THE CYBER TOOLCHAIN

NO. 122

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 27, 2026 · EVERY DAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

INFORMATION DENSITY

everything



BRIEFED | EVERYTHING

Every tool in the issue, in full.

all 138 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 33 tools

▾ // THE BRIEF HIDE

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

PROTECT

Stop a secrets store from handing out credentials it never should have

vim ? help

OpenBao now denies SSH wildcard identity templating and public OIDC clients presenting client credentials by default, and closes a path-traversal hole in Transit key imports — three ways a misread config turned into broader access than intended, now off unless you ask. ToolHive tightens auth and token handling in `thv run` and vMCP and adds backend-scoped vMCP authorization, so an MCP client's token buys access to one backend rather than everything the proxy fronts.

openbao · ToolHive

PROTECT

Keep an agent account from reaching past the servers it was scoped to

Runlayer locks server scoping for federated agent accounts, masks cookies in approval prompts, and caps connector redirects and MCP request sizes — the gateway stops being a pass-through for whatever the agent asks for. GitHub's assisted permissions for Copilot agents go GA across CLI, app and VS Code, so permission grants happen at the point of use instead of being pre-granted broadly.

Runlayer · GitHub

RESPOND

Get AD attack-path findings out of the console and into the remediation queue

BloodHound adds CSV export from the findings table and bulk ingestion of Priority Zone rules, so defining which tier-zero assets matter stops being a one-at-a-time click job and findings can go to whoever owns the fix. DefectDojo's product-level export API and bulk Jira unlinking cover the other end — pulling a whole product's findings programmatically, and detaching stale tickets without touching each one.

bloodhound · defectdojo

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Keygraph shannon

Feedback ⓘ v3.5.0 NOTES ↗ RANK

Shannon is an AI pentester for web applications and APIs. It analyzes your source code, identifies attack vectors, and executes real exploits to prove vulnerabilities before they reach production.

> Shannon adds an experimental CLI probe for verifying PHP-to-shell upgrade attack paths before they reach production.

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

01 **Experimental PHP-to-shell upgrade probe** NEW 75

Adds an experimental `--exp-php2shell` flag to the CLI, an upgrade-verification probe for testing PHP-to-shell attack paths. Can be run via `npx @keygraph/shannon@latest --exp-php2shell` to validate shell-upload exploit paths before they reach production.

Run the experimental PHP-to-shell upgrade-verification probe against a target during a pentest to validate shell-upload exploit paths before they reach production.

```
$ npx @keygraph/shannon@latest --exp-php2shell
```

— Named flag and runnable command, but mechanism of probe not detailed. v3.5.0

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

instaloader

Feedback ⓘ v4.16 NOTES ↗ RANK

Download pictures (or videos) along with their captions and other metadata from Instagram.

> Instaloader v4.16 made a small usability tweak to highlight retrieval logging.

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

01 **Progress counter for highlight retrieval** IMPROVED 30

The Retrieving highlights log line now shows a progress counter, giving better visibility into highlight retrieval progress.

– Describes the change but no mechanism or example v4.16

WAS THIS USEFUL? 👍 👎

tookie-osint

Feedback ⓘ v4.2.2.3 NOTES ↗ RANK

Tookie is a advanced OSINT information gathering tool that finds social media accounts based on inputs.

> Tookie-osint added a `--tor` flag to route username scans through the Tor network for anonymised lookups.

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

01 **Tor-routed username scans** NEW 70

A new `--tor` flag routes username scans through the Tor network for anonymised lookups, e.g. `python3 brib.py --tor <username>`.

Run a username scan anonymously through Tor to avoid exposing your source IP to target sites

```
$ python3 brib.py --tor <username>
```

– Names exact flag and runnable command, but no mechanism detail. v4.2.2.3

WAS THIS USEFUL? 👍 👎

◆ Post-Exploitation

SpecterOps BloodHound

Feedback ⓘ v9.9.0 NOTES ↗ RANK

Six Degrees of Domain Admin

> BloodHound v9.9.0 adds CSV export from the findings table, bulk Priority Zone rule ingestion, and DCA-migrated post-processing for 20+ AD attack paths, alongside new null-aware API filtering, extension-based saved queries, jobs API/configuration additions, and Explorer view enhancements; it also removes the legacy ingest parser and zone node feature flag and raises the OpenHound version requirement for `use_raw_object_id`.

