field support** NEW 50

Adds `provides` mapping to the Dependency model to comply with the CycloneDX 1.6 spec.

— Names exact field and spec version but no usage example cyclonedx-core-java-13.3.0

02 **SPDX license list and mapping updates** IMPROVED 35

Bumps the bundled SPDX license list to version 3.29.0, expanding license identifier coverage, and improves license string mappings for more accurate SPDX identifier resolution.

— Names version bump but mapping improvement is vague cyclonedx-core-java-13.3.0

WAS THIS USEFUL? 👍 👎

zarf

Feedback ⓘ v0.88.0 NOTES ↗ RANK

The Airgap Native Package Manager for Kubernetes

> Zarf's v0.88.0 and v0.87.0 releases extend `v1beta1` package format support across the dev CLI toolchain and overhaul package signing into a decoupled, per-component workflow, while shipping several breaking changes to package layout, state storage, and internal APIs.

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

01 v1beta1 package format support across dev CLI IMPROVED

78

Extends `zarf dev lint` to validate, `zarf dev generate-schema` to emit JSON schemas for, `zarf dev inspect` to support, and `zarf dev find-images` to support `v1beta1` packages and component configs, with `v1beta1` commands now publicly visible in the CLI behind an opt-in feature flag. Adds `v1beta1` init config support for cluster initialization with JSON Schema validation for init chart values, honors explicit image sources in `v1beta1` packages, adds component import validation for `v1beta1` packages, and stores multiple API versions under `v1beta1` in deployed packages for forward-compatible state tracking. Various previously-undefined `v1beta1` edge-case behaviours have also been cleaned up, which may change behaviour for packages relying on them.

— Names every command and surface but lacks a worked example v0.88.0 v0.87.0

02 `--components` filter for `zarf dev inspect values-files`

NEW

75

Adds `--components` flag to `zarf dev inspect values-files` to filter inspection by specific components, e.g. `zarf dev inspect values-files --components=my-component ./zarf.yaml`.

Inspect Helm values files for only specific components in a package definition, rather than all components.

```
$ zarf dev inspect values-files --components=my-component
./zarf.yaml
```

— Exact flag plus a runnable example command v0.87.0

03 `--skip-arch-check` flag for package deploy NEW

70

Adds `--skip-arch-check` to `zarf package deploy` to bypass architecture validation when deploying packages to mismatched hosts.

— Exact flag and command given, no usage example v0.88.0

04 **Component `services` spec changed to named object** BREAKING

65

Component `services` in the package spec changes from a list to a named object (map); existing `zarf.yaml` definitions using a list for `services` must be updated.

— States before/after but no migration example v0.88.0

05 **Component signing and decoupled signing workflow** BREAKING

65

Adds component sign and verify capability, enabling individual components to be signed and their signatures verified, and moves package signing out of `zarf package publish` into a dedicated signing step, decoupling publish and sign workflows; callers that expected `publish` to sign must now invoke signing separately.

— Describes mechanism and workflow change but no exact command v0.87.0

THINNER COVERAGE BELOW

06 **Internal package type and signing API changes** BREAKING

54

`SignBlobOptions` is now passed as a pointer, so code calling the signing API with a value type will fail to compile; the internal package type is switched from `v1alpha1.ZarfPackage` to `api.Package` throughout the codebase, requiring updates to any code depending on `v1alpha1.ZarfPackage` directly.

— Names both Go API changes, developer-facing only v0.87.0

07 **Checksum verification for file sources** NEW

50

Adds `sha256` and `sha512` checksum verification for file sources in package definitions.

— Names the hashes but not how to configure them v0.88.0

08 **Symlinks materialized as real files on package creation**

50

BREAKING

Package creation now materializes symlinks as real files instead of preserving them as symlinks; packages that relied on symlink preservation will produce different layout content.

— Clear before/after but no config to adjust v0.88.0

09 **File path isolation enforced in package create** BREAKING

50

File paths in `zarf package create` are now isolated and enforced to the layout directory; references outside the layout root will be rejected.

– Names the command and new restriction, no remediation steps v0.88.0

10 **Artifact server disabled by default** BREAKING

49

The artifact server is now disabled by default, reducing attack surface in deployments that do not require it; deployments relying on it must now explicitly enable it.

– States the change but not how to re-enable it v0.87.0

11 **Images and repositories removed from package secret** BREAKING

47

Images and repositories are no longer stored in the package secret during connected deploys; existing tooling that reads those fields from the secret will find them absent.

– Clear before/after, no replacement location named v0.87.0

12 **Explicit git server mode field in Zarf state** NEW

45

Adds an explicit git server mode field to Zarf state, allowing the git server mode to be tracked and queried independently.

– Explains purpose but not the field name or location v0.87.0

13 **Helm pod-template labeling scoped to built-in workload kinds**

45

BREAKING

Helm pod-template labeling now applies only to built-in workload kinds; custom resource pod templates that previously received labels will no longer be labeled.

– Describes scope change but no workaround given v0.87.0

14 **Validation for `api.Package` definitions** NEW

38

Adds validation for `api.Package` to catch malformed package definitions earlier in the workflow.

– Names the type but not what checks run v0.87.0

└──▷ **BREAKING ON UPGRADE**

! Symlinks in package source directories are now materialized as real files on package creation; packages that relied on symlink preservation will produce different layout content.

- ! File paths in `zarf package create` are now isolated and enforced to the layout directory; references outside the layout root will be rejected.
- ! Component `services` is now a named object (map) rather than a list; existing `zarf.yaml` definitions using a list for `services` must be updated.
- ! `v1beta1` edge-case behaviours have been cleaned up; packages relying on previously-undefined `v1beta1` semantics may behave differently.
- ! Images and repositories are no longer stored in the package secret during connected deploys; existing tooling that reads those fields from the secret will find them absent.
- ! Helm pod-template labeling now applies only to built-in workload kinds; custom resource pod templates that previously received labels will no longer be labeled.
- ! The artifact server is disabled by default; deployments relying on it must now explicitly enable it.
- ! `SignBlobOptions` is now passed as a pointer; code calling the signing API with a value type will fail to compile.
- ! Package signing is no longer performed as part of `zarf package publish`; callers that expected publish to sign must now invoke signing separately.
- ! The internal package type is switched from `v1alpha1.ZarfPackage` to `api.Package` throughout the codebase; any code depending on `v1alpha1.ZarfPackage` directly must be updated.

WAS THIS USEFUL? 👍 👎

◆ WAF, Bot & API Protection

Fastly

Feedback ⓘ

Automatically enable CVE virtual patches (2026-10-22) BLOG ↗ RANK

Fastly is an edge cloud platform offering content delivery network, load balancing, edge computing, and web application security services for developers and enterprises.

> Fastly shipped automatic CVE virtual patching for workspace protection and added configurable API endpoint targeting to its beta Terraform provider.

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

01 Configurable API endpoint in Terraform provider NEW

85

The Fastly Terraform Provider (beta) 0.2.5 adds an `api_endpoint` provider argument and a `FASTLY_API_URL` environment variable to target a non-default Fastly API endpoint (HTTPS only), useful for staging or private Fastly deployments.

Point the Fastly Terraform provider at a non-default API endpoint — useful for staging environments or private Fastly deployments.

hcl

```hcl
provider "fastly" {
  api_endpoint = "https://api.staging.fastly.com"
}
```

Override the API endpoint via environment variable without changing Terraform config — useful in CI pipelines targeting different Fastly environments.

```
$ export FASTLY_API_URL=https://api.staging.fastly.com
terraform apply
```

— Names exact flag, env var, and runnable config/shell examples

Fastly Terraform Provider (beta) 0.2.5 …

THINNER COVERAGE BELOW

02 Automatic CVE virtual patching NEW

50

A new Automatic CVE Virtual Patching setting in the Fastly control panel auto-applies all future CVE virtual patches to a workspace as soon as they are released, removing the need for manual intervention per CVE.

Turn on automatic CVE virtual patching so your workspace is protected by new CVE patches the moment Fastly releases them, without manual intervention.

> 📍 In the Fastly control panel, go to your workspace settings and enable `Automatic CVE Virtual Patching`.

— Names the setting and UI path but no scope or limits

`Automatically enable CVE virtual patche…`

WAS THIS USEFUL? 👍 👎

safeline

Feedback ⓘ v9.4.3 NOTES ↗ RANK

SafeLine is a self-hosted WAF(Web Application Firewall) / reverse proxy to protect your web apps from attacks and exploits.

> SafeLine v9.4.3 adds a non-blocking Monitor Mode for Deny Rules, stabilizes certificate fingerprints across ACME renewals, and introduces quota handling when downgrading to Personal/Lite plans.

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

01 **Monitor Mode for Deny Rules** NEW 60

Deny Rules can now be set to Monitor Mode, which logs matched requests instead of blocking them, letting admins observe traffic patterns and tune rules before enforcing blocks in production. Configured via Protection > Allow & Deny Rules in the SafeLine console by setting a rule's action to Monitor Mode.

Set a Deny Rule to Monitor Mode to observe traffic patterns and tune rules before enforcing blocks in production.

> 📍 In the SafeLine console, go to Protection › Allow & Deny Rules, open or create a Deny Rule, and set its action to Monitor Mode to log matched requests without blocking them.

— Names exact UI path and rule action to enable it v9.4.3

THINNER COVERAGE BELOW

02 **Application quota notice on plan downgrade** NEW 45

When downgrading to the Personal or Lite plan, SafeLine now shows an application quota notice; applications exceeding the new limit are disabled while their data and configurations are retained.

— Describes behavior but no steps to act on it v9.4.3

03 **Stable key reuse on ACME auto-renewal** IMPROVED 40

ACME auto-renewal now reuses the existing private key on certificate renewal, keeping the public key fingerprint stable across renewals.

— Explains mechanism but no user-facing action v9.4.3

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

burp-ai-agent

Feedback ⓘ v1.1.0 NOTES ↗ RANK

Burp Suite extension that adds built-in MCP tooling, AI-assisted analysis, privacy controls, passive and active scanning and more

> burp-ai-agent v1.1.0 splits its build into separate BApp Store and full JARs with checksum and dependency-list verification, adds a Suggested targets panel for passive-to-active scan handoff, and removes OkHttp in favor of Burp's networking layer, which breaks older drop-in backend JARs and resets scanner on/off state on every load.

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

01 Separate BApp and full-build JAR tasks with verification

95

BREAKING

Adds `./gradlew bappJar` and `./gradlew fullJar` build tasks, producing `build/libs/custom-ai-agent-bapp.jar` (BApp Store build) and `build/libs/Custom-AI-Agent-full-<version>.jar` (GitHub full build) from separate source roots; adds a `bom.json` to each release listing the full build's runtime dependencies, including every dependency the BApp Store build packages, for dependency version verification; and adds per-release `.sha256` files in `sha256sum` format, verifiable with `sha256sum -c <asset>.sha256`. `./gradlew shadowJar` and the `-PstoreBuild` flag are removed; `shadowJar` now fails with a message naming the two replacement tasks.

— Names exact gradle tasks, file paths and verification command v1.1.0

02 Suggested targets panel for passive-to-active handoff NEW

88

Adds a Suggested targets section to the AI Active Scanner Queue dialog, listing passive findings by Suggested At, Vulnerability, Injection Point, URL, and Confidence, with Queue N suggested targets and Queue selected actions and a cap of 500 entries per session. A new opt-in Auto-queue passive findings for active testing setting on the AI Active Scanner tab (full build only) defaults OFF on every install, including upgrades from 1.0.1 where auto-queue had been enabled.

— Clear UI path and fields, no runnable command v1.1.0

03 Drop-in backend JARs and removal of OkHttp BREAKING

88

The full build loads drop-in backend JARs from `~/.burp-ai-agent/backends`; the BApp Store build never creates, lists, or loads that directory. OkHttp is removed from both JARs, so every AI request now goes through Burp's networking layer — drop-in backend JARs built against the OkHttp classes or HTTP helpers bundled in earlier versions no longer load in the full build, with the failure written to Burp's error output and no backend from that directory registered on that load.

— Names config path and explains breakage mechanism precisely v1.1.0

---

04 **AI Active Scanner no longer auto-resumes, applies settings live**

75

BREAKING

Both the passive and active AI scanners now start OFF at every extension load instead of resuming their saved on/off state, including upgrades from 1.0.1 that had a scanner enabled. The AI Active Scan confirmation dialog now reads 'This will send active test payloads to N target(s) on H host(s).', and AI Active Scanner settings (spinners, risk level, scan mode, checkboxes) and Save now apply to the next target without restarting the scanner, preserving pending Collaborator checks and baselines.

— Describes before/after behaviour but no direct config surface v1.1.0

---

05 **Status pill health checks without extra requests** IMPROVED

70

The status pill updates from each chat reply in both builds, showing AI: OK, AI disabled in Burp, or AI: Offline with error detail in the tooltip, without an extra request. The BApp Store build performs health checks on demand only (no startup check and no timer), showing Not checked · click to test until triggered by Save, a backend change, a pill click, Test connection, or the first chat message.

— Describes exact states and triggers but no command v1.1.0

---

06 **Automatic Ollama context window discovery** IMPROVED

65

Ollama queries `/api/show` before the first message to discover the model's context window when the setting is left at its default of 8192, reusing the result for the connection.

— Names API call and default value, limited scope v1.1.0

THINNER COVERAGE BELOW

---

07 **Non-blocking chat, site-map, and connection test actions**

IMPROVED

35

Chat send, site-map actions, and Test connection now run in the background, no longer blocking Burp's UI.

— Brief statement with no mechanism or named surface v1.1.0

└──▷ BREAKING ON UPGRADE

! `./gradlew shadowJar` and the `-PstoreBuild` flag are removed; `shadowJar` now fails with a message naming the two replacement tasks (`bappJar` and `fullJar`).

! The AI passive and active scanners no longer resume their saved on/off state at extension load — both start OFF in every install, including upgrades from 1.0.1 that had a scanner enabled.

! An install that had auto-queue enabled in 1.0.1 will find it turned OFF after upgrading to 1.1.0.

! Drop-in backend JARs in `~/.burp-ai-agent/backends` built against the OkHttp classes or HTTP helpers bundled in earlier versions no longer load in the full build; the failure is written to Burp's error output and no drop-in backend from that directory is registered on that load.

! OkHttp is removed from both JARs; every AI request now goes through Burp's networking layer.

WAS THIS USEFUL? 👍 👎

sitadel

Feedback ⓘ v2.0.0 NOTES ↗ RANK

Web Application Security Scanner

> Sitadel v2.0.0 is a major expansion adding a real-time TUI dashboard, new discovery modules (OpenAPI, WebSocket, exposed-file, HTML form), a dedicated JWT/token attack module, new vulnerability checks for XXE/SSRF/SSTI and more, and richer finding enrichment with OWASP/CWE/WSTG mapping.

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

01 **JWT and token security testing module** NEW 75

New module tests JWT/token handling for `alg=none` acceptance, RS256-to-HS256 algorithm confusion, `kid` parameter abuse, and weak signing secrets.

— Names specific attack techniques, no command example v2.0.0

02 **Real-time TUI dashboard during scans** NEW 70

Adds a `--tui` flag that enables an optional real-time terminal UI dashboard while a scan runs.

— Names the exact flag but no further UI detail v2.0.0

03 **New discovery modules for APIs, WebSockets, files and forms** NEW 60

Adds API-first discovery via OpenAPI/Swagger specs with JSON/XML body injection testing, WebSocket discovery with a Cross-Site WebSocket Hijacking (CSWSH) attack module, exposed-file discovery with passive secret-leak scanning of responses, and HTML form discovery with POST/body parameter injection testing.

— Names four modules and their techniques, no flags given v2.0.0

THINNER COVERAGE BELOW

04 **New vulnerability checks: XXE, SSRF, SSTI and more** NEW 50

Adds in-band XXE (XML external entity) file-disclosure detection, in-band SSRF detection, security-headers audit and CORS misconfiguration checks, and open redirect, SSTI, and path-traversal/LFI checks.

— Names several checks but no mechanism or config detail v2.0.0

05 **Enriched findings with confidence and compliance mapping** IMPROVED 50

Findings now include deduplication, confidence scores, evidence, remediation guidance, and OWASP/CWE/WSTG mapping.

— Lists enrichment fields but no output format or flag v2.0.0

06 **HTTP rate limiting and scope restriction controls** NEW

40

Adds controls for HTTP rate limiting and restricting scans to in-scope hosts.

— Describes the controls but no config keys or flags named v2.0.0

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

pwndbg

Feedback ⓘ 2026.10.22 NOTES ↗ CODE ↗ RANK

Exploit Development and Reverse Engineering with GDB & LLDB Made Easy

> pwndbg 2026.10.22 adds full m68k architecture support, Objective-C class inspection, Android Bionic libc detection, Linux 7.0 SLUB sheaf/barn visibility, a setrlimit command, SIMD register display, and several LLDB launch/targeting options.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 Objective-C class inspection commands NEW

88

`objc show-class -n <ClassName>` displays Objective-C class metadata in pseudo-ObjC syntax, including method implementations and selector addresses, with `-a <address>` and `-m <module>` lookups. `objc list-classes` lists Objective-C classes in the main module, where `-a` includes all modules except the DYLD shared cache and `-A` includes it.

Inspect an Objective-C class during macOS/iOS binary analysis to see method selectors and implementation addresses without needing a separate class-dump tool.

```
$ objc show-class -n NSURLSession
```

— Full command set with flags and runnable example 2026.10.22

02 setrlimit command for debuggee resource limits NEW

81

`setrlimit` sets resource limits on the debuggee — e.g. `setrlimit nofile 1 2` sets soft and hard open-file limits, and `setrlimit core infinite` sets both core-file limits to unlimited.

Remove the core-dump size limit so a crashing debugee actually writes a full core file for post-mortem analysis.

```
$ setrlimit core infinite
```

— Exact command and runnable example given 2026.10.22

03 Opcode byte display in nearpc NEW

78

`nearpc --opcode` / `nearpc -o` shows opcode bytes sized to the longest instruction in the listing, or `nearpc -o 4` for a fixed-width column, without changing the `nearpc-num-opcode-bytes` setting.

— Exact flags and config key named, no usage example 2026.10.22

04 **LLDB launch and targeting options** NEW 70

Adds the LLDB `target create -r` flag for specifying a remote executable path and `process launch -w` for setting the working directory, plus `target.run-args` support so launch now respects pre-configured run arguments.

— Names flags and config key but no example 2026.10.22

05 **SIMD and flags-register decoding in regs** IMPROVED 62

`regs` now displays SIMD registers (e.g. `regs xmm0`) and decodes individual flags-register bits (e.g. `regs eflags`).

— Named examples but no deeper mechanism explained 2026.10.22

06 **Deprecated resolve-heap-via-heuristic config key** DEPRECATED 62

The `resolve-heap-via-heuristic` configuration key is deprecated and has no effect — pwndbg now always tries symbols first and then heuristics regardless of this setting.

— Names exact config key and new fallback behavior 2026.10.22

THINNER COVERAGE BELOW

07 **Android Bionic libc detection** NEW 58

Adds Android Bionic libc detection, surfaced via the `libc` / `libcinfo` command showing libc and loader mappings.

