THE CYBER TOOLCHAIN

NO. 035

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 4, 2026 · EVERY DAY

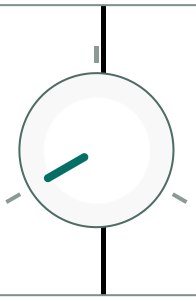
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 40 tools

OFFENSIVE SECURITY

vim ? help

◆ AI-Powered Offensive Security

Dreadnode

Feedback ⓘ 2 RELEASES · SEEN 2026-08-04 NOTES ↗

> Dreadnode shipped a major AI red-teaming expansion — a `dn airt run-classifier` CLI for traditional-ML attacks (evasion, extraction, membership inference, model inversion) and an ATLAS multi-agent attack framework — alongside an agent memory UI, session monitoring overhaul, five data-integration connectors, and air-gapped hub sync commands.

└──▷ WHAT SHIPPED · 27 FEATURES most completely described first ? what's the number?

01 **`--project-memory-preload-limit` flag for session memory preload**

IMPROVED 90

New `--project-memory-preload-limit` flag sets the number of memories preloaded at session start (range: 1–200), replacing a hardcoded default of 20.

Increase the number of project memories available at session start beyond the default of 20 to give agents broader context on large engagements.

```
$ dn ... --project-memory-preload-limit 100
```

— Exact flag, range, default, and example command

Traditional-ML red teaming suite — evas…

02 **`dn airt run-classifier` CLI for traditional-ML attacks** NEW

85

New `dn airt run-classifier` CLI command runs classic-ML attacks — evasion, extraction, membership inference, and model inversion — directly against a predict endpoint without SDK code.

— Exact CLI command and target named, no runnable example given

Traditional-ML red teaming suite — evas…

03 **Traditional-ML attack suite in AIRT** NEW 85

Adds four traditional-ML attack families (evasion, extraction, membership inference, model inversion) to AIRT, covering 20+ named attacks — including equation-solving, Jacobian, copycat, knockoff, PWWS, BAE, TextBugger, threshold, and label-only — across tabular, image, and text modalities, with per-finding metrics, adversarial trajectory traces, and MITRE ATLAS/SAIF/NIST compliance mappings. Model inversion attacks

reconstruct class-representative inputs via black-box queries to surface privacy risks, completing the fourth traditional-ML attack category.

— Extensive named attacks and mappings, but no direct invocation

Traditional-ML red teaming suite — evas…

04 **`/report-bug` command in TUI** NEW 85

Adds `/report-bug` command (or `r` in console) to the TUI, assembling a privacy-reviewed diagnostic bundle and submitting it without losing logs after a crash.

File a bug report from the TUI immediately after a crash without losing session logs.

```
$ /report-bug
```

— Exact command plus example demonstrating usage Agent memory, triage, and session monit…

05 **`dn evaluation create --agent` flag pins evaluation agent**

NEW 80

New `dn evaluation create --agent <name>` flag specifies which agent runs in an evaluation; eval output reports which agent actually ran.

Pin a specific agent to an evaluation run so results are reproducible and the run report identifies which agent executed.

```
$ dn evaluation create --agent <name>
```

— Exact flag with runnable example command Traditional-ML red teaming suite — evas…

06 **`dn capability sync` and `dn task sync` for air-gapped hub loading**

NEW 80

Adds `dn capability sync` and `dn task sync` CLI commands to pull hub content from SaaS (or another on-prem) into a local directory and push it to a self-hosted instance, enabling air-gapped on-prem loading.

— Named CLI commands with clear directionality, no example given

Agent memory, triage, and session monit…

07 **Agent Output triage workspace with export and deep-linking**

IMPROVED 75

Agent Output now functions as a severity-first triage workspace — searches, sorts, and filters an engagement's findings; each finding can be exported to CSV, deep-linked, or

copied into a report.

— Names CSV export and actions, no exact endpoint/path

Traditional-ML red teaming suite — evas…

08 **AST-based JS/TS analysis via `ast-grep`** NEW 75

Adds AST-based JS/TS static analysis via `ast-grep` in the web-security runtime, removing the need for semgrep, CodeQL, or a jxscout license.

— Names tool and replaced alternatives, no exact invocation

Agent memory, triage, and session monit…

09 **ATLAS multi-agent attack framework** NEW 70

Adds the ATLAS multi-agent attack framework for running adaptive campaigns against deployed multi-agent environments across three attack surfaces, with executed tool calls captured per-trial and surfaced in findings across the UI, API, and SDK.

— Named framework and surfaces, no exact command given

Traditional-ML red teaming suite — evas…

10 **Richer privacy-attack metrics in findings** IMPROVED 70

Extraction findings include KL divergence and accuracy retention; membership inference includes attack accuracy, precision, and recall; model inversion includes reconstruction quality.

— Specific metrics named per attack type, no action needed

Traditional-ML red teaming suite — evas…

11 **`llm-wiki` and `analysis-ledger` capabilities** NEW 70

Adds the `llm-wiki` capability (compiles source material into linked wiki pages) and `analysis-ledger` capability (maintains a resumable code-analysis frontier with evidence-bearing claims).

— Named capabilities with described mechanism, no usage example

Agent memory, triage, and session monit…

12 **Session monitoring redesign with activity tracking and transcript pagination**

IMPROVED 70

Session monitoring redesigned: adds a Last Msg relative-activity column, highlights sessions active in the last 15 minutes, and adds forward/backward transcript pagination with tailing mode.

— Names UI columns and modes, no exact path Agent memory, triage, and session monit…

13 **Web-security initial-recon skill via ProjectDiscovery toolchain**

NEW 65

New web-security initial-recon skill provides structured reconnaissance for domain, wildcard, ASN, IP-list, and CIDR entry points using the ProjectDiscovery toolchain.

— Names toolchain and entry-point types, no exact command

Traditional-ML red teaming suite — evas…

14 **Five self-contained AIRT demo notebooks** NEW 65

Five self-contained AIRT demo notebooks cover model evasion, extraction, membership inference, jailbreaks, and multi-agent attacks — runnable via `dn login` with no provider keys required.

— Names entry command, notebooks not individually detailed

Traditional-ML red teaming suite — evas…

15 **Sessions group-by control (None / Workflow / Capability / Agent / Source)**

IMPROVED 65

Sessions group-by control (None / Workflow / Capability / Agent / Source) is now available to all orgs, with workflow runs collapsed by default.

— Lists exact options, no config key Agent memory, triage, and session monit…

16 **Automatic recovery of Agent Output items from span backups**

IMPROVED 65

Agent Output items from failed `report_item` POSTs are automatically recovered from span backups when a session is frozen.

— Names failure mode and recovery mechanism, no user action needed

Agent memory, triage, and session monit…

17 **Memories page under Agents** NEW 60

New Memories page under Agents lets operators view, inspect, and manage project memory entries, including version history and diffs.

— Describes UI page and capabilities, no config path

Agent memory, triage, and session monit…

18 **Clickable stat-card filters and sortable Updated column in Agent Outputs**

IMPROVED 60

Stat cards for Needs review and Critical + High on the Agent Outputs page are now clickable filter toggles; the items grid gains a sortable Updated column.

— Names UI elements, no deeper mechanism Agent memory, triage, and session monit…

THINNER COVERAGE BELOW

19 **Executive/CISO summary panel on assessment view** NEW

55

New executive/CISO summary panel on the assessment view shows attack-surface posture, query cost to compromise, risk score, and a plain-English headline for stakeholder reporting.

— Lists panel contents, no navigation path given Traditional-ML red teaming suite — evas…

20 **Runtime component health recording in evaluations** NEW

55

Evaluation runtime now records the health of all runtime components (broken tools, skills, policies) at run time, so degraded capabilities are visible after the fact.

— Describes mechanism but no named config/API surface

Traditional-ML red teaming suite — evas…

21 **'Inactive' filter in Agent Sessions** NEW 55

Adds an Inactive filter to Agent Sessions to isolate sessions idle past 15 minutes but not yet archived.

— Describes threshold and UI location, no exact path given

Agent memory, triage, and session monit…

22 **Five data-integration connectors (Atlassian, Azure DevOps, GitHub, GitLab, Linear)**

NEW 50

Adds five data-integration connectors — Atlassian, Azure DevOps, GitHub, GitLab, and Linear — as platform capabilities.

— Names five connectors but no config or usage detail

Agent memory, triage, and session monit…

23 **Finding triage (status and notes) in Agent Output detail panel**

NEW 50

Enables finding triage (status and notes) directly from a finding's detail panel in Agent Output.

— Describes UI action, no example given Agent memory, triage, and session monit…

24 **Table view and unified toolbar on Hub browse pages** IMPROVED

50

Models and Datasets Hub browse pages add a table view alongside cards, with a unified toolbar offering separate sort and filter controls across all four Hub browse pages.

— Describes UI change, lacks navigation path Agent memory, triage, and session monit…

25 **Organization picker in new workspace dialog** NEW 50

New workspace dialog includes an organization picker to create a workspace in any org without switching context.

— Simple UI addition, minimal detail Agent memory, triage, and session monit…

26 **Dedicated media gallery for session tool outputs** NEW

45

New dedicated media gallery collects all images and other media logged by a session's tools in one place for reviewers.

— Brief description, minimal specificity Traditional-ML red teaming suite — evas…

27 **AI Red Teaming Learning Guide documentation** NEW 45

New AI Red Teaming Learning Guide adds eight documentation pages covering all four traditional-ML attack categories and generative AI attacks, with runnable SDK, TUI, and CLI examples.

— Describes doc scope, no specific commands listed

Traditional-ML red teaming suite — evas…

WAS THIS USEFUL?  

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ 1 RELEASE · SEEN 2026-08-04

NOTES ↗

Horizon3 NodeZero is an autonomous security tool that continuously identifies and exploits vulnerabilities in networks to assess organizational risk.

> NodeZero's biggest update centers on Rapid Response's new asset-level exploitability view and the introduction of Tripwires honeytoken decoys, alongside a rebuilt External Pentest Scopes model, a Horizon3-hosted MCP Server with OAuth 2.1 auth, and a broad set of MSP, Vulnerability Management Hub, and integration improvements.

└──▷ WHAT SHIPPED · 19 FEATURES most completely described first ? what's the number?

01 **Rapid Response asset-level exploitability** IMPROVED

83

Rapid Response now surfaces personalized asset-level exploitability status, remediation guidance, and verification feedback to distill vulnerability noise. The Rapid Response Details page shows an Asset Breakdown, CVE or H3 IDs, and a Relevant Assets table with filtering, Action Items and Fix Actions guidance, and a Take Action button offering options to run a Rapid Response Test, add a note, create a ticket, or run an external or internal test.

- UI path named but no API or flag given. snapshot-20260804

02 **Tripwires honeytoken decoys** NEW

75

Tripwires deploys silent decoy credentials — real credentials generated by Horizon3 that blend into the environment — to assets flagged as high-risk by NodeZero autonomous pentest results; tokens stay silent until triggered, then alert Horizon3's cloud infrastructure. Tripwires now also supports CSV export, credential rotation, and improved Active Directory validation.

- Names export and rotation features but no config path given. snapshot-20260804

03 **`No Longer Found` status in VMH** NEW

65

A new No Longer Found status was added to the Vulnerability Management Hub (VMH), with enhanced sorting and filtering by Host, Weakness ID, and Downstream Impact.

- Names exact status and filters but no API. snapshot-20260804

04 **External Pentest rebuilt with Scopes**

60

BREAKING

External Pentest was rebuilt, replacing Asset Groups with reusable Scopes, improving bulk asset management and uploads, and enhancing support for large external attack surfaces.

— Names replacement of Asset Groups but no migration steps. snapshot-20260804

05 **Shorter asset-return window on client deletion** BREAKING

60

Client deletion now returns scanned asset allocation to the parent within 30 days, down from a previous 60-day hold.

— Exact timeframe change stated, no further action needed. snapshot-20260804

THINNER COVERAGE BELOW

06 **Runner Command History and Credentials tabs** IMPROVED

55

Runner visibility was upgraded with Command History, Credentials, and Past Pentests tabs.

— Names three tabs but no navigation detail. snapshot-20260804

07 **Sankey diagram pentest summaries with CISA KEV badges**

IMPROVED

55

Pentest summaries were redesigned with Sankey diagrams carrying CISA KEV badges to highlight actively exploited vulnerabilities.

— Describes visualization but no navigation path. snapshot-20260804

08 **Horizon3-hosted MCP Server** NEW

50

A new MCP Server hosted by Horizon3 provides OAuth 2.1 authentication and multi-user support; FedRAMP High customers can now access the Horizon3-hosted MCP Server.

— No endpoint or setup steps given. snapshot-20260804

09 **Real-Time View redesign** IMPROVED

50

Real-Time View (RTV) was redesigned with a visual timeline, event-log summaries, and granular event details with flexible filtering to help distinguish NodeZero testing activity from real attacks.

— Describes UI changes without exact navigation steps. snapshot-20260804

10 **Linux privilege escalation implants and NetBIOS support** NEW

50

Linux privilege escalation implants and NetBIOS support were added for Active Directory operations.

— Names capabilities but no usage detail. snapshot-20260804

11 **MSP license and access management pane** IMPROVED 45

MSP and license-management improvements let parent accounts manage clients' product access, asset allocations, and custom branding from a single pane.

— General description, no specific settings named. snapshot-20260804

12 **Automated weakness retrieval for Cortex XSOAR** NEW 45

Automated weakness retrieval was added for the Palo Alto Cortex XSOAR integration.

— Names integration only, no setup steps. snapshot-20260804

13 **Co-branded report downloads and branding enforcement** IMPROVED 45

Co-Branding now supports downloading co-branded sample reports in groups and enforcing parent branding across child accounts.

— Describes features but no exact controls named. snapshot-20260804

14 **Threat Actor Intelligence tier expansion** IMPROVED 40

Threat Actor Intelligence, normally an Elite-tier feature, was made available to Pro and Core accounts for a limited time.

— Notes tier access but no duration or scope specifics. snapshot-20260804

15 **NodeZero Federal external attack support** NEW 35

NodeZero Federal external attack support was enabled across all federal environments.

— States rollout scope but no mechanism detail. snapshot-20260804

16 **Account names in multi-tenant email alerts** IMPROVED 30

Multi-tenant email notifications now include account names to help MSPs and Enterprises triage alerts across subclients.

— Simple field addition, no config path given. snapshot-20260804

17 **Cross-domain credential injection** NEW 30

Cross-domain credential injection was added for more realistic modeling and validation of multi-domain attacks.

— Names capability without mechanism detail. snapshot-20260804

18 **Subclients draw from parent asset pools** IMPROVED 30

Subclients can now draw from parent asset pools for more flexible MSP/MSSP resource management.

— General description, no configuration steps. snapshot-20260804

19 **Expanded API for subclient management** IMPROVED 25

Expanded API support was added for subclient management.

— No endpoint names or examples given. snapshot-20260804

WAS THIS USEFUL? 👍 👎

Rapid7 Metasploit Pro

Feedback ⓘ

August 3, 2026: 5.1.0-2026073001 NOTES ↗

Metasploit Pro is a commercial penetration testing framework for identifying, exploiting, and validating security vulnerabilities in systems.

Metasploit Pro 5.1.0-2026073001 adds MALLEABLEC2 HTTP(S) shaping, 8 new MCP tools, RISC-V payloads, expanded CertificateTrace/KerberosTicketTrace, and an HTTP-to-SMB NTLM relay module.

└──▷ **TRY IT**

Shape Meterpreter HTTPS traffic using a Malleable C2 profile to blend into expected network baselines during an engagement.

```
$ use exploit/multi/handler
set PAYLOAD windows/x64/meterpreter/reverse_https
set LHOST 10.0.0.1
set LPORT 443
set MALLEABLEC2 /opt/profiles/jquery.profile
run
```

Start the MCP server with dangerous module-execution actions enabled so MCP tool integrations can run and check modules via RPC.

```
$ msfmcpd --enable-dangerous-actions
```

Adds `MALLEABLEC2` option to all Meterpreter HTTP(S) payloads (Windows, Java, Python, PHP, Linux) to load a standard Malleable C2 profile and reshape HTTP(S) traffic.

Adds `--enable-dangerous-actions` flag to the `msfmcpd` CLI to opt in to dangerous actions across the eight new MCP tools that wrap existing RPC endpoints for module and session control.

Adds eight new MCP tools wrapping existing RPC endpoints for module and session control, backed by strict per-tool input validation.

Adds `run_uuid` job tracking across all module types when executed through RPC endpoints to support MCP tool integrations.

Extends `CertificateTrace` peer cert tracing (using the existing `off` / `metadata` /

`full` enum) to LDAP over TLS and RDP.

+18 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Portkey

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

Portkey is an AI gateway providing observability, security, and reliability features for LLM applications with multi-model support.

> Portkey's March 2026 update adds vault-backed Secret References for credential management, weekly and endpoint-scoped rate limits, cross-cloud GCS log delivery via GCP Workload Identity Federation, and four new guardrail integrations alongside expanded provider support for DeepInfra, DeepSeek, Bedrock, Vertex AI, and Azure AI Foundry.

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

01 Vault-backed Secret References for credentials NEW

90

Credentials stored in AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault can be mapped to integrations and virtual keys via `secret_mappings`, with Portkey fetching values at runtime instead of storing them directly.

Keep LLM provider credentials out of Portkey by pointing a virtual key at a secret already stored in AWS Secrets Manager.

```json
{
  "virtual_key": "my-openai-key",
  "secret_mappings": {
    "api_key": "arn:aws:secretsmanager:us-east-
1:123456789012:secret:prod/openai-api-key"
  }
}
```

— Named config field, three vault providers, and runnable JSON example. snapshot-20260804

02 Cross-cloud GCS log delivery via GCP Workload Identity Federation

NEW

90

Adds `GCP_WIF_AUDIENCE` and `GCP_WIF_SERVICE_ACCOUNT_EMAIL` environment variables to enable GCP Workload Identity Federation (`gcs_assume` style) for GCS log delivery when the gateway runs in AWS, eliminating static cross-cloud credentials.