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

01 **Jobs API and configuration additions** NEW 70

Adds `profiles create` endpoint (`BED-9573`) and `profiles update` endpoint (`BED-9574`) to the jobs API, adds secret store config for jobs (`BED-9852`), adds an experiment storage configuration option, and adds a job history table to the UI.

— Names every jobs-area endpoint and config addition with ticket IDs v9.9.0

02 **Legacy ingest parser removed for CHOW streaming validation** 60

BREAKING

The legacy ingest parser, its types, and its schemas have been removed and replaced with CHOW-based streaming validation; integrations relying on the legacy ingest format will break.

— Explains before/after and impact, no migration steps given v9.9.0

THINNER COVERAGE BELOW

03 **Extension-based saved queries** NEW 55

Adds `extension` as a valid scope value for the list queries endpoint, enabling retrieval of extension-based saved queries, and surfaces these queries in the UI where they are read-only (non-modifiable).

— Names the API scope value and UI behavior, no endpoint path given v9.9.0

04 **Explorer graph view enhancements** IMPROVED 55

Adds graph direction and spacing controls to the Explorer view, renders non-traversable Explorer edges as dashed lines for visual distinction, and displays associated node labels in the Explore entity panel.

— Describes three concrete UI behaviors, navigation only implied v9.9.0

05 **Null-aware filtering on API endpoints** NEW — 53

Adds `is_null` and `is_not_null` query parameters to API endpoints, enabling filtering on whether a field is or is not null.

— Names exact query parameters but not which endpoints support them v9.9.0

06 **OpenHound version requirement raised for use_raw_object_id** — 50

BREAKING

The `use_raw_object_id` flag now requires OpenHound (OH) version >v0.3.0.

— Names flag and exact version threshold, clear upgrade trigger v9.9.0

07 **CSV export from findings table** NEW — 45

Adds CSV export from the findings table in the UI.

— Names format and surface but no mechanism or limits v9.9.0

08 **Zone node feature flag removed** BREAKING — 45

The zone node feature flag has been removed; zone node behavior is now controlled by Enterprise entitlement only.

— Names the flag and new control mechanism, no config path v9.9.0

09 **DCA-migrated post-processing for AD attack paths** IMPROVED — 30

Post-processing for over 20 AD attack paths has been migrated to DCA-based processing.

— Names the count and DCA migration but no behavioral detail v9.9.0

10 **Enriched logging for extension upserts** IMPROVED — 25

Adds enriched logging for extension upsert operations.

— Bare statement with no detail on log content or access v9.9.0

11 **Bulk Priority Zone rule ingestion** NEW — 25

BloodHound v9.9.0 adds bulk Priority Zone rule ingestion.

— Mentioned only in summary with no mechanism or interface detail v9.9.0

└──▷ **BREAKING ON UPGRADE**

! The legacy ingest parser, its types, and its schemas have been removed and replaced with CHOW-based streaming validation — any integrations relying on the legacy ingest format will break.

! The zone node feature flag has been removed; zone node behavior is now controlled by Enterprise entitlement only.

! The `use_raw_object_id` flag now requires OpenHound (OH) version >v0.3.0.

WAS THIS USEFUL?  

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Semgrep

Feedback ⓘ snapshot-20261027 NOTES ↗ RANK

Get started with Semgrep to help you catch, flag, and fix real vulnerabilities before they ship.

> Semgrep's latest snapshot adds a dedicated related-issues API endpoint, direct Anthropic connectivity for its Multimodal AI features, PowerShell support in Guardian, and a string of dashboard/Policies filters, while also removing QL/CodeQL support and tightening rate limits on public deployment endpoints.

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

01 **Related issues API endpoint for findings** NEW 82

Adds an experimental `GET /api/deployments/{deployment_id}/issues/{issue_id}/related-issues` endpoint that returns a finding's related issues; the existing `relatedIssues` field on finding responses is now deprecated in favor of this endpoint.