— Names command but no example output 2026.10.22

08 **Per-CPU sheaf and per-node barn info in slab info** IMPROVED 53

`slab info` now displays per-CPU sheaves and per-node barns for Linux 7.0+ SLUB allocator internals.

— Names command and kernel version but no example 2026.10.22

09 **m68k architecture support** NEW 52

Adds m68k architecture support including Capstone disassembly, Unicorn emulation, branch target resolution, and instruction annotations.

— Describes scope but no command or example 2026.10.22

10 **Fallback support for unknown architectures** IMPROVED

38

Pwndbg no longer refuses to start on unknown architectures — it falls back to the underlying debugger's disassembler with reduced context and no emulation.

— Behavioural change described, no named surface 2026.10.22

└──▷ **ALSO FROM THESE RELEASES**

└──▷ **BREAKING ON UPGRADE**

! The `resolve-heap-via-heuristic` configuration key is deprecated and has no effect — pwndbg now always tries symbols first and then heuristics regardless of this setting.

WAS THIS USEFUL? 👍 👎

ntoseye

Feedback ⓘ v0.51.0 NOTES ↗ RANK

WinDbg-like kernel debugger for Windows, from Linux and macOS

> ntoseye's biggest window yet: full Windows hypervisor debugging (partition/VTL inspection, EPT walking, hypercall and VM-exit breakpoints), live breakpoints and stepping inside Hyper-V/Sandbox guests, a wide batch of kernel inspection commands (user-mode minidumps, registry, NDIS, StorPort, virtio, WDF, power, APIC), a native Tern REPL UI, and a ground-up Python SDK overhaul with typed records and a redesigned Stop hierarchy.

└──▷ **WHAT SHIPPED · 34 FEATURES** most completely described first ? what's the number?

01 Hypercall and VM-exit inspection and breakpoints NEW

100

`!hvcalls` lists the hypervisor's hypercall table with TLFS names, rep/variable-header flags, input/output sizes, and handlers (`-a` includes unimplemented codes) (SDK: dbg.hypercalls()). `!hvbp <call> [partition [vp]]` breaks on a hypercall only from a given caller, with a condition evaluated on the caller's registers and memory, e.g. $pqwo(rdx) for slow-call input (SDK: Breakpoints.add_hypercall(), `Breakpoint.hypercall` , Cpu.hypercall_caller()). `!hvexit <reason> [partition [vp]]` stops on a VM exit by reason (`cpuid` , `rdmsr` , `wrmsr` , `ept_violation` , or a number) from a given caller, with a condition on the caller's registers at the exit, e.g. `!hvexit wrmsr if @rcx==0x6e0` (SDK: Breakpoints.add_exit(), `Breakpoint.vm_exit`). `!hvcall` at a hypervisor stop shows the full hypercall input, including XMM fast calls, rep lists, and extended GVA flush ranges with addresses, page counts, and page sizes.

Break on a specific hypercall only when issued from a particular Hyper-V partition and VP, with a condition on the caller's input — useful for tracing guest-to-hypervisor communication without flooding on every hit.

```
$ !hvbp HvCallSignalEvent 4 2 if $pqwo(rdx)
```

Stop execution whenever a guest writes to `IA32_TSC_DEADLINE` via `wrmsr` , scoped to a specific partition and VP — useful for diagnosing timer-related hypervisor exits.

```
$ !hvexit wrmsr 4 2 if @rcx==0x6e0
```

— Named commands with conditions and two runnable examples v0.45.0

02 Hypervisor partition and VTL state inspection NEW

98

`!hvpartitions` shows the hypervisor's partitions as a tree with privileges and `!hvvps` lists a partition's virtual processors and the host processors running them (SDK: dbg.hypervisor_partitions(), `VirtualProcessor.processors`). `!hvr [partition vp [vtl]]` displays the registers of any hypervisor VP, including those no processor currently runs, such as a WSL2 VP blocked in `HLT` (SDK: VirtualProcessor.registers()). `.partition <id>` inspects the Windows guest of a hypervisor partition (e.g. a Windows Sandbox) in place of the target, read-only, using its kernel symbols, with `.partition 1` returning to the target (SDK: Debugger.select_partition(), `Debugger.partition`). `.vtlcxr 1` selects the hypervisor-saved state for VTL1 so `k`, `r`, `u`, and expressions inspect the secure kernel where it left off, with `.vtlcxr` returning to VTL0; the SDK adds `SavedVtlState.general_registers` and `may_be_stale`, marking stale states '(may be one exit behind)'. Hypervisor stack unwinding uses `hvix64.exe` unwind data (importable with `.fetchimage /f <file>` or Symbols.import_image()), with prologue-based fallback marked `[prolog]`, and Cpu.backtrace() gains a `vtl=` parameter. `~` now also accepts vCPU IDs (`~p01.03s`, `~p01.03k`, `~p01.03r`), case-insensitive, in addition to processor numbers.

Inspect a Windows Sandbox guest's kernel in-place — run `lm`, `!process`, `dt`, `u`, and `k` against its kernel and symbols without leaving the current session.

```
$ .partition 4
```

— Deep named surface with mechanism and a runnable example v0.45.0

03 Python SDK typed records and Stop hierarchy overhaul BREAKING

95

`Stop` is now a class hierarchy per kind (`Stop.Reboot` → `Stop.Reload`, attach is `Stop.Entry`), with stop.hit(bp) replacing `bp in stop.breakpoints`. Exceptions.set() now takes (code, *, mode=None, command=None, disposition=None) and returns `ExceptionPolicy`, with new set_module_event('ld'|'ud', module=None, *, mode=None, command=None). `ListEnd` is renamed `ListTermination` with `ListTermination.Cycle` / `ListTermination.Corrupt` subclasses. A stop's `thread` / `process` are now `ThreadInfo` / `ProcessInfo` records and `stop.cpu` is an `int`. dbg.run_to()/dbg.cont() are removed in favor of run(address=...), step(until=address), step_over(until=address), and run(disposition='handled'). Results are frozen records (use attribute access and to_dict() instead of `record['field']`); all options are keyword-only, address args are named `address`, and Memory.read() takes `size` (was `n`). Fields with fixed values are `Literal` snake-case names. Renames: `LoadedModule` → `ModuleInfo`, `ThreadSummary` → `ThreadInfo`,

`MemoryRegion` → `MemoryRegionInfo`; `Cpu.id` → `Cpu.number`, Cpu.info()→Cpu.cpuinfo(); Regions.at()→Regions.get(); `Watchpoint.access` / `.length` → `Breakpoint.watch_access` / `watch_length`; `UserThread` time fields moved into `UserThread.times`; Memory.describe() is refused on user-mode minidumps in favor of UserDump.address_report().

— Extensive named renames and signature changes, but migration is manual `v0.50.0`

04 **Crash-dump blackbox and ETW session decoding** NEW 93

`.enumtag`, `!blackboxbsd`, `!blackboxntfs`, `!blackboxpnp`, `!blackboxwinlogon`, and `!blackboxpci` list and decode bugcheck-callback tagged data in a crash dump; `!wdfkd.wdfcrashdump` shows the KMDF In-Flight Recorder log; `!wmitrace` crash-dump support reads ETW sessions Windows saves in a minidump (e.g. EventLog-System, WiFiSession) and decodes their events. The SDK exposes `version().dump.tagged_blocks`, inspect.read_tagged(), inspect.wdf_crash_log(), inspect.wdf_crash_drivers(), inspect.etw_crash_sessions(), and inspect.etw_crash_events().

In a post-mortem crash dump, list bugcheck tagged blocks and decode Windows blackboxes to triage the stop reason.

```
$ .enumtag; !blackboxbsd; !blackboxpnp; !wdfkd.wdfcrashdump
```

Read ETW sessions saved in a minidump to review EventLog-System events captured at crash time — detail WinDbg cannot provide.

```
$ !wmitrace.logdump EventLog-System
```

— Full command set plus two runnable triage examples `v0.48.0`

05 **Live breakpoints and stepping inside Hyper-V/Sandbox guests**

NEW 93

In a `.partition` view (experimental, `gdb` backend), `ba` sets a hardware breakpoint or data watchpoint scoped to that guest's VPs, with `/c` selecting a VP and `/p`, `/t`, and conditions reading that guest. `g`, `t`, `p`, `gu`, and `g <address>` run, step, step-over, step-out, and run-to-address inside the guest's thread; `eb`, `ed`, and other memory-write commands change the guest's memory, and `r` changes registers of a VP stopped at a breakpoint. The Python SDK adds breakpoints.add(..., hardware=True), step(), step_over(), step_out(), and run_to() after select_partition() for scripted guest debugging, with `Breakpoint.partition` naming the target guest.

— Full command and SDK surface for guest debugging named `v0.46.0`

06 **dx expression and debugger data-model command** IMPROVED

90

`dx` shows typed expressions and fields as WinDbg lays them out, supporting casts with C type names, typed pseudo-registers (`@$proc->UniqueProcessId`), and bitfield layout. It was extended to read WinDbg's debugger data model: `Debugger.Sessions` , `@$cursession` , `@$curprocess` , `@$curthread` , `KernelObject` , `Io.Handles` , LINQ-style queries with lambdas (`Where` , `Select` , `OrderBy` , `SelectMany` , `First` , `Any` , `Count`), `Debugger.Utility.Collections.FromListEntry` , grid output via `dx -g` , and format modifiers `, d` and `, <count>` .

Query the live kernel for lsass.exe using the debugger data model, equivalent to a WinDbg dx expression.

```
$ dx @$cursession.Processes.Where(p => p.Name ==
"lsass.exe").Select(p => new { Name = p.Name, Id = p.Id })
```

— Named query surface with a runnable dx example v0.49.0 v0.48.0

---

07 **User-mode minidump analysis via --dump** NEW

90

`--dump` opens Windows user-mode minidumps (AMD64 and ARM64) from procdump, Task Manager, and Windows Error Reporting, with `lm` , `~` , `.ecxr` , `.exr -1` , `k` , `!analyze` , `!teb` , `!peb` , `!gle` , `!dlls` , `!handle` , `!address` , `!vprot` , `!heap` , `.time` , `vertarget` , `dt` , `x` , `ln` , `dx` , `$teb` , and `$peb` reading captured memory; the SDK exposes results via `dbg.user_dump` and inspect.triage().

— Huge command surface named, no usage example given v0.49.0

---

08 **StorPort adapter, unit, and log inspection** NEW

90

`!storagekd.storadapter` and `!storagekd.storunit` show StorPort adapters and units (viostor, vioscsi, storahci) with state, frozen/paused/busy queue, and each miniport request's IRP and SRB. `!storagekd.storloglist` , `!storagekd.storlogirp` , `!storagekd.storlogsrb` , `!storagekd.storsrb` , and `!storagekd.storclass` add adapter event-log listing, individual request-entry and SRB decoding, and class devices with in-flight requests; the SDK exposes storport_adapters(), storport_unit(), storport_log(), srb(), and class_devices().

— Area extended across two releases with every command named v0.49.0 v0.48.0

---

09 **NDIS miniport, filter, and protocol inspection** NEW

90

`!ndiskd.miniports` , `!ndiskd.miniport` , and `!ndiskd.minidriver` show guest network miniports (e.g. NetKVM) including driver, state, link, adapter context, filter stack, bindings, and pending OID request. `!ndiskd.filter` , `!ndiskd.filterdriver` , `!ndiskd.protocol` , `!ndiskd.oid` , and `!ndiskd.nbl`

(with `-data` for frame bytes) add inspection of filter modules, filter drivers, protocol drivers, pending OID requests, and `NET_BUFFER_LIST`s. The SDK exposes dbg.inspect.ndis_miniports() and other `ndis_*` methods including ndis_nbl_data().

— Area extended across two releases with every command named `v0.49.0` `v0.48.0`

---

10 **Hypervisor memory and VMCS inspection** NEW 90

`!hvept <gpa>` (or `-v <address>` for a root virtual address) walks each VTL's extended page tables to show host address, access, page size, and memory type, while `!hveptdiff` lists guest physical ranges where VTL0 and VTL1 map differently (SDK: HypervisorVtl.translate(), VirtualProcessor.ept_differences()). `!hvvmcs [-msr|-io]` displays every field of a VTL's Enlightened VMCS, the MSRs it intercepts or passes (named, e.g. `IA32_EFER`), and its intercepted I/O ports (SDK: HypervisorVtl.vmcs_fields(), msr_intercepts(), io_intercepts()). `!hvd` reads and `!hvu` disassembles a guest partition's memory (virtual 4-level long-mode or `-p` physical) (SDK: HypervisorVtl.read(), translate_virtual(), disassemble()).

— Named commands and SDK methods, no worked example `v0.45.0`

---

11 **New SDK methods for previously uncovered commands** NEW 85

Adds SDK methods for commands not previously scriptable: dbg.inspect.hang(), decode_bugcheck(), blackbox_bsd() and other blackbox_*() reports, session_processes(), vtop(), set_global_flags(), virtio_driver(), vring_raw(), etw_save(), etw_crash_save(), etw_crash_session(), Memory.pte(), Process.inspect(), decode_peb(), loader_modules(), Thread.decode_teb(), UserDump.address_report(), Breakpoints.list(), and commands(). Every SDK handle (`Process`, `Thread`, `Module`, `Cpu`, `Heap`, `Driver`, `Device`, `Breakpoint`, `Trustlet`, `UserThread`, and the hypervisor handle) now exposes its record as `.info`, each record field as a property, to_dict(), and `==` / `hash` equality scoped to the named object and its target.

Inspect a specific process's PEB and loaded modules from a Python script during a live debug session.

python

```
import ntoseye
dbg = ntoseye.attach()
for proc in dbg.processes:
    if proc.info.name == 'lsass.exe':
        print(proc.decode_peb())
        print(proc.loader_modules())
        break
```

Read a thread's TEB by passing a thread handle as the address argument, using the new expression-or-handle addressing.

```python
import ntoseye
dbg = ntoseye.attach()
for thread in dbg.threads:
    print(thread.decode_teb())
    break
```

Analyze a bugcheck from a crash dump and retrieve the full ETW crash session log via the new SDK methods.

```python
import ntoseye
dbg = ntoseye.open('C:/dumps/MEMORY.DMP')
print(dbg.decode_bugcheck())
print(dbg.etw_crash_session())
```

— Named methods plus runnable Python examples `v0.50.0`

---

12 **Windows registry inspection from kernel memory** NEW

85

`!reg hivelist`, `!reg querykey` (alias `!reg q`), `!reg keyinfo`, and `!reg cellindex` read the Windows registry from hives in kernel memory across all backends and crash dumps, with `-v <value>` for a single value and tab-completion of key paths; the SDK gains inspect.registry_hives(), inspect.registry_key(), inspect.registry_value(), inspect.registry_key_at(), and inspect.registry_cell().

— Named subcommands and SDK methods, no worked example `v0.49.0`

---

13 **Full-screen Tern code/memory browser added then removed**

DEPRECATED

85

v0.47.0 added the `browse` command (also reachable via F2 in Tern), a full-screen code and memory browser paging through disassembly and memory as hex, colored bytes, an inspector, or pointers with symbols, with `g` to navigate to an expression, `/` to find text/bytes/a pointer, and `b` to set or clear breakpoints in code or write watchpoints in memory. v0.48.0 added stepping inside it with WinDbg keys (F10, F11, Shift+F11, F5, F7 to run to cursor, Esc to break), changed-register tinting, caller-frame navigation with `[` / `]`, mixed source/disassembly mode with a private PDB, and a memory-as-type view where `t` renders `nt!_EPROCESS` or any PDB type in `dt` columns with nested structures, pointer following, and list-link walking, with `e` writing values down to individual bits of a flags word. v0.49.0 removed the `browse` command and its F2 key entirely, directing users to `u`, `db` / `dq`, `dt`, `k`, and `r` at the prompt, or the disassembly and memory views of an editor over the DAP adapter.

Open the full-screen memory browser to inspect raw kernel memory at a known address and set a write watchpoint, without leaving the REPL.

```
$ browse
```

— Full lifecycle traced across three releases with exact keys v0.49.0 v0.48.0 v0.47.0

14 **Recursive and prepended symbol path loading** IMPROVED

83

`--sympath-append <directory>` prepends local PDBs (plain files or symbol-store layout) to the symbol path at startup. Both `--sympath-append` and `.sympath+` now scan subdirectories recursively (eight levels deep), so a mounted virtio-win ISO's driver PDBs (e.g. `viostor/w11/amd64/viostor.pdb`) load without copying them out; `.reload` re-scans the directory for PDBs added since the session started.

— Exact flags and concrete scan depth given v0.51.0 v0.48.0

15 **Module unload exception events (sxe/sxn/sxd/sxi ud)** BREAKING

83

`sxe` / `sxn` / `sxd` / `sxi` `ud[:<module>]` stop on, report, or ignore kernel module unloads (mirroring `ld` for loads), with `-c` to run commands at the stop; DAP reports the stop reason `module unload`. In the SDK, `Exceptions.module_loads` is replaced by `Exceptions.module_events` and the stop type is `Stop.ModuleUnload`; existing code referencing `module_loads` must be updated.

— New stop type plus a required SDK rename v0.45.0

16 **KMDF request, I/O target, and object inspection** NEW

78

`!wdfkd.wdfrequest`, `!wdfkd.wdfiotarget`, and `!wdfkd.wdfobject` inspect KMDF requests (IRP, queue, cancel callback), I/O targets (state, device, pending requests), and arbitrary KMDF objects by handle or `FxObject` address; the SDK exposes inspect.wdf_request(), inspect.wdf_io_target(), and inspect.wdf_object().

— Named commands and fields covering three KMDF surfaces v0.49.0

17 **Local APIC and I/O APIC inspection** NEW 78

`!apic [processor]` shows a processor's local APIC (ID, priorities, LVT, timer, ICR, vectors in service and requested) in xAPIC or x2APIC mode over KD and in kernel dumps, and `!ioapic` shows each I/O APIC and PIC line with vector, IRQL, trigger, polarity, and target; the SDK exposes cpu.apic() and interrupt_controllers().

— Named commands with field-level detail v0.49.0

18 **Virtio guest device and queue inspection** NEW 75

`!virtio` lists guest virtio devices and queue states (held, unclaimed, moved) and names a stuck queue; `!vring` shows a queue's buffers and decodes SCSI, network, or GPU requests with device answers. The SDK exposes inspect.virtio() and inspect.vring().

— Named commands and behavior, no worked example v0.48.0

19 **Guest DbgPrint output streaming** NEW 75

Guest debug output (`DbgPrint`) is supported on `gdb` and `memory` backends: the REPL prints it live, MCP results carry it as `[dbgprint]` , and DAP shows it in its console; `!dbgprint` dumps the full kernel buffer on all backends including crash dumps.

— Named surfaces across REPL, MCP, DAP and a crash-dump command v0.48.0

20 **WPP recorder in-flight log decoding** NEW 73

`!rcdrkd.rcdrloglist` and `!rcdrkd.rcdrlogdump` read WPP recorder in-flight logs for drivers such as usbxhci and pci, formatted from TMF annotations in loaded PDBs; the SDK exposes rcdr_drivers(), rcdr_logs(), and rcdr_log().