Route GCS log writes through GCP Workload Identity Federation when the Portkey gateway runs in AWS, avoiding long-lived GCP credentials in your AWS environment.

```
$
GCP_WIF_AUDIENCE=//iam.googleapis.com/projects/123456/locations/global/workloadIdentityPools/my-pool/providers/my-provider \
GCP_WIF_SERVICE_ACCOUNT_EMAIL=portkey-log-writer@my-project.iam.gserviceaccount.com \
docker run portkey/gateway ...
```

— Exact env vars plus a runnable docker command. snapshot-20260804

03 **Zscaler AI Guard guardrail integration** NEW 80

Applies Zscaler Detection Policies to LLM inputs and outputs using `beforeRequestHook` and `afterRequestHook` hooks, requiring a `policyId` and supporting a configurable `timeout` (default 10000 ms).

— Named hooks and fields but no worked example. snapshot-20260804

04 **Vertex AI enterprise web search grounding** NEW 80

Adds Vertex AI enterprise web search grounding via `enterpriseWebSearch` / `enterprise_web_search` with separate cost attribution, Portkey metadata mapping to Vertex resource labels, and AWS-to-Vertex access via `GCP_WIF_AUDIENCE` and `GCP_WIF_SERVICE_ACCOUNT_EMAIL`.

— Names params and env vars but lacks a worked example. snapshot-20260804

05 **DeepSeek streaming and reasoning options** IMPROVED 75

Adds `stream_options` support for DeepSeek `deepseek-chat`, and maps `reasoning_effort` to thinking mode for `deepseek-reasoner`, returning `reasoning_content` in streams.

— Precise param and model names, but no example call. snapshot-20260804

06 **Weekly and endpoint-scoped rate limits** NEW 70

Adds `rpw` (requests-per-week) budget and usage windows so rate limits can align with weekly planning cycles, alongside endpoint-scoped limits that let chat completions, embeddings, and admin routes carry independent caps.

— Names the config unit but no runnable example of setting it. snapshot-20260804

07 **Akto Agentic Security guardrail integration** NEW 65

Scans LLM inputs and outputs for prompt injection and sensitive data leakage via hooks with a configurable `timeout` (default 5000 ms).

— Names the mechanism and default timeout, no example given. snapshot-20260804

08 **DeepInfra provider support** NEW 65

Adds DeepInfra provider support with tool calling via `tools`, `tool_choice`, and `parallel_tool_calls`, plus completions and embeddings endpoints.

— Names parameters and endpoints, no runnable example. snapshot-20260804

09 **Azure AI Foundry rerank support for Cohere models** NEW 65

Adds Azure AI Foundry rerank support for Cohere rerank models (e.g. `cohere.Cohere-rerank-v4.0-pro`), with the gateway stripping the `cohere.` prefix for the provider.

— Names model and prefix-stripping behaviour, no example call. snapshot-20260804

10 **OpenTelemetry GenAI semantic spans updated to semconv 1.40.0** IMPROVED 65

OpenTelemetry GenAI semantic spans now follow semconv 1.40.0 for inference and embeddings, with OTEL exporter support for guardrail flows and custom resource attributes.

— Names the spec version and scope but no config example. snapshot-20260804

THINNER COVERAGE BELOW

11 **Custom host for Bedrock guardrail plugin** IMPROVED 55

Adds `customHost` configuration to the Bedrock guardrail plugin so evaluations can target private or regional Bedrock-compatible endpoints instead of the default public Bedrock URLs.

— Names the config field but no example of use. snapshot-20260804

12 **Archived workspaces in analytics and saved filters by slug** IMPROVED 45

Organization admins and owners can include archived workspaces in graphs and summaries; saved filters now accept workspace slugs alongside IDs.

— Describes the change with no config example. snapshot-20260804

13 **Required metadata key-value guardrails** NEW 40

Adds required metadata key–value guardrails that block requests at the gateway before they reach the model if required metadata fields are missing or invalid.

— Describes behaviour only, no named field or example. snapshot-20260804

14 **Bedrock batch inference for embeddings** IMPROVED 40

Extends Bedrock batch inference to support embeddings in addition to chat completions.

— Brief statement with no mechanism or example. snapshot-20260804

WAS THIS USEFUL? 👍 👎

Cisco skill-scanner

Feedback ⓘ 1 RELEASE · 2026-08-03

NOTES ↗ CODE ↗

The skill-scanner is a CLI security scanner for AI agent skills.

> skill-scanner shipped a broad set of new capabilities this window: OSV.dev dependency vulnerability scanning, a demote-only LLM adjudicator for false positives, a pre-commit hook with incremental CI scanning, a `scan-repo` subcommand for scanning GitHub repositories directly, simultaneous multi-format reporting, and the ATR/PromptGuard rule packs alongside expanded static-analysis coverage.

└──▷ **WHAT SHIPPED · 22 FEATURES** most completely described first ? what's the number?

01 **Pre-commit hook with incremental CI scanning** NEW 93

`skill-scanner-pre-commit --install` installs a Git pre-commit hook that maps each changed path to its nearest parent `SKILL.md`, scans each affected skill once, and blocks commits at the configured severity; `--scan-all` scans every configured skill instead of only affected ones. Incremental CI scans are supported via `pre-commit run skill-scanner --from-ref "$BASE_SHA" --to-ref "$HEAD_SHA"`, scanning only files changed between two revisions.

Install the pre-commit hook so commits are blocked when changed skills exceed the configured severity threshold.

```
$ skill-scanner-pre-commit --install
```

Run incremental skill scans in CI between a base and head commit, limiting work to only the changed files in a PR.

```
$ pre-commit run skill-scanner --from-ref "$BASE_SHA" --to-ref
"$HEAD_SHA"
```

— Names exact commands and flags with clear behaviour. 2.0.13

02 **Adjudicator demote-only false-positive review** NEW 90

`--adjudicate` flag / `adjudicator.enabled` and `adjudicator.min_fp_confidence` policy keys enable a demote-only LLM review that can reduce deterministic HIGH/CRITICAL literal-regex false positives to INFO without ever promoting a finding; disabled by default and records decisions for audit.

`SKILL_SCANNER_ADJUDICATOR_LLM_MODEL` overrides the model used specifically for this path (e.g. `vertex_ai/gemini-2.5-pro`).

Enable adjudication with a confidence threshold to automatically demote likely false positives while keeping an audit trail.

```yaml
adjudicator:
  enabled: true
  min_fp_confidence: 4
```

Review deterministic HIGH/CRITICAL regex matches with LLM adjudication before committing — reduces false-positive noise without a full LLM enrichment pass.

```
$ skill-scanner scan ./my-skill --adjudicate
```

Enable the adjudicator alongside LLM enrichment and the meta-analyzer to reduce false positives before results are promoted.

```
$ skill-scanner scan ./skill --use-llm --enable-meta --adjudicate
```

— Named flags, config keys, env var, and exact behaviour. 2.0.13

---

03 **build_analyzers factory for programmatic analyzer assembly**

NEW 88

`build_analyzers` in `skill_scanner.core.analyzer_factory` assembles optional analyzers (`use_behavioral`, `use_llm`, `use_osv`, `llm_consensus_runs`) programmatically, matching CLI and API parity; the demote-only adjudicator pass can be enabled via `policy.adjudicator.enabled = True` before constructing `SkillScanner`.

Assemble all optional analyzers including behavioral, LLM, and OSV checks with 3 consensus runs — matching what the CLI and API produce.

```python
from skill_scanner.core.analyzer_factory import build_analyzers
analyzers = build_analyzers(
    policy,
    use_behavioral=True,
    use_llm=True,
    use_osv=True,
    llm_consensus_runs=3,
)
scanner = SkillScanner(analyzers=analyzers, policy=policy)
```

Enable the demote-only adjudication pass before scanning to apply adjudicator logic programmatically.

```python
policy.adjudicator.enabled = True
scanner = SkillScanner(analyzers=analyzers, policy=policy)
```

— Exact module path, function signature, and runnable code. product docs

04 **ATR and PromptGuard rule packs** IMPROVED 88

`--rule-packs` selects named packs; the `atr` (Cisco AI Threat Research) pack expanded from 314 to 712 rules (pack version 3.5.6), and the `promptguard` pack ships 26 rules covering PII, provider secrets, and LLM/Markdown exfiltration. `--rule-packs list` enumerates all bundled packs.

Run both the ATR and PromptGuard rule packs together to maximise detection coverage in a single scan.

```
$ skill-scanner scan ./skill --rule-packs atr promptguard
```

Run the full analyzer stack — behavioral, OSV vulnerability DB, LLM semantic, and metadata — with ATR and PromptGuard rule packs to maximize detection coverage on a local skill.

```
$ skill-scanner scan /path/to/skill --use-behavioral --use-osv --use-llm --enable-meta --rule-packs atr promptguard
```

— Exact rule counts, pack names, and version given. 2.0.13

05 **Simultaneous multi-format report output** NEW 86

Repeating `--format` (e.g. `--format json --format sarif`) produces multiple report formats in a single scan run, paired with per-format output path flags `--output-json`, `--output-sarif`, `--output-markdown`, `--output-html`, and `--output-table`, eliminating the need to scan twice for different report formats. `--render-markdown` / `--no-render-markdown` override terminal Markdown rendering.

Run a full local skill scan and emit both a machine-readable JSON report and a SARIF file for upload to a code-scanning platform in one pass.

```
$ skill-scanner scan ./my-skill --format json --format sarif \
  --output-json results.json --output-sarif results.sarif
```

Produce both JSON and SARIF reports from a single scan pass — ideal for feeding results simultaneously into a SIEM and a code-review tool like GitHub Advanced Security.

```
$ skill-scanner scan /path/to/skill --format json --format sarif
--output-json results.json --output-sarif results.sarif
```

— All flags named with runnable command shown. product docs

06 **scan-repo subcommand for GitHub repositories** NEW 85

New `scan-repo` subcommand clones a public GitHub repository (accepts `owner/repo` shorthand or a full HTTPS URL) into a temp directory, recursively discovers and scans every skill in it, then removes the clone; `--no-recursive` disables recursive discovery (recursive is on by default).

Scan a public GitHub repo for skills, including OSV dependency advisories and overlap checks, without cloning manually.

```
$ skill-scanner scan-repo https://github.com/owner/repo --use-osv
--check-overlap
```

Scan a remote GitHub repository for all skills without cloning it locally — useful in CI pipelines that operate against upstream repos.

```
$ skill-scanner scan-repo owner/repo
```

— Runnable command and exact flags given. product docs

07 **OSV.dev dependency vulnerability scanning** NEW 80

`--use-osv` flag / `use_osv` config field queries the OSV.dev advisory database for vulnerabilities in exactly pinned Python dependencies, requiring network access but no API key or additional package.

Enable OSV.dev vulnerable dependency checking and trigger specificity analysis alongside LLM consensus across 3 runs.

```yaml
use_osv: true
use_trigger: true
llm_consensus_runs: 3
```

Scan a public GitHub repo for skills, including OSV dependency advisories and overlap checks, without cloning manually.

```
$ skill-scanner scan-repo https://github.com/owner/repo --use-osv
--check-overlap
```

Check pinned Python dependencies against the OSV database without any API key as part of a CI pipeline.

```
$ skill-scanner scan ./skill --use-osv
```

Scan a skill's dependencies against the OSV advisory database to catch known vulnerabilities in pinned or unpinned packages.

```
$ skill-scanner scan /path/to/skill --use-osv
```

— Named flag/config with exact scope and requirements. 2.0.13

08 **OpenAI-compatible LLM endpoint routing** NEW 78

An `openai-compatible` LLM provider shortcut plus `SKILL_SCANNER_LLM_PROVIDER` (set to `openai`), `SKILL_SCANNER_LLM_BASE_URL` to point at any OpenAI-compatible base URL, `SKILL_SCANNER_LLM_USER` to inject a raw Chat Completions `user` field, and `SKILL_SCANNER_LLM_TEMPERATURE` (or `none` to omit sampling temperature) configure routing to OpenAI-compatible endpoints.

— Four named env vars with exact effects, no example shown. product docs

09 **Trusted reference domains for LLM policy** NEW 72

`llm_analysis.trusted_reference_domains` config key lists organization-controlled domains whose transitive-trust and supply-chain findings are automatically

demoted to LOW severity; described elsewhere as allowlisting domains that should not be flagged as suspicious during semantic analysis.

Restrict LLM transitive-trust findings to only your internal registries and package hosts so external references still trigger full-severity alerts.

```yaml
llm_analysis:
  trusted_reference_domains:
    - git.example.com
    - packages.example.com
```

— Named config key with clear effect and example. `2.0.13`

---

10 **Expanded CLI allowlist for reusable GitHub Actions workflow**

`IMPROVED` 68

The reusable workflow's CLI allowlist now includes `--rule-packs`, `--use-aidefense`, `--use-trigger`, and `--enable-meta`, plus LLM limits and repository-local policy/taxonomy paths; unknown flags are rejected before the scan runs. `--use-osv` and `--adjudicate` are not accepted by the reusable workflow and require the self-hosted pattern instead.

— Names exact flags allowed/rejected, no full workflow example. `product docs`

---

11 **LLM consensus runs for adjudication** `NEW` 64

`llm_consensus_runs` integer config field (default `1`) controls the number of LLM runs used for majority consensus adjudication, and is also exposed programmatically through `build_analyzers`.

Enable OSV.dev vulnerable dependency checking and trigger specificity analysis alongside LLM consensus across 3 runs.

```yaml
use_osv: true
use_trigger: true
llm_consensus_runs: 3
```

Assemble all optional analyzers including behavioral, LLM, and OSV checks with 3 consensus runs — matching what the CLI and API produce.

```python
from skill_scanner.core.analyzer_factory import build_analyzers
analyzers = build_analyzers(
    policy,
    use_behavioral=True,
    use_llm=True,
    use_osv=True,
    llm_consensus_runs=3,
)
scanner = SkillScanner(analyzers=analyzers, policy=policy)
```

— Config key and default given, mechanism thin. product docs

12 **LLM token usage in scan results** NEW 62

Scan results include an aggregate `llm_usage` field with `input_tokens`, `output_tokens`, and `total_tokens` when the configured provider reports usage, across LLM, meta, and adjudication calls, enabling cost tracking and auditing; also exposed as a `dict[str, int]` field in the Python SDK.

— Named field with exact sub-fields but no usage command. 2.0.13

THINNER COVERAGE BELOW

13 **Ambient Vertex AI and AWS credential support** NEW 59

Supports Vertex AI ambient credentials via `GOOGLE_APPLICATION_CREDENTIALS` or Application Default Credentials, including Workload Identity, with no extra configuration; also supports ambient AWS/IAM credentials for Bedrock with no extra configuration.

— Named credential mechanisms, no setup command given. product docs

14 **Unpinned-dependency detection across 9 ecosystems** IMPROVED 59

The static analyzer flags unpinned dependencies across Python, Node/npm, Ruby, Go, Rust, Java/Maven, .NET/NuGet, PHP/Composer, and Swift manifests, surfacing supply-chain risk in skill packages.

— Lists concrete ecosystems but no flag or example. 2.0.13

15 **Custom taxonomy and threat mapping file paths** NEW 57

`--taxonomy PATH` overrides the Cisco AI taxonomy JSON/YAML file, and `--threat-mapping PATH` overrides the threat mapping JSON file.

— Named flags with file formats, no example command. product docs

16 **LLM response token budget override** NEW 55

`--llm-max-tokens N` overrides the LLM response-token budget (default: 8192).

— Exact flag and default given, no usage example. product docs

17 **Suspicious-URL detection expanded to more tunneling domains**

IMPROVED 55

Suspicious-URL detection in YAML, JSON, TOML, INI, and environment/config files expanded to include `bore.pub`, `serveo.net`, and `localtunnel.me` alongside existing ngrok domains.

— Names exact domains and file formats, no flag shown. 2.0.13

18 **Unicode smuggling and jailbreak pattern detection** NEW

44

Adds Unicode Tag Block (ASCII-smuggling) payload detection, generic jailbreak pattern detection, and improved OpenAI/Anthropic project-key detection to core coverage.

— Names detection types but no flag or config to invoke. product docs

19 **Build provenance attestations and SBOM for releases** NEW

43

GitHub release artifacts ship with build-provenance attestations and an SBOM, enabling supply-chain verification and software-inventory retention for controlled environments.

— Describes artifact contents but no retrieval steps. 2.0.13

20 **Trigger specificity analysis** NEW 40

`use_trigger` config field / `--use-trigger` flag enables trigger specificity analysis.

Enable OSV.dev vulnerable dependency checking and trigger specificity analysis alongside LLM consensus across 3 runs.

yaml

```yaml
use_osv: true
use_trigger: true
llm_consensus_runs: 3
```

— Named flag but no explanation of mechanism. product docs

21 **JavaScript/TypeScript analyzability scoring and symlink discovery**

IMPROVED

38

Adds JavaScript and TypeScript content to analyzability scoring, plus recursive discovery of symlinked skill directories.

— Brief description, no example or flag given. product docs

22 **LLM analysis budget and meta multiplier controls** NEW

32

`llm_analysis` policy section covers prompt/output token budgets and a meta multiplier for controlling LLM analysis behavior.

— Named section but no concrete keys or values. product docs

└──▷ ALSO FROM THESE RELEASES

Run a full CI scan with LLM analysis and emit SARIF output, while failing the build on any high-severity finding and tracking token spend in the JSON result.

```
$ skill-scanner scan-all ./skills --recursive --use-llm --use-osv --fail-on-severity high --format sarif --output results.sarif
```

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ 3 RELEASES · 2026-07-30 → 2026-08-03

BLOG ↗

StepSecurity is a security platform that hardens GitHub Actions workflows and CI/CD pipelines by reducing attack surface and preventing supply chain compromises.