Fetch related issues for a specific finding to surface correlated alerts without relying on the deprecated `relatedIssues` field in the finding response.

```
$ curl -s -H 'Authorization: Bearer <token>' \

'https://semgrep.dev/api/deployments/<deployment_id>/issues/<issue_id>/related-issues'
```

— Exact endpoint and deprecated field named, with runnable curl example.

snapshot-20261027

02 **CSV export scoped to selected finding groups** IMPROVED 75

Enables CSV export scoped to selected finding groups in the Findings view, with the existing 10,000-issue limit now applying to the selection rather than the full filter result; if nothing is selected, active filters determine the export.

Export only a selected subset of findings to CSV when more than 10,000 issues match your filters — select the relevant groups first, then export.

> 📍 In Semgrep AppSec Platform, go to Findings, select the finding groups you want, then choose Export › CSV. The export applies to your selection; if nothing is selected, your active filters determine the export (10,000-issue limit applies to the selection).

— Exact limit, scope rule, and UI steps for exporting are given. snapshot-20261027

03 **QL/CodeQL language support removed** BREAKING 75

Semgrep now rejects `--lang ql` and the `ql` language key in rules, and no longer parses `.ql` or `.qll` files — QL (CodeQL query language) support has been removed entirely.

— Names the exact flag, language key, and file extensions dropped. snapshot-20261027

04 **Rate limiting on public deployment endpoints** BREAKING 70

The public deployment endpoints `/api/v2/deployments/public/{id}` and `/api/v2/deployments/public/search` now count every request, including repeated and concurrent ones, against rate limits, returning HTTP 429 when exceeded.

— Names both endpoints and exact response code for the new limit. snapshot-20261027

05 **Filters and incident column on Supply Chain Advisories page** NEW 62

Adds Language, Severity, and Advisory type filters to the Supply Chain Advisories page, with an Incident column added to the advisories table.

— Names all three filters and the new column explicitly. snapshot-20261027

06 **Finding type filter on dashboards** NEW 60

Adds a Finding type filter to dashboards, enabling separate views of AI-powered and Rule-based Semgrep Code findings in charts, drilldowns, and PDF exports.

— Names filter and affected views but no exact navigation path. snapshot-20261027

07 **False positive rate column in Policies rules table** NEW 60

Adds a sortable False positive rate column to the rules table in Policies, covering Semgrep Code and Semgrep Secrets rules, with drill-down into ignored findings.

— Names the column and scope, mechanism partially explained. snapshot-20261027

THINNER COVERAGE BELOW

08 **Direct Anthropic connection for Multimodal AI** NEW 55

Adds direct Anthropic connection support in Semgrep Multimodal AI settings via Manual mode, used for in-app features such as Autotriage and Autofix.

— Names the setting and features it powers, no exact setup steps. snapshot-20261027

09 **Expanded Python dependency coverage in Supply Chain SBOM**

IMPROVED

55

Expands Semgrep Supply Chain dependency inventory and SBOM to include every package declared in a Python `requirements.txt` file, not only packages pinned with `==`.

— Clear before/after behavior but no user action required or shown. snapshot-20261027

---

10 **Required Base model field for Azure OpenAI setups** BREAKING

55

Adds a Base model required field to Azure OpenAI settings, specifying the model your Azure deployment runs; new Azure OpenAI setups must populate this field before the connection can be configured.

— Names the required field and when it applies, no UI path given. snapshot-20261027

---

11 **Secret value redaction in legacy issue details endpoint**

55

BREAKING

The legacy issue details endpoint now returns `null` for `code_snippet` on Semgrep Secrets findings instead of the secret value.

— Names the field and endpoint behavior change precisely. snapshot-20261027

---

12 **Reviewer-approved exception notes as Semgrep memories** NEW

50

Adds reviewer-approved exception notes as Semgrep memories, applying to Semgrep Code, AI-powered detection, and Semgrep Secrets findings.