— Named commands with concrete driver examples v0.49.0

21 **display and undisplay expression watch command** NEW 73

`display <expression>` shows an expression's value at every stop, marked when it changed; `display` alone lists active displays and `undisplay` removes them.

Track a key variable across every breakpoint stop without re-typing it — useful when stepping through a driver routine to watch a pointer change.

```
$ display nt!PsInitialSystemProcess
display
undisplay
```

— Named command with a runnable example v0.47.0

22 **REPL and session behavior improvements** IMPROVED 70

Ctrl+C during the SDK's run() or wait() now breaks in and leaves the target halted instead of leaving it running. `--plain-repl` scripts now continue after a command fails, reporting the error, matching interactive REPL behavior. `.thread .` selects the current thread. `ib`, `iw`, and `id` are refused on a dump and while the target runs, since reading an I/O port can change the device. The REPL now refuses unrecognized words a command does not take (e.g. `vertarget x`, `k 3 extra`) instead of ignoring them.

— Several named commands/flags, each thin on its own v0.50.0

---

23 **System and device power state inspection** NEW 70

`!poaction`, `!podev`, and `!popolicy` show the system power action in progress, per-device power state and IRPs, and the system power policy; the SDK exposes inspect.power_action() and inspect.device_power().

— Named commands, thin on mechanism beyond that v0.49.0

---

24 **Structured notices for exceptions and session events** BREAKING 65

An exception passing a notify filter (`sxn`) is now reported as a notice on every host: in the REPL, via dbg.notices() in the SDK, and over MCP and DAP, and its `-c` command runs. Debugger.notices() now returns Notice records with a `level` field (`"info"` or `"warning"`) and a `text` field, instead of plain strings.

— Behavior and return-type change named, no code example v0.50.0 v0.48.0

---

25 **Tern prompt completion and navigation improvements** IMPROVED 65

The Tern prompt gains Tab completions narrowing in a popup, live line coloring as you type, Ctrl+Z undo, paste-as-one-command, Alt+P palette searching commands/recent lines/symbols/types/processes, and Ctrl+R opening the palette on history.

— Several UX additions named, no mechanism beyond description v0.47.0

---

26 **lm a shows module containing an address** NEW 65

`lm a <address>` shows only the module containing a given address, resolving kernel modules for kernel addresses even while a process is selected (matching WinDbg behaviour).

Find which kernel module owns an address while a user-mode process is selected, without switching context.

```
$ lm a fffff804`12340000
```

— Named flag with a runnable example v0.46.0

27 **Structured operand fields on DisassembledInstruction** NEW

62

`DisassembledInstruction` gains `length`, `mnemonic`, and `operands` fields, and a structured `Operand` type describing register, memory, immediate, or branch operands, so scripts no longer need to parse `asm` strings.

— Named fields and type, no worked example v0.45.0

---

28 **MCP output tool for paging long results** NEW

60

The MCP `output` tool pages or searches results longer than ~24,000 characters (e.g. `x nt!*`) without re-running the command.

— Named tool and threshold, minimal further detail v0.48.0

---

29 **TERN_TSP environment variable disables native rendering** NEW

60

Setting `TERN_TSP=0` disables native Tern rendering and keeps plain text output; Tern versions too old to draw the views fall back to text with a warning.

— Exact env var and fallback behavior named v0.47.0

THINNER COVERAGE BELOW

30 **Changed-register highlighting on repeated stops** NEW

50

After a stop on the same thread as the previous stop, changed registers are highlighted in yellow in the register grid; in Tern the folded register summary names them (e.g. `registers rsp rip changed`) and shows their new values in amber.

— Visual behavior described, nothing to run v0.47.0

---

31 **Registers.get() with default fallback** NEW

50

The SDK gains Registers.get(name, default=None) to return a register's value or `default` when the register file (e.g. an unwound frame) does not hold it.

— Single named method, minimal further detail v0.45.0

---

32 **Driver source tree mapping via --srcpath** NEW

45

`--srcpath <directory>` maps a driver's source tree from session start (repeatable), equivalent to `.srcpath+`.

— Named flag but one line, no mechanism detail v0.49.0

---

33 **PDB integer field typing matches C types** BREAKING

45

In the SDK, `dt`, and DAP variables, a PDB field declared `int` or `unsigned int` now has type `INT` or `UINT` instead of `LONG` or `ULONG`.

— Named type change but no migration guidance `v0.48.0`

---

34 **MCP command tool drops format argument** `BREAKING` 45

MCP's `command` tool no longer accepts a `format` argument and rejects calls that pass it; callers expecting typed results must switch to the Python SDK.

— Clear removal but brief migration note only `v0.45.0`

└──▷ BREAKING ON UPGRADE

! `Stop` is now a class hierarchy with a subclass per kind; `Stop.Reboot` is renamed to `Stop.Reload`; stop.hit(bp) replaces `bp in stop.breakpoints`; a stop found at attach is `Stop.Entry` instead of `Stop.Interrupt`.

! Exceptions.set() signature changed: it now takes (code, *, `mode=None`, `command=None`, `disposition=None`) and returns `ExceptionPolicy`; set('ld') and set('ud:<module>') raise `ValueError`. New set_module_event('ld' | 'ud', module=None, *, mode=None, command=None) replaces the old module-event overloads.

! `ListEnd` is renamed to `ListTermination`, now a subclass hierarchy: `ListTermination.Cycle` (with `address`), `ListTermination.Corrupt` (with `error`), each with `diagnostic`.

! A stop's `thread` and `process` are now `ThreadInfo` and `ProcessInfo` records; `stop.cpu` is the processor number (an `int`), not a handle.

! dbg.run_to() and dbg.cont() are removed; use run(address=...), step(until=address), step_over(until=address), or run(disposition='handled') instead. step() and step_over() take `timeout` first and `until=` by keyword.

! Results are now frozen value records, no longer mappings; use attribute access instead of `record['field']`, and to_dict() instead of keys(), items(), len(), or iteration. The `Record` class is removed; name-keyed maps are plain `dict`s.

! All SDK method options are now keyword-only; address arguments are named `address` everywhere (was `addr` on `Memory` and Type.at()); Memory.read() takes `size` (was `n`).

! Fields and arguments with a fixed set of values are now typed as `Literal` of snake-case names (e.g. `'kernel_module'`, `'ntstatus'`, `'no_pdb'`, `'pde'`, `'nonpaged'`).

! Many SDK classes are renamed: `LoadedModule` → `ModuleInfo`, `ThreadSummary` → `ThreadInfo`, `MemoryRegion` → `MemoryRegionInfo` (and `MemoryRegion` is now a user-mode minidump region); several fields moved.

! Collections now use uniform iteration as typed `Iterator[...]`; the `ProcessIterator`-style classes are removed. Regions.at(address) is now Regions.get(address); Threads.at(address) returns `None` instead of raising when no thread has that `_ETHREAD`.

! dbg.breakpoints.watch() returns a `Breakpoint`; `Watchpoint.access` and `.length` are now `Breakpoint.watch_access` and `Breakpoint.watch_length`; `Breakpoint.address` is `None` for unresolved breakpoints.

! `Cpu.id` is replaced by `Cpu.number`; Cpu.info() (`!cpuinfo`) is renamed to Cpu.cpuinfo(); `Cpu.info` is now the `VcpuInfo` record. `VcpuInfo` has `number` instead of `id`. `dbg.cpus` is a plain sequence indexed by `int`.

! `UserThread`'s thread-info stream fields (`start_address`, `create_time`, `exit_time`, `kernel_time`, `user_time`, `exit_status`, `affinity`) are now in `UserThread.times` (a `UserThreadTimes` or `None`); `UserThread.stack` is now a `CapturedStack` instead of a tuple; identity is `tid` (the `id` text is removed).

! Memory.describe(address) is refused on a user-mode minidump; use UserDump.address_report() instead. Inspect.global_flags() now reads flags; set_global_flags(change) sets them.

! The REPL now refuses unrecognized words a command does not take (e.g. `vertarget x`, `k 3 extra`) instead of ignoring them.

! The `browse` command and its F2 key in Tern are removed; use `u`, `db`/`dq`, `dt`, `k`, and `r` at the prompt, or the disassembly and memory views of an editor over the DAP adapter.

! In the SDK, Debugger.notices() now returns `Notice` records with a `level` field (`"info"` or `"warning"`) and a `text` field, instead of plain strings.

! In the SDK, `dt`, and DAP variables, a PDB field declared `int` or `unsigned int` now has type `INT` or `UINT` instead of `LONG` or `ULONG`.

! `Exceptions.module_loads` is replaced by `Exceptions.module_events`; any SDK code referencing `module_loads` must be updated.

! MCP's `command` tool no longer accepts a `format` argument and rejects calls that pass it; callers expecting typed results must switch to the Python SDK.

WAS THIS USEFUL? 👍 👎

ipatool

Feedback ⓘ v2.7.0 NOTES ↗ RANK

Command-line tool that allows you to search for iOS, iPadOS, tvOS, watchOS, visionOS, and macOS apps on the App Store, and download .ipa or macOS .pkg app packages.

> ipatool v2.7.0 adds an MCP server mode for AI agent integration, phone-number sign-in for Apple Accounts in China and India, and expands app downloads to delisted tvOS apps, standalone watchOS apps, and legacy macOS installer packages.

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

01 MCP server subcommand for AI agents NEW

88

Adds the `mcp` subcommand to serve App Store tools (search, download, list-purchases, list-versions, get-version-metadata, purchase) over MCP via stdio, allowing AI agents such as Claude Desktop to search and download App Store apps on ipatool's behalf.

Wire ipatool as an MCP stdio server so an AI agent (e.g. Claude Desktop) can search and download App Store apps on your behalf.

json

```
{
  "mcpServers": {
    "ipatool": {
      "command": "ipatool",
      "args": ["mcp"]
    }
  }
}
```

Start the MCP stdio server manually to verify it launches correctly before configuring a client.

```
$ ipatool auth login && ipatool mcp
```

— Names subcommand, tools exposed, and runnable config/command examples v2.7.0

THINNER COVERAGE BELOW

02 Expanded download coverage for tvOS, watchOS, macOS NEW

50

Adds support for downloading delisted tvOS apps by numeric app ID, downloading standalone watchOS apps, and retrieving legacy installer packages (`.pkg`) for old

macOS app versions.

— Names three download targets but no commands or flags given v2.7.0

03 **Phone-number sign-in for Apple Accounts** NEW 35

Adds phone-number sign-in support for Apple Accounts in mainland China and India.

— States what changed but no flag or command shown v2.7.0

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

kubeshark

Feedback ⓘ v53.6.0 NOTES ↗ RANK

eBPF-powered network observability for Kubernetes. Indexes L4/L7 traffic with full K8s context, decrypts TLS without keys.

> Kubeshark v53.6.0 adds private-CA support for S3-compatible snapshot storage, a new `trace` log level across the Hub, worker, and dissector, detailed step-by-step authorization diagnostics for REST/MCP/Connect-RPC and OIDC, and fixes the MCP `list_api_calls` endpoint to respect snapshot filters.

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

01 **Private-CA S3 snapshot connectivity** NEW 90

Adds `tap.snapshots.cloud.s3.insecureSkipVerify` config key in `~/.kubeshark/config.yaml` to let the Hub connect to S3-compatible endpoints signed by a private CA, enabling cloud snapshots that previously failed the startup bucket check.

Enable cloud snapshots against an internal S3-compatible store whose certificate is signed by a private CA.

yaml

```yaml
tap:
  snapshots:
    cloud:
      s3:
        insecureSkipVerify: true
```

— Names exact config key and gives runnable YAML example. v53.6.0

02 **Trace log level across Hub, worker, dissector** NEW 85

Adds `logLevel: trace` support in the Hub, worker, and dissector; the Hub now propagates its log level to spawned dissection jobs, which previously ran at their own default log level.

Turn on maximum verbosity across the Hub, worker, and dissector to diagnose OIDC auth failures or snapshot errors without a debugger.

yaml

```yaml
logLevel: trace
```

— Exact config key with example and clear before/after behavior. v53.6.0

03 **Step-by-step authorization diagnostics** IMPROVED 75

Every REST, MCP, and Connect-RPC denial now logs the route, required capability, resolved role, and effective capability set; the OIDC pipeline logs discovery, login, token verification, and role resolution; snapshot creation records request details and per-worker results, distinguishing three previously identical 'no PCAP files found' causes — all logged at `trace` with actionable outcomes surfaced at `warn`.

— Detailed mechanism but no standalone command beyond enabling trace. v53.6.0

04 **MCP list_api_calls snapshot filtering fix** IMPROVED 70

Improves MCP `list_api_calls` so that `snapshot_id` and `dissection_id` are honoured on the GET endpoint, preventing snapshot queries from falling through to live worker stores.

— Names exact endpoint and fields but lacks a usage example. v53.6.0

WAS THIS USEFUL? 👍 👎

ecapture

Feedback ⓘ v2.8.2 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 v2.8.2 adds Android executables to its release archives, making direct downloads available for Android devices.

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

01 **Android binaries in release archives** NEW 45

Release archives now include Android executable binaries for x86_64 and aarch64, enabling direct download and use on Android devices without building from source.

— Names platforms and archive change but no exact filenames or commands v2.8.2

NEW SUPPORTED PLATFORM Linux/LoongArch64 platform support

> v2.8.0 NOTES ↗ STABLE
>
> Adds `text`, `keylog`, and `pcapng` capture modes to the `gnutls` module, matching OpenSSL module parity.
>
> └──▷ **GET THIS VERSION**
>
> ```
> $ git clone --branch v2.8.0 https://github.com/gojue/ecapture.git
> # already have the repo? check out this version:
> $ git checkout v2.8.0
> ```

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

CISA ScubaGear

Feedback ⓘ v2.0.0 NOTES ↗ RANK

Automation to assess the state of your M365 tenant against CISA's baselines

> ScubaGear v2.0.0 replaces the retired Microsoft Defender baseline with a new Security Suite baseline, adds a risky-AI-agent-blocking policy for Entra ID, a `Invoke-SCuBADiff` comparison command, and severity scoring and CSV export for Risky Apps, alongside new SharePoint and Entra ID policies and reporting improvements.

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

01 Security Suite baseline replaces Defender baseline

BREAKING

83

Adds a new `MS.SECURITYSUITE` baseline (Groups 1–8) covering advanced threat protection — including email attachment blocking, post-delivery scanning, malware quarantine, DLP for PII, unified audit logging, alert enablement, and anti-spam/anti-phishing controls — replacing the retired Microsoft Defender baseline (`MS.DEFENDER.1.1v1` , `MS.DEFENDER.1.2v1` , `MS.DEFENDER.1.3v1` , `MS.DEFENDER.1.4v1` , `MS.DEFENDER.1.5v1` , `MS.DEFENDER.4.1v2` , `MS.DEFENDER.5.1v1` , `MS.DEFENDER.6.1v1`). Any configuration or automation referencing the old policy IDs will no longer find them.

Run a full assessment with the new Security Suite baseline included, producing HTML, JSON, and CSV outputs for BOD 25-01 submission.

```
$ Invoke-SCuBA -ConfigFilePath 'C:\scuba\config.yaml' -
Organization 'example.onmicrosoft.com' -ProductNames *
```

— Names all retired and replacement policy groups, no migration command given. v2.0.0

02 New baseline policy checks across platforms

NEW

75

Adds new baseline policies `MS.AAD.9.1v1` (requiring risky AI agents to be blocked in Entra ID, needing Entra ID P2 plus Microsoft Agent 365 or M365 E7, and not applicable to GCC High/DoD tenants), `MS.AAD.5.5v1` , `MS.AAD.5.6v1` , and `MS.AAD.5.7v1` for application management, `MS.SHAREPOINT.1.3v1` , `MS.SHAREPOINT.3.1v1` , `MS.SHAREPOINT.3.2v1` , and updated `MS.SHAREPOINT.3.3v2` for external sharing domain restrictions, 30-day Anyone link expiration, view-only link permissions, and reauthentication limits, plus new Power BI policy checks.

— Names policy ID and licensing prerequisites, no run command shown. Names all four policy IDs and what each enforces. Names policy IDs but not what each one checks. Single line with

no policy IDs or mechanism named.

v2.0.0

03 **Invoke-SCuBADiff comparison command** NEW 75

Adds the `Invoke-SCuBADiff` command to diff two `ScubaResults.json` runs, enabling change-over-time comparison of assessment outputs.

— Names exact command but no usage example or output shown. v2.0.0

04 **Unconditional MFA evaluation for MS.AAD.3.2v2** BREAKING 65

`MS.AAD.3.2v2` now evaluates independently of `MS.AAD.3.1v1` — tenants that previously relied on the conditional 'if phishing-resistant MFA has not been enforced' logic will now be evaluated for MFA enforcement unconditionally.

— Full before/after logic with policy IDs, no remediation step given. v2.0.0

05 **SPF soft fail accepted by MS.EXO.2.2v3** BREAKING 60

`MS.EXO.2.2v3` now accepts SPF soft fail (`~all`) as compliant; configurations that enforced hard fail only will pass where they may have previously been used to detect deviations.

— Names policy ID and exact SPF syntax change affected. v2.0.0

THINNER COVERAGE BELOW

06 **CSV output and severity scoring for Risky Apps** IMPROVED 45

Adds a CSV output file and severity scoring to the Risky Apps module, providing spreadsheet-compatible data and risk-based prioritization for remediation.

— Describes format and scoring but no field names or thresholds. v2.0.0

07 **HTML report identification improvements** IMPROVED 40

Adds tenant domain name to product-specific HTML reports, a Privileged User Table, and new indicators across all HTML reports for easier identification.

— Lists additions but no field names or UI path. v2.0.0

08 **Near-miss detection for Entra ID exclusions** NEW 40

Adds near-miss detection for Entra ID (AAD) exclusions, surfacing conditional access policy exclusions that almost match risky patterns.

— Explains concept but no policy IDs or worked example. v2.0.0

09 **Automatic policy pass configuration analysis** NEW 20

Adds configuration analysis for automatic policy pass detection, reducing manual review burden.

— States benefit but not which policies or how it works. v2.0.0

└──▷ BREAKING ON UPGRADE

! The Microsoft Defender baseline (`MS.DEFENDER.1.1v1` , `MS.DEFENDER.1.2v1` , `MS.DEFENDER.1.3v1` , `MS.DEFENDER.1.4v1` , `MS.DEFENDER.1.5v1` , `MS.DEFENDER.4.1v2` , `MS.DEFENDER.5.1v1` , `MS.DEFENDER.6.1v1`) has been retired and replaced by the new `MS.SECURITYSUITE` baseline; any configuration or automation referencing these policy IDs will no longer find them.

! `MS.AAD.3.2v2` now evaluates independently of `MS.AAD.3.1v1` — tenants that previously relied on the conditional 'if phishing-resistant MFA has not been enforced' logic will now be evaluated for MFA enforcement unconditionally.

! `MS.EXO.2.2v3` now accepts SPF soft fail (`~all`) as compliant; configurations that enforced hard fail only will pass where they may have previously been used to detect deviations.