> StepSecurity's biggest update this window is full NuGet (.NET) ecosystem support spanning GitHub Checks, Secure Registry, and OSS Package Search, alongside Dev Machine Guard/Device Policy improvements that add npm/pnpm/Yarn/bun registry enforcement, VS Code private marketplace policies, and MDM drift verification.

└──▷ WHAT SHIPPED · 13 FEATURES most completely described first ? what's the number?

01 NuGet package proxying and policy controls in Secure Registry

NEW 90

Secure Registry now serves NuGet (Beta) at `https://registry.stepsecurity.io/dotnet/index.json` (V3 protocol only, resources under `/dotnet/v3/`), routing all NuGet metadata requests and package downloads through Secure Registry for policy evaluation. Policy controls available at launch are Cooldown Period, Compromised Packages, and Custom Block List (Typosquatting Protection is Coming Soon). Supports three integration paths — JFrog Artifactory, Sonatype Nexus Repository, and Direct (with a Package manager selector switching setup instructions between .NET CLI, NuGet CLI, and Visual Studio) — and NuGet now appears in the Registry selector on both the Policy and Setup Guide tabs, configured independently per ecosystem.

— Exact endpoint, protocol path and integration options given

NuGet Ecosystem Support in GitHub Check…

02 Package config policy for npm/pnpm/Yarn/bun registries

IMPROVED 88

Adds `package_config` device policy support: the agent converges a StepSecurity-owned block in `~/.npmrc`, bounded by `BEGIN` / `END` markers, to route npm, pnpm, Yarn classic, and bun through the tenant's Secure Registry using a per-device `::dev:-suffixed` authentication token; absent policy is a no-op with snapshot rollback, and surrounding configuration is left untouched.

— Names exact config file, markers, and rollback mechanism

Device Policy: Package Config Policies,… Dev Machine Guard v1.15.0 (2026-08-03)

03 Verify-only vs self-healing enforcement with MDM drift detection

IMPROVED

83

Introduces a verify-only enforcement channel that reads the OS-managed VS Code policy (Windows registry, macOS managed preferences, `/etc/vscode/policy.json`) or effective `~/.npmrc` and reports observed values and drift without writing. MDM verification performs field-by-field drift diffs, reporting one of eight states (`MDM verified`, `MDM drift`, `Not deployed`, `Verification failed`, and others) with desired vs. observed values per field. Every profile now declares an explicit enforcement type — `Verify only` (MDM-enforced) or `Self healing` (agent-enforced) — and switching an existing profile's type migrates all assigned devices, which report `Pending` until each agent checks in.

— Named policy file paths and explicit state values given

Device Policy: Package Config Policies,… | Dev Machine Guard v1.15.0 (2026-08-03)

---

04 **NuGetAudit unsupported, NU1905 workaround** BREAKING

80

Secure Registry does not yet proxy the VulnerabilityInfo resource for NuGet, causing audit-enabled restores to emit `NU1905`; the documented workaround is to pass `-p:NuGetAudit=false` or exclude the Secure Registry source from audit sources, e.g. `dotnet restore -p:NuGetAudit=false`.

Perform a NuGet restore without triggering NU1905 warnings while NuGetAudit support is pending in Secure Registry.

```
$ dotnet restore -p:NuGetAudit=false
```

— Named error code and exact runnable workaround command

NuGet Ecosystem Support in GitHub Check…

---

05 **VS Code private marketplace gallery URL enforcement** NEW

75

Writes `extensions.gallery.serviceUrl` into VS Code user-scope `settings.json` atomically alongside `extensions.allowed` when a policy sets a gallery URL; removal is ownership-gated so developer-configured URLs are never deleted. IDE extension policies now support an optional gallery URL that points assigned devices at a self-hosted extension gallery instead of the public Visual Studio Marketplace, and this gallery URL is itself verified for drift.

— Names exact settings.json fields and behavior | Device Policy: Package Config Policies,… | Dev Machine Guard v1.15.0 (2026-08-03)

---

06 **NuGet ecosystem in OSS Package Search** NEW

60

NuGet is now selectable from the Package ecosystem toggle in OSS Package Search; packages are indexed from `*.csproj`, `*.vbproj`, `*.fsproj`, `packages.config`, `packages.lock.json`, `Directory.Packages.props`, `Directory.Build.props`, and `*.deps.json` files.

— Named file types indexed and UI toggle NuGet Ecosystem Support in GitHub Check…

THINNER COVERAGE BELOW

07 **MDM deployment scripts for package config policies** NEW

55

Generates audit, remediation, and uninstall script sets per platform for MDM-delivered package config policies, designed for deployment into Intune Remediations detection/remediation slots or an Iru Custom Script library.

— Names deployment targets but no exact script contents

Device Policy: Package Config Policies,…

08 **NuGet Package Compromised Updates check** NEW 55

Adds a `NuGet Package Compromised Updates` PR Content Security control to GitHub Checks that blocks pull requests introducing or updating a NuGet dependency flagged in StepSecurity's compromised-packages database; can be set to Required or Optional.

— Named check and its two config states NuGet Ecosystem Support in GitHub Check…

09 **NuGet Package Cooldown check** NEW 55

Adds a `NuGet Package Cooldown` PR Content Security control to GitHub Checks that blocks pull requests adding recently published NuGet versions until they age beyond a configured cooldown window; can be set to Required or Optional.

— Named check and its two config states NuGet Ecosystem Support in GitHub Check…

10 **Homebrew detection outside PATH** IMPROVED 50

Detects Homebrew installed outside `PATH` by resolving `brew` from `/opt/homebrew`, `/usr/local`, and `/home/linuxbrew` and reading version metadata from Homebrew's repository instead of invoking `brew --version`, completing formula and cask inventory on shell-less scans.

— Named install paths but no user-facing action Dev Machine Guard v1.15.0 (2026-08-03)

11 **UTF-8 BOM support for VS Code settings.json** IMPROVED

35

Makes VS Code `settings.json` files beginning with a UTF-8 byte order mark (written by PowerShell 5.1 `Set-Content -Encoding UTF8` or `utf8bom -configured` editors) enforceable.

— Names encoding sources but no reader action Dev Machine Guard v1.15.0 (2026-08-03)

12 **NuGet filter in Policy Evaluations audit log** IMPROVED

35

Policy Evaluations audit log now includes `.NET` in the Ecosystem filter, enabling narrowing of audit log entries to NuGet requests.

— Names the filter field but little else NuGet Ecosystem Support in GitHub Check…

13 **Per-category device policy ownership tracking** IMPROVED

30

Records device policy ownership for every category in a single locked state file, preventing one category's policy clear from dropping another category's ownership record.

— Describes fix briefly without naming the state file

Dev Machine Guard v1.15.0 (2026-08-03)

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Veracode

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

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

> Veracode shipped a new Custom PDF REST API, expanded Pipeline Scan support to artifacts over 200 MB, and added a GitHub app uninstallation workflow, alongside CLI and GitHub Workflow Integration updates.

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

01 **GitHub Workflow Integration runner image update** IMPROVED 75

`GitHub Workflow Integration v2.1.1` updates the default runner image to `scm-packaging-5.0.0`, adding Alpine v3.24.0, Ubuntu v26.04, .NET v10, PipEnv v2026.6.2, libicu v78, Node.js v22.22.1, and Python v3.14.4.

– Lists exact image name and every bundled tool version changelog-20260804-ea32fe1a

THINNER COVERAGE BELOW

02 **Custom PDF REST API** NEW 50

A new Custom PDF REST API enables programmatic generation of custom PDF reports.

– Names the API but not exact endpoint path or method changelog-20260804-ea32fe1a

03 **Pipeline Scan support for artifacts over 200 MB** IMPROVED 50

Pipeline Scan now supports packaged artifacts larger than 200 MB, with no changes required to existing integrations.

– Concrete size limit given but no action needed to adopt changelog-20260804-ea32fe1a

04 **CLI batch fix processing** NEW 50

`Veracode CLI v2.51.2` adds enhanced batch fix processing, enabling remediation of a large number of vulnerabilities in a single batch operation.

– Names CLI version and capability but no command syntax changelog-20260804-ea32fe1a

05 **GitHub app uninstallation workflow** NEW 40

A new uninstallation workflow on the Veracode Platform lets you uninstall the Veracode app from GitHub organizations.

– Describes a UI workflow without exact navigation steps changelog-20260804-ea32fe1a

06 **CLI policy enforcement consistency** IMPROVED 35

`Veracode CLI v2.51.2` improves policy enforcement to ensure vulnerability severity thresholds are consistently applied.

— Vague on what changed in enforcement mechanism `changelog-20260804-ea32fe1a`

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

WPScan

Feedback ⓘ 1 RELEASE · 2026-07-23 NOTES ↗ CODE ↗

WPScan WordPress security scanner. Written for security professionals and blog maintainers to test the security of their WordPress websites. Contact us via contact@wpscan.com

> WPScan v4.1.0 introduces an opt-in local vulnerability database mode for enterprise customers, bypassing online API rate limits.

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

01 **Local enterprise vulnerability database mode** NEW 95

The new `--enterprise-db-token TOKEN` flag (also settable via the `WPSCAN_ENTERPRISE_DB_TOKEN` environment variable or the `enterprise_db_token` config-file key) lets WPScan scan against a local vulnerability database dump instead of the online API, with no per-request limits. When set, WPScan downloads `plugins.json.gz`, `themes.json.gz`, and `wordpresses.json.gz` from `enterprise-data.wpscan.org`, authenticated via the `X-DB-JSON-AUTH` header, and reads vulnerability data locally; API status is then reported as plan `enterprise` with `Unlimited` requests remaining. `--enterprise-db-token` and `--api-token` are mutually exclusive (including via their environment variables) — supplying both raises an error — and the token is redacted from logged command lines.

Scan a WordPress site using a local enterprise vulnerability database dump instead of online API calls, avoiding per-request limits.

```
$ wpscan --url https://target.tld --enterprise-db-token <your-enterprise-token> --enumerate vp,vt
```

Store the enterprise token in a config file so it's applied automatically to every scan without passing it on the CLI.

```yaml
cli_options:
  enterprise_db_token: "<your-enterprise-token>"
  url: "https://target.tld"
```

— Names flag, env var, config key, endpoint files, header, and behaviour v4.1.0

WAS THIS USEFUL? 👍 👎

◆ Digital Forensics & Incident Response

Yamato Security Hayabusa

Feedback ⓘ

1 RELEASE · 2026-08-03 NOTES ↗ CODE ↗

Hayabusa (隼) is a sigma-based threat hunting and fast forensics timeline generator for Windows event logs.

> Hayabusa 4.0.0 merges its CSV/JSON timeline commands into a single `dfir-timeline` command with rich output-format, GeoIP, HTML-report, sorting and deduplication options, while standardizing all long flag names to lowercase and expanding RDP/PowerShell detection coverage in `logon-summary` and the new `extract-base64` subcommand; it also adds a Linux ARM64 musl binary.

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

01 extract-base64 subcommand for Base64 detection NEW

93

New `extract-base64` subcommand scans `.evtx` or JSON logs for Base64-encoded content and saves results to a CSV file via `-o <FILE>`; it scans PowerShell event IDs `4100` / `4102` (Microsoft-Windows-PowerShell/Operational and PowerShellCore/Operational) across `ContextInfo` and `Payload` fields, and classic Windows PowerShell EIDs `403` / `600` via the `EventData.Data` blob.

Extract Base64-encoded content from a directory of EVTX logs and save to CSV for offline analysis.

```
$ hayabusa.exe extract-base64 -d ../hayabusa-sample-evtx -o
base64-extracted.csv
```

— Named subcommand, event IDs and fields, runnable example v4.0.0

02 Unified dfir-timeline command replaces csv/json-timeline

92

BREAKING

New `dfir-timeline` command replaces `csv-timeline` and `json-timeline`; output format is selected via `-t, --output-type`: `csv` (default), `json`, or `jsonl` (case-insensitive, e.g. `-t JSONL`). The `-t` short flag is reassigned from `--threads` to `--output-type` (use `--threads <N>` for thread count), the old `-L, --JSONL-output` flag is removed in favor of `-t jsonl`, CSV-only flags (`-M, --multiline`, `-S, --tab-separator`, `-R, --remove-duplicate-data`) now error if combined with a non-CSV `-t` value, and the scan wizard is enabled by default.

Output JSONL for bulk import into the Elastic Stack — faster and smaller than JSON, with case-insensitive format flag.

```
$ hayabusa.exe dfir-timeline -d .\hayabusa-sample-evtx -t JSONL -
o results.jsonl
```

— Command, flag values and behaviour change fully specified `v4.0.0`

---

03 **Duplicate data and detection removal flags** NEW 88

`-R, --remove-duplicate-data` replaces duplicate field data with `"DUP"` in `dfir-timeline` output (requires `--sort`); `-X, --remove-duplicate-detections` strips duplicate detections from `dfir-timeline` (requires `--sort`); `-X, --remove-duplicate-records` deduplicates event records in `eid-metrics` and `logon-summary` (default: disabled).

Speed up scanning on large collections by restricting to common Event IDs, then deduplicate detections for a cleaner timeline.

```
$ hayabusa.exe dfir-timeline -E -s -X -d .\hayabusa-sample-evtx -
o results.csv
```

— Multiple named flags with exact behaviour and prerequisites `product docs`

---

04 **RDP session detection in logon-summary** IMPROVED 85

`logon-summary` now counts RDP session events: Security EIDs `4778` / `4779` (session reconnect/disconnect with client workstation name and IP) and Terminal Services events RemoteConnectionManager/Operational `1149` and LocalSessionManager/Operational `25`, covering RDP activity even when `4624` events have been overwritten.

— Specific event IDs and log channels named, no CLI action needed `v4.0.0`

---

05 **Lowercase standardization of long option names** BREAKING 83

All long option names are renamed to lowercase: `--geo-ip`, `--HTML-report` → `--html-report`, `--UTC` → `--utc`, `--ISO-8601` → `--iso-8601`, `--RFC-2822` / `--RFC-3339` → `--rfc-2822` / `--rfc-3339`, `--JSON-input` / `--JSON-output` → `--json-input` / `--json-output`, `--EID-filter` → `--eid-filter`; short flags are unchanged and scripts using the old uppercase forms will break.

— Exact old/new flag names given; no migration script `v4.0.0`

---

06 **GeoIP enrichment for dfir-timeline** NEW 81

`-G, --geo-ip <MAXMIND-DB-DIR>` enriches IP addresses in `dfir-timeline` output with GeoIP data (ASN, city, country).

Enrich a DFIR timeline with GeoIP context (ASN, city, country) for IP addresses — useful when triaging lateral movement or C2 traffic.

```
$ hayabusa.exe dfir-timeline -d .\hayabusa-sample-evtx -G
.\maxmind-db -o results.csv
```

— Named flag with argument and exact enriched fields `v4.0.0`

---

07 **Common-EID filter for faster scanning** NEW 78

`-E, --eid-filter` restricts scanning to common Event IDs listed in `./rules/config/target_event_IDs.txt` for faster performance; added to `dfir-timeline` and `pivot-keywords-list`.

Speed up scanning on large collections by restricting to common Event IDs, then deduplicate detections for a cleaner timeline.

```
$ hayabusa.exe dfir-timeline -E -s -X -d .\hayabusa-sample-evtx -
o results.csv
```

— Flag and config file path named; performance benefit stated generically `product docs`

---

08 **Additional timestamp format flags** NEW 78

Adds `--european-time`, `-O, --iso-8601` (always UTC), `--rfc-2822`, `--rfc-3339`, `-U, --utc`, `--us-military-time`, and `--us-time` timestamp format flags across `dfir-timeline`, `eid-metrics`, `extract-base64`, `log-metrics`, `logon-summary`, and `search`.

Collect log-file metadata from a directory, scoping to only Security and System channels, with UTC timestamps.

```
$ hayabusa.exe log-metrics -d ../logs --include-channel
Security,System -U -o log-metrics.csv
```

— Full list of named format flags with commands they apply to `product docs`

---

09 **Channel and filename filters for log-metrics** NEW 78

`--exclude-channel <CHANNEL...>`, `--exclude-filename <FILE...>`, `--include-channel <CHANNEL...>`, and `--include-filename <FILE...>` flags added to `log-metrics` for fine-grained channel and file filtering.

Collect log-file metadata from a directory, scoping to only Security and System channels, with UTC timestamps.

```
$ hayabusa.exe log-metrics -d ../logs --include-channel
Security,System -U -o log-metrics.csv
```

— Four named flags with runnable filtering example `product docs`

---

10 **Sort and terminal timeline visualization** NEW 72

`-s, --sort` sorts `dfir-timeline` results before saving (note: significantly increases memory usage); `-T, --visualize-timeline` outputs an event-frequency timeline in the terminal (requires unicode terminal support and `--sort`).

— Flags and prerequisites given but visual output not shown `product docs`

---

11 **CSV-only multiline and tab-separator flags** NEW 72

`-M, --multiline` separates event field information by newline characters (CSV only), added to `dfir-timeline`, `log-metrics`, and `search`; `-S, --tab-separator` separates fields by tabs (CSV only) in `dfir-timeline`; combining either with a non-CSV `-t` value produces an error.

— Named flags with scope and error condition specified `product docs`

---

12 **HTML results summary report** NEW 67

`-H, --html-report <FILE>` saves a Results Summary HTML report alongside the `dfir-timeline` output.

— Flag and file argument given, but no report contents detail `v4.0.0`

---

13 **Linux ARM64 musl release binary** NEW 65

New Linux aarch64 (ARM64) musl release binary (`hayabusa-<ver>-lin-aarch64-musl`) built with a proper aarch64 musl cross toolchain.

— Exact binary naming pattern given, download path implied `v4.0.0`

---

14 **JSON log input for dfir-timeline** NEW 64

`-J, --json-input` scans JSON-formatted logs (`.json` or `.jsonl`) instead of `.evtx` in `dfir-timeline`.