— Names the finding types covered but no mechanism or usage path. snapshot-20261027

---

13 **xAI provider renamed to SpaceXAI with fixed endpoint** BREAKING

45

The xAI provider is now called `SpaceXAI` and always uses `https://api.x.ai/v1`; the endpoint field no longer appears in its settings.

— Names the rename and fixed URL but limited further explanation. snapshot-20261027

---

14 **PowerShell support in Semgrep Guardian** NEW

40

Adds PowerShell support to Semgrep Guardian, which now runs hooks on both PowerShell and Bash commands.

— States the addition but no command or config example. snapshot-20261027

└──▷ BREAKING ON UPGRADE

! Semgrep now rejects `--lang ql` and the `ql` language key in rules, and no longer parses `.ql` or `.qll` files — QL (CodeQL query language) support has been removed.

! The `relatedIssues` field on finding responses is deprecated; use `GET /api/deployments/{deployment_id}/issues/{issue_id}/related-issues` instead.

! The legacy issue details endpoint now returns `null` for `code_snippet` on Semgrep Secrets findings instead of the secret value.

! The xAI provider is now called `SpaceXAI` and always uses `https://api.x.ai/v1`; the endpoint field no longer appears in its settings.

! New Azure OpenAI setups require a `Base model` field to be filled in before the connection can be configured.

! The public deployment endpoints `/api/v2/deployments/public/{id}` and `/api/v2/deployments/public/search` now count every request (including repeated and concurrent ones) against rate limits, returning HTTP 429 when exceeded.

WAS THIS USEFUL? 👍 👎

GitHub Feedback ⓘ

GitHub Copilot assisted permissions generally available (2026-10-26) BLOG ↗

RANK

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

> GitHub shipped GA assisted permissions for Copilot agents across CLI, app, and VS Code, alongside general availability of the GPT-6.1 Terra model across nearly every Copilot surface.

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

01 **Assisted permissions mode for Copilot agents** NEW 85

Assisted permission mode reaches general availability, letting an LLM auto-approve clearly safe commands, file-system access, and URL opens instead of prompting for every action. It's available in the Copilot CLI via the `/permissions` command (select `assisted`), in the GitHub Copilot app via Tool permissions in the mode picker, and in Visual Studio Code via the permissions picker. Enterprise administrators get managed settings to block Allow All while keeping Assisted available, to disable Assisted separately, and to set the default permission mode for new sessions.

Switch to assisted permissions in the Copilot CLI so safe commands run without interrupting your workflow.

```
$ /permissions
# Then select: assisted
```

Enable assisted permissions in VS Code's Copilot permissions picker to reduce approval prompts during agentic tasks.

📍 Open the permissions picker in VS Code › Select Assisted permissions

— Names exact command, mode, and three surfaces plus admin controls

GitHub Copilot assisted permissions gen…

02 **GPT-6.1 Terra model in Copilot** NEW 70

GPT-6.1 Terra is generally available for agentic coding and terminal work, selectable in the Copilot model picker across Visual Studio Code, Visual Studio, Copilot CLI, GitHub Copilot cloud agent, GitHub Copilot app, github.com, GitHub Mobile (iOS and Android),

JetBrains IDEs, Xcode, and Eclipse. Copilot Enterprise and Business administrators can manage access via the model policy in Copilot settings; new models are auto-enabled by default unless the global default is off or the model is explicitly disabled.

Restrict GPT-6.1 Terra from being used org-wide until your team has evaluated it — useful for Enterprise admins who want controlled rollout.

> 📍 In GitHub Copilot settings, go to Policies › Model policy, find GPT-6.1 Terra, and set it to `Disabled` to prevent automatic enablement for your organization.

— Lists all platforms and admin policy, but no technical model detail

GPT-6.1 Terra in GitHub Copilot (2026-1…

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

openbao

Feedback ⓘ v2.7.2 NOTES ↗ RANK

OpenBao is a software solution to manage, store, and distribute sensitive data including secrets, certificates, and keys.

> OpenBao's v2.6.5 and v2.7.2 releases tighten several security defaults — denying SSH wildcard identity templating and public OIDC clients with client credentials by default, and closing a path-traversal issue in Transit key imports — while also adding PKI external-key UI fields.