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ Secure SaaS - October 22, 2026 BLOG ↗

Sysdig Secure is a cloud security platform that provides runtime threat detection, vulnerability management, and compliance monitoring for containers, Kubernetes, and cloud workloads.

CLI Scanner extends risk acceptance to standalone binaries and non-package components; Registry Scanner adds standalone binary scanning and license detection (Technical Preview).

Extends risk acceptance to every component type in the CLI Scanner, so accepted risks now apply to vulnerabilities in standalone binaries and other non-package components.

NEW SUPPORTED PLATFORM Adds standalone binary scanning to the Registry Scanner, enabling detection of applications not installed by a package manager in distroless and scratch images.

NEW DETECTION Adds license detection for packages in Registry Scanner scanned images, available in Technical Preview.

└──▷ BREAKING ON UPGRADE

! Registry Scanner v0.13.0: images that previously showed no findings (such as distroless and scratch images) may show new findings after upgrading due to standalone binary scanning being enabled.

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

cicd-sensor

Feedback ⓘ releases/v0.0.48 NOTES ↗

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

Agent now emits `job_event_channel_closed`, `job_drain_in_progress`, and `job_drain_completed` log records at job finalize to expose drain status.

└──▷ **GET THIS VERSION**

```
$ git clone --branch releases/v0.0.48 https://github.com/cicd-sensor/cicd-sensor.git
# already have the repo? check out this version:
$ git checkout releases/v0.0.48
```

NEW DETECTION @subql/common npm compromise IOCs

NEW DETECTION tensorlake npm compromise IOCs

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks

Feedback ⓘ

Box Audit Logs managed ingestion connector (Beta) (2026-10-23) BLOG ↗ RANK

Databricks is a cloud-based data lakehouse platform for data engineering, analytics, machine learning, and AI workloads, built on Apache Spark.

> Databricks brought Genie Code tasks in Lakeflow Jobs to general availability (with default-on rollout for compliance-profile workspaces coming in November 2026) and promoted the Lakeflow pipelines system tables to GA for account-wide pipeline auditing.

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

01 **Lakeflow pipelines system tables GA** NEW 60

The `lakeflow` system tables `pipelines` and `pipeline_update_timeline` are now generally available, recording pipeline definitions and update history across the account for audit purposes.

– Names exact table names but no query example given.

Lakeflow pipelines system tables are no…

THINNER COVERAGE BELOW

02 **Genie Code tasks in Lakeflow Jobs** NEW 55

A new Genie Code task type in Lakeflow Jobs runs a prompt autonomously as part of a job pipeline, reads upstream task outputs, calls tools as needed, and returns a link to the resulting Genie Code chat. This task type is now generally available, and will be enabled by default for workspaces with the compliance security profile enabled starting in November 2026.

– Describes behavior and rollout timing but no setup steps or config shown.

Genie Code tasks in Lakeflow Jobs are n… Genie Code tasks in Lakeflow Jobs are a…

NEW INTEGRATION Box Audit Logs managed ingestion connector (Beta)

NEW SUPPORTED PLATFORM DataFrame caching on serverless compute

└──▷ **ALSO FROM THESE RELEASES**

Box Audit Logs managed ingestion connector (Beta) (2026-10-23) BLOG ↗

STABLE

DataFrame caching on serverless compute is now in Public Preview (2026-10-23)

BLOG ↗ LATEST STABLE

DataFrame caching via df.cache() and df.persist() is now in Public Preview on serverless compute, enabled by default for most workspaces.

Enables df.cache() and df.persist() on serverless compute to reuse expensive DataFrame results within a session.

Workspace admins can disable the preview from the Previews page.

WAS THIS USEFUL?  

◆ Security Data Pipeline

Cribl

Feedback ⓘ changelog · Oct 22, 2026 BLOG ↗ RANK

Cribl is the AI Platform for Telemetry, providing open, flexible infrastructure that helps enterprises manage, connect, and analyze their data across IT, security, and observability environments.

> Cribl shipped network-isolation controls for Azure Storage Locations in Cribl Lake, letting both management calls and Search query traffic stay inside the customer's Azure network.

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

01 **Outpost support for Azure Storage Locations** NEW 87

Cribl Outpost support routes all Azure management calls (container creation, lifecycle and retention policies, validation) for an Azure Storage Location through an Outpost deployed inside your Azure vNet, rather than from the Cribl control plane. It is opt-in per storage location and requires no migration for existing Azure Storage Locations.

Bind an Azure Storage Location to an Outpost Group so all Azure management calls run from inside your vNet — required for regulated environments where public control-plane access to storage is blocked.

📍 In Cribl.Cloud, go to Lake › Storage Locations › [your Azure Storage Location] › Edit. Under Outpost, select your deployed Outpost Group and supply the storage account's Azure Resource ID, then save. Cribl validates the path end-to-end and rolls back automatically on misconfiguration.

— Exact UI path and config steps given, but no API/CLI surface 2026-10-22-43392606

02 **Private Endpoint support for Cribl Search on Azure** NEW 86

Azure Private Endpoint support lets Cribl Search query traffic to Azure Blob Storage travel over the Microsoft backbone instead of the public internet, via a connection you approve in the Azure portal and can revoke at any time. Like Outpost support, it is opt-in per storage location with no migration required for existing Azure Storage Locations.

Enable Private Endpoint support on an Azure Storage Location so Cribl Search query traffic never leaves the Microsoft backbone — satisfying 'no public internet on the read path' audit requirements.

📍 In Cribl.Cloud, go to Lake › Storage Locations › [your Azure Storage Location] › Edit. Under Private Endpoint, enable the option and save. Then, in the Azure portal,

> locate the pending private endpoint connection request on your storage account and approve it. No bytes flow until approval.

— Clear UI path and approval flow, no numeric limits given `2026-10-22-43392606`

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Helm Charts v0.2.0: Support slow node startups (2026-10-22) BLOG ↗

RANK

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

> Tenzir shipped Helm chart reliability fixes, organization-level machine credentials (service tokens and federated OIDC) for the Platform, and a Skills v4.0.0 update adding Codex/Claude Code plugin support and a restructured OCSF skill, while removing package management and developer skills from the public collection.

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

01 **Machine credentials for organizations** NEW 68

Adds organization service tokens to cover automation use cases such as the Terraform provider, with immediate invalidation on revocation, and adds federated credentials for workload identity, letting workloads such as GitHub Actions authenticate using their own OIDC tokens instead of a stored secret.

— Names mechanisms (Terraform, GitHub Actions, OIDC) but no exact API/flag

Tenzir Platform v1.41.0: Machine creden…

THINNER COVERAGE BELOW

02 **OCSF skill rebuilt on structured schema data** IMPROVED 54

Rebuilds the OCSF skill on structured YAML schema data sourced from the official OCSF compiler, replacing the previous schema representation, and adds OCSF release notes surfacing inside the OCSF skill.

— Explains source and change but no concrete command to try it

Tenzir Skills v4.0.0: Plugin support an…

03 Startup probe and version labels in Helm charts IMPROVED

53

Adds an optional startup probe to give nodes with large partition catalogs time to initialize before Kubernetes liveness checks begin, and Kubernetes version labels now reflect the configured image tag.

— Describes behaviour but no exact config key or flag given

Tenzir Helm Charts v0.2.0: Support slow…

04 Codex and Claude Code plugin support in Tenzir Skills NEW

40

Adds Codex plugin support, enabling Tenzir Skills to run inside the Codex environment, and adds Claude Code plugin support, enabling Tenzir Skills to run inside Claude Code.

— Names two integrations but no setup detail Tenzir Skills v4.0.0: Plugin support an…

05 Package management and developer skills removed BREAKING

25

The package management skill has been removed from the public collection, and the developer skills have been removed from the public collection.

— States removal only, no migration guidance Tenzir Skills v4.0.0: Plugin support an…

└──▷ **BREAKING ON UPGRADE**

! The package management skill has been removed from the public collection.

! The developer skills have been removed from the public collection.

WAS THIS USEFUL?  

Nebulock Agentic Threat Hunting Framework

Feedback ⓘ v0.23.0 NOTES ↗ RANK

ATHF is a framework for agentic threat hunting - building systems that can remember, learn, and act with increasing autonomy.

> ATHF's biggest change this window is a five-verdict ladder with evidence and provenance gates enforced directly in hunt findings, paired with a new `/gates` skill that scores findings for promotion to detection engineering; the release window also adds a hunt_type category and raises the minimum Python version to 3.11.

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

01 Five-verdict ladder with evidence and provenance gates NEW

90

Hunt `findings` and `ruled_out` entries now carry a `verdict` field with five explicit values: `confirmed`, `suspected`, `attempted_not_vulnerable`, `benign`, `inconclusive`. An evidence gate requires `verdict: confirmed` entries to carry a `confirmation` mapping with a supported `method`, a `produced_by` producer declared in `.athfconfig.yaml` whose capabilities include that method and reach beyond query-only access, an `attested_by` person distinct from the producer, and a non-empty `detail`; telemetry-only or query-only producers cap an entry at `suspected`. A routing rule enforces that `attempted_not_vulnerable`, `benign`, and `inconclusive` verdicts live only in `ruled_out`, and that `attempted_not_vulnerable` entries name the control that held.

— Names every field, value and config key, with full gate mechanism v0.23.0

02 /gates skill for hunt-to-detection promotion scoring NEW

90

Adds a `/gates` skill with methodology at `.claude/skills/gates/SKILL.md`, output schema at `.claude/skills/gates/AGENT_MEMORY_SCHEMA.md`, and worked examples under `.claude/skills/gates/examples/` for every verdict type. GATES scoring produces five verdicts — PROMOTE, CONDITIONAL, HOLD, TIME-BOX, RECURRING HUNT — evaluated across five dimensions (Generalizability, Accuracy, Telemetry, Exposure-tested, Sustainability); PROMOTE and prerequisite-complete CONDITIONAL verdicts emit YAML for ADEF to ingest, bridging hunt findings to production detection engineering.

— Full file paths, verdict set, dimensions and output format given v0.22.0

03 workspace_root parameter in hunt validation functions

IMPROVED 70

Adds a `workspace_root` parameter to validate_hunt_file(), hunt_file_warnings(), and `HuntParser` so producer-registry lookups use the workspace root, closing a validation/aggregation divergence in nested `hunts/` directories.

— Names exact functions and the bug it fixes v0.23.0

---

04 **Verdict tallies in hunt stats, metrics, and MCP validation**

IMPROVED 65

Verdict tallies and precision pairs are now surfaced in `athf hunt stats` and `athf metrics`, and the MCP `athf_hunt_validate` tool now returns warnings alongside errors.

— Names the commands and MCP tool but no example output v0.23.0

THINNER COVERAGE BELOW

05 **Minimum Python version raised to 3.11** BREAKING 50

The framework's minimum supported Python version is raised to 3.11; installations on Python 3.10 and below will no longer work.

— Clear before/after but no migration steps given v0.21.0

---

06 **hunt_type category with stats, list, validate, metrics support**

NEW 35

Adds a hunt_type category to the framework, with support across `stats`, `list`, `validate`, and `metrics`.

— Named but no commands, fields or examples given v0.21.0

---

07 **Hypothesis-generator runnable from Next Steps output** IMPROVED 30

`hypothesis-generator` can now be run directly from the printed Next Steps output.

— Single sentence with no further mechanism v0.21.0

└──▷ **BREAKING ON UPGRADE**

! Python minimum version raised to 3.11 (was lower); installations on Python 3.10 and below will no longer work.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Patch Management

Ivanti Neurons for Patch Management

Feedback ⓘ

Version 2026.4 October | NOTES ↗ | RANK

Ivanti neurons for patch management is a prioritize and patch vulnerabilities based on active risk exposure, patch reliability and device compliance with Ivanti's cloud-based patch management solution.

> Ivanti Neurons for Patch Management adds an AI-powered summary tile to help prioritize patch deployment issues without manual review.

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

01 **AI Risk Summary tile on Deployment History** NEW 55

A new AI Risk Summary tile on the Deployment History page uses AI to summarize and prioritize patch deployment status, surfacing successful deployments, issues needing attention, and urgent conditions without manual data-grid review. This feature is in Technical Preview.

— Names exact UI tile and location but no deeper mechanism detail. Version 2026.4 October

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

syzkaller

Feedback ⓘ changes since 2026-10-22 CODE ↗

syzkaller is an unsupervised coverage-guided kernel fuzzer

syzkaller adds `LLMAdvisor` and `AllowPartialResults` to its aflow agent framework, plus a new `linux-subsystems-update` agent skill and a refactored security assessment flow.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-10-22
https://github.com/google/syzkaller.git
# already have the repo? check out this version:
$ git checkout commits-2026-10-22
```

Adds `LLMAdvisor` to the aflow framework, usable via the `Advisor` field on `LLMAgent`, to provide iterative evaluation guidance during agent runs.

Adds `AllowPartialResults` field to `LLMAgent` so agents can return partial outputs when iteration limits are reached rather than failing entirely.

Adds `linux-subsystems-update` agent skill across the aflow system.

Splits the aflow security assessment flow into focused, discrete steps.

UPDATED SUPPORTED PLATFORM Linux subsystem extraction rules updated for v7.3

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

Tenable Nessus

Feedback ⓘ Tenable Nessus10.12.6 (2026-10-22)

NOTES ↗ RANK

Tenable Nessus is a vulnerability scanner that identifies security weaknesses, misconfigurations, and missing patches across networks, servers, and applications.

> Nessus 10.12.6 lifts the Amazon Linux restriction on `--aws-scanner` linking and drops support for several legacy operating systems.

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

01 **AWS scanner linking beyond Amazon Linux** IMPROVED 70

Linking a Tenable One Vulnerability Management-managed scanner using the `--aws-scanner` flag from an AWS-hosted scanner is no longer restricted to Amazon Linux.

— Names exact flag but no further mechanism detail Tenable Nessus10.12.6 (2026-10-22)

02 **Legacy OS support removed** BREAKING 65

Nessus drops support for Windows Server 2012 and 2012 R2, Ubuntu 16.04, RHEL and derived distributions version 8 minor versions before 8.10, and macOS 14.

— Lists exact dropped OS versions but no migration path

Tenable Nessus10.12.6 (2026-10-22)

└──▷ **BREAKING ON UPGRADE**

! Support for Windows Server 2012 and 2012 R2 is removed — scanners running on those OSes will no longer be supported.

! Support for Ubuntu 16.04 is removed.

! Support for RHEL and derived distributions version 8 minor versions before 8.10 is removed.

! Support for macOS 14 is removed.

WAS THIS USEFUL? 👍 👎

OXO

Feedback ⓘ v2.19.0 NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> OXO's major addition this window is scan snapshotting with pause/resume and warm-restart support for local scans, alongside expanded AI provider support and small protocol- and messaging-level additions to the scanning pipeline.

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

01 Scan snapshot, pause, and resume for local scans NEW

80

Adds pause and resume support for local scans, letting practitioners halt and continue in-progress scans without losing state. Captures and persists the state of a local scan universe to a local file, and saves the full scan definition to disk so it can be restarted later without re-specifying parameters. Adds snapshot bucket options to the `oxo scanner` command so agents receive snapshot storage configuration, captures the full MQ (RabbitMQ) and Redis state of a running universe into that bucket, and supports restoring a scan snapshot before starting universe agents for a warm restart of interrupted scans.

– Names command and backing stores but no exact flags given v2.19.0 v2.18.0 v2.15.0

02 AI provider infrastructure and new model providers NEW

60

Adds shared AI model and provider construction infrastructure (`ostorlab.ai`) for use across OXO agents, then expands supported providers to Mistral, Qwen, and Ollama, and finally adds a generic `openai_compatible` provider enabling self-hosted or third-party OpenAI-compatible models.

– Names module and all providers but no usage steps v2.14.0 v2.14.1

THINNER COVERAGE BELOW

03 vulnerability_target field in vulnerability message NEW

55

Adds `vulnerability_target` field to the vulnerability message and mixin, enabling agents to specify the exact target associated with a reported vulnerability.

– Names field and message/mixin but no usage example v2.11.0

04 Experimental release channels for agent installation NEW

35

Adds experimental release channel support for agent installation and fetching, enabling use of pre-release agent builds.

— States capability but no flag or command shown v2.15.0

---

05 **TCP and UDP open-port discovery with new TCP message type**

NEW

35

Introduces a new TCP message type and adds open-port reporting for both TCP and UDP protocols.

— Names protocols and message type but no mechanism or command v2.16.0

WAS THIS USEFUL? 👍 👎

gvm-libs

Feedback ⓘ v23.15.0 NOTES ↗ RANK

Greenbone Vulnerability Management Libraries

> gvm-libs added CA certificate support to its LDAP and Active Directory connectors, updated the Agent Controller client for a revised response format, and introduced encryption and anonymization options for agent support bundles.

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

01 CA certificate support in LDAP and AD connectors NEW

60

The `gvmldap` and `ad_connector` interfaces now accept CA certificate data directly, supporting more flexible TLS configurations for LDAP and Active Directory connections respectively.

– Names both connectors but no config syntax shown v23.15.0

THINNER COVERAGE BELOW

02 Agent Controller client handles revised response format

IMPROVED

40

The `agent_controller` client was updated to handle the revised agent response format returned by the Agent Controller.

– States change but not the format details v23.14.0

03 Encryption and anonymization for agent support bundles NEW

25

Adds encryption and anonymization options for agent support bundles.

– Bare feature statement, no mechanism or flags given v23.13.0

WAS THIS USEFUL? 👍 👎

openvas-scanner

Feedback ⓘ v23.52.0 NOTES ↗ RANK

This repository contains the scanner component for Greenbone Community Edition.

> OpenVAS Scanner added a new performance-monitoring API endpoint and rewrote its HTTP API on axum, tightening request validation and routing semantics in ways that break existing integrations.

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

01 Axum-based HTTP API rewrite tightens routing and validation

83

BREAKING

The HTTP API was replaced with an axum-based implementation. Routes that accept a JSON-encoded payload now require the `Content-Type: application/json` header, rejecting requests without it; routes with a trailing `/` are now treated as distinct and return 404 instead of matching the non-trailing route; and the `enable-get-scans` option under `[endpoints]` now correctly defaults to `false` as documented, whereas it was previously silently ignored and defaulted to `true`.

— Names config key and header but no migration tooling given v23.51.0

02 Scanner performance metrics endpoint NEW

60

Adds a `GET /health/performance` endpoint to the REST API for querying scanner performance metrics.

— Names exact endpoint but no detail on returned metrics v23.52.0

└──▷ **BREAKING ON UPGRADE**

! Routes that accept a JSON payload MUST now include `Content-Type: application/json`; requests lacking this header will be rejected.

! Routes with a trailing `/` now return 404 instead of matching the corresponding non-trailing route — any client or integration calling a URL with a trailing slash will break.

! The `enable-get-scans` option under `[endpoints]` now defaults to `false`; GET scan endpoints that were previously accessible without explicit configuration will stop responding after upgrade.

WAS THIS USEFUL? 👍 👎

gvmd

Feedback ⓘ v26.42.0 NOTES ↗ RANK

Greenbone Vulnerability Manager - The database backend for the Greenbone Community Edition

> gvmd shipped two releases adding scan config details to report XML, support for executable report format files, and validation of keys and certificates on credential modification.

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

01 **Scan config info in report XML** IMPROVED 35

Report XML output now includes scan config information, giving more context about which configuration produced a given report.

— Names the output format but no fields or schema detail v26.42.0

02 **Executable report format files** NEW 35

gvmd now supports executable report format files, enabling custom executable-based report generators.

— Describes capability but no file structure or setup steps v26.41.0

03 **Credential key/certificate validation** IMPROVED 30

Keys and certificates are now validated when modifying credentials, hardening credential store mutations against invalid input.

— States behaviour change but no mechanism or error details v26.41.0

WAS THIS USEFUL? 👍 👎

vulnerablecode

Feedback ⓘ v40.2.0 NOTES ↗ RANK

A free and open vulnerabilities database and the packages they impact. And the tools to aggregate and correlate these vulnerabilities. Sponsored by NLnet https://nlnet.nl/project/vulnerabilitydatabase/ for https://www.aboutcode.org/ Chat at https://gitter.im/aboutcode-org/vulnerablecode Docs at https://vulnerablecode.readthedocs.org/

> VulnerableCode's latest releases expand the advisory API with new identification fields and curation workflows, add CWE weakness and mitigation curation, and remove a legacy advisory endpoint.

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

01 **Retrieve endpoint removed from PackageAdvisoriesViewSet** 65

BREAKING

The retrieve endpoint is removed from `PackageAdvisoriesViewSet`; any client calling that endpoint will receive an error after upgrading.

— Names exact viewset and consequence for callers v40.2.0

02 **`datasource_id` field in advisory API responses** NEW 62

Advisory API responses now include a `datasource_id` field, letting callers identify which data source produced each advisory.

— Names the exact response field but no further detail v40.2.0

THINNER COVERAGE BELOW

03 **AVID lookup and todo curation in advisory API** NEW 55

The API now supports advisory `todo` lookup by AVID identifier, and API responses expose an advisory `todo` count along with curating advisories.

— Names AVID and todo field but lacks endpoint path v40.1.0

04 **CWE weakness and extra mitigation curation** NEW 41

Adds support for CWE weakness curation and extra mitigation curation on advisories.

— Names CWE and mitigation curation but no mechanism v40.1.0

05 **Automated Docker image publishing workflow** NEW 40

Adds a workflow to publish the VulnerableCode Docker image automatically.

— Describes CI workflow but no trigger or registry detail v40.1.0

06 **Pipeline-run cleanup job** NEW 38

Adds config and a job to automatically clean up old pipeline runs, keeping the database tidy over time.

— No named config key or schedule given v40.2.0

07 **Federated advisory data schema** NEW 25

Adds a schema for federated advisory data.

— Bare mention, no structure or usage detail v40.1.0

└──▷ **BREAKING ON UPGRADE**

! The retrieve endpoint is removed from `PackageAdvisoriesViewSet`; any client calling that endpoint will receive an error after upgrading.

WAS THIS USEFUL? 👍 👎

gsa

Feedback ⓘ v28.5.0 NOTES ↗ RANK

Greenbone Security Assistant - The web frontend for the Greenbone Community Edition

> GSA v28.5.0 improves how NVT technical information is displayed and adds dimming hover effects across dashboard charts.

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

01 Markdown-rendered NVT technical information with anchor links

IMPROVED 45

NVT technical information is now rendered using React-markdown, and the `DetailsPage` supports scrolling to anchors. Result views also gained technical information links, alongside updated localization files.

— Names component and page but no exact flags or commands v28.5.0

02 Dimming hover effects across dashboard chart types IMPROVED

45

Adds hover effects with dimming transitions to `BarChart`, `BubbleChart`, `WordCloudChart`, `ScheduleChart`, and `DonutChart` to improve dashboard interactivity.

— Names all five chart components but no usage detail v28.5.0

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Didit

Feedback ⓘ snapshot-20261023 NOTES ↗ RANK

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

> Didit's session deletion API gained an async job workflow for large batch deletes, a filtered select-all delete option, and visibility into child verifications removed during business-verification cleanup.

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

01 **Async batch-delete job workflow for sessions** NEW 85

Batch Delete Sessions returns HTTP `202` when more than 20 sessions match a filter, and callers poll Get Session Deletion Job until the job reports `COMPLETED` or `FAILED`.

Purge all failed sessions in bulk and poll until the background job finishes — useful for large compliance-driven data-deletion runs.