— Flag and accepted formats given, no further behaviour `product docs`

---

15 **Configurable thread count** NEW 64

`--threads <NUMBER>` controls thread count (default: optimal for performance), added to `dfir-timeline`, `eid-metrics`, `extract-base64`, `log-metrics`, `logon-summary`, `pivot-keywords-list`, and `search`.

– Flag and default named across commands, no benchmark given product docs

16 **Checksum validation flag across commands** NEW 63

`-V, --validate-checksums` enables checksum validation, added to `dfir-timeline`, `computer-metrics`, `eid-metrics`, `extract-base64`, `log-metrics`, `logon-summary`, `pivot-keywords-list`, and `search`.

– Flag named across commands but validation mechanism unexplained product docs

17 **JSON/JSONL output for search command** NEW 60

`-J, --json-output` and `-L, --jsonl-output` flags save `search` results in JSON or JSONL format.

– Named flags and formats, no further mechanism product docs

THINNER COVERAGE BELOW

18 **First/Last logon columns in logon-summary** IMPROVED 55

Added First Logon / Last Logon columns to the `logon-summary` successful-logons table and First Attempt / Last Attempt columns to the failed-logons table.

– Column names given but no mechanism or example v4.0.0

19 **Deterministic results-summary ordering** IMPROVED 50

Results-summary output is now deterministically ordered: equal-count rows are tie-broken by channel and event ID, and by logon grouping key, eliminating run-to-run diff noise.

– Mechanism described but no flag or config to act on product docs

20 **config-critical-systems subcommand** NEW 35

New `config-critical-systems` subcommand, accepting an `<INPUT>` argument and `[OPTIONS]`, for configuring critical systems.

– Only name and generic argument shape given, no behaviour product docs

└──▷ **BREAKING ON UPGRADE**

! The `-L, --JSONL-output` flag is removed; replace with `-t jsonl` on the `dfir-timeline` command.

! The `-t` short flag no longer means `--threads`; it now means `--output-type`. Scripts passing `-t <number>` for thread count must switch to `--threads <number>`.

! Long option names are renamed to lowercase: `--geo-ip`, `--UTC` → `--utc`, `--HTML-report` → `--html-report`, `--JSON-input` → `--json-input`, `--JSON-output` → `--json-output`, `--EID-filter` → `--eid-filter`; any scripts using the old casing will break.

- ! CSV-only output options now error when used with `-t json` or `-t jsonl`; previously they were silently accepted.
- ! The `csv-timeline` and `json-timeline` commands are replaced by `dfir-timeline`; scripts invoking either old command must be updated.
- ! The `-t` short flag is reassigned from `--threads` to `--output-type`; any use of `-t <N>` to set thread count must switch to `--threads <N>`.
- ! The `json-timeline` flag `-L, --JSONL-output` is removed; use `dfir-timeline -t jsonl` instead.
- ! Long-form flags renamed: `--GeoIP` → `--geo-ip`, `--HTML-report` → `--html-report`, `--UTC` → `--utc`, `--ISO-8601` → `--iso-8601`, `--RFC-2822` → `--rfc-2822`, `--RFC-3339` → `--rfc-3339`, `--JSON-input` → `--json-input`, `--JSON-output` → `--json-output`, `--EID-filter` → `--eid-filter`; any scripts using the old uppercase forms will break.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

ILSpy

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

.NET Decompiler with support for PDB generation, ReadyToRun, Metadata (&more) - cross-platform!

> ILSpy v11.0-rc adds two new ilspycmd inspection modes, a way to attach the GUI to a running process, and several decompiler-view and metadata-panel navigation improvements.

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

01 **ilspycmd --dump-table metadata dump mode** NEW 72

Adds `--dump-table` metadata dump mode to `ilspycmd` for inspecting raw metadata tables, e.g. `ilspycmd --dump-table MyAssembly.dll`.

Dump a raw metadata table from an assembly for inspection without launching the GUI.

```
$ ilspycmd --dump-table MyAssembly.dll
```

— Names flag and gives runnable command, mechanism not detailed further. v11.0-rc

02 **ilspycmd --member single-member decompilation** NEW 72

Adds `--member` flag to `ilspycmd` for single-member decompilation, e.g. `ilspycmd --member "MyNamespace.MyClass.MyMethod" MyAssembly.dll`, useful for quick diffs in CI.

Decompile a single method or type member to stdout — useful for quick diffs in CI.

```
$ ilspycmd --member "MyNamespace.MyClass.MyMethod" MyAssembly.dll
```

— Named flag with runnable CLI example and stated use case. v11.0-rc

THINNER COVERAGE BELOW

03 **Ctrl+Click member reference navigation** NEW 55

Ctrl+Click navigates member references while a plain click highlights them in the editor.

— Exact interaction described, immediately usable, no further scope. v11.0-rc

04 **Publicized IProjectFileWriter for programmatic writing**

IMPROVED

50

Publicizes `IProjectFileWriter` implementations for programmatic project-file writing.

— Names interface but no usage example or code sample. v11.0-rc

05 **WholeProjectDecompiler explicit constructor** NEW

50

Adds a `WholeProjectDecompiler` constructor with no default parameter values for explicit API use.

— Names class and constructor change but no call-site example. v11.0-rc

06 **Full-width metadata row details panel** IMPROVED

45

Metadata row details panel now uses a full-width layout with a decompiler-styled editor for text blobs.

— Describes UI redesign but no exact navigation path given. v11.0-rc

07 **FieldList/MethodList values for memberless types** NEW

45

Show stored `FieldList` / `MethodList` values for memberless types in the metadata view.

— Names concrete metadata fields but no usage steps. v11.0-rc

08 **Open from running process dialog** NEW

40

New Open from running process dialog lets you attach ILSpy directly to a live process.

— Names UI dialog but no navigation steps or mechanism given. v11.0-rc

09 **Overrides analysis and this/base navigation** NEW

40

Adds overrides analysis and `this` / `base` keyword navigation.

— Named keywords but mechanism and scope unexplained. v11.0-rc

10 **Hover tooltips for dynamic member accesses** NEW

35

Shows hover tooltips for dynamic member accesses in the decompiler view.

— Describes behaviour only, no detail on trigger scope or format. v11.0-rc

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Sublime Security

Feedback ⓘ 1 RELEASE · 2026-08-03 BLOG ↗

Sublime Security is an email security platform that detects and prevents advanced threats like phishing, malware, and business email compromise attacks.

> Sublime Security introduces Tailored Phishing Simulations (Beta), letting teams run simulation campaigns prioritized by attack patterns already detected in their own environment, with multi-template campaigns and cross-campaign tracking.

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

01 **Tailored Phishing Simulations (Beta)** NEW 63

Run phishing simulation campaigns prioritized by attack patterns Sublime is already detecting in your environment. Supports multiple templates per campaign with staggered delivery across recipients, and tracks click rates, report rates, and repeat clickers in a cross-campaign dashboard with recipient-level detail per campaign.

— Describes behavior and metrics but no config or API surface named.

Platform v.1.88 Tailored Phishing Simul…

WAS THIS USEFUL? 👍 👎

Fastly

Feedback ⓘ CLI v15.5.0 (2026-08-04) 2 RELEASES · 2026-08-04

BLOG ↗ LATEST STABLE

Fastly is a content delivery network and edge computing platform that accelerates website performance and provides security services globally.

Fastly CLI v15.5.0 adds Python support for Compute init and new Notification Service Integrations management.

Adds Python language support to `compute/init`, enabling Python-based Compute service scaffolding.

Adds support for Notification Service Integrations via the `integrations` commands.

▾ 1 more release in this issue · 2026-08-04

> Fastly Terraform Provider 9.5.0 (2026-08-04) BLOG ↗ STABLE
>
> Fastly Terraform Provider 9.5.0 adds five new integration types to `fastly_integration`: Datadog, Jira Issue, JSM, OpsGenie, and Splunk On-Call.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion

Feedback ⓘ 2 RELEASES · SEEN 2026-08-04 NOTES ↗

Plerion is a cloud security platform that enables continuous vulnerability discovery and remediation across cloud infrastructure and applications.

> Plerion added organization-wide GCP onboarding with keyless federation and automatic project discovery, native Terraform-based AWS onboarding, and expanded Kubernetes workload scanning with SBOM generation and in-image secret detection.

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

01 **AWS onboarding via native Terraform template** NEW 75

A new Terraform template provisions the IAM role and policies as native Terraform resources and registers the AWS integration automatically, requiring no CloudFormation stack or manual console step. It offers a choice of running workload scanning inside your own account or through a Plerion-managed service account.

— Names Terraform resources and account-scanning choice, no template path

Onboard AWS accounts with Terraform:

02 **SBOM generation for Kubernetes container images** NEW 70

Workload scanning now builds a software bill of materials (SBOM) for every container image across Kubernetes Deployments, DaemonSets, StatefulSets, Pods, Jobs, and CronJobs, surfaced on each asset's SBOM tab.

— Lists workload types and named SBOM tab UI location

Onboard AWS accounts with Terraform:

03 **GCP organization-wide onboarding via keyless federation** NEW 60

Onboards an entire Google Cloud organization from a single generated command using keyless access via Workload Identity Federation, with no service account keys required. An equivalent Terraform template is also provided for organizations preferring infrastructure-as-code onboarding.

— Names WIF mechanism and Terraform option but no exact command

Connect your entire GCP organization at…

04 **Secret scanning inside container images** NEW 60

Secret scanning detects API keys, tokens, and passwords inside container images and reports them with severity on the asset's Sensitive data tab, so practitioners know what

to rotate first.

— Names detected secret types and Sensitive data tab location

Onboard AWS accounts with Terraform:

THINNER COVERAGE BELOW

05 **Organization-level resource scanning scope** NEW 50

Scans organization-level resources including organization details, organization-level IAM policy, folders, organization policies, and VPC Service Controls service perimeters.

— Lists concrete resource types but no way to act on it

Connect your entire GCP organization at…

06 **Configurable workload scan types** IMPROVED 45

Scan types (SBOM, secret scanning) can be selected at cluster-connect time or changed at any point from the integration's workload scan settings.

— Names the settings location but no exact field or key

Onboard AWS accounts with Terraform:

07 **Automatic discovery of GCP org projects** NEW 40

Automatically discovers every project in the GCP organization, including projects created after initial onboarding, with the option to include all projects automatically or select them individually.

— Describes behaviour but no config surface named

Connect your entire GCP organization at…

08 **Revoke GCP org access via single resource deletion** NEW 25

Revoke all Plerion access to the GCP organization by deleting a single resource.

— Mentions an action but doesn't name the resource

Connect your entire GCP organization at…

WAS THIS USEFUL?  

◆ Cloud Security Posture

Aikido Security

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

Aikido Security is a developer-focused platform that identifies and remediates vulnerabilities in code, dependencies, and infrastructure.

> Aikido Security introduces a code-based Data Security Posture Management (DSPM) capability, the Data Exposure Audit, that traces sensitive data flows and flags exposure risks entirely from source code without connecting to production data stores.

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

01 **Data Exposure Audit for code-based DSPM** NEW 70

The new Data Exposure Audit uses LLM agents to trace sensitive data from sources (request bodies, uploads, DB reads, config, application logic) to sinks (logs, third-party APIs, LLM prompts, caches, queues, storage) entirely from code, with no production data store connection required. It detects credential and token leakage in URLs, query strings, redirects, and clone URLs; API over-exposure where endpoints return more fields than the caller is entitled to or lack tenant ownership checks; unsafe logging and exports of unredacted PII and secrets in logs, traces, caches, and dumps; over-broad third-party and AI data sharing to vendors, analytics, LLM prompts, and vector stores; weak storage patterns such as regulated data stored unencrypted or protected only with base64, MD5, or SHA-1; and incomplete deletion flows that leave downstream records, caches, shadow copies, staging data, or snapshots intact. Each Data Exposure Audit run consumes credits from the Aikido Wallet.

— Explains mechanism and detection categories but gives no UI path or command.

New - DPSM: Find data exposure risks wi…

WAS THIS USEFUL? 👍 👎

Prowler

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗ CODE ↗

Prowler is the Open Cloud Security Platform trusted by thousands to automate security and compliance in any cloud environment. With thousands of ready-to-use checks and compliance frameworks, delivers real-time, customizable monitoring and seamless integrations, making cloud security simple, scalable, cost-effective for organizations of any size.

> Prowler 5.37.0 adds a new Huawei Cloud provider, 11 AWS Nitro Enclaves CSPM checks, four privilege-escalation Attack Path queries, and a 32-tool Prowler MCP Cloud Server namespace, while Prowler Cloud's Lighthouse AI becomes context-aware and can drive the full MCP tool surface under RBAC.

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

01 **Huawei Cloud provider support** NEW 95

Prowler CLI adds a new `huaweicloud` provider authenticated via `HUAWEICLOUD_ACCESS_KEY_ID` and `HUAWEICLOUD_SECRET_ACCESS_KEY` environment variables, running `prowler huaweicloud` across 25 checks covering CTS, ECS, ELB, EVS, IAM, KMS, OBS, RDS, VPC, and WAF services, plus the CIS Huawei Cloud Foundations Benchmark 1.0 compliance framework.

Scan a Huawei Cloud account for the first time using IAM access key credentials.

```
$ HUAWEICLOUD_ACCESS_KEY_ID="your-access-key-id"
HUAWEICLOUD_SECRET_ACCESS_KEY="your-secret-access-key" prowler
huaweicloud
```

Scan a Huawei Cloud account for the first time using IAM user access key credentials.

```
$ export HUAWEICLOUD_ACCESS_KEY_ID="your-access-key-id"
export HUAWEICLOUD_SECRET_ACCESS_KEY="your-secret-access-key"
prowler huaweicloud
```

— Names env vars, command, exact check count and services. 5.37.0

02 **Prowler MCP Cloud Server 32-tool namespace** NEW 90

Prowler MCP Cloud Server at `https://mcp.prowler.com/mcp` adds 32 tools in the `prowler_cloud_*` namespace covering Alerts (rule creation/management, dry-run, fired-alert history), Findings Triage (status and notes), Scan Scheduling (daily, interval, weekly, or monthly per-provider or bulk), and Scan Configurations (reusable check/compliance selections attached to providers).

— Names endpoint, tool namespace, and exact scheduling options. 5.37.0

03 **AWS Nitro Enclaves CSPM checks** NEW 85

Adds 11 new checks: five `ec2_confidential_workload_host_*` checks (IMDSv2 enforcement, public IP exposure, unrestricted ingress, exposed vsock proxy ports, stopped hosts) and six `kms_key_enclave_*` checks (attestation enforcement, bypass, deployment binding, debug-mode, PCR mismatches, unknown enclave images), mapped across 23 compliance frameworks including NIST 800-53 Rev 5, PCI-DSS v4.0, ISO 27001:2022, SOC 2, HIPAA, and MITRE ATT&CK.

— Names all checks and frameworks; no invocation example. 5.37.0

04 **Context-aware Lighthouse AI chat panel** IMPROVED 75

Lighthouse AI in Prowler Cloud gains a context-aware chat panel that automatically attaches the current app page and any open finding or resource metadata to each message — shown as a chip in the composer — so queries like explain this resolve against what is on screen, and Lighthouse can now invoke all Prowler MCP Server tool families (scan configurations, scheduling, finding triage, alert rules, integrations, users, roles), gated by RBAC.

— Describes context mechanism and RBAC gating in detail. 5.37.0

05 **AWS privilege-escalation Attack Path queries** NEW 70

Adds four new Attack Paths detection queries: `STS-002` (cross-account role trust), `STS-003` (wildcard role trust), `IAM-022` (user permissions-boundary removal), and `SSO-001` (IAM Identity Center permission-set escalation).

— Named query IDs given, but no run instructions. 5.37.0

06 **MCP tool families for Integrations, Users, Roles** NEW 70

Prowler MCP Server (Cloud and self-hosted Local) gains three tool families: Integrations (full lifecycle for Amazon S3, AWS Security Hub, and Jira — create, update, re-check, delete, and dispatch findings as Jira work items), Users (list tenant users and identify

the authenticated user), and Roles (browse RBAC roles, inspect capabilities, and set a user's role).

— Named tool families and actions, no call syntax given. 5.37.0

07 **Cross-provider compliance grouping and PDF report** IMPROVED 70

Compliance section now groups results for all providers of the same type under a single framework card with per-provider breakdown, findings drill-down, and a combined executive PDF report; the tabs are renamed to Single Scan and Multiple Scans, with Multiple Scans as the new default landing tab.

— Names tab labels and default, no config key given. 5.37.0

08 **CodeCommit hardcoded secrets check** NEW 65

Adds `codecommit_repository_no_secrets` check under a new `codecommit` service, scanning files at the tip of each repository's default branch for hardcoded secrets.

— Named check and service, no config example given. 5.37.0

09 **OCSF MITRE ATT&CK enrichment** IMPROVED 65

OCSF detection finding output now populates `finding_info.analytic` with the Prowler check rule and `finding_info.attacks` with MITRE ATT&CK technique and tactic objects, for findings that carry MITRE ATT&CK compliance metadata.

— Named OCSF fields, but no sample output shown. 5.37.0

10 **Guided GCP organization onboarding flow** IMPROVED 65

GCP Organization onboarding in Prowler Cloud is now a single guided flow: provide an organization-level credential, browse the full folder/project hierarchy in a selection tree, set aliases, test the connection, and register each selected project as a provider in one operation.

— Describes flow steps but no exact UI path. 5.37.0

THINNER COVERAGE BELOW

11 **EC2 stopped-instance age check** NEW 55

Adds `ec2_instance_stopped_older_than_specific_days` check to detect EC2 instances stopped longer than a configurable number of days (default 30).

— Named check with configurable default, single source. product docs

12 **Glue Data Catalog secrets check** NEW 45

Adds `glue_catalog_connection_no_secrets` check to detect secrets in Glue Data Catalog connection properties.