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

01 Wildcard control for SSH identity templates

BREAKING

90

Adds `allow_wildcards_in_identity_templates` parameter to `secrets/ssh` controlling whether wildcards (`*`) are permitted in SSH certificate identity templating. Wildcards are now denied by default, so existing roles relying on them will be blocked unless the parameter is explicitly enabled, e.g. `bao write ssh/roles/my-role allow_wildcards_in_identity_templates=true` .

Re-enable wildcard identity templating on an SSH role that requires it, after upgrading to v2.6.5 where wildcards are denied by default.

```
$ bao write ssh/roles/my-role
allow_wildcards_in_identity_templates=true
```

— Names exact parameter, endpoint, default change and runnable remediation command.

v2.7.2 v2.6.5

THINNER COVERAGE BELOW

02 OIDC public client restriction with client credentials

57

BREAKING

`client_type=public` is now denied when `client_credentials=true` is set in the OIDC provider (`secrets/identity`); configurations combining both settings will be rejected on upgrade.

— Names exact fields and endpoint but gives no fix command. v2.7.2 v2.6.5

03 Path traversal fix in transit key imports

BREAKING

50

Imported key backups in `secrets/transit` that relied on path traversal to target arbitrary keys will no longer succeed.

— Names affected endpoint and vulnerability but no remediation steps. v2.6.5

---

04 **UI fields for PKI external keys** NEW

25

Adds UI fields to configure external keys in PKI.

— Bare mention with no detail on fields or workflow. v2.7.2

└──▷ **BREAKING ON UPGRADE**

! Wildcards (`*`) in SSH certificate identity templating are now denied by default in `secrets/ssh` ; existing roles relying on wildcard identity templates will be blocked unless `allow_wildcards_in_identity_templates` is explicitly enabled.

! `client_type=public` is now denied when `client_credentials=true` in the OIDC provider (`secrets/identity`); configurations combining both settings will be rejected on upgrade.

! Wildcards (`*`) in SSH certificate identity templating are now denied by default in `secrets/ssh` ; set `allow_wildcards_in_identity_templates=true` on affected roles to restore prior behavior.

! `client_type=public` is now rejected when `client_credentials=true` is set in the `secrets/identity` OIDC provider; configurations combining these two settings will fail.

! Imported key backups in `secrets/transit` that relied on path traversal to target arbitrary keys will no longer succeed.

WAS THIS USEFUL? 👍 👎

infisical

Feedback ⓘ v0.167.4 NOTES ↗ RANK

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

> Infisical v0.167.4 makes batched secret editing in Secrets Overview the default behavior and repositions the quick-add row for faster entry.

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

01 Always-on batched changes in Secrets Overview IMPROVED

50

Batched secret changes are now always-on in the Secrets Overview, removing the need to opt in. The quick-add row has also been moved to the top of the Secrets Overview table for faster secret entry.

— Describes UI behavior change but no mechanism detail v0.167.4

UPDATED INTEGRATION Cloudflare Workers and Pages secret sync — variable type selection

v0.167.3 NOTES ↗ STABLE

Adds a v2 batch commit API, bulk secret-approval reviews, PAM auto-approvals, Agent Vault member-provided credential fields, and group access in the org users table.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.167.3
https://github.com/Infisical/infisical.git
# already have the repo? check out this version:
$ git checkout v0.167.3
```

Adds a private v2 batch commit API route for atomically committing batched secret and folder changes.

Adds bulk review of change requests in secret approvals, letting reviewers approve or reject multiple requests at once.

Adds auto-approvals for PAM (Privileged Access Management) workflows.

Adds member-provided credential fields with personal values and variables in Agent Vault.

Shows group access in the organization users table.

WAS THIS USEFUL?  

◆ WAF, Bot & API Protection

Fastly

Feedback ⓘ CLI v17.0.0 (2026-10-26) BLOG ↗ RANK

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

> Fastly CLI v17.0.0 adds a flag to automate virtual patch enablement in NGWAF workspaces, while also shipping two breaking changes: a removed flag on AI Runtime Control key creation and a changed JSON output shape across several list commands.

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

01 Automated Virtual Patching flag for NGWAF workspaces NEW

83

The CLI adds `--autoEnableNewVirtualPatches` to `ngwaf workspace create` and `ngwaf workspace update`, letting new virtual patches be applied to a workspace automatically rather than requiring manual enablement; the setting is also surfaced as output in `ngwaf workspace get`.

Enable Automated Virtual Patching when creating a new NGWAF workspace so new virtual patches are applied automatically without manual intervention.