```
$ # Step 1: Kick off the filtered select-all deletion
curl -X POST https://api.didit.me/management-api/sessions/batch-delete \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"delete_all": true, "filter_status": "FAILED"}'
# Returns 202 with a deletion job ID

# Step 2: Poll until COMPLETED or FAILED
curl https://api.didit.me/management-api/sessions/get-deletion-job/<job_id> \
  -H 'Authorization: Bearer <token>'
```

— Named endpoints, HTTP status, job states, and a runnable curl example snapshot-20261023

02 **Filtered select-all deletion for sessions** NEW 80

Batch Delete Sessions now accepts `delete_all: true` to delete every session matching a filter set, narrowed with `filter_` `-prefixed` parameters (e.g. `filter_status`, `filter_date_from`) mirroring the List Sessions filters. Invalid or unreadable filter values are rejected before any deletion begins.

Purge all failed sessions in bulk and poll until the background job finishes — useful for large compliance-driven data-deletion runs.

```
$ # Step 1: Kick off the filtered select-all deletion
curl -X POST https://api.didit.me/management-api/sessions/batch-delete \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"delete_all": true, "filter_status": "FAILED"}'
# Returns 202 with a deletion job ID

# Step 2: Poll until COMPLETED or FAILED
curl https://api.didit.me/management-api/sessions/get-deletion-job/<job_id> \
  -H 'Authorization: Bearer <token>'
```

— Names exact flag and filter params with a working example `snapshot-20261023`

THINNER COVERAGE BELOW

03 **Child verification counts in deletion results** NEW

45

Deletion results now include `deleted_child_verifications`, reporting how many key-person and company verifications were deleted as children of a business verification.

— Names the field but no usage example or endpoint shown `snapshot-20261023`

WAS THIS USEFUL? 👍 👎

cerbos

Feedback ⓘ v0.57.0 NOTES ↗ RANK

Cerbos is an open-core authorization management platform for authorizing every identity and governing every action across applications, gateways, workloads, and AI agents.

> Cerbos v0.57.0 focuses on internal reliability and performance: it fixes Hub audit log sync gaps, tightens query plan accuracy, speeds up policy evaluation, and adds a Homebrew cask for cerbosctl.

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

01 **Hub audit log sync reliability and throughput** IMPROVED 48

Hub audit log sync now resumes correctly after the last shipped marker, preventing missed entries, and syncs audit log batches outside the Badger read transaction to improve throughput.

– Explains before/after behavior but no user-facing config exposed v0.57.0

02 **Query plan `has` emission accuracy** IMPROVED 48

Improves query plan accuracy by emitting `has` in query plans only when it guards an expression on the same attribute.

– Names exact mechanism but is an automatic internal fix v0.57.0

03 **cerbosctl Homebrew cask** NEW 37

Adds a `cerbosctl` Homebrew cask for easier installation via Homebrew.

– Names the install mechanism but no exact command given v0.57.0

04 **RuleTable glob pre-filtering for performance** IMPROVED 29

Pre-filters globs in RuleTable to optimize policy evaluation performance.

– Names component but no benchmark or scope given v0.57.0

05 **Shutdown event includes API activity** IMPROVED 25

The shutdown event now includes API activity for more complete observability.

– Bare description of an observability tweak v0.57.0

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

Application Profiles are available in Early Access (2026-10-22) BLOG ↗ RANK

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

> Auth0 shipped two Early Access features for securing AI agents and MCP clients — centralized Application Profiles and Client ID Metadata Document (CIMD) dynamic registration — alongside a new manual unblock mode for Brute-Force Protection to stop email scanners from auto-consuming recovery links.

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

01 Application Profiles for DCIMD applications NEW 70

Application Profiles (Early Access) centralize security settings for DCIMD applications used by AI agents and MCP clients from a single configuration point, covering refresh token lifetimes, leeway, and DPoP. They also configure B2B behaviour — whether an application serves individuals, organizations, or both — plus login experience settings, and manage refresh token policies, DPoP, and Organizations for dynamically registered CIMD clients.

Configure Application Profiles for DCIMD apps via the Auth0 dashboard to enforce uniform refresh token and DPoP settings across all AI agent-managed applications.

> 📍 1. In the Auth0 dashboard, navigate to Applications section. 2. Open Application Profiles . 3. Set default values for refresh token lifetimes, leeway, and DPoP. 4. Configure B2B behaviour (individuals / organizations / both) and login experience. 5. Save — settings apply to all DCIMD applications at runtime.

— UI steps given; settings named but no API/config key

Application Profiles are available in E… CIMD with Dynamic Registration in Early…

02 CIMD with dynamic registration for AI agents NEW 70

Client ID Metadata Document (CIMD) with dynamic registration (Early Access) lets MCP clients and AI agents present an HTTPS metadata URL during the OAuth authorization flow and be discovered and registered just-in-time, with no administrator pre-registration required. Auth0 automatically picks up changes to a client's metadata and public signing keys without manual syncing, and also supports an optional manual pre-registration flow where an administrator submits the client metadata URL, reviews it, and approves the client before creation.

— Mechanism well described but no concrete command or endpoint

CIMD with Dynamic Registration in Early…

03 Manual unblock mode for Brute-Force Protection NEW

70

Brute-Force Protection gains an `Unblock Behavior` setting with `Automatic` (default) and `Manual` modes; Manual mode presents a confirmation page so headless email scanners (e.g. Microsoft 365 Defender) cannot silently consume the recovery link before the user opens it. It is configured via Auth0 Dashboard > Security > Attack Protection > Brute-Force Protection > Unblock Behavior.

Enable Manual unblock mode so Microsoft 365 Defender and similar link-scanning gateways cannot auto-consume brute-force recovery links.

> 📍 1. Go to Auth0 Dashboard › Security › Attack Protection › Brute-Force Protection. 2. Scroll to `Unblock Behavior`. 3. Select `Manual`. 4. Click `Save`.

— Exact setting name and dashboard path given, clear starting point

`Configurable Unblock Mode for Brute For…`

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Castle

Feedback ⓘ Identify AI agents NOTES ↗ RANK

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

> Castle now identifies AI agents and browser automation tools in traffic, surfacing details about the agent in its response.

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

01 **AI agent identification in traffic responses** NEW 55

Castle adds a new top-level object in its API response identifying AI agents and browser automation tools, including the agent's name, operator, category, and detection method. Detection includes support for verified Web Bot Auth signatures.

– Names response fields and detection method but no endpoint or schema shown.

Identify AI agents

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ

Ory Network, Ory Hydra, Ory Kratos, Ory Keto v26.4.5 released NOTES ↗ RANK

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

> Ory's single release this window adds actor identity to admin API audit events across Identities, OAuth2, and Permissions.

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

01 **Actor identity in admin API audit events** IMPROVED 45

Events emitted by Ory Identities, Ory OAuth2, and Ory Permissions when handling authenticated admin API requests now include actor identity, enabling audit trails that show who made each admin action.

– Explains what changed and scope but no API/field names given.

Ory Network, Ory Hydra, Ory Kratos, Ory…

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ Access profile types (2026-10-20)

BLOG ↗ RANK

ConductorOne is an identity governance platform that automates workforce access management and deploys AI agents to streamline access reviews and provisioning.

> ConductorOne introduced distinct access profile types — Catalog, Role, and Hybrid — giving admins explicit control over whether access is self-service, automatically granted by membership, or a hybrid of both.

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

01 Catalog, Role, and Hybrid access profile types NEW

85

Access profiles now come in three types: Catalog profiles are a collection of entitlements users can self-service request from the catalog; Role profiles automatically grant entitlements based on a membership condition, with membership itself grantable, revocable, and requestable like any entitlement, but access is all-or-nothing with no piecemeal grants; Hybrid profiles are the original ConductorOne behaviour — entitlements requestable or auto-granted on a condition — now explicitly named, and all existing profiles are automatically migrated to this type. Users are prompted to choose a type on profile creation and can change the type later from the profile's details page.

Create a Role profile to automatically provision all entitlements for a job function — no piecemeal access, membership can be requested or revoked like any other entitlement.

> 📍 In the C1 console, go to Access Profiles › Create Profile, select Role as the type, define a membership condition (e.g. department = Engineering), and assign the entitlements the role should grant.

Convert an existing Hybrid profile to a Catalog profile so employees can self-service request only the specific entitlements they need.

> 📍 In the C1 console, open the Access Profile's details page, click Edit, change the type from Hybrid to Catalog, and save.

— Names three types, mechanisms, migration behaviour, and exact UI path

Access profile types (2026-10-20)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

mock-oauth2-server

Feedback ⓘ 6.1.0 NOTES ↗ RANK

Mock OAuth2/OpenID Connect server for JVM tests and Docker Compose. No security workarounds needed.

> mock-oauth2-server 6.1.0 is a spec-compliance release that corrects OAuth 2.0/OIDC error codes, discovery document contents, and logging behaviour — all of it breaking for tests that assert on the old values.

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

01 **Client error logging rewritten to single-line INFO with debug stack traces**

75

BREAKING

Client errors are now logged at `INFO` as a single line containing method, URL, error code, and description, instead of at `ERROR` with a stack trace; server errors remain at `ERROR`. Setting `LOG_LEVEL=DEBUG` restores stack traces for these client errors. Log-level assertions or alerting rules keyed on `ERROR` for client errors will no longer trigger.

— Exact env var and log format given, actionable config change 6.1.0

02 **Discovery document advertises only configured signing algorithm**

65

BREAKING

`id_token_signing_alg_values_supported` in the discovery document now lists only the configured algorithm (`RS256` by default) instead of every algorithm the server knows. Tests asserting the full algorithm list will need updating.

— Names field and default value, clear before/after 6.1.0

03 **Correct error code for unknown grant_type** BREAKING

62

An unknown `grant_type` now returns `unsupported_grant_type` instead of `invalid_grant`, matching RFC 6749 section 5.2. Tests asserting the old `invalid_grant` code for this case will break.

— Names exact codes and RFC section but no migration snippet 6.1.0

THINNER COVERAGE BELOW

04 **error_description casing preserved in error responses**

45

BREAKING

`error_description` in error responses now preserves its original casing instead of being lowercased by the server. Tests asserting on lowercase descriptions will need updating.

— Clear before/after but no field path or example given `6.1.0`

└──▷ **BREAKING ON UPGRADE**

! `error_description` fields in error responses are no longer lowercased — tests asserting on lowercase descriptions will need updating.

! Unknown `grant_type` values now return `unsupported_grant_type` instead of `invalid_grant` — tests asserting the old error code will break.

! `id_token_signing_alg_values_supported` in the discovery document now lists only the configured algorithm instead of all known algorithms — tests asserting the full list will need updating.

! Client errors are now logged at `INFO` (single line) instead of `ERROR` (with stack trace) — log-level assertions or alerting rules keyed on `ERROR` for client errors will no longer trigger.

WAS THIS USEFUL? 👍 👎

openfga

Feedback ⓘ v1.22.0 NOTES ↗ RANK

A high-performance and flexible authorization/permission engine built for developers and inspired by Google Zanzibar.

> OpenFGA v1.22.0 adds MySQL read-replica support for scaling read-heavy workloads, alongside breaking internal API changes to the Go datastore interfaces that implement it.

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

01 MySQL read-replica support NEW

95

OpenFGA now supports primary/secondary MySQL datastores via the `--datastore-secondary-uri` flag (or `OPENFGA_DATASTORE_SECONDARY_URI` environment variable). Reads that don't require `HIGHER_CONSISTENCY` are routed to the replica while all writes go to the primary; `GetStore`, `ListStores`, `ReadAssertions`, `ReadAuthorizationModel`, `ReadAuthorizationModels`, `FindLatestAuthorizationModel`, and `ReadChanges` always hit the replica when one is configured.

Start OpenFGA with a MySQL read replica so that read-heavy workloads are offloaded from the primary.

```
$ openfga run \
  --datastore-engine mysql \
  --datastore-uri 'user:pass@tcp(primary-host:3306)/openfga' \
  --datastore-secondary-uri 'user:pass@tcp(replica-
host:3306)/openfga'
```

Set the secondary URI via environment variable in a container deployment to route replica reads without changing the command line.

```
$ OPENFGA_DATASTORE_SECONDARY_URI='user:pass@tcp(replica-
host:3306)/openfga'
```

— Names exact flag, env var, and affected methods with a runnable command. v1.22.0

02 Breaking changes to MySQL datastore internals BREAKING

75

`mysql.NewWithDB` gained a required `secondaryDB *sql.DB` parameter, so callers with no replica must now pass `nil` explicitly.

`sqlcommon.ReadAuthorizationModel` and `sqlcommon.FindLatestAuthorizationModel` were removed from `sqlcommon` and moved into per-backend implementations, breaking any code that imports them from `sqlcommon`.

— Names exact functions and packages affected but no migration command. v1.22.0

└──▷ **BREAKING ON UPGRADE**

! `mysql.NewWithDB` gained a required `secondaryDB *sql.DB` parameter — callers with no replica must pass `nil` explicitly.

! `sqlcommon.ReadAuthorizationModel` and `sqlcommon.FindLatestAuthorizationModel` were removed from `sqlcommon` and moved to per-backend implementations; any code importing them from `sqlcommon` will break.

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

bytebase

Feedback ⓘ 3.24.0 NOTES ↗

Database governance built for humans and agents — controlling changes and access across every major database.

MCP read-write agents can now manage access requests end-to-end; self-hosted admin recovery runs without downtime; `ListPlans` gains release and task-status filters.

└──▷ GET THIS VERSION

```
$ git clone --branch 3.24.0 https://github.com/bytebase/bytebase.git
# already have the repo? check out this version:
$ git checkout 3.24.0
```

Extends MCP read-write mode so a connected agent can request just-in-time access, approve or reject issues and access requests, toggle project just-in-time access, grant or revoke masking exemptions, and revoke access grants — all bounded by the signed-in user's own permissions.

Adds self-hosted admin recovery: re-enable password sign-in or reset a user's password while Bytebase keeps running, with no restart required, scriptable for unattended use.

Adds `release` and `has_release` filter parameters to `ListPlans` API, plus `task_status` filter, so GitOps plans can be distinguished from console-created plans.

Adds TLS enforcement for SQL Server and Oracle connections: Verify server certificate now checks the certificate chain and host name, and Oracle TLS connects over TCPS without a wallet.

Adds a 10-second connect timeout for Oracle connections, with `CONNECTION TIMEOUT` available under Extra Parameters to override (use `0` for no limit).

+1 more

NEW SUPPORTED PLATFORM BigQuery SQL review

└──▷ BREAKING ON UPGRADE

! HA mode: stop every replica before starting 3.24.0 — rolling upgrades are not supported. A replica still on an older version saves access grants that 3.24.0 cannot match, and those grants stop working for queries and exports.

! SQL Server: `verify_tls_certificate = true` on a data source is turned off during upgrade; re-enable it via Terraform only after confirming the certificate verifies.

! Oracle: `use_ssl = true` on a data source is turned off during upgrade, and stored CA and client certificates are removed; re-enable via Terraform only after confirming the port serves TCPS.

WAS THIS USEFUL? 👍 👎

◆ Backup & Recovery

Veeam Kasten

Feedback ⓘ 9.0.9 NOTES ↗ RANK

Veeam Kasten is a Kubernetes-native platform for backing up, restoring, and migrating containerized applications and their persistent data.

> Kasten K10 9.0.9 adds compatibility with Veeam Backup & Replication 13.2 and enforces stricter minimum values on Helm timeout parameters.

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

01 Minimum values enforced on Helm timeout parameters

BREAKING

90

Helm parameters `timeout.backup`, `timeout.restore`, `timeout.retire`, `timeout.blueprintResourceHooks`, `timeout.checkRepoPodReady`, `timeout.statsPodReady`, `timeout.efsRestorePodReady`, `timeout.workerPodReady`, and `timeout.imageStreamImageResolveSec` now require a value of at least `1`; `timeout.jobWait` must be empty or at least `1`. Installs or upgrades that supply `0` or a negative value will fail chart validation, so existing Helm values files must be checked before upgrading to 9.0.9.

Ensure all timeout Helm parameters meet the new minimum-of-1 requirement when upgrading to 9.0.9 to avoid install/upgrade failures.