— Named check only, thin description. 5.37.0

13 **SageMaker endpoint KMS encryption check** NEW 45

Adds `sagemaker_endpoint_config_kms_encryption_enabled` check to verify SageMaker endpoint configurations use a KMS key for storage volume encryption.

— Named check only, thin description. 5.37.0

14 **Attack Paths query links to Prowler Hub** IMPROVED 40

Attack Paths query info panel now links every query to its page on Prowler Hub.

— Single sentence, minimal mechanism detail. 5.37.0

WAS THIS USEFUL?  

◆ Container & Image Security

Inspektor Gadget

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

CODE ↗

Inspektor Gadget is a set of tools and framework for data collection and system inspection on Kubernetes clusters and Linux hosts using eBPF

> Inspektor Gadget v0.55.0 adds a GPU telemetry bridge daemon, map pinning and non-destructive iteration for the eBPF operator, three new eBPF program types, and OTel/Python stack symbolization for perf_event-based gadgets.

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

01 **OTel/Python stack symbolization for perf_event gadgets** NEW 65

Adds OTel/Python stack symbolization for `perf_event -based` gadgets such as `profile_cpu`, including aggregation of identical stacks.

— Names the affected gadget and stack aggregation behaviour. v0.55.0

02 **eBPF map pinning support** NEW 60

The eBPF operator now supports map pinning via `LIBBPF_PIN_BY_NAME` in gadget maps, letting maps persist and be referenced across gadget invocations.

— Names the pinning key but gives no usage walkthrough. v0.55.0

03 **Non-destructive eBPF map iteration** NEW 60

A new `iter/bpf_map_elem` program type combined with `GADGET_ITER_TARGET_MAP` enables non-destructive iteration over eBPF maps in the eBPF operator.

— Names the iterator type and target-map setting, no example given. v0.55.0

THINNER COVERAGE BELOW

04 **Three new eBPF program types supported** NEW 45

The eBPF operator adds support for `ebpf.SockOps`, `ebpf.SkSKB`, and `ebpf.SkMsg` eBPF program types.

— Only lists the three type names, no behaviour detail. v0.55.0

05 **GPU telemetry bridge daemon** NEW 30

Adds a `gpu-ebpf-bridge` daemon for collecting GPU telemetry.

— Mentioned only in the release summary, no mechanism given. v0.55.0

WAS THIS USEFUL? 👍 👎

Docker Compose

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

Define and run multi-container applications with Docker

> Docker Compose v5.4.0 introduces a new reconciliation model for volumes and networks, adding first-class lifecycle tracking for both resource types during deployment.

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

01 Volume recreation modelling in reconciliation plan NEW

40

Compose's reconciliation plan now models volume recreation, giving finer control over when volumes are recreated during `up`.

— Names `up` but no flag or config detail on control. v5.4.0

02 Network lifecycle modelling in reconciliation plan NEW

33

Compose's reconciliation plan now tracks network create/remove operations as first-class steps in the deployment plan, enabling more precise network lifecycle management.

— Describes mechanism but no concrete flag or command shown. v5.4.0

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

NVIDIA OpenShell

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

OpenShell is the safe, private runtime for autonomous AI agents.

> OpenShell v0.0.97 adds gateway callback listener negotiation for its compute layer and introduces a prototype Bazel build path with OpenTelemetry crate targets.

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

01 **Gateway callback listener negotiation** NEW 40

Adds gateway callback listener negotiation for the compute layer, enabling dynamic coordination between sandboxes and the gateway control plane.

— Describes mechanism but no config/API surface named v0.0.97

02 **Prototype Bazel build system** NEW 35

Introduces a prototype Bazel build system, including OpenTelemetry crate targets, for an alternative build path.

— Names Bazel and OpenTelemetry but no exact commands given v0.0.97

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Huntress

Feedback ⓘ 1 RELEASE · SEEN 2026-08-04 CHANGELOG ↗

Huntress is a managed security platform that provides threat detection, incident response, and endpoint protection for small and mid-sized businesses.

> Huntress expanded Managed ITDR with a new API for IP-based access rules, smarter handling of trusted corporate devices, automatic onboarding escalation pacing, and Microsoft usage-location inference to cut false escalations.

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

01 **API endpoint for IP-based ITDR access rules** NEW 65

A new API endpoint lets partners create IP address-based allow/block rules for Managed ITDR Unwanted Access, accepting a single IP, a CIDR range, or a list of IPs in one call.

– Names endpoint capability and input formats but no exact path/method

changelog-20260804-eba7c64e

02 **Trusted-device suppression for corporate login escalations**

IMPROVED 65

Logins from corporate-managed devices no longer trigger location and VPN escalations by default for new accounts; existing partners can configure which device trust levels qualify — compliant, domain-joined, or registered — per account or per organization.

– Explains behavior and configurable trust levels, no exact config path

changelog-20260804-eba7c64e

THINNER COVERAGE BELOW

03 **Microsoft Usage Location Inference for identity escalations**

NEW 50

Huntress infers an identity's home country from the tenant when Entra usage location is absent, eliminating Unexpected Country and 'login without usage location' false escalations with no partner configuration needed.

– Names mechanism and specific escalation types but no config surface

changelog-20260804-eba7c64e

04 **Automatic onboarding escalation pacing** NEW 35

Managed ITDR now automatically paces escalation flow during a new tenant's first week to prevent onboarding noise spikes; this is on by default with no configuration required.

– Describes behavior but no configurable surface or metrics changelog-20260804-eba7c64e

WAS THIS USEFUL? 👍 👎

Sophos Core Agent

Feedback ⓘ 1 RELEASE · SEEN 2026-08-04

NOTES ↗

Sophos Core Agent is an endpoint protection platform component providing malware detection, threat response, and security monitoring across managed devices.

> Sophos Core Agent 2025.2.1.709 adds more granular CredGuard exclusion handling, DNS-based trusted-network detection for ZTNA to cut latency on local resources, and expanded removable storage device support.

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

01 **ZTNA trusted-network detection via DNS** NEW 55

ZTNA agents can now detect trusted networks via DNS configuration and bypass traffic interception, reducing latency for local resources such as CIFS and RDP (WINEP-57828).

— Explains mechanism and benefit but no setup instructions

Version 2025.2.1.709Windows 10 64-bit a…

02 **Process-level CredGuard exclusion in HitmanPro.Alert** NEW 40

Adds process-level exclusion capability in `HitmanPro.Alert` to support more specific exclusion of CredGuard detections (WINEP-56443).

— Names mechanism and ticket but no config steps Version 2025.2.1.709Windows 10 64-bit a…

03 **Kingston IronKey added to Secure Removable Storage** IMPROVED 35

Adds Kingston IronKey peripheral to the Secure Removable Storage classification in Sophos Endpoint Defense (WINEP-60202).

— Names specific device and feature but minor, no action needed

Version 2025.2.1.709Windows 10 64-bit a…

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

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

Tines is a no-code automation platform for security teams to build workflows that connect tools and respond to threats.

> Tines 42.3.0 shipped dynamic SMTP configuration, native JSON support in records, cross-story find-and-replace, and AI-generated story names and descriptions.

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

01 **Find and replace across stories** NEW 33

Find and replace content across stories for bulk editing workflows.

– Names the feature and use case but no exact UI steps. changelog-20260804-2193b6d8

02 **Dynamic SMTP management** IMPROVED 31

Manage SMTP settings dynamically without requiring static configuration changes.

– No config keys or UI path given, just a description. changelog-20260804-2193b6d8

03 **JSON support in records** NEW 31

Store and work with JSON natively inside Tines records.

– States the capability but no schema or field detail. changelog-20260804-2193b6d8

04 **AI-generated story names and descriptions** NEW 28

AI-generated story names and descriptions to accelerate documentation.

– Describes benefit but not how or where triggered. changelog-20260804-2193b6d8

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

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

Tenzir is a security data platform that ingests, transforms, and analyzes security telemetry and threat intelligence data.

> Tenzir Ship v1.12.0 makes release publishing safer and more flexible, adding a tag-only publish mode and pre-publish confirmation of all git actions.

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

01 **Publish release tag without GitHub release** NEW 50

Adds an option to publish a release tag without creating a GitHub release, for cases where another repository owns the release or downstream builds must validate the tag first.

— Explains use case but no flag or command name given

Tenzir Ship v1.12.0: Safer and more fle…

02 **Pre-publish confirmation of commits, tags, pushes** IMPROVED 30

Confirms every commit, tag, and push before making any changes during release publishing.

— Describes behavior but no config or flag named Tenzir Ship v1.12.0: Safer and more fle…

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

Open Cyber Threat Intelligence Platform

> OpenCTI's biggest shift this window is supply-chain transparency via per-release CycloneDX SBOM assets, alongside updated Map view styling, a stop-operation control for data sanity tasks, and tighter HTTP header handling in the Express layer.

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

01 **CycloneDX SBOM asset per release** NEW 85

Each release now ships a CycloneDX JSON SBOM file named `sbom-opencti-platform-<release version>.cdx.json`, generated by the `.github/workflows/cd-build-rolling-images-on-merge.yml` workflow. One SBOM asset is produced per Docker image base (standard, FIPS, and UBI9), with a nested component structure covering the platform, frontend, GraphQL API, and worker sub-components, enabling component verification, compliance reviews, and supply chain transparency.

— Names exact file pattern and generating workflow, but no consumption steps.

product docs

THINNER COVERAGE BELOW

02 **Dark and light mode Map view styles** IMPROVED 55

OpenCTI Map gains updated styling via two new style files, `filigran-dark3.json` for dark mode and `filigran-light3.json` for light mode, for the OpenCTI Map view.

— Names the two style files but not how they're applied. product docs

03 **Stop control for data sanity tasks** NEW 15

Adds a stop-operation control for data sanity tasks, per the release summary; no further mechanism or UI location was described.

— Only named in the summary, no bullet detail or surface given. 7.260803.0

04 **HTTP header management improvements** IMPROVED 15

Improves HTTP header management in the Express layer.

— No specifics on which headers or what changed. 7.260803.0

WAS THIS USEFUL? 👍 👎

◆ Managed Detection & Response (MDR)

Arctic Wolf

Feedback ⓘ 1 RELEASE · 2026-08-03 BLOG ↗

Arctic Wolf is a cloud-native security operations platform providing managed detection, response, and threat hunting services.

> Arctic Wolf introduced a bundled cyber resilience offering that combines security operations, exposure management, endpoint protection, and incident response with up to $3 million in incident warranty coverage.

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

01 **Integrated cyber resilience offering with incident warranty**

NEW 50

A new comprehensive cyber resilience offering bundles security operations, exposure management, endpoint protection, and incident response into a single integrated package, and includes up to $3 million in warranty protection coverage for cybersecurity incidents.

— Names bundled components and warranty figure but no operational detail or how to enroll.

Arctic Wolf Launches Comprehensive Cybe…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Intruder

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

Intruder is a cloud-based vulnerability scanner that automatically discovers and tests systems for security weaknesses.

> Intruder launched AI Pentesting, an on-demand penetration testing capability with whitebox code-repository integration that produces full PDF reports.

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

01 AI Pentesting with whitebox code-repo integration NEW

76

A new AI Pentesting capability is launched from the portal via left nav → AI Pentesting → Start new pentest → AI Pentest option. It supports whitebox testing through code repository integration, covering access control, business logic, and auth bypass issues that require multi-step chaining, and produces full PDF reports containing an executive summary, severity ratings, reproduction steps, and remediation guidance. Practitioners can define custom entry points, credentials, and application context to scope the test to their real environment.

Kick off an on-demand AI pentest ahead of an audit or security questionnaire without waiting weeks for a traditional engagement.

> 📍 In the portal, go to AI Pentesting in the left nav → click Start new pentest → select the AI Pentest option, then connect your code repository and define entry points, credentials, and application context to scope the test.

— Exact UI navigation given but no API/config surface or limits.

changelog-20260804-4f8b84fe

WAS THIS USEFUL? 👍 👎

Censys

Feedback ⓘ 1 RELEASE · 2026-08-03 BLOG ↗

Censys is a search engine for discovering and analyzing internet-connected devices, networks, and certificates for security research.

> Censys added OAuth authentication to its cencli command line tool and expanded platform detection with 3 new ASM risk fingerprints and 12 new software fingerprints.

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

01 **OAuth login for cencli via censys auth login** NEW 77

Adds OAuth authentication support to the `cencli` command line tool, triggered via `censys auth login`, which redirects to an authentication page to select an account and approve the connection. OAuth connections can be managed through the Connected Applications page in the Platform web console.

— Names exact command and console page for management. August 3, 2026 (2026-08-03)

02 **Three new ASM risk fingerprints** NEW 69

Adds 3 new ASM risk fingerprints: Exposed json-server Instance (unauthenticated CRUD REST API exposure), Exposed PTC Windchill Application (PLM platform reachable from public internet), and Exposed Express Default Page (unconfigured Express.js scaffold in production). New risk types may be disabled by default and are reviewable under risk type configuration in the ASM web console.

— Named risk types with UI path but no exact config key. August 3, 2026 (2026-08-03)

THINNER COVERAGE BELOW

03 **12 new software fingerprints on the Platform** NEW 42

Adds 12 new software fingerprints to the Platform, covering Bitwarden Server, json-server, nPerf Server, Alinto SOGo, TileServer GL, FMSoft uniGUI, Vaultwarden, PTC Windchill, DevExpress, OpenJS Foundation Express, SearXNG, and Vaadin.

— Lists software names but no mechanism or usage path. August 3, 2026 (2026-08-03)

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

defectdojo

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo 3.2 overhauls its Vulnerability model with per-ID type classification, multi-CWE support, and KEV/EPSS enrichment, adds 34+ new scanner parsers, and begins deprecating six API-pull parsers, Tool Type/Configuration, and dbbackup ahead of removal in 3.5.0.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01 Multiple CWEs per finding via Finding_CWE table NEW

85

Adds a `Finding_CWE` relationship table (migration `0284_finding_cwe`) so a finding can carry multiple CWEs; the primary CWE stays on `Finding.cwe` while additional CWEs are accessible via `finding.cwes`. An optional `manage.py migrate_cwe` command idempotently reconciles `Finding_CWE` rows from current `Finding.cwe` values after a parser upgrade.

After upgrading a parser that changes how CWEs are derived, re-sync Finding_CWE rows without re-importing scan data.

```
$ ./manage.py migrate_cwe
```

— Exact runnable management command plus schema detail. product docs

02 Hash_code recompute required for Burp Suite DAST Scan

75

BREAKING

Existing Burp Suite DAST Scan findings will not match new imports until hash codes are recomputed; reimporting before running the recompute will incorrectly close existing findings as absent and create duplicates. Recommended fix is `./manage.py dedupe --parser "Burp Suite DAST Scan" --hash_code_only` followed by `./manage.py dedupe --parser "Burp Suite DAST Scan" --dedupe_only`.

Recompute hash codes for Burp Suite DAST Scan findings before your next reimport to avoid false closures and duplicates.

```
$ ./manage.py dedupe --parser "Burp Suite DAST Scan" --
hash_code_only
./manage.py dedupe --parser "Burp Suite DAST Scan" --dedupe_only
```

— Exact remediation commands with flags provided. product docs

03 **Deprecation of Tool Type / Tool Configuration and related endpoints**

DEPRECATED 75

Deprecates Tool Type and Tool Configuration features along with the API endpoints `/api/v2/tool_types/`, `/api/v2/tool_configurations/`, and `/api/v2/product_api_scan_configurations/`; affected endpoints now return an X-Deprecated response header.

— Exact endpoint paths and header named, no removal date beyond 3.5.0. 3.2.0

04 **Per-vulnerability-ID type classification** NEW 70

Adds a `vulnerability_id_type` column to `Vulnerability_Id`, auto-detected from the identifier's prefix (`CVE`, `GHSA`, `RUSTSEC`, etc.), indexed, and backfilled by migration `0282_backfill_vulnerability_id_type`. The column is `NULL` when no non-numeric prefix exists.

— Names column, migration id, examples; migration runs automatically. product docs

05 **Per-parser deduplication algorithm overrides** IMPROVED 70

Adds `DEDUPLICATION_ALGORITHM_PER_PARSER` entries for Snyk Code Scan, Cycognito Scan, and `n0s1 Scanner` so deduplication now uses their configured `HASHCODE_FIELDS_PER_SCANNER` fields instead of the legacy algorithm.

— Names config keys and affected parsers, no usage step shown. product docs

06 **Deprecation of dbbackup/dbrestore management commands**

DEPRECATED 70

Deprecates `manage.py dbbackup` / `dbrestore` (the `django-dbbackup` integration) with removal in DefectDojo 3.5.0; recommends native `pg_dump` / `pg_restore` instead.

— Names deprecated and replacement commands with removal version. 3.2.0

07 **Deprecation of six API-based pull parsers** DEPRECATED 65

Deprecates the `BlackDuck API`, `Bugcrowd API Import`, `Cobalt.io API Import`, `Edgescan Scan`, `SonarQube API Import`, and `Vulners` pull parsers, with removal scheduled for DefectDojo 3.5.0 (November 2026).

— All six parsers and removal version named, no migration path given. 3.2.0

08 **Stable hash_code for AWS Security Hub findings** IMPROVED

60

The AWS Security Hub parser now sorts resource ARNs before joining them into the finding description, making `hash_code` stable across re-imports for findings with multiple resources; previously the hash was randomized per import, affecting false-positive history, risk acceptance copies, and similar-findings features tied to old hashes.

— Explains before/after behaviour but gives no remediation command. product docs

THINNER COVERAGE BELOW

09 **Deterministic component ordering in Yarn Audit parser**

IMPROVED

55

The Yarn Audit parser (Yarn 2) now sorts dependents and tree versions before joining them into `component_version` and the finding description, eliminating UI reshuffling between imports.

— Names parser and fields changed, no action needed. product docs

10 **Unique constraint on finding/vulnerability-id pairs** NEW

50

Adds a unique constraint on (finding, vulnerability_id) via migration `0283_unique_finding_vulnerability_id`, consolidating any pre-existing duplicate rows.