```
$ fastly ngwaf workspace create --autoEnableNewVirtualPatches
```

Opt an existing NGWAF workspace into Automated Virtual Patching after upgrade.

```
$ fastly ngwaf workspace update --autoEnableNewVirtualPatches
```

— Names exact flag and three commands with runnable examples CLI v17.0.0 (2026-10-26)

THINNER COVERAGE BELOW

02 JSON output changed to bare array for several list commands

58

BREAKING

`--json` output for `ai-runtime-control usage-metrics list`, `ngwaf workspace list`, `ngwaf rule list`, and `ngwaf workspace rule list` is now a bare array instead of an object with `data` and `meta` fields.

— Names all four affected commands and the exact shape change CLI v17.0.0 (2026-10-26)

03 `--user-id` flag removed from `ai-runtime-control key create`

45

BREAKING

The `--user-id` flag is no longer accepted by `ai-runtime-control key create` because it was removed from the underlying API.

— Names the flag and command but gives no migration path CLI v17.0.0 (2026-10-26)

└──▷ **BREAKING ON UPGRADE**

! `ai-runtime-control key create` no longer accepts the `--user-id` flag, which was removed from the API.

! `--json` output for `ai-runtime-control usage-metrics list`, `ngwaf workspace list`, `ngwaf rule list`, and `ngwaf workspace rule list` is now a bare array instead of an object with `data` and `meta` fields.

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

StepSecurity

Feedback ⓘ

Secure Registry: Inventory and Statistics (Beta) (2026-10-26) BLOG ↗ RANK

StepSecurity is a security platform that hardens CI/CD pipelines, such as GitHub Actions, by detecting supply-chain risks, enforcing least-privilege runners, and monitoring workflow tampering.

> StepSecurity added two beta pages to Secure Registry — Inventory for package visibility and Statistics for request analytics — giving teams drill-down views into what packages are being pulled and how registry controls are acting on requests.

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

01 **Secure Registry Inventory page** NEW 80

The new Inventory page lists every package requested through Secure Registry with its last pulled version, request counts, licenses, and known vulnerabilities. It supports search by package name, SPDX license identifier, or advisory ID/alias to surface affected versions, filtering by ecosystem, vulnerability severity, minimum CVSS score, version, and license status, and a per-package detail view showing first-seen date, last puller, advisories with fixed versions and upgrade guidance, and the last five allowed pulls.

Identify every package version your teams have pulled that is affected by a specific advisory, so you can prioritise remediation.

> 📍 In Secure Registry, go to Inventory › search by advisory ID or alias › review the listed package versions, their fixed versions, and the upgrade guidance shown in each package detail view.

Audit all packages requested under a specific open-source license by searching with its SPDX identifier (e.g. `GPL-3.0`) in Inventory.

> 📍 In Secure Registry, go to Inventory › enter the SPDX license identifier (e.g. `GPL-3.0`) in the search box › apply License Status filter to review every requested version under that license.

— Names exact filters and fields; UI navigation given but no API/flag.

Secure Registry: Inventory and Statisti…

02 **Secure Registry Statistics page** NEW 63

The new `Statistics` page shows requests handled over a selected time range, broken down into served-as-is, modified, and blocked, with a daily request history and per-control counts of modified and blocked requests.

Review how many requests your controls blocked versus served over the past week to validate policy effectiveness.

> 📍 In Secure Registry, go to Statistics › select the desired time range › inspect the daily request history chart and the per-control breakdown of modified and blocked request counts.

— Describes breakdown categories and UI path but no deeper mechanism.

Secure Registry: Inventory and Statisti…

WAS THIS USEFUL? 👍 👎

Endor Labs

Feedback