```
$ helm upgrade k10 kasten/k10 --namespace kasten-io \
  --set timeout.backup=180 \
  --set timeout.restore=300 \
  --set timeout.retire=60 \
  --set timeout.blueprintResourceHooks=30 \
  --set timeout.checkRepoPodReady=30 \
  --set timeout.statsPodReady=30 \
  --set timeout.efsRestorePodReady=30 \
  --set timeout.workerPodReady=30 \
  --set timeout.imageStreamImageResolveSec=30 \
  --set timeout.jobWait=120
```

— Names every affected flag and gives a runnable upgrade command 9.0.9

THINNER COVERAGE BELOW

02 Veeam Backup & Replication 13.2 compatibility

IMPROVED

35

Adds support for the updated data mover certification exchange required for compatibility with Veeam Backup & Replication 13.2 and later.

— Names the compatibility target but no mechanism detail beyond cert exchange `9.0.9`

└──▷ **BREAKING ON UPGRADE**

! Helm parameters `timeout.backup`, `timeout.restore`, `timeout.retire`, `timeout.blueprintResourceHooks`, `timeout.checkRepoPodReady`, `timeout.statsPodReady`, `timeout.efsRestorePodReady`, `timeout.workerPodReady`, and `timeout.imageStreamImageResolveSec` now require a value of at least `1`; `timeout.jobWait` must be empty or at least `1`. Installs or upgrades that supply `0` or a negative value will fail chart validation.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

amneziawg-installer

Feedback ⓘ v6.1.0 NOTES ↗ RANK

1 VPN server on a fresh Ubuntu or Debian VPS with one command. Made for networks that block WireGuard.

> amneziawg-installer now defaults new installs to AmneziaWG 3.1 with header encryption and an automatic/explicit 2.0 fallback path, and adds fine-grained control over the obfuscation I1 packet plus privilege diagnostics and documentation matrices.

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

01 Custom I1 concealment packet via --i1 flag NEW

98

The `--i1` flag lets operators pin a custom I1 concealment packet instead of accepting the generated default, e.g. `--i1='<b 0xHEX><r 16>'`, validated immediately against accepted tags (`<b 0xHEX>` , `<r N>` , `<rc N>` , `<rd N>` , `<t>` ; packet ≤ 1024 bytes); the `<c>` tag is rejected because it prevents mobile profile loading. The value is pinned in `awgsetup_cfg.init` so reinstalls and 2.0 fallbacks preserve it, while `--no-cps` , `--preset` , `--mobile` , `--jc` , `--jmin` , or `--jmax` drop it with a warning. On a running server, `--force --i1=...` updates only I1 and leaves other parameters unchanged, with issued profiles receiving the new value after `regen` .

Set a specific I1 concealment packet on a fresh install to control obfuscation fingerprinting rather than accepting the generated default.

```
$ sudo bash ./install_amneziawg.sh --i1='<b 0xdeadbeef><r 16>'
```

Update only the I1 packet on a running server without changing any other obfuscation parameters, then regenerate client profiles to distribute the new value.

```
$ sudo bash ./install_amneziawg.sh --force --i1='<b 0xdeadbeef>
<rc 32>'
```

— Exact flag syntax, validation rules and update behaviour given v6.1.0

02 AmneziaWG 3.1 default with 2.0 protocol fallback BREAKING

95

New installations now default to AmneziaWG 3.1 with header encryption (`HeaderProtectionKey` , unique per installation) and `ContentPaddingAddition =`

`32-128` padding. The `--protocol=2.0` flag forces installation of AmneziaWG 2.0 for clients that don't support 3.1 (routers on stock firmware, Hiddify, older apps); new `AWG_PROTOCOL_SOURCE` (`default` or `explicit`) and `AWG_PROTOCOL_FALLBACK` (reason code) keys are written to `awgsetup_cfg.init`. Automatic fallback to 2.0 triggers on kernels older than 6.7 (stock Debian 12), non-x86_64 architectures including ARM, or installed `awg` utilities/module that don't support the 3.x line. Switching a running server between generations requires `--uninstall` followed by a fresh install, and all client profiles must be re-issued.

Install an AmneziaWG 2.0 server when your clients include stock-firmware routers or Hiddify that do not support 3.1.

```
$ sudo bash ./install_amneziawg.sh --protocol=2.0
```

— Full mechanism, triggers, config keys and migration path given v6.0.0

03 **Kernel module privilege warning in diagnostics** IMPROVED

70

The `check` and `diagnose` commands now warn when the kernel module is accessible without root via user namespace privilege, and name the `sysctl` setting (`kernel.unprivileged_userns_clone = 0`) that closes that path. A new Kernel module and user namespaces section was added to the ADVANCED documentation.

— Names commands and exact sysctl key but no runnable example v6.0.0

THINNER COVERAGE BELOW

04 **Client and support compatibility matrices in docs** NEW

45

ADVANCED documentation now includes a client compatibility matrix built from `docs/clients-matrix.json` (with evidence source and check date per row) and a support matrix built from `docs/support-matrix.json`.

— Names the data files but is purely informational, no action v6.0.0

└──▷ **BREAKING ON UPGRADE**

! New installations now default to AmneziaWG 3.1; clients without 3.1 support (stock-firmware routers, Hiddify, outdated apps) will fail to connect silently — re-install with `--protocol=2.0` to serve them.

! Switching a running server between protocol generations (3.1 to 2.0 or vice versa) requires `--uninstall` followed by a fresh installation; all client profiles must be re-issued.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 6.55.0 (2026-10-22) BLOG ↗ RANK

RegScale is a compliance automation platform that helps organizations manage security controls, risk, and regulatory frameworks like NIST, FedRAMP, and SOC 2 through continuous monitoring.

> RegScale CLI 6.55.0 introduces configurable approval routing for DevSecOps Change records and adds Python 3.15 support.

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

01 Configurable approval routing for Change records NEW

86

Adds the `change.approval_by_type` config key and `--approval-mode` CLI flag to route Change records directly to completion, to policy approval, or always to the CCB, configurable per change type. This routing applies to the new DevSecOps `release` and `deploy` Change record types introduced in the same release.

Force all DevSecOps deploy Change records to route through the CCB rather than auto-completing, using the new per-type approval config.

yaml

```yaml
change:
  approval_by_type:
    deploy: ccb
    release: policy
```

— Names config key, flag, and record types with a working example CLI 6.55.0 (2026-10-22)

THINNER COVERAGE BELOW

02 Python 3.15 support IMPROVED

25

Adds Python 3.15 support with updated CI test matrices and wheel bounds.

— Bare statement of version support, no further detail CLI 6.55.0 (2026-10-22)

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.34.0 (2026-10-23) BLOG ↗ RANK

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

> Runlayer shipped an idempotent Auto-sync API for connectors, plugins and skills, audit logging and cascading cancellation for agent runs, org-isolated evaluation replays, policy-managed Chrome installs for AI Watch, and artifact search plus Slack/Teams comment notifications in the Launcher.

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

01 Auto-sync API for connectors, plugins, skills NEW

85

Adds an Auto-sync API to read, set, or remove a connector, plugin, or skill's device auto-sync setting by its type and ID, using an idempotent `PUT`.

Idempotently enable auto-sync for a specific connector by type and ID without risking duplicate writes.

```
$ curl -X PUT https://api.runlayer.com/v2/auto-sync/<type>/<id> \
  -H 'Authorization: Bearer <api-key>' \
  -H 'Content-Type: application/json' \
  -d '{"auto_sync": true}'
```

— Runnable PUT endpoint example with exact request shape

1.33.0 - Artifact notifications and res…

02 AI Watch Chrome extension install via policy NEW

80

Adds a Chrome extension installed by your policy toggle in AI Watch browser extension settings, letting you install AI Watch via Google Admin, Chrome Enterprise Core, or MDM instead of having AI Watch write its own Chrome policy — preventing it from overwriting your policy and removing extensions like Okta.

Use your existing Chrome fleet policy to deploy AI Watch without letting Runlayer overwrite or shadow your other managed extensions (e.g. Okta).

> 📍 In AI Watch browser extension settings, enable Chrome extension installed by your policy, then add the Runlayer entry to your Google Admin or MDM Chrome extension policy. AI Watch will write no Chrome policy of its own, and the device health check will confirm your policy installed the extension.

— Names the toggle and gives exact deployment steps 1.34.0 (2026-10-23)

03 **Deprecation header on v1 API endpoints** DEPRECATED 75

Adds a Deprecation response header to API-key calls made to v1 endpoints, making it easy to identify scripts that still depend on unsupported v1 automation paths.

Audit which of your automation scripts still call unsupported v1 endpoints by inspecting the Deprecation header returned on any v1 API-key request.

```
$ curl -si -H 'Authorization: Bearer <api-key>'
https://api.runlayer.com/v1/<endpoint> | grep -i deprecation
```

— Runnable curl to detect the header, clear migration signal

1.33.0 - Artifact notifications and res…

THINNER COVERAGE BELOW

04 **Go reliability improvements** IMPROVED 55

Go now retries failed lookups via alternate routes, checks whether a failed write landed before retrying (preventing duplicates), cites the source of each number in its answers, and executes the next step rather than offering it.

— Mechanism-rich description but no command or config to invoke

1.33.0 - Artifact notifications and res…

05 **Audit logging for agent run cancellations** NEW 50

Adds audit log entries when an agent run is stopped via the API, web app, or Slack, recording who stopped it and whether the sandbox halted.

— Names channels and fields logged but no UI path given 1.34.0 (2026-10-23)

06 **Cascading cancellation of session agent runs** IMPROVED 50

Stopping a session now also cancels all Runlayer Agents runs inside it, including runs you do not own, and only reports Terminated once the agent's runtime confirms.

— Describes behavior change but no trigger command shown 1.34.0 (2026-10-23)

07 **Artifact search in Launcher** NEW 50

Adds artifact search to the Launcher — find accessible artifacts by title or filename, open them directly, and return to recently opened files.

— Names search fields but no exact navigation steps 1.34.0 (2026-10-23)

08 **Artifact comment notifications via Slack and Teams** NEW

45

Artifact comment notifications via Slack or Teams: mentions, replies, and comments on owned artifacts now trigger a message from the Runlayer assistant linking directly to the artifact.

— Names channels and triggers but no config step shown

1.33.0 - Artifact notifications and res…

09 **Org isolation for evaluation replays** IMPROVED 40

Evaluation replays are now isolated to your organization — a replay can no longer list or call a connector belonging to another organization.

— States the fix but no reproduction or scope detail 1.34.0 (2026-10-23)

WAS THIS USEFUL? 👍 👎

jadx-ai-mcp

Feedback ⓘ v6.5.0 NOTES ↗ RANK

Plugin for JADX to integrate MCP server

> jadx-ai-mcp v6.5.0 adds class/method/field annotation tools with new persistence endpoints, slims Kotlin metadata out of decompiled source by default, ships an OpenAI-compatible function-calling client, and fixes null-argument validation across several tools.

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

01 **Kotlin metadata slimming on source-returning tools** BREAKING 91

`get_class_source`, `get_method_by_name`, `fetch_current_class`, `get_main_activity_class`, and `get_main_application_classes_code` gain a `slim=True` default parameter that strips `@Metadata` and `@DebugMetadata` Kotlin compiler annotations, shrinking returned source by 25-50%; pass `slim=False` to restore the original, unslimmed output. Workflows depending on full annotation output must now explicitly set `slim=False`.

— Lists exact tools, flag, stripped annotations and size impact v6.5.0

02 **OpenAI-compatible function-calling client** NEW 90

New `examples/openai_compatible_client.py` runs against the server's `--http` mode and converts all 34 MCP tools into OpenAI function-calling tools, driving any `POST /v1/chat/completions` `-compatible` provider (OpenRouter, LiteLLM, vLLM, Ollama, OmniRoute) via `--base-url`, `--api-key`, `--model`, `OPENAI_*` environment variables, `--list-tools`, `--max-steps`, and `--max-tool-chars`.

— Full command-line surface and supported providers named v6.5.0

03 **Comment annotation tools for decompiled code** NEW 86

Adds `add_comment` and `list_comments` MCP tools (bringing the total to 34), backed by new plugin `GET /add-comment` and `GET /list-comments` endpoints in `CommentRoutes`, letting agents annotate decompiled classes, methods, fields, or individual lines with comments that persist in the JADX project file.

— Names both new tools and both endpoints with persistence behaviour v6.5.0

04 **Explicit null support for optional tool arguments** IMPROVED 72

Nine optional arguments — `get_method_by_name.method_signature`, `rename_method.class_name`, `rename_method.method_signature`,

`rename_variable.reg`, `rename_variable.ssa`, `add_comment.method_name`, `add_comment.method_signature`, `add_comment.field_name`, and `add_comment.line` — now accept explicit `null`, so clients that send null for use the default no longer fail validation.

— Lists all nine affected arguments and the fixed behaviour v6.5.0

└──▷ **BREAKING ON UPGRADE**

! Tools `get_class_source`, `get_method_by_name`, `fetch_current_class`, `get_main_activity_class`, and `get_main_application_classes_code` now strip Kotlin `@Metadata` and `@DebugMetadata` annotations by default (`slim=True`); workflows that depended on full annotation output must pass `slim=False` to restore previous behaviour.

WAS THIS USEFUL? 👍 👎

wazuh-mcp-server

Feedback ⓘ v5.0.0 NOTES ↗ RANK

Production-grade MCP server for Wazuh SIEM — 55 security tools for alert triage, threat hunting, vulnerability management, compliance (PCI DSS, GDPR, HIPAA, NIST CSF, ISO 27001) and active response. Connect Claude or any LLM to your SOC.

> Wazuh MCP Server v5.0.0 adds OIDC sign-in, per-tool scope annotations, environment-based toolset filtering and a hardened default security posture (self-signed certs rejected, write actions require confirmation, strict input schemas), alongside a string of breaking changes to existing deployments.

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

01 OIDC sign-in for MCP sessions

BREAKING — 90

Adds `OAUTH_IDP_ISSUER` to enable OpenID Connect sign-in via Entra ID, Google Workspace, Okta or any OIDC provider, with RS256/JWKS token verification, tenant/domain/user allow-lists (e.g. `OAUTH_IDP_ALLOWED_USERS`) and group-to-scope mapping. Breaking: `/oauth/authorize` no longer auto-approves requests — users must sign in with their `wazuh_` API key; all existing bearer tokens are invalidated and must be re-minted after upgrading; combining `OAUTH_ENABLE_DCR` with `OAUTH_IDP_ISSUER` is now refused at startup.

Authenticate OAuth users through Okta so every MCP action is tied to a real corporate identity and scoped by the user's Wazuh API key.

```ini
OAUTH_IDP_ISSUER=https://your-org.okta.com
OAUTH_IDP_ALLOWED_USERS=alice@example.com,bob@example.com
```

— Full auth mechanism plus concrete breaking behaviour and config keys v5.0.0

02 Private CA trust for Wazuh Manager TLS

BREAKING — 80

Adds `WAZUH_CA_BUNDLE` (and a per-cluster `ca_bundle` option) so a private CA can be verified instead of disabling TLS checks; the bundle is validated at startup and passed to httpx as an SSL context. `WAZUH_ALLOW_SELF_SIGNED` now defaults to `false`, so existing deployments using the Manager's stock self-signed certificate will fail TLS verification on upgrade until the certificate is reissued or the flag is set explicitly.

Trust a private CA when the Wazuh Manager uses an internally issued certificate, avoiding the need to disable TLS verification.

```ini
WAZUH_CA_BUNDLE=/etc/ssl/certs/internal-ca.pem
WAZUH_ALLOW_SELF_SIGNED=false
```

— Mechanism, default-change impact and config keys all specified v5.0.0

03 **Toolset and tool filtering via environment variables** NEW

75

Adds `WAZUH_TOOLSETS` and `WAZUH_DISABLED_TOOLS` environment variables to limit which tools appear in `tools/list` and are callable via `tools/call`; unknown tool names now fail at startup instead of being silently ignored.

Restrict an LLM agent to read-only SOC investigation tools to reduce the token budget and eliminate accidental write-tool selection.