— Migration id given but no reader action required. product docs

11 **KEV/EPSS enrichment columns on findings** NEW

50

Introduces per-vulnerability KEV/EPSS enrichment columns on the Vulnerability entity, surfacing exploit and prioritization data directly on findings.

— Names KEV/EPSS and entity but no field names or UI path. 3.2.0

12 **New parsers for composer, pnpm, dotnet, and mix_audit** NEW

45

Adds parsers for `composer audit`, `pnpm audit`, `dotnet vulnerable packages`, and `mix_audit` output formats.

— Tool names given but no ingestion detail or command shown. 3.2.0

13 **SPDX, CSAF 2.0, and OpenVEX interchange parsers** NEW

45

Adds SPDX, CSAF 2.0, and OpenVEX interchange-format parsers for software bill-of-materials and VEX ingestion.

— Formats named but no field or endpoint detail given. 3.2.0

14 **Vulnerability entity normalization for cross-finding correlation**

IMPROVED

40

Normalizes `Vulnerability_Id` into a dedicated Vulnerability entity with ordered references, enabling cross-finding vulnerability correlation.

— Describes restructuring but no concrete usage or query surface. 3.2.0

15 **30 new CLI tool parsers across scan categories** NEW

30

Adds 30 new command-line cyber tool parsers spanning SAST, IaC, host hardening, web/network discovery, and malware categories.

— Count and categories given, no tool names or specifics. 3.2.0

└──▷ **BREAKING ON UPGRADE**

! Existing Burp Suite DAST Scan findings will not match new imports until hash codes are recomputed; reimporting before running `manage.py dedupe --parser "Burp Suite DAST Scan" --hash_code_only` will close existing findings as absent and create duplicates.

! AWS Security Hub Scan findings that carry two or more resources have a new stable `hash_code`; the previous value was random per import, so false-positive history, risk acceptance copies, and similar-findings features linked to old hashes are affected.

! API-based pull parsers, Tool Type/Configuration, and `dbbackup` are deprecated as of 3.2.0 with end-of-life set at 3.5.0; integrations relying on these surfaces should begin migration now.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Netskope

Feedback ⓘ 5 RELEASES · SEEN 2026-08-04 NOTES ↗

Netskope is a cloud security platform that inspects and controls data moving to cloud applications and services.

> Netskope shipped GA for Next-Gen API Data Protection across Box, Cisco Webex, ServiceNow and Slack Enterprise, released Cloud Exchange v5.1 with a new Unified Cloud Risk Exchange and full ticket-lifecycle orchestration, and introduced an AI agent that auto-generates Netskope Private Access app segments and policies, alongside smaller updates to RBI, SWG and NPA Browser Access.

└──▷ **WHAT SHIPPED · 22 FEATURES** most completely described first ? what's the number?

01 **CE v5.1 deployment and migration restrictions** BREAKING

80

CE v5.1 GA does not support the CE `Small` stack sizing, so deployments on the Small stack must meet new minimum system requirements before upgrading; migration from CE v5.0.x (HA) to CE v5.1 (Standalone) is not supported. Prior versions reach end of life on set dates: CE v3.x EOL/EOS 31 January 2024, CE v4.0.x and v4.1 EOL/EOS 31 July 2024, and CE v4.2 and v5.0.0 EOL/EOS 31 January 2025.

— Gives exact dates and migration constraints for action

Netskope Cloud Exchange (CE) v5.1 is NO…

02 **Unified Cloud Risk Exchange (CRE)** NEW

75

Adds Unified Cloud Risk Exchange (CRE) providing user and application risk context for adaptive zero trust policy enforcement, with integrations for Microsoft Entra, CrowdStrike Identity Protect, CrowdStrike Cloud Security (CNAPP), Okta, and Wiz.

— Names integrations and purpose but no setup steps

Netskope Cloud Exchange (CE) v5.1 is NO…

03 **Cloud Ticket Orchestrator v2.0.0 full lifecycle management**

IMPROVED

70

Upgrades Cloud Ticket Orchestrator (CTO) to support complete ticket/incident lifecycle management between Netskope and third-party ITSM tools including ServiceNow and Jira (`Ticket Orchestrator v2.0.0 Plugin`).

— Names plugin version and integrations, no config steps

Netskope Cloud Exchange (CE) v5.1 is NO…

04 **Netskope URL Lookup API (Controlled GA)** NEW 60

Adds Netskope URL Lookup API (Controlled GA) enabling programmatic URL classification queries.

— Named API but no endpoint, method, or auth detail

Netskope Cloud 119.0.0 Documentation Av…

05 **IoC Source Labeling in Cloud Threat Exchange** NEW 60

Adds IoC Source Labeling in Cloud Threat Exchange (CTE) to identify the origin of IoCs shared between Netskope and third-party Threat Intel Providers (TIPs) including CrowdStrike, SentinelOne, MISP, and Anomali.

— Names TIPs but no field or API to act on it Netskope Cloud Exchange (CE) v5.1 is NO…

06 **IoC Retraction in Cloud Threat Exchange** NEW 60

Adds IoC Retraction in Cloud Threat Exchange (CTE) to remove or roll back IoCs shared between Netskope and third-party TIPs such as CrowdStrike, SentinelOne, MISP, and Anomali.

— Names TIPs but no procedure for retraction Netskope Cloud Exchange (CE) v5.1 is NO…

THINNER COVERAGE BELOW

07 **Transaction Events Field - Format 4 (Beta)** NEW 55

Adds `Transaction Events Field - Format 4` (Beta) for Netskope Secure Web Gateway transaction event logging.

— Names exact field format, no usage instructions given

Netskope Cloud 119.0.0 Documentation Av…

08 **Multiple IdP support for NPA Browser Access** IMPROVED 55

Adds support for Multiple IDPs with Browser Access (Beta) in NPA; this later reached general availability, enabling more than one identity provider per deployment.

— Explains deployment scope but no configuration steps

Netskope Cloud 121.0.0 Documentation Av… Netskope Cloud 119.0.0 Documentation Av…

09 **Endpoint Events & Alerts visibility in Cloud Log Shipper** NEW 55

Adds Endpoint (Device & Content) Events & Alerts visibility to Cloud Log Shipper (CLS) via the `Log Shipper v2.0.0 Plugin`.

— Names plugin version but no field or query detail

Netskope Cloud Exchange (CE) v5.1 is NO…

10 **Next-Gen API Data Protection GA for major SaaS apps** NEW

50

General Availability of Next Generation API Data Protection for Box, Cisco Webex, ServiceNow, and Slack Enterprise.

– Names apps covered but no usage or config detail

Netskope Cloud 119.0.0 Documentation Av…

11 **Private Access Tunnel Browser Access with URI Auth Bypass** NEW

45

Adds Private Access Tunnel Browser Access with URI Auth Bypass (Controlled GA) for Netskope Private Access.

– Named feature but no configuration or use-case detail

Netskope Cloud 119.0.0 Documentation Av…

12 **AI agent for NPA app segment and policy generation** NEW

45

Introduces an AI agent within Netskope Private Access (NPA) that automatically generates app segments and policies.

– Describes capability but no mechanism or limits given

Netskope Cloud 140.0.0 Documentation Av…

13 **Microsoft Purview Information Protection write action GA** NEW

40

General Availability of Microsoft Purview Information Protection Write Capability Action in Next Generation API Data Protection.

– Names the action but no usage steps Netskope Cloud 119.0.0 Documentation Av…

14 **HTTP/2 tenant compatibility with RBI (Beta)** NEW

40

Adds HTTP/2 Enabled Tenant Compatibility with RBI (Beta) in Netskope Secure Web Gateway.

– Names the compatibility mode but no rollout mechanics

Netskope Cloud 119.0.0 Documentation Av…

15 **XRBI SKU scope and third-party auth flow isolation** NEW

40

Adds XRBI SKU Scope and isolation of third-party auth flows in Remote Browser Isolation (RBI).

– Names SKU and mechanism but no operational detail

Netskope Cloud 122.0.0 Documentation Av…

16 **Automatic file classification using Box Labels** NEW

35

Adds Automatic Classification of Files using Box Labels in Next Generation API Data Protection.

– Single bullet with no mechanism detail Netskope Cloud 119.0.0 Documentation Av…

17 **Browser Access for RUH1 in NPA** NEW 35

Adds Browser Access for RUH1 in Netskope Private Access.

– Names region code but no further detail Netskope Cloud 122.0.0 Documentation Av…

18 **WebTX v2.0.0 Plugin for Cloud Log Shipper** NEW 35

Introduces new `WebTX v2.0.0 Plugin` for Cloud Log Shipper.

– Named plugin version but no functional description

Netskope Cloud Exchange (CE) v5.1 is NO…

19 **NGWeb migration for NPA Publishers** IMPROVED 30

Adds NGWeb migration for Netskope Private Access (NPA) Publishers.

– Thin bullet with no migration mechanics described

Netskope Cloud 119.0.0 Documentation Av…

20 **New inline file constraints in NG SWG** NEW 30

Adds new Inline File Constraints to the Netskope Secure Web Gateway (NG SWG).

– No specifics on what constraints were added Netskope Cloud 122.0.0 Documentation Av…

21 **Netskope Tenant Plugin for Cloud Exchange** NEW 30

Introduces new Netskope Tenant Plugin for Cloud Exchange.

– Named plugin but no functional description Netskope Cloud Exchange (CE) v5.1 is NO…

22 **Cloud TAP general availability** IMPROVED 25

Promotes Cloud TAP to general availability (GA).

– No description of what Cloud TAP does Netskope Cloud 121.0.0 Documentation Av…

└──▷ **BREAKING ON UPGRADE**

! CE v5.1 GA does NOT support the CE `Small` stack sizing; deployments using the Small stack must meet the new minimum system requirements before upgrading.

! Migration from CE v5.0.x (HA) to CE v5.1 (Standalone) deployment is NOT supported.

! CE v3.x reached EOL/EOS on 31 January 2024; CE v4.0.x and CE v4.1 reached EOL/EOS on 31 July 2024; CE v4.2 and CE v5.0.0 will reach EOL/EOS on 31 January 2025.

WAS THIS USEFUL? 👍 👎

Varonis

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

Varonis is a data security platform that detects and responds to threats targeting sensitive data across cloud and on-premises environments.

> Varonis Atlas introduces Agent IBAC (Intent-Based Access Control), a new capability to detect AI agent intent drift and automatically quarantine rogue agent identities in real time.

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

01 **Agent Intent-Based Access Control (IBAC)** NEW 48

New Agent Intent-Based Access Control (IBAC) in Varonis Atlas validates that an AI agent's reasoning, tool use, and data access stay consistent with its original instructions, alerting on or blocking out-of-policy actions. Intent drift detection monitors every agent action across sessions and provides a complete audit trail for every agent interaction.

.
.
.
.

— Names the feature and mechanism but no config or API surface given

Detect agent intent drift and rogue act…

02 **Quarantine rogue agent identities** NEW 38

When an agent crosses policy boundaries, Atlas can quarantine the agent's identity and block all subsequent actions for a defined window.

— Describes response action but no exact trigger or duration values

Detect agent intent drift and rogue act…

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

SailPoint

Feedback ⓘ 3 RELEASES · 2026-07-22 → 2026-08-03 BLOG ↗

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

> SailPoint launched Agentic Fabric, a new governance and protection product for AI agents and non-human identities spanning 32+ frameworks, and separately shipped rule-based auto-discovery of privileged entitlements and free-form ad hoc values for Access Model Metadata attributes.

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

01 Ad hoc values for Access Model Metadata attributes NEW

95

A new Allow Ad Hoc Values option on custom metadata attributes lets admins assign free-form values up to 100 characters, with no pre-defined list, to roles, access profiles, and entitlements. A companion Allow Multivalued Assignment option permits multiple values per attribute (irreversible once enabled). Capacity rises to up to 5,000 values per custom attribute and up to 250 custom metadata attributes per access object type, roles/profiles/entitlements become searchable by ad hoc value, and previously entered ad hoc values persist even if the option is later disabled.

Configure an existing AMM attribute to accept free-form project names or application IDs sourced from a CMDB, without maintaining a static list.

📍 In the ISC console, go to Admin → Access Model Metadata, open the target attribute, enable Allow Ad Hoc Values, and save. Then open any role, access profile, or entitlement and assign values directly in the metadata field.

— Named config options, exact limits, and a runnable configuration path.

New Capability: Access Model Metadata A…

02 Agentic Fabric workspace and unified NHI registry NEW

75

A new Agentic Fabric tab in the top navigation (MIS and AIS tenants) opens a dedicated governance workspace for AI agents and non-human identities, with new Applications and Accounts left-nav items consolidating machine accounts (IAM roles) and machine/application identities into one view. It is backed by a unified NHI registry for discovering and governing agents across enterprise, endpoint, and browser surfaces, with an embedded Identity Graph for visualizing relationships between agents, identities, and resources.

Review all machine accounts and NHIs in one place to assess IAM role exposure across your environment.

> 📍 In the ISC console, open the Agentic Fabric tab in the top navigation, then select Accounts from the left nav to view all machine accounts (IAM roles) and machine/application identities.

— Named tabs and nav items with exact navigation, but no API or flag.

New Product: SailPoint Agentic Fabric (…

---

03 **Agent provisioning approval routing** NEW 75

Approval routing for agent provisioning (activation and deactivation) is configurable under `Global → System Settings → Approval Settings → Agent Requests`, letting admins require approval before agents are provisioned or deprovisioned.

Navigate to agent approval routing to configure who must approve activation or deactivation of AI agents before provisioning proceeds.

> 📍 In the ISC console, go to Global → System Settings → Approval Settings → Agent Requests and configure approval routing for agent provisioning operations.

— Exact config path plus a runnable navigation example.ed

New Product: SailPoint Agentic Fabric (…

---

04 **Multi-framework and cross-surface agent discovery** NEW 65

Agentic Fabric supports 32+ agent frameworks from day one across three surfaces: endpoint agents (Claude Code, Claude Desktop, Cursor, Cline, Windsurf, GitHub Copilot, OpenAI Codex, Gemini CLI, and others), enterprise agents (Amazon Bedrock, Amazon Bedrock AgentCore, Anthropic Claude Console, Databricks Mosaic AI, Google Vertex AI, Microsoft Copilot Studio, Microsoft Foundry, Salesforce Agentforce, ServiceNow AI Platform, Snowflake Cortex AI, N8N, WIZ, and custom agents via BYO schema/API), and browser agents (Zendesk, MS Office Copilot 365, Cursor, Dify, Claude, OpenAI). Endpoint and browser discovery capabilities detect agents running on managed devices and in browsers.

— Extensive named framework list but no way to invoke it directly.

New Product: SailPoint Agentic Fabric (…

---

05 **Owner correlation and lifecycle controls for AI agents** NEW 60

Agentic Fabric adds owner correlation and lifecycle controls for AI agents, with an optional default owner mapping configurable in Datasets and Resources for automated ownership assignment.

— Names config location but gives no exact steps or field names.

New Product: SailPoint Agentic Fabric (…

06 **AWS IAM roles integration via AWS SaaS connector** NEW

60

AWS IAM Roles / Datasets & Resources integration is available via the AWS SaaS connector, with automatic migration for existing MIS and AIS customers who already have AWS configured.

— Names the connector and migration behavior but no setup steps.

New Product: SailPoint Agentic Fabric (…

THINNER COVERAGE BELOW

07 **Auto-discovery of privileged entitlements** NEW 55

SailPoint surfaces recommendations for potentially privileged entitlements based on aggregate privilege classification pattern analysis, covering access outside vendor-published lists or custom criteria. Recommendations appear on existing entitlement and source pages, where administrators can accept or reject each one and assign a privilege level, or delegate the decision to entitlement owners, source owners, governance groups, or other subject matter experts via the existing approvals flow; all classification decisions are kept in an auditable record.

— Names mechanism and workflow but no exact UI path or config key.

New Capability: Auto-Discovery of Privi…

08 **AI agent threat protection capabilities** NEW 35

As part of the Protect capability set, Agentic Fabric adds CLI threat detection and remediation for endpoint agents, prompt and response security monitoring, and agent authorization controls.

— Capabilities named but no mechanism, limits, or usage path given.

New Product: SailPoint Agentic Fabric (…

09 **Agent audit and compliance reporting** NEW 20

Agentic Fabric adds agent audit and compliance reporting for audit-ready governance of AI agent activity.

— Only a one-line description with no detail on scope or format.

New Product: SailPoint Agentic Fabric (…

WAS THIS USEFUL?  

Auth0 Feedback ⓘ 1 RELEASE · 2026-08-04 BLOG ↗

Auth0 is a cloud-based identity and access management platform providing authentication, authorization, and user management solutions for applications.

> Auth0 expanded Organizations Search with new filter types and the ability to combine multiple filters, plus shareable filtered views.

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

01 **Advanced filtering in Organizations Search** NEW 80

The Organizations Search interface now supports filtering by Name, Display name, ID, Metadata (custom key-value pairs), Third-party client access (allow or block), and App entitlement active status. Up to 5 filters can be combined simultaneously using AND/OR logic, and filter pills persist in the URL, enabling shareable and bookmarkable filtered Organization views.

.
.
.
.

Find all organizations where a specific metadata key is set and third-party client access is blocked — useful for auditing tenant isolation policies.

> 📍 In the Auth0 Dashboard, go to Organizations › search bar, add a Metadata filter (e.g. key: tier, value: enterprise), then add a Third-party client access filter set to block, and combine them with AND logic.

— Names exact filter types, combination logic, and UI path with example

Organizations Search Expands with Advan…

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Termius

Feedback ⓘ 9.42.2 (2026-08-03) BLOG ↗

Termius is a cross-platform SSH client and terminal emulator for secure remote server access and management.

Termius 9.42.2 adds host context-menu actions to open or add selected hosts to workspaces.