```
$ WAZUH_TOOLSETS=investigation
WAZUH_DISABLED_TOOLS=wazuh_firewall_drop,wazuh_host_deny docker
run --env-file .env ghcr.io/gensecaihq/wazuh-mcp-server:5.0.0
```

— Names exact env vars and MCP endpoints with a runnable example v5.0.0

04 **Active-response scope controls for Manager and fleet** NEW

65

Adds `WAZUH_ALLOW_MANAGER_AR=true` to permit active-response actions targeting agent `000` (the Manager) or restarting the Manager, and `WAZUH_ALLOW_FLEET_AR=true` to permit fleet-wide active-response blocks.

— Both flags and their exact scope named, no example shown v5.0.0

05 **Configurable rate limiting on MCP endpoints** IMPROVED

65

Applies `RATE_LIMIT_REQUESTS` / `RATE_LIMIT_WINDOW` to the `/mcp` and `/` endpoints, which previously used hard-coded constants.

— Before/after and exact endpoints and config keys named v5.0.0

06 **Startup validation of configuration values** BREAKING

65

Startup now refuses invalid configuration such as `ENVIRONMENT=prod` typos, `RATE_LIMIT_REQUESTS=0`, bad `SESSION_TTL_SECONDS`, or malformed `clusters.json`, whereas previously these were silently misread.

— Specific invalid-value examples given, clear migration note v5.0.0

07 **Bounded OAuth session store with LRU eviction** NEW

60

Adds `MAX_SESSIONS` and `MAX_SESSIONS_PER_PRINCIPAL` to cap the in-memory OAuth session store, with LRU eviction of expired sessions.

— Mechanism and config keys named, no concrete limits given v5.0.0

08 **Per-tool capability annotations for clients** NEW 60

Adds MCP tool annotations (`readOnlyHint`, `destructiveHint`, `idempotentHint`, `openWorldHint`) on every tool, derived from scope, so clients and gateways can gate approvals.

— Named annotation fields but no client-side action shown v5.0.0

THINNER COVERAGE BELOW

09 **Quarantine directory allow/deny prefixes** NEW 55

Adds `WAZUH_QUARANTINE_DENY_PREFIXES` and `WAZUH_QUARANTINE_ALLOW_PREFIXES` to extend or restrict the directories the `wazuh_quarantine_file` tool may operate on.

— Named config keys but no usage example or default values given v5.0.0

10 **Mandatory confirmation for write tools** BREAKING 55

Adds `WAZUH_REQUIRE_ACTION_CONFIRMATION` to control whether write tools require `confirm=true`; it defaults to on in production, so callers that do not send the flag are now refused.

— Clear config key and behaviour but no code example v5.0.0

11 **`/sse` endpoint replaced by `/mcp`** BREAKING 55

`/sse` now returns 410 Gone; clients must use `/mcp` instead.

— Exact endpoint and status code, simple migration action v5.0.0

12 **Strict tool input schemas reject unknown arguments** BREAKING

> Tool input schemas now set `additionalProperties: false`; clients sending unknown tool arguments that were previously silently ignored will now receive errors.
>
> — Named schema flag and behaviour change, no example given `v5.0.0`

└──▷ **BREAKING ON UPGRADE**

- ! `WAZUH_ALLOW_SELF_SIGNED` now defaults to `false`; existing deployments using the Manager's stock self-signed certificate will fail TLS verification on upgrade until the certificate is reissued or `WAZUH_ALLOW_SELF_SIGNED=true` is set explicitly.
- ! `/oauth/authorize` no longer auto-approves requests — users must now sign in with their `wazuh_` API key; any automation relying on the previous auto-approval flow will break.
- ! All bearer tokens must be re-minted after upgrading; existing tokens are invalidated.
- ! Write tools now require `confirm=true` by default in production (`WAZUH_REQUIRE_ACTION_CONFIRMATION`); callers that do not send this flag will be refused.
- ! Startup now refuses invalid configuration values (e.g. `ENVIRONMENT=prod` typos, `RATE_LIMIT_REQUESTS=0`, bad `SESSION_TTL_SECONDS`, malformed `clusters.json`); previously these were silently misread.
- ! `OAUTH_ENABLE_DCR` combined with `OAUTH_IDP_ISSUER` is refused at startup.
- ! `/sse` now returns 410 Gone; clients must use `/mcp` instead.
- ! Tool input schemas now set `additionalProperties: false`; clients sending unknown tool arguments that were previously silently ignored will now receive errors.

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Agentic Infrastructure

microsandbox

Feedback ⓘ v0.7.9 NOTES ↗ RANK

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

> microsandbox v0.7.9 ships a mix of SDK and CLI additions: configurable backend restrictions with new `context` commands, authenticated HTTP CONNECT proxy egress, JSON-formatted policy denials, SDK-exposed sandbox exit outcomes, live label edits, and a `modify` command for changing published ports on running sandboxes.

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

01 **Port modify command for running sandboxes** NEW 60

Supports adding and removing published ports via the `modify` command on already-running sandboxes.

— Names the exact CLI command and its effect v0.7.9

THINNER COVERAGE BELOW

02 **Authenticated HTTP CONNECT proxy egress** NEW 50

Adds HTTP CONNECT proxy basic authentication support, allowing sandbox network egress through authenticated proxies.

— Names the exact proxy mechanism but no config detail v0.7.9

03 **JSON-formatted policy denials** IMPROVED 50

Network requests blocked by policy now return JSON-formatted denial responses by default.

— Clear behavior change, no schema or example shown v0.7.9

04 **Sandbox exit outcomes in SDK** NEW 40

Exposes retained sandbox exit outcomes through the SDK so callers can inspect how a sandbox terminated.

— No method name or data shape given v0.7.9

05 **Live label edits via SDK** NEW 35

Applies label modifications live on already-running sandboxes via the SDK.

— Describes effect but not the SDK call used v0.7.9

06 **Backend restriction and context commands** NEW 32

Adds `context` commands and the ability to restrict allowed backends via configuration.

— Two capabilities bundled with no config key or command syntax given v0.7.9

WAS THIS USEFUL?  

GOVERN

◆ AI Red Teaming & Testing

ifixai

Feedback ⓘ v4.3.0 NOTES ↗ RANK

Independent Auditing of AI Agents. Run by human or the agent itself, to answer the most crucial question in the AI Agent Economy.

> ifixai v4.3.0 brings chat-native auditing to 15 inspections, adds fixture-authoring and governance-testing commands, and tightens judge-budget, validation and CLI semantics in ways that break compatibility with 4.2.x runs.

└──▷ **WHAT SHIPPED · 13 FEATURES** most completely described first ? what's the number?

01 **Stricter `--judge-budget` accounting** BREAKING 95

`--judge-budget` now counts every judge call including retries, fallback models, and ensemble judges in Full mode. `--judge-budget 0` is now a runaway guard of 10x the estimate (minimum 2,000 calls) rather than unlimited, and a run that hits its ceiling or whose judge fails stops early and exits 1 (or 2 when a mandatory minimum that finished failed).

Cap judge spend on a long CI run so a runaway retry storm cannot exhaust your quota.

```
$ ifixai run --fixture my-agent.yaml --full-run --judge-budget
500
```

— Exact flag, thresholds, exit codes, and a runnable example. v4.3.0

02 **`ifixai validate` evidence floor checks and run-time refusal** 86

BREAKING

`ifixai validate` now checks evidence floors, refusing any fixture that would leave an inspection without a verdict, and accepts the same `-b`, `-c`, `--suite`, `--strategic`, and `--full-run` flags as `ifixai run`. `ifixai run` itself now refuses, before any call, a fixture that cannot give a selected inspection a verdict — run `ifixai validate` to see what is missing.

— Names both commands, shared flags, and the new refusal behavior. v4.3.0

03 **Chat-native auditing for 15 inspections** BREAKING 85

Enables 15 inspections — `B01`, `B02`, `B04`, `B23`, `C02`, `C05`, `C11`, `M02`, `M03`, `M07`, `P01`, `P08`, `V08`, `X04`, and `X11` — to audit agents over a plain chat endpoint

with no adapter required. These are now decided by agent replies rather than hooks, so scores from these checks are not comparable with 4.2.x runs.

— Names all 15 check codes and the scoring-compatibility break. v4.3.0

04 **`ifixai windtunnel` command for governance testing** NEW

80

Adds `ifixai windtunnel` with subcommands `controls`, `run`, `sweep`, and `sensitivity` to measure whether governance defects are caught by the suite and report false-positive rates; it runs offline without changing grades.

— Names command, all four subcommands, and offline behavior. v4.3.0

05 **CLI flag and script removals** DEPRECATED

75

`--governance` is now ignored (governance rules must be placed in the fixture's `governance:` block); `--profile` has been removed and now gives an error (use `--mode standard` or `--mode full`); and the record/replay bridge, the unused embedding classifier, and the `ifixai.examples` scripts have been removed.

— Names every removed flag and component plus replacement. v4.3.0

06 **`ifixai author` command for fixture generation** NEW

70

Adds `ifixai author` to generate a fixture from a short JSON agent profile; `--example` writes a starter profile and `--force` allows overwriting an existing file.

— Exact command and flags, no deeper mechanism given. v4.3.0

07 **Optional `--api-key` via provider env var** IMPROVED

70

`--api-key` is now optional: when omitted, the SUT key is read from the provider's environment variable (e.g. `OPENAI_API_KEY`).

— Names flag and env-var fallback but no further mechanism. v4.3.0

08 **V11 Guilt Tripping check** NEW

70

Adds V11 Guilt Tripping (exploratory, report-only): 40 four-turn companion chats checking the agent never harms a struggling user, belittles them, or delays pointing them to a human.

— Specific scenario count and criteria, but no CLI surface given. v4.3.0

09 **Guided fixture flow in Claude Code / Codex plugin** IMPROVED

70

Updates the Claude Code / Codex plugin and `ifixai install` skill to a new guided flow: discover agent, ask at most two questions, author and validate fixture, show `--`

`dry-run` cost, then run only after confirmation.

— Describes the step sequence and `--dry-run` cost check. v4.3.0

THINNER COVERAGE BELOW

10 **Python API `run_single` raises on not-applicable checks** 50

BREAKING

`run_single` now raises `ValueError` (with reason) for a check that is not applicable.

— Names the API change but gives no further usage detail. v4.3.0

11 **Mock demo agent behavior change** BREAKING 50

The mock demo agent now intentionally fails `B01`, `B12`, `B16`, and `P01`, receives grade D, and exits 2.

— Specific codes and exit code but minor demo-only scope. v4.3.0

12 **`not_applicable` field in reports** NEW 45

Surfaces a `not_applicable` field in reports explaining why a check was set aside before scoring.

— Names the field but no example of its output. v4.3.0

13 **MiniMax default model update** IMPROVED 30

Sets MiniMax default model to MiniMax-M3.1-Flash.

— Single config default with no further context. v4.3.0

└──▷ BREAKING ON UPGRADE

! The 15 chat-probed inspections (B01, B02, B04, B23, C02, C05, C11, M02, M03, M07, P01, P08, V08, X04, X11) are now decided by agent replies, not hooks — scores are not comparable with 4.2.x runs.

! `ifixai run` now refuses, before any call, a fixture that cannot give a selected inspection a verdict; run `ifixai validate` to see what is missing.

! `--judge-budget 0` is now a runaway guard (10x estimate, minimum 2,000 calls), not unlimited; a run that hits its ceiling or whose judge fails stops early and exits 1 (or 2 when a mandatory minimum that finished failed).

! `--governance` is now ignored; governance rules must be placed in the fixture's `governance:` block.

! `--profile` has been removed and now gives an error; use `--mode standard` or `--mode full`.

! The record/replay bridge, the unused embedding classifier, and the `ifixai.examples` scripts have been removed.

! Python API: `run_single` now raises `ValueError` (with reason) for a check that is not applicable.

! The mock demo agent now intentionally fails B01, B12, B16, and P01, receives grade D, and exits 2.

WAS THIS USEFUL? 👍 👎

◆ AI/LLM Security

tenuo

Feedback ⓘ v0.3.3 NOTES ↗ RANK

Task-scoped authorization for AI agents. Cryptographic warrants constrain tools and arguments, prevent privilege escalation at every delegation hop, and produce signed evidence of what was allowed or denied.

> Across three releases Tenuo added delegated warrant-chain support spanning Temporal workflows and five agent-framework adapters, centralized proof and chain verification in a shared Rust core, and tightened MCP argument bounds, OpenAI guard() restrictions, and URL-safety matching — alongside new path_glob() path helpers, a public enforce_tool_call API, and an installable agent-authorization skill. Several changes are breaking: health probes move ports, BoundWarrant validation now requires an explicit trust anchor, and GrantBuilder tool scoping becomes additive.

└──▷ **WHAT SHIPPED · 22 FEATURES** most completely described first ? what's the number?

01 **path_glob()/pathGlob() safe path helpers** NEW 88

Adds path_glob() (Python) and pathGlob() (TypeScript) to compose `Subpath` with a glob that decides after normalization, replacing the unsafe bare `Pattern` which let `*` cross `/` boundaries.

Scope a file-reading tool to a specific directory subtree without allowing `*` to escape across `/` path boundaries.

python

```python
from tenuo import Capability, path_glob, SigningKey, configure,
mint_sync

configure(issuer_key=SigningKey.generate(), dev_mode=True,
audit_log=False)

authority = Capability(
    "read_file",
    path=path_glob("/data/*"),
)

with mint_sync(authority):
    # '/data/report.csv' is allowed; '/data/../secrets/key' is
rejected after normalization
    result = read_file("/data/report.csv")
```

— Named functions with a runnable scoping example v0.3.1

02 **Temporal workflow delegation and warrant-chain enforcement** 82

BREAKING

Adds `warrant_chain=` on Temporal workflow-start helpers so delegated warrants reach workers that trust only the root, with workflow-scoped capabilities via the reserved `_workflow_namespace` and `_workflow_id` argument names. Temporal signal and update allowlists now apply during the workflow's first activation; a missing or wrong-key proof of possession is reported as `POP_VERIFICATION_FAILED` instead of `CONSTRAINT_VIOLATED`; replaying a 0.3.2 history whose run failed on an off-list signal now reports nondeterminism; and `temporalio>=1.33` is required.

— Names exact params, error codes, and a version requirement v0.3.3 v0.3.2

03 **Installable agent-authorization skill** NEW 81

Adds an installable `tenuo-agent-authorization` skill via `npx skills add tenuo-ai/tenuo --skill tenuo-agent-authorization`, with automated release re-pin, to place Tenuo enforcement at effect boundaries and generate allow and deny tests.

Install the agent-authorization skill so a coding agent can place Tenuo enforcement at effect boundaries and generate both allow and deny tests automatically.

```
$ npx skills add tenuo-ai/tenuo --skill tenuo-agent-authorization
```

— Exact install command and stated purpose v0.3.1

04 **GrantBuilder tool scoping now additive** BREAKING 81

GrantBuilder.tool() / tools() are now additive and raise after inherit_all(); the narrowing idiom is retain_tool() / retain_tools().

— Exact methods and the migration idiom named v0.3.1

05 **MCP integration hardening: argument bounds, closed-world denial, and SDK 2.x middleware**

78

BREAKING

Adds `TenuoServerMiddleware` for MCP SDK 2.x. Enforces MCP argument JSON size bounds of 256 KiB encoded JSON, 64 KiB decoded string/key bytes, 4,096 values, 256 entries per container, and depth 9, with strict JSON parsing that rejects repeated argument keys. MCP client tools now deny calls when no warrant is in scope (closed-world verification). Mixed 0.3.1/0.3.2 MCP peers reject proofs containing `null` or integral floats such as `1.0` with `invalid_pop`.

Use closed-world MCP verification so any tool call without a warrant in scope is denied.

```python
from tenuo.mcp import MCPVerifier

verifier = MCPVerifier(config=closed_world_config)
```

— Concrete numeric limits and a runnable closed-world example v0.3.3 v0.3.2

06 **Delegated warrant chains across agent framework adapters**

76

BREAKING

Adds delegated warrant-chain support across `CrewAI`, Google ADK, `AutoGen`, `OpenAI`, and `LangChain` adapters. Extends OpenAI's guard() to verify tool calls from async clients and from legacy `function_call` responses, and extends AutoGen stream verification to cover every buffered tool call regardless of finish reason. Google ADK `TenuoPlugin` callbacks are now async and keyword-only to match `BasePlugin`, so synchronous or positional callback signatures will break.

— Names five frameworks and the exact ADK callback break v0.3.2 v0.3.3

07 **OpenAI guard() configuration restrictions** BREAKING

75

responses.create(stream=True) under guard() now raises `CFG_004`, and `TenuoToolGuardrail` in `input_guardrails` raises `CFG_006`; use guard_agent(), which requires `openai-agents>=0.3.3`.

— Exact error codes and the replacement API given v0.3.3

08 **A2A audience binding opt-in** BREAKING

75

A2A `require_audience` now defaults to False; set `require_audience=True` or `TENUO_A2A_REQUIRE_AUDIENCE=true` to keep audience binding.

Opt back in to A2A audience binding after the default changed to False.

```env
TENUO_A2A_REQUIRE_AUDIENCE=true
```

— Exact flag and env var to restore prior behavior v0.3.3

09 **Rust-core chain verification and proof construction** IMPROVED

75

Centralizes holder-side proof construction, approval request-hash verification, and receipt-chain verification in the Rust core, sharing one `_meta.tenuo` envelope with conformance vectors across Rust, Python, and TypeScript. Chain verification now requires a depth-0 root without a parent hash, enforces issuance ordering along the chain, offers an optional lifetime ceiling, and fails closed on an empty All([]).

— Mechanism and rules named, no direct code example v0.3.2 v0.3.1

10 **Health and status routes moved to dedicated listener** BREAKING

75

`/health`, `/healthz`, `/ready`, and `/status` now live on a dedicated listener (port 9091), preventing Envoy HTTP ext_authz from interpreting unauthenticated health responses as ALLOW decisions; authorizer health probes must move from port 9090 to port 9091 since they no longer bypass authorization on the main listener.

— Exact ports and routes named with the operational reason v0.3.2

11 **BoundWarrant validation requires explicit trust anchor**

75

BREAKING

BoundWarrant.validate() and headers() now require a trust anchor via `trusted_roots=`; previously they built their verifier from the warrant's own issuer, so every warrant validated.

— Exact parameter and the prior insecure behavior explained v0.3.1

12 **enforce_tool_call public API export** NEW 73

Exports `tenuo.enforce_tool_call`, `enforce_tool_call_async`, and `EnforcementResult` as first-class public API for checking a tool invocation against an in-scope warrant without the `@guard` decorator.

Call `enforce_tool_call` directly to check a tool invocation against an in-scope warrant without using the `@guard` decorator.

```python
from tenuo import enforce_tool_call, EnforcementResult

result: EnforcementResult = enforce_tool_call(
    tool="scale_cluster",
    arguments={"cluster": "staging-web", "replicas": 3},
)
if not result.allowed:
    raise PermissionError(result.reason)
```

— Named exports with a runnable usage example v0.3.1

13 **Rust guard denial logging now opt-in** BREAKING 72

Rust guards no longer print denials to stderr by default; call .denial_reporting(DenialReporting::Error) to retain the prior stderr behavior.

Opt back in to stderr denial logging in a Rust guard after the default changed to silent in v0.3.2.