Adds the ability to open selected hosts in a new workspace or add them to an open one directly from the Hosts screen context menu.

Makes app updates less disruptive.

└─▷ **BREAKING ON UPGRADE**

! MultiKey is being sunset and will no longer be available after October 15, 2026; migrate to SSH ID or another authentication method before that date.

WAS THIS USEFUL? 👍 👎

tailscale

Feedback ⓘ 1 RELEASE · 2026-08-03 BLOG ↗

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale v1.102.1 adds three new CLI commands, per-connection Serve metrics, a macOS deep-linking protocol handler, several iOS/tvOS client upgrades, and performance and certificate-renewal improvements, while dropping support for deprecated 4via6 MagicDNS name formats.

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

01 **`tailscale get` CLI command** NEW 65

The new `tailscale get` command returns the current node's preferences, letting a user inspect its configuration directly from the command line.

Inspect the current node's configured preferences from the command line.

```
$ tailscale get
```

— Named command with runnable example, minimal mechanism detail.

Tailscale v1.102.1 (2026-08-03)

02 **`tailscale whoami` CLI command** NEW 65

The new `tailscale whoami` command displays information about the current user and device, allowing identification of the authenticated session without opening the GUI.

Identify the currently authenticated user and device without opening the GUI.

```
$ tailscale whoami
```

— Named command with runnable example, minimal mechanism detail.

Tailscale v1.102.1 (2026-08-03)

03 **`tailscale service list` CLI command** NEW 65

The new `tailscale service list` command lists Tailscale Services visible to the current node, useful for auditing what is exposed on the tailnet.

List all Tailscale Services visible to this node to audit what's being exposed on the tailnet.

```
$ tailscale service list
```

— Named command with runnable example, minimal mechanism detail.

Tailscale v1.102.1 (2026-08-03)

04 **Tailscale Serve traffic metrics** NEW 60

New `tailscaled_serve_outbound_bytes_total` and `tailscaled_serve_inbound_bytes_total` client metrics report bytes sent to and received from peers on Tailscale Serve connections.

— Named metrics but no scrape/config path given. Tailscale v1.102.1 (2026-08-03)

THINNER COVERAGE BELOW

05 **macOS deep-linking protocol handler** NEW 45

A new macOS protocol handler supports deep-linking to devices, exit nodes, and settings panels directly within the application window.

— Describes capability but no URL scheme or invocation shown.

Tailscale v1.102.1 (2026-08-03)

06 **Proactive TLS certificate auto-renewal on idle servers**

IMPROVED 45

TLS certificates on idle servers now proactively auto-renew even in the absence of traffic, with warnings issued when no valid cached certificate exists.

— Explains behaviour and warning condition, no config surface.

Tailscale v1.102.1 (2026-08-03)

07 **Exit node subtitles and filtering in Shortcuts** IMPROVED 45

Exit node names now appear as subtitles in Shortcuts on macOS and iOS and can be used to filter the displayed exit node list.

— Names integration and behaviour but no exact steps. Tailscale v1.102.1 (2026-08-03)

08 **iOS status widget exit node toggle** NEW 40

The new iOS status widget supports toggling the active exit node inline, without opening the full app.

— Named UI surface but no exact widget setup steps. Tailscale v1.102.1 (2026-08-03)

09 **TLS-ALPN-01 for Funnel domains** IMPROVED 40

Tailscale Funnel domains switch to TLS-ALPN-01 for faster HTTPS certificate renewals.

– Names protocol and product area but no config or command.

Tailscale v1.102.1 (2026-08-03)

10 **iPad and wide iPhone split-plane layout** NEW 35

New split-plane layout on iPad (portrait and landscape) and wide iPhones (landscape) for the iOS app improves the interface on larger screens.

– UI change named but no navigation path given. Tailscale v1.102.1 (2026-08-03)

11 **Constant-time node membership updates** IMPROVED 35

Node additions and removals are now processed in constant time, significantly reducing CPU usage on large tailnets.

– Explains mechanism and benefit but no metrics or scope named.

Tailscale v1.102.1 (2026-08-03)

12 **Synology RMv7 binaries for NAS models** IMPROVED 35

Synology RMv7 binaries with software floating point are now produced for certain NAS models, replacing older ARMv5 binaries.

– Names architecture but not which specific NAS models. Tailscale v1.102.1 (2026-08-03)

13 **Removal of deprecated 4via6 MagicDNS name formats** BREAKING 35

Deprecated formats for 4via6 MagicDNS names are no longer available in this release.

– States removal but no migration path or old format shown.

Tailscale v1.102.1 (2026-08-03)

14 **ARM32 assembly crypto routines** IMPROVED 30

New assembly crypto routines for 32-bit ARM on Linux improve cryptographic performance on those platforms.

– Names platform but no benchmark or config detail. Tailscale v1.102.1 (2026-08-03)

15 **tvOS remote management via web client** NEW 30

New Tailscale web client support enables remote management of Apple TV devices on tvOS.

– States capability without mechanism or access path. Tailscale v1.102.1 (2026-08-03)

16 **Linux userspace TUN throughput optimizations** IMPROVED 20

Userspace TUN optimizations on Linux improve throughput and performance.

– Generic statement with no benchmark or mechanism detail.

Tailscale v1.102.1 (2026-08-03)

└──▷ **BREAKING ON UPGRADE**

! Deprecated formats for 4via6 MagicDNS names are no longer available.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ **Integrated & IT Risk Management**

CISO Assistant

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 150+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

> CISO Assistant's v3.20.7 release lets users rename qualifications to risk categories and enriches the context available to in-window AI chat sessions.

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

01 **Rename qualifications to risk categories** IMPROVED 30

CISO Assistant now allows renaming of qualifications to risk categories.

– Names the change but no mechanism or UI path given v3.20.7

02 **Expanded context for AI chat** IMPROVED 20

The in-window AI chat sessions now receive more context.

– Vague on what context was added or how v3.20.7

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Cursor

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

Cursor is an AI-powered code editor built on VS Code that helps developers write, edit, and debug code faster using artificial intelligence.

> Cursor added five new Google Workspace plugins — Gmail, Drive, Calendar, Docs, and Sheets — giving coding agents direct read/write access to Google services from inside the editor.

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

01 **Google Drive plugin for agents** NEW 50

A new Google Drive plugin lets agents search files and folders, open and download content, and create and organize files from within Cursor. It can be installed from the Cursor Marketplace or the Customize page in Cursor.

— Names actions and install path, no API or flag detail Google Workspace Plugins

02 **Gmail plugin for agents** NEW 50

A new Gmail plugin lets agents search and read mail, draft and send messages, apply labels, and manage threads. It can be installed from the Cursor Marketplace or the Customize page in Cursor.

— Names actions and install path, no deeper mechanism given Google Workspace Plugins

03 **Google Calendar plugin for agents** NEW 50

A new Google Calendar plugin lets agents read schedules, create and update events, and find free time. It can be installed from the Cursor Marketplace or the Customize page in Cursor.

— Names actions and install path, lacks further detail Google Workspace Plugins

04 **Google Docs plugin for agents** NEW 50

A new Google Docs plugin lets agents open and read documents, write and edit text, and create new docs. It can be installed from the Cursor Marketplace or the Customize page in Cursor.

— Names actions and install path, no further mechanism Google Workspace Plugins

05 **Google Sheets plugin for agents** NEW 50

A new Google Sheets plugin lets agents read ranges and find values, update cells, and create and edit spreadsheets. It can be installed from the Cursor Marketplace or the Customize page in Cursor.

— Names actions and install path, lacks further specifics Google Workspace Plugins

WAS THIS USEFUL? 👍 👎

Claude Code

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

Claude Code is an agentic coding tool that lives in your terminal, understands your codebase, and helps you code faster by executing routine tasks, explaining complex code, and handling git workflows - all through natural language commands.

> Claude Code v2.1.221 adds sandbox credential masking, a VSCode Focus view, and a prompt-audit subcommand for the claude-api skill, alongside smarter plugin lifecycle handling, background-session git behavior, and a batch of smaller fixes across Windows startup, Gateway validation, and Stats reporting.

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

01 **Sandbox credential masking mode** NEW 95

Adds `mode: "mask"` for sandbox credential files on Linux and WSL — sandboxed commands read a sentinel copy (whole file or spans captured by an `extract` regex) while the sandbox proxy substitutes the real value on egress; on macOS this falls back to `deny`.

Configure a sandbox credential file to use masking on Linux/WSL so sandboxed commands see a sentinel value while the proxy substitutes the real credential on egress.

yaml

```yaml
mode: "mask"
extract: "(sk-[A-Za-z0-9]+)"
```

— Named config key, mechanism, and platform behaviour with example v2.1.221

02 **Focus view in VSCode extension** NEW 80

Adds Focus view: a chat-menu toggle that hides tool activity behind an expandable per-turn summary with a live running-tool indicator, toggled with Ctrl+Alt+F or the Claude Code: Toggle Focus view command.

— Exact shortcut and command name given, no screenshot v2.1.221

03 **prompt-audit subcommand for claude-api skill** NEW 70

Adds `prompt-audit` subcommand to the `claude-api` skill for auditing prompts and tool descriptions for patterns written for older models.

Audit your prompts and tool descriptions for patterns written for older models before migrating to Claude 4.5-generation.

```
$ claude-api prompt-audit
```

— Named command with runnable example, mechanism unstated v2.1.221

04 **Marketplace name validation warnings** IMPROVED 70

Adds warnings to `claude plugin validate` when a marketplace or plugin name would be rejected by Claude Desktop's managed marketplace sync.

Validate a plugin and catch names that would be rejected by Claude Desktop's managed marketplace sync.

```
$ claude plugin validate
```

— Named command with example, limited on trigger conditions v2.1.221

05 **Windows startup avoids PowerShell spawn** IMPROVED 65

Improves Windows startup by reading process creation times via a native `kernel32` call instead of spawning PowerShell, so endpoint security tools that gate `powershell.exe` no longer prompt.

— Names mechanism and API but no user action needed v2.1.221

06 **/status shows session kind** IMPROVED 60

Changes `/status` to show the session kind: `interactive`, or a background job that is `attached` or `unattended`.

— Named command and states, no example shown v2.1.221

07 **Background sessions commit and push work** IMPROVED 60

Changes background sessions to commit and push to preserve work, open a draft PR only when the task calls for one, follow `CLAUDE.md` git instructions, and always report where the work lives.

— Names CLAUDE.md and behaviour, no exact trigger rules v2.1.221

08 **Gateway model field rejects non-string** BREAKING 60

Changes Gateway `model` field validation to reject non-string values with a 400 instead of forwarding them.

— Names field, status code, and endpoint behaviour v2.1.221

THINNER COVERAGE BELOW

09 **/fork creates its own worktree** IMPROVED 55

Changes sessions forked with `/fork` to create a new worktree of their own instead of working in the original session's checkout.

— Named command, behaviour change clear, no example v2.1.221

10 **Plugin skills path accepts dot** IMPROVED 55

Changes plugins to accept `"."` as a `skills` path, with the root-level `SKILL.md` validation error now suggesting using the plugin root.

— Named config field, edge-case detail, no example v2.1.221

11 **/plugin install retries stale marketplace catalog** IMPROVED 55

Changes `/plugin install` to refresh a stale marketplace catalog and retry before reporting a plugin not found.

— Named command, behaviour explained, no example v2.1.221

12 **Plugins activate immediately when safe** IMPROVED 55

Changes plugins installed from `/plugin` to activate immediately when safe, instead of always requiring `/reload-plugins`.

— Named commands, condition of 'safe' unspecified v2.1.221

13 **Stats panel counts cache tokens** IMPROVED 55

Improves Stats panel to count cache tokens in its token totals, with a breakdown by input, output, cache read, and cache write.

— Names breakdown categories, no navigation path given v2.1.221

14 **Tool search re-enabled on Vertex AI** IMPROVED 45

Re-enables tool search on Google Vertex AI for Claude 4.5-generation and newer models.

— Names platform and model generation, thin on mechanism v2.1.221

15 **Emoji autocomplete alternate shortcodes** IMPROVED 45

Changes emoji autocomplete to accept common alternate shortcodes like `:thumbsup:`, `:thumbsdown:`, and `:love:`.

— Names shortcodes but minor cosmetic feature v2.1.221

16 **Fast mode reports credit exhaustion** IMPROVED 35

Changes fast mode to report on the stream when usage credits run out mid-session, instead of failing silently.

— Behaviour described, no surface name or example v2.1.221

17 **Monitor reports empty watch output** IMPROVED 35

Changes Monitor: a watch that exits without producing any output now says so instead of reporting `stream ended`.

— Minor message change, thin detail v2.1.221

WAS THIS USEFUL? 👍 👎

OpenHands

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

OpenHands: AI-Driven Development

> OpenHands shipped a preview desktop app for the Agent Canvas, a new plugin inspection view, per-purpose LLM profile selection, and backend improvements including live agent activity in chat and a domain-neutral extension-manifest host.

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

01 **Agent Canvas desktop app preview** NEW 95

A preview desktop app for macOS (Apple silicon) and Windows bundles Node.js and required runtimes so no prerequisites or open terminal are needed; download `Agent-Canvas-<version>-arm64.dmg` (macOS) or `Agent-Canvas-Setup-<version>.exe` (Windows) from the OpenHands releases page. The app auto-starts its local backend on launch — select Show details during startup to view and copy the live startup log for troubleshooting — and supports in-app update checking and version display via `Settings > Application`. It is also reachable via a Try the desktop preview link.

— Names exact installer files, download source, and troubleshooting UI path product docs

02 **LLM profile selector for conversation titles** NEW 70

A dedicated LLM profile selector for conversation title generation is available under `Settings > Application > Conversation titles`, with an Automatic option to use the active local profile or a specific saved profile.

— Exact settings path and option names given product docs

03 **Plugin detail view in agent canvas** NEW 60

The agent canvas gains a plugin detail view listing bundled skills (including command-derived skills) with their icons, names, and descriptions, plus an expandable directory tree with inline syntax-highlighted file viewing — select a file to open it, select again to close.

— Describes mechanism but no exact navigation path given product docs

THINNER COVERAGE BELOW

04 **Compact backend chooser in UI** IMPROVED 35

A compact Cloud vs Agent-server backend chooser is available in the UI for faster backend selection.

— Names the two backend options but no exact path v1.9.0

05 **Domain-neutral extension-manifest host** NEW 30

A domain-neutral extension-manifest host has been introduced, enabling extensions to be served without domain-specific coupling.

— No endpoint, config key, or usage detail provided v1.9.0

06 **Live agent activity in chat view** NEW 25

Live agent activity is now shown directly in the chat view so users can monitor what the agent is doing in real time.

— No named UI surface or mechanism beyond brief description v1.9.0

WAS THIS USEFUL?  

AGENT

◆ AI Agent Frameworks

Hermes

Feedback ⓘ 1 RELEASE · 2026-08-03 NOTES ↗ CODE ↗

The agent that grows with you

> Hermes shipped a complete onboarding/install CLI (portal login, model picker, gateway setup, doctor, desktop launch), a wave of in-session power commands (shell escape, prompt stash, `/init`, `/diff`, `/context`, `/focus`), streaming conversational voice with barge-in and wake-word detection, an Agent-to-Agent v1.0 plugin, signed outbound webhooks, a grounded-citations skill, and a plugin SDK, alongside deeper context-compression controls and measurable performance fixes.

└──▷ **WHAT SHIPPED · 45 FEATURES** most completely described first ? what's the number?

01 **`!<command>` shell mode** NEW 96

Prefixing a line with `!` (e.g. `!git status`, `!pytest -x tests/cli`) runs it as a shell command without invoking the model — no API call, no tokens, no latency, and no history entry. Dangerous commands (e.g. `rm -rf`) still go through the approval prompt, non-zero exits print `! exited <code>`, and shell mode is CLI-only — Gateway platforms (Discord, Telegram, Slack) and cron runs ignore it.

Instantly run a shell command from inside a Hermes CLI session without consuming a model turn.

```
$ !git status
```

— Exact syntax, edge cases, and a runnable example. v2026.8.3

02 **Ctrl+S prompt stash** NEW 94

Ctrl+S parks the current composer draft and clears the input so you can send something else first; pressing Ctrl+S again on an empty composer restores the draft (with cursor position and attached images). Repeated presses build a stack — with two or more stashed drafts, Ctrl+S opens a browse panel (↑ / ↓ to navigate, Enter to restore, D to discard, Esc / Ctrl+S to close), and a 📌 N badge in the status bar shows the number of parked drafts.

— Full keybinding behaviour and UI feedback described. product docs

03 **Compression timeout caps** NEW 93

`context_timeout_seconds` caps the inactivity budget for in-agent `compress_context` operations (conversation loop, preflight compaction, manual

compress), settable to `0` to disable; `context_total_ceiling_seconds` sets an absolute pre-commit ceiling on in-agent compression even while summary tokens are still streaming, clamped to at least `context_timeout_seconds`.



Prevent a slow or hung summary model from stalling a session indefinitely by setting both a streaming ceiling and an inactivity cutoff.

```yaml
context_timeout_seconds: 30
context_total_ceiling_seconds: 90
```

— Two named keys with defaults, clamping, and example. product docs

04 **hermes setup --portal one-command onboarding** NEW 88

`hermes setup --portal` logs in via the Nous Portal, sets Nous as the LLM provider, and enables the Tool Gateway (web search, image generation, TTS, cloud browser) in a single command.

Fastest onboarding: authenticate with Nous Portal, set the provider, and enable the full Tool Gateway (search, image gen, TTS, cloud browser) in a single step.

```
$ hermes setup --portal
```

— Exact command with full onboarding scope given. product docs

05 **Install layout detection and deployment modes** NEW 87

The installer auto-detects install method (git installer, Docker, or NixOS) and surfaces matching `hermes update` commands and an environment summary via `hermes`

`doctor`. It supports a root-mode FHS layout (`/usr/local/lib/hermes-agent/`, `/usr/local/bin/hermes` via `sudo curl … | sudo bash`) for shared machines, a per-user layout (code at `~/.hermes/hermes-agent/`, binary symlink at `~/.local/bin/hermes`), and graceful degradation for non-sudo/system-service-user installs — skipping `--with-deps`, installing Chromium into the user's local Playwright cache, and printing the admin command needed for system libraries.