```rust
guard.denial_reporting(DenialReporting::Error)
```

— Exact method call to restore prior behavior, with example v0.3.2

14 **UrlSafe pattern matching hardening** BREAKING 70

`UrlSafe` `allow_domains` now applies to IP literals, so `https://8.8.8.8/` is denied unless that IP is listed. A malformed allow or deny pattern (`"*"`, `"https://x"`, `"example.com:443"`, a bare IPv6 address) in `UrlSafe` denies every call, fails chain verification, and makes the Python UrlSafe() constructor raise.

— Concrete before/after examples, no direct remediation command v0.3.3

15 **Control plane status and audit resilience** NEW 68

Adds `ControlPlaneClient.status` to expose connection state and audit delivery counters (Rust callers use `start_heartbeat_loop_with_status`), plus audit-event buffering through control plane outages that is drained on shutdown.

— Named methods given, no runnable example shown v0.3.3

16 **Denial evidence redaction and Explorer links** IMPROVED 65

Strips warrant tokens and argument values from denial text, logs, and audit reasons across `why_denied`, FastAPI, OpenAI, and `check_constraint`. Adds `why_denied(...).explorer_url` for a pre-filled Explorer link; `WhyDenied.suggestion` no longer embeds the warrant.

— Named surfaces given, usage path is clear but no runnable example v0.3.3

17 **TypeScript SDK: one-use approvals, session export, and runnable examples**

NEW 65

Adds one-use approvals for in-process tools in the TypeScript SDK, which apply only when a `nonceStore` is passed, plus Session.toWireString() and runnable quickstart, `BullMQ`, and `LangChain.js` examples.

— Named methods and gating condition, examples referenced but not shown v0.3.3

18 **TypeScript constraint expressions fail closed on invalid types**

65

BREAKING

TypeScript (`@tenuo/core`, WASM) constraint expressions now fail closed on unknown options and wrong-typed values; `{ kind: 'urlSafe', domains: [...] }` no longer compiles to an unrestricted `urlSafe`.

— Clear before/after example, no migration command v0.3.1

19 **FastAPI argument typing for proof-of-possession** BREAKING 60

FastAPI default path and query arguments are now typed from the endpoint signature; proof-of-possession clients must sign typed values (e.g. `{"limit": 5}`, not `"5"`).

— Concrete example of the typing change, limited mechanism detail v0.3.3

THINNER COVERAGE BELOW

20 **Filesystem jail for Rust guard** NEW — 57

Adds an optional `filesystem` feature on the Rust guard to open an authorized `Subpath` argument through a pinned jail.

— Feature flag and mechanism named, no example v0.3.3

21 **ARM64 Linux Python wheel** NEW — 35

Adds a native `manylinux_2_28_aarch64` Python wheel for ARM64 Linux targets.

— Bare packaging change, no further detail v0.3.2

22 **Capability checks enforced before human approval handlers run** IMPROVED — 25

Calls outside a warrant's capabilities are now denied before any human approval is requested, and Temporal workers verify the chain and proof of possession before calling an approval handler.

— Behavior described but no named API or config v0.3.2 v0.3.3

└──▷ BREAKING ON UPGRADE

! `UrlSafe` `allow_domains` now applies to IP literals: `https://8.8.8.8/` is denied unless that IP is listed.

! A malformed allow or deny pattern (`"*"` , `"https://x"` , `"example.com:443"` , a bare IPv6 address) in `UrlSafe` denies every call, fails chain verification, and makes the Python UrlSafe() constructor raise.

! A2A `require_audience` now defaults to `False` ; set `require_audience=True` or `TENUO_A2A_REQUIRE_AUDIENCE=true` to keep audience binding.

! FastAPI default path and query arguments are typed from the endpoint signature; proof-of-possession clients must sign typed values (e.g. `{"limit": 5}` , not `"5"`).

! responses.create(stream=True) under guard() raises `CFG_004` .

! `TenuoToolGuardrail` in `input_guardrails` raises `CFG_006` ; use guard_agent(), which requires `openai-agents>=0.3.3` .

! A missing or wrong-key PoP is now reported as `POP_VERIFICATION_FAILED` instead of `CONSTRAINT_VIOLATED` in Temporal.

! `_workflow_id` and `_workflow_namespace` are reserved argument names in Temporal and cannot be used as tool arguments.

! Replaying a 0.3.2 Temporal history whose run failed on an off-list signal reports nondeterminism.

! `temporalio>=1.33` is required.

! `WhyDenied.suggestion` no longer embeds the warrant; use `why_denied(...).explorer_url` for the pre-filled Explorer link.

! TypeScript (beta) inputs are stricter; one-use approvals apply only when a `nonceStore` is passed.

- ! Mixed 0.3.1/0.3.2 MCP peers reject proofs containing `null` or integral floats such as `1.0` with `invalid_pop`; upgrade MCP clients and servers together when calls may contain those argument shapes.
- ! Authorizer health probes must move from port 9090 to port 9091; `/health`, `/healthz`, `/ready`, and `/status` no longer bypass authorization on the main listener.
- ! Google ADK `TenuoPlugin` callbacks are now async and keyword-only to match `BasePlugin`; synchronous or positional callback signatures will break.
- ! Rust guards no longer print denials to stderr by default; code relying on stderr denial output must add .denial_reporting(DenialReporting::Error) to retain the prior behavior.
- ! BoundWarrant.validate() and headers() now require a trust anchor via `trusted_roots=`; they previously built their verifier from the warrant's own issuer, so every warrant validated.
- ! GrantBuilder.tool() / tools() are now additive and raise after inherit_all(); the narrowing idiom is retain_tool() / retain_tools().

WAS THIS USEFUL? 👍 👎

ToolHive

Feedback ⓘ v0.56.0 NOTES ↗ RANK

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive v0.56.0 strengthens its auth server with TLS and SPIFFE support, adds declarative DCR client limits with automatic expiry, and ships several smaller security and networking hardening changes.

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

01 **DCR client caps and expiry in CRD** NEW 58

ToolHive now exposes a DCR client cap and register rate limit as CRD fields, letting Kubernetes operators declaratively limit Dynamic Client Registration (DCR) clients; unused DCR clients are automatically expired after 24 hours.

– Mechanism and timing given but exact CRD field names not specified. v0.56.0

02 **TLS and SPIFFE support for auth server** NEW 45

Adds a TLS listener and SPIFFE middleware to the auth server, and starts the auth server in vMCP.

– Names the components added but no config or flag details. v0.56.0

03 **401 response for sessions without upstream token** IMPROVED 45

Returns HTTP 401 when a token session holds no upstream access token.

– Concrete status code and condition named, no endpoint specified. v0.56.0

04 **Decode IPv4-mapped and SIIT IPv6 addresses** IMPROVED 40

Decodes IPv4 addresses embedded in IPv4-compatible and SIIT (IPv4-mapped) IPv6 addresses.

– Specific address formats named but no usage context. v0.56.0

05 **Skip OAuth consent for pre-authorized clients** NEW 30

Lets embedders skip the OAuth consent flow for pre-authorized clients.

– Behavior described but no mechanism or API named. v0.56.0

06 **Native Go fuzz targets for parsers** NEW 25

Adds native Go fuzz targets for untrusted parsers.

— Bare mention of fuzzing addition with no scope detail. v0.56.0

07 **OSS Scanner audit configuration** NEW 10

Adds OSS Scanner audit configuration.

— Only a bare name is given, no further detail. v0.56.0

WAS THIS USEFUL? 👍 👎

coi

Feedback ⓘ v0.15.0 NOTES ↗ RANK

Give the agent a machine. Just not yours.

> coi's first tracked release (v0.15.0) ships a full CLI for running AI coding agents inside isolated Incus containers—shell access, command execution, image building, reusable profiles, and lifecycle management—paired with kernel-level active defense, nftables-based network isolation, and supply-chain guardrails.

└──▷ **WHAT SHIPPED · 10 FEATURES** most completely described first ? what's the number?

01 Reusable agent profiles and tool configuration NEW

93

Adds `coi profile create <name>` and `coi profile list` for managing reusable named container profiles, plus `coi schema profile` and `coi validate profile` for schema-checking profile configs. A `--profile <name>` flag on `coi shell` applies a reusable profile (`config.toml`) specifying image, tool, resource limits, mounts, network mode, and build scripts, with profile inheritance via the `inherits` key in `~/.coi/profiles/<name>/config.toml` . `[tool] permission_mode` set to `bypass` or `interactive` in `~/.coi/config.toml` or `./.coi/config.toml` controls agent autonomy, and six AI coding tools— `claude` , `codex` , `opencode` , `opencode2` , `pi` , and `omp` —are selectable via `[tool] name` .

Lock an agent session to a hardened profile with restricted egress so untrusted code cannot reach arbitrary hosts.

```
$ coi shell --profile hardened
```

— Names every config key, flag, inheritance path and tool option. v0.15.0

02 **Observability commands: monitor, audit, top** NEW 80

Adds `coi monitor` real-time security dashboard for in-session threat visibility, `coi audit` to stream the JSONL threat-event log (pipeable to a SIEM or `jq`), and `coi top` for per-container CPU, memory, and IO metrics resolved to workspace and alias.

Stream the JSONL threat-event log into jq to filter HIGH and CRITICAL events for SIEM ingestion or incident triage.

```
$ coi audit | jq 'select(.severity == "HIGH" or .severity == "CRITICAL")'
```

— Three named commands with formats and a runnable jq example. v0.15.0

03 **Command execution in sandbox** NEW 75

Adds `coi run -- <command>` to execute arbitrary commands in the sandbox with streamed output and propagated exit codes, and `coi run --prompt-name <name>` for headless, unattended agent runs from a predefined prompt, suitable for cron and CI.

Run an agent headlessly against a nightly predefined prompt from cron, with exit code propagation for CI failure detection.

```
$ coi run --prompt-name nightly
```

— Clear mechanism and a runnable cron/CI example. v0.15.0

04 **Network egress isolation modes** NEW 71

Supports `[network] mode` set to `open`, `restricted`, or `allowlist` in `config.toml`, backed by nftables egress control with DNS pinning and per-host port scoping.

— Names config key and enforcement mechanism but no direct command. v0.15.0

05 **Kernel-level active defense** NEW 65

Active defense auto-pauses the container on HIGH-severity events and auto-kills on CRITICAL-severity events, based on kernel-level detection of reverse shells, C2, exfiltration, and DNS tunneling.

— Clear thresholds and threat types, but behavior is automatic, not invoked. v0.15.0

06 **Container lifecycle management commands** NEW 65

Adds `coi list --all` to show active containers and saved sessions, `coi attach` to reconnect to a running session, `coi shutdown` and `coi kill` to stop or force-kill containers, and `coi clean` to remove stopped containers and orphaned resources.

— Five named commands, no usage examples given. v0.15.0

07 **Sandbox shell session command** NEW 60

Adds `coi shell` to drop into an isolated Incus system container with the project mounted at `/workspace`, correct file ownership, and no host credentials exposed.

— Exact command and mount path named, no extended example. v0.15.0

THINNER COVERAGE BELOW

08 **Supply-chain guard protections** NEW 55

Supply-chain guards mount git hooks and IDE configs read-only, enforce protected paths, and block direct commits and pushes to `main`.

— Named mechanisms but automatic, no configurable surface given. v0.15.0

09 **Base container image build command** NEW 45

Adds `coi build` to construct the base container image with Docker, systemd, and package-manager access pre-installed.

— Thin one-line description of what gets pre-installed. v0.15.0

10 **Setup diagnostics command** NEW 23

Adds `coi health` to diagnose the local coi setup.

— Bare name with no detail on what it checks. v0.15.0

WAS THIS USEFUL?  

◆ AI Runtime Defense & Guardrails

doberman-core

Feedback ⓘ v0.19.0 NOTES ↗ RANK

Your AI's guard dog. Doberman sits at runtime, gating every input, output and tool call to stop unsafe or unintended actions before they execute.

> Doberman v0.19.0 adds an alert-only ambient monitor for scoring non-inline agent activity, plus a dry-run mode for log pruning, stronger git branch protection, and secret masking in approval prompts.

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

01 Alert-only ambient activity monitor NEW

85

New `doberman monitor run` and `doberman monitor status` subcommands score non-inline agent activity through the same risk engine used for inline gating, but operate in alert-only mode and never enforce or block.

Run the ambient monitor to observe agent activity without blocking anything — useful for baselining a new project before enabling enforcement.

```
$ doberman monitor run
```

— Named subcommands and runnable example, mechanism and scope stated v0.19.0

02 Dry-run mode for decision-log pruning NEW

80

Adds a `--dry-run` flag to `doberman decision-log-prune` that lists what a prune would delete without committing the deletion, letting operators verify retention policy before data is removed.

Preview which decision-log entries would be removed before committing a prune, so you can verify retention policy before data is gone.

```
$ doberman decision-log-prune --dry-run
```

— Exact flag and runnable command, clear before/after behaviour v0.19.0

03 Doctor check for disabled Claude Code hooks IMPROVED

65

Extends `doberman doctor` to fail when Claude Code's `disableAllHooks` setting turns every hook off, catching a configuration state that would silently disable enforcement.

— Names the config key and command but no run example v0.19.0

04 **Protected branch-delete blocking** IMPROVED 60

Blocks protected-branch deletion via `git push --delete` or the `:<branch>` push syntax, using the same protection mechanism already applied to force-pushes.

— Names exact git syntaxes blocked but no command to invoke v0.19.0

THINNER COVERAGE BELOW

05 **Secret masking in approval notifications** IMPROVED 40

Masks bearer tokens and password flags in the approval prompt and phone push notifications so sensitive values aren't exposed when a human reviews a pending action.

— States what is masked but no mechanism or location detail v0.19.0

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ

Obsidian 1.15.0 Desktop (Early access) (2026-10-22) BLOG ↗ RANK

Obsidian is a note-taking app that lets you write and organize Markdown files stored locally in a vault.

> Obsidian 1.15.0 replaces its bundled Mermaid diagram engine with the new Pele renderer for themed, zoomable diagrams, while the mobile app adds voice-note capture with live transcription, PDF export, and Siri/Spotlight Quick Capture integration.

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

01 **Mermaid diagrams rendered via Pele engine** BREAKING

85

Replaces the bundled Mermaid rendering engine with the new Pele library, supporting Mermaid `12.1.0` syntax, theme-matched colors and fonts, responsive sizing, and zoom/pan for scaled-down diagrams, with no code-execution sandbox required. Adds `--mermaid-*` CSS variables for colors, corner radius, and font weights to style diagrams, exposes `window.pele` for new code once loadMermaid() resolves, and updates loadMermaid() and `window.mermaid` to draw with Pele, supporting initialize(), render(), parse(), and mermaidAPI.getConfig(), setConfig(), render(). Breaking: the Show Mermaid diagrams in notes setting has been removed since diagrams always render via Pele now, and SVG output carries Pele's `.pele-*` classes instead of Mermaid's, breaking plugins relying on the old class names.

— Full mechanism, CSS hooks and API surface named, plus breaking impact

Obsidian 1.15.0 Desktop (Early access) …

02 **Quick Capture enhancements** NEW

80

Adds voice note recording and live transcription to iOS Quick Capture, including from the Lock Screen, with recordings saveable to your vault (requires iOS 26); adds Quick Capture to Spotlight actions and Siri, enabling invocation via phrases like Open Quick Capture in Obsidian; and setting Note Title to Automatic in Quick Capture uses the first 80 characters of the capture as the note title.

Capture a voice note hands-free from the iOS Lock Screen and have it transcribed and saved directly to your vault.

> 📍 On your iOS 26 device Lock Screen, tap the Quick Capture widget › tap the microphone icon › speak your note › tap Stop. The transcription is saved

automatically to your vault.

Open Quick Capture instantly without unlocking your phone by asking Siri.

📍 Say: Hey Siri, Open Quick Capture in Obsidian.

— Exact steps given with platform requirement and save behaviour Exact invocation phrase given as a runnable action Named setting and exact character limit given

Obsidian 1.15.0 Mobile (Early access) (…

03 **Vim mode command and regex additions** IMPROVED 70

Updates Vim mode with new commands including `g?`, `:put`, search patterns in `:` ranges (e.g. `:/foo/,/bar/d`), Vim-style regexes via `:set nopcre`, and `<A-…>` mappings on macOS.

— Multiple exact Vim commands and flags named Obsidian 1.15.0 Desktop (Early access) …

04 **Filter reordering and limits in Bases** IMPROVED 65

Adds dragging to reorder filters in Bases and a per-filter menu to add filters or move them between All views and This view; indicates when a result limit hides matches in search (e.g. Showing 10 of 40); and allows formula fields to be resized up to the width of the window.

— Three concrete UI behaviours named with example wording

Obsidian 1.15.0 Desktop (Early access) …

05 **URI/CLI commands to append or prepend headings** NEW 60

Adds URI/CLI commands to append or prepend content to a heading, with control over line breaks around the appended or prepended content.

— Names the capability but not exact command syntax

Obsidian 1.15.0 Desktop (Early access) …

06 **Accent-insensitive fuzzy and simple search** IMPROVED 60

Updates `prepareFuzzySearch` and `prepareSimpleSearch` to ignore accents on Latin and Greek letters, so a query like `cafe` matches `café`.

— Named functions and example match behaviour given

Obsidian 1.15.0 Desktop (Early access) …

THINNER COVERAGE BELOW

07 **Closing formatting classes in Live Preview** NEW 55

Adds `cm-formatting-*-end` class (e.g. `cm-formatting-strong-end`) to closing formatting tokens in Live Preview, matching the existing `cm-formatting-link-end` and `cm-formatting-highlight-end`.

— Named CSS classes but narrow theming use case `Obsidian 1.15.0 Desktop (Early access) …`

---

08 **Lazy-loaded PIXI graph library** BREAKING 55

Loads the `PIXI` global on first graph open instead of at start-up, and exposes loadPixi() for code that needs it earlier; code that accessed `PIXI` at startup must now call loadPixi() first.

— Names the function and migration step for plugin developers

`Obsidian 1.15.0 Desktop (Early access) …`

---

09 **PDF export for notes on mobile** NEW 40

Adds PDF export for notes on mobile.

Export a finished note as a PDF from your phone to share or archive it.

> 📍 Open the note › tap the '...' menu › select Export as PDF.

— Clear UI steps given but no format options described

`Obsidian 1.15.0 Mobile (Early access) (…`

---

10 **Create Sync vault from Manage Vaults** IMPROVED 40

Manage Vaults can now create a new Obsidian Sync vault from the same screen.

— Names the screen but no further mechanism `Obsidian 1.15.0 Mobile (Early access) (…`

---

11 **Tab dragging between tab groups** IMPROVED 30

Enables tab dragging between tab groups in the Tab switcher.

— Bare description with UI location only `Obsidian 1.15.0 Mobile (Early access) (…`

---

12 **Faster community plugin startup** IMPROVED 10

Community plugins now load faster when Obsidian starts.

— No mechanism, numbers or scope given `Obsidian 1.15.0 Desktop (Early access) …`

└──▷ BREAKING ON UPGRADE

! The Show Mermaid diagrams in notes setting has been removed; diagrams now always render via Pele.

! Obsidian no longer bundles Mermaid: loadMermaid() and `window.mermaid` now draw with Pele and their SVG output carries Pele's `.pele-*` classes, not Mermaid's — plugins relying on Mermaid's SVG class names will break.

! The `PIXI` global is no longer loaded at start-up; code that accessed it at startup must now call loadPixi() first.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ announced · Oct 1, 2026 PRESS ↗ RANK

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

> Cloudflare launched Basin, a generally available serverless data platform built on Apache Iceberg and R2 for ingesting, storing, cataloging, and querying analytics data without infrastructure management.

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

01 **Basin serverless Iceberg data platform** NEW 65

Cloudflare Basin is a generally available serverless data platform built on Apache Iceberg and Cloudflare R2 that lets teams ingest, store, catalog, and query analytical data without provisioning or managing infrastructure. It ships with named components — Basin Pipelines, Catalog, and SQL — intended to replace complex setups such as AWS S3 and Athena with a unified, serverless architecture. Because all data is stored in the open Apache Iceberg format, any compatible tool can read it directly with no copying, converting, or vendor lock-in, and Cloudflare eliminates data egress fees when querying across different clouds, platforms, or regions.

— Names components and format but no commands or UI path given `launch-20261001-1074b374`

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe Feedback on this issue