— Concrete file paths and commands for each layout. product docs

06 **`--skip-browser` install flag** NEW 83

The install script supports `--skip-browser` to bypass Playwright/Chromium installation for headless deployments, e.g. `curl -fsSL https://hermes-agent.nousresearch.com/install.sh | bash -s -- --skip-browser`.

Install Hermes on a headless server or unprivileged service account where browser automation is not needed, skipping the Playwright/Chromium download entirely.

```
$ curl -fsSL https://hermes-agent.nousresearch.com/install.sh | bash -s -- --skip-browser
```

— Exact flag and full runnable install command. product docs

07 **`agent.session_stall_timeout` watchdog** NEW 83

`agent.session_stall_timeout` (default `0`, disabled) is a notify-only gateway watchdog that logs a WARNING and sends a one-shot user notification when a session with a pending inbound follow-up has been idle beyond the threshold.

Enable the session stall watchdog so users in busy multi-session deployments are notified when the agent appears wedged, without killing the turn.

```yaml
agent:
  session_stall_timeout: 120  # seconds; 0 disables
```

— Named key, default value, and exact behaviour. product docs

08 **`hermes config set/get/check/migrate` commands** NEW 81

`hermes config set` and `hermes config get` read and write individual config values (e.g. `hermes config set OPENROUTER_API_KEY your_key`); `hermes config check` and `hermes config migrate` inspect and migrate configuration after updates.

— Exact commands with a concrete example value. product docs

09 **`terminal.font_family` config key** NEW 78

`terminal.font_family` sets the embedded terminal font, accepting a locally installed family name or CSS font stack with JetBrains Mono as fallback; also editable at Settings → Appearance → Terminal Font.

— Named config key plus equivalent UI path. product docs

10 **max_concurrency cap for auxiliary LLM calls** NEW 78

`max_concurrency` caps auxiliary compression LLM calls and title-generation calls to prevent retry bursts from overwhelming rate-limited providers across multi-session deployments such as Discord/Telegram channels and multiple terminals.

Cap auxiliary LLM concurrency to avoid amplifying retries during a provider incident across many simultaneous sessions.

```yaml
# under your compression / title-generation auxiliary config
max_concurrency: 2
```

— Named key, scope, and example config. product docs

11 **Signed outbound webhooks** NEW 68

Signed outbound webhooks push HMAC-verified lifecycle events (session activity, turn completions, tool events) to any registered HTTP endpoint — no polling required.

— Signing method and event types named, no registration endpoint given. v2026.8.3

12 **hermes doctor diagnostics** NEW 67

`hermes doctor` surfaces environment diagnostics, the detected install method, and actionable fixes for a broken install.

Diagnose a broken install — surfaces missing dependencies, detected install method, and exact remediation steps.

```
$ hermes doctor
```

— Named command with runnable example. product docs

13 **Automatic dependency provisioning** NEW 65

The installer automatically provisions all dependencies — Python 3.11 (via uv), Node.js v22, ripgrep, and ffmpeg — with no manual prerequisite setup required beyond `git` (and `curl` / `xz-utils` on Linux).

— Named dependency versions, no reader action needed. product docs

14 **latency status-bar field** NEW 65

`latency` is an opt-in status-bar field showing wall-clock turn duration (e.g. `22s` / `1m05s`) alongside the existing default set of `["model", "context_pct", "cwd"]`.

— Named field and format, no toggle command given. product docs

15 **/init AGENTS.md generator** NEW 65

`/init` scans a project and generates or updates an `AGENTS.md` file.

Bootstrap a new project by having Hermes scan it and generate an AGENTS.md with its conventions and structure.

```
$ /init
```

— Named command and output file, runnable example. v2026.8.3

16 **Voice note support across messaging platforms** NEW 65

Voice note support is added on WhatsApp, Feishu, DingTalk, LINE, QQ, Photon, and Weixin, with auto-TTS replies delivered in the platform-appropriate codec (opus with captions).

— Named platforms and codec, no enabling steps. v2026.8.3

17 **grounded-citations skill** NEW 65

The `grounded-citations` skill matches quotes against actual page text, links citations to exact evidence, and includes a fact-checking mode for any document or claim.

— Named skill and mechanism, no invocation shown. v2026.8.3

18 **Plugin SDK with ctx.download** NEW 65

A plugin SDK adds `ctx.download` for handing users files, floating pane placement, and multiple GUI windows; it includes Kanban as the founding plugin.

— Named API surface and example plugin, no docs link. v2026.8.3

19 **hermes import-agent migration command** NEW 64

`hermes import-agent` migrates a Claude Code or Codex CLI setup into Hermes in one command.

— Named command and source tools, mapping unspecified. v2026.8.3

20 **hermes gateway setup for messaging platforms** NEW 62

`hermes gateway setup` connects messaging platforms such as Telegram, Discord, Slack, and others to Hermes.

— Command and platforms named, no config detail. product docs

21 **hermes tools command with STT category** IMPROVED 62

`hermes tools` configures which tools are enabled; a later release adds a dedicated STT configuration category to this command with GUI toggles and dashboard dropdowns.

— Command and new category named, UI mechanics only sketched. v2026.8.3

22 **/personality none command** NEW 62

`/personality none` (also accepts `default` or `neutral`) removes any active personality overlay and returns to the default.

— Exact command and accepted aliases given. product docs

23 **Streaming conversational voice mode** NEW 62

Streaming conversational voice mode adds clause-by-clause TTS, barge-in interruption, and busy-aware silence detection across the CLI, desktop, and gateway adapters.

— Mechanism described, no enabling command named. v2026.8.3

24 **hermes model interactive provider selection** NEW 60

`hermes model` interactively lets you choose an LLM provider and model, including custom endpoints for local or self-hosted models.

— Command named, selection flow not detailed. product docs

25 **On-device wake-word detection** NEW 60

On-device open-vocabulary wake-word detection adds multi-profile voice routing and a `stop` keyword to end voice chat on any surface — no audio leaves the device while waiting.

— Mechanism and privacy behaviour given, no config surface. v2026.8.3

26 **Desktop artifacts with sandboxed preview** NEW 60

Desktop artifacts adds versioned cards with sandboxed live preview in a right-rail viewer so generated HTML/apps run safely beside the chat.

— UI mechanism described, no exact navigation given. v2026.8.3

27 **Faster cold start and update times** IMPROVED 60

`hermes -w` cold start is reduced from ~14 s to ~1.8 s, and `hermes update` no-ops are 2–6 s faster.

— Concrete before/after numbers and named command. v2026.8.3

THINNER COVERAGE BELOW

28 **Per-model compression-threshold configuration** IMPROVED

59

Compression-threshold configuration is now settable per-model and in absolute tokens, with a guaranteed N-user-message tail so recent conversation always survives pruning.

— Mechanism named, no exact config key given. v2026.8.3

29 **Approval circuit breaker and docker/podman gate** NEW

59

A consecutive-denial circuit breaker stops misbehaving approval loops, plus a new approval gate specifically for docker/podman daemon-redirect commands.

— Named command class, no threshold or flag given. v2026.8.3

30 **Self-diagnosis for patch/write_file/search tools** NEW

59

Adds `patch` self-diagnosis for already-applied edits and whitespace mismatches, `write_file` on-disk content verification, and near-miss probing for searches that match nothing.

— Named tools and failure modes, no invocation detail. v2026.8.3

31 **/context window breakdown** NEW

58

`/context` breaks down exactly what is filling the context window.

Inspect what is consuming your context window before a long autonomous run to decide whether to compress.

```
$ /context
```

— Command named with runnable example, thin mechanism. v2026.8.3

32 **Nix/NixOS support** NEW

57

Best-effort Nix/NixOS support is available via a Nix flake, a declarative NixOS module, and an optional container mode, documented in `nix-setup.md`.

— Named artifacts, setup steps not included here. product docs

33 **/diff change viewer** NEW

56

`/diff` shows staged, all, or session changes from any surface.

— Command named with scope options, no example. v2026.8.3

34 **hermes approvals suggest allowlist mining** NEW 56

`hermes approvals suggest` mines approval history into allowlist proposals.

— Command named, output format not specified. v2026.8.3

35 **Tool-calling iteration limit raised to 500** IMPROVED 56

The default tool-calling iteration limit is raised from 90 to 500, enabling long autonomous runs that previously hit an artificial ceiling.

— Concrete before/after numbers, no adjustable config key given. v2026.8.3

36 **A2A v1.0 plugin** NEW 56

A bundled Agent-to-Agent (A2A v1.0) plugin lets Hermes discover, communicate with, and be driven by other A2A-compatible agents.

— Protocol version named, no setup steps given. v2026.8.3

37 **gpt-transcribe STT provider** NEW 52

OpenAI `gpt-transcribe` is added as a supported STT provider with unified language resolution.

— Named provider, no configuration steps given. v2026.8.3

38 **Mid-turn redirect** NEW 52

Typing a correction while the agent is working redirects the active turn, preserving work in flight and the original prompt.

— Behaviour explained, no trigger command named. v2026.8.3

39 **Context compression engine improvements** IMPROVED 52

Context compression for large-window models adds proactive tool-result pruning, per-turn micro-compaction, and progress-aware timeouts.

— Named mechanisms, no config surface or numbers. v2026.8.3

40 **/focus reduced-output view** NEW 51

`/focus` provides a reduced-output view with hidden-line recovery.

— Command named, hidden-line recovery mechanism unclear. v2026.8.3

41 **hermes desktop launch command** NEW 50

`hermes desktop` launches the desktop app after a CLI-only install.

— Simple named command, little further detail. product docs

42 **Prompt caching for tool schemas on Anthropic** IMPROVED

46

Prompt caching now covers tool schemas on native Anthropic without history loss.

— Named provider and scope, no mechanism detail. v2026.8.3

43 **Android (Termux) install target** NEW 45

Android is supported as an install target via Termux, using the standard curl installer.

— Platform named, no Termux-specific steps given. product docs

44 **Global-hotkey quick-entry window** NEW 45

A global-hotkey quick-entry window captures a thought into any session from anywhere in the OS.

— Feature named, no exact hotkey given. v2026.8.3

45 **Truncated terminal output spill-to-file** NEW 45

Truncated terminal output now spills to file so the agent can read back output that exceeded the display limit.

— Mechanism named, no file path or limit given. v2026.8.3

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Agentic Infrastructure

Hugging Face

Feedback ⓘ 1 RELEASE · SEEN 2026-08-04 CHANGELOG ↗

Hugging Face is a platform providing pre-trained machine learning models, datasets, and tools for natural language processing and computer vision.

> Hugging Face added a credential revocation endpoint and schema updates to its OAuth/user APIs, plus label-based filtering for Jobs pages.

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

01 **Label-based filtering on Jobs pages** NEW 60

Jobs pages for users and organizations now support filtering by label, with the most-used labels shown as clickable chips displaying job counts, plus a free-form `key=value` input field for filtering by any label.

— Describes UI mechanism and input format but no code example changelog-20260804-42376c35

THINNER COVERAGE BELOW

02 **Credential revocation endpoint** NEW 55

New endpoint `POST /api/credentials/revoke` was added to the API surface.

— Names endpoint and method but no request/response detail 0.0.1

03 **OAuth userinfo response schema changed** IMPROVED 45

The response schema changed for both `GET /oauth/userinfo` and `POST /oauth/userinfo`.

— Endpoints named but no detail on what changed 0.0.1

04 **User overview endpoint schema changed** IMPROVED 45

The response schema changed for `GET /api/users/{username}/overview`.

— Endpoint named but no detail on schema change itself 0.0.1

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Promptfoo

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

Promptfoo is a testing framework for evaluating and benchmarking language model prompts and AI applications across multiple providers.

> Promptfoo's snapshot-20260804 release adds an adaptive rate-limit scheduler, a Transformers.js provider for local inference, several new gateway/video providers, telecom and RAG-attribution red team plugins, multi-input red team scanning, and a batch of CLI, config and UI refinements.

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

01 **Multiple `--env-file` flags for eval** NEW 75

Supports specifying multiple `--env-file` flags to load multiple environment files in a single run, e.g. `promptfoo eval --env-file .env.base --env-file .env.staging`.

Load multiple environment files (e.g. base secrets plus environment-specific overrides) in a single eval run.

```
$ promptfoo eval --env-file .env.base --env-file .env.staging
```

— Runnable multi-flag command example provided. snapshot-20260804

02 **`promptfoo logs` CLI command** NEW 70

Adds `promptfoo logs` command for viewing log files directly from the CLI without navigating log directories.

View recent promptfoo logs directly from the CLI without navigating log directories.

```
$ promptfoo logs
```

— Runnable command example given, no further mechanism detail. snapshot-20260804

03 **xAI Voice provider enhancements** IMPROVED 65

Adds `apiBaseUrl`, `websocketUrl` override, and function call support to the xAI Voice provider.

— Named config fields but no usage example. snapshot-20260804

04 **`numTests` config key caps generated test cases** NEW

60

Adds `numTests` config key to cap the number of test cases generated per red team strategy.

— Config key named but file path/example missing. snapshot-20260804

05 **Transformers.js local inference provider** NEW 60

New Transformers.js provider that runs Hugging Face models locally in Node.js or the browser.

— Named provider and runtimes, no config example. snapshot-20260804

06 **AWS Bedrock Video provider** NEW 60

New AWS Bedrock Video provider supporting Nova Reel and Luma Ray 2 video generation.

— Named models supported, no config example. snapshot-20260804

07 **Azure AI Foundry Video provider** NEW 60

New Azure AI Foundry Video provider for Sora video generation via Azure.

— Named model and platform, no config example. snapshot-20260804

08 **OpenAI Codex provider collaboration mode and tracing** IMPROVED 60

Adds `collaboration_mode` support and integrated tracing to the OpenAI Codex provider.

— Named config field and tracing capability, no walkthrough. snapshot-20260804

THINNER COVERAGE BELOW

09 **`--share`/`--no-share` flags for Model Audit** NEW 55

Adds `--share` and `--no-share` flags to Model Audit for controlling cloud sharing of audit results.

— Flags named but no example command or default. snapshot-20260804

10 **Adaptive rate limit scheduler** NEW 55

New adaptive rate limit scheduler that automatically adjusts concurrency based on provider rate limits and response headers.

— Mechanism described but no config surface named. snapshot-20260804

11 **`-d`/`--description` flag for redteam generate** NEW 50

Adds `-d` / `--description` flag to `redteam generate` commands.

– Flag named, no description of resulting behaviour. snapshot-20260804

12 **`$VAR` syntax in file paths** NEW 50

Adds `$VAR` syntax support for environment variables in file paths.

– Syntax named, no example path shown. snapshot-20260804

13 **`word-count` assertion type** NEW 50

Adds `word-count` assertion type for validating response word counts.

– Named assertion type, no config example. snapshot-20260804

14 **`__count` variable for derived metrics** NEW 50

Adds `__count` variable for use in derived metrics when computing averages.

– Variable named, usage context not detailed. snapshot-20260804

15 **`metadata.sessionId` column in eval tables** IMPROVED 50

Surfaces `metadata.sessionId` as a variable column in eval tables and exports.

– Field named, no export format specifics given. snapshot-20260804

16 **Vercel AI Gateway provider** NEW 50

New Vercel AI Gateway provider for routing requests through Vercel's AI gateway.

– Provider named, no setup detail given. snapshot-20260804

17 **Cloudflare AI Gateway provider** NEW 50

New Cloudflare AI Gateway provider for routing requests through Cloudflare's AI gateway.

– Provider named, no setup detail given. snapshot-20260804

18 **HTTP provider native session endpoint support** IMPROVED 50

Adds native session endpoint support to the HTTP provider for stateful conversations.

– Feature named, no endpoint schema given. snapshot-20260804

19 **RAG Source Attribution red team plugin** NEW 50

New RAG Source Attribution red team plugin to test whether RAG systems properly attribute sources.

– Purpose explained, no plugin ID or config shown. snapshot-20260804

20 **Fork PR and comment-triggered scans in Code Scanning** NEW

50

Adds Fork PR support and comment-triggered scans to Code Scanning.

— Capability named, no trigger syntax or setup shown. snapshot-20260804

21 **`-j` flag for Python worker pool concurrency** IMPROVED

45

Adds `-j` flag propagation to Python worker pools for concurrency control.

— Flag named, mechanism and range not explained. snapshot-20260804

22 **Telecom red team plugin suite** NEW

45

New Telecom red team plugin suite for telecommunications AI systems.

— Plugin suite named, no individual plugins listed. snapshot-20260804

23 **User-rated filter in eval results** NEW

45

Adds a user-rated filter to show only manually rated eval results.

— UI filter named, navigation path not given. snapshot-20260804

24 **Provider config hover in eval results UI** NEW

45

Adds provider config hover in the eval results UI to view provider configuration details.

— UI feature named, no exact navigation path. snapshot-20260804

25 **Test-level prompts and providers filters** NEW

40

Adds a test-level prompts filter and providers filter for per-test-case scoping in configuration.

— Filters named generically, no config key shown. snapshot-20260804

26 **Per-test structured output configuration** NEW

40

Supports per-test structured output configuration at the test case level.

— Capability described only in prose, no key named. snapshot-20260804

27 **Multi-input variable red team scanning** IMPROVED

40

Red team scans now support multiple input variables for testing systems with complex input structures.

— Described generically, no config key or example. snapshot-20260804

28 **Early scan stop on plugin generation failure** NEW

40

Adds early scan stop behavior when plugins fail to generate test cases.

— Behaviour named, no threshold or flag given. snapshot-20260804

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 4, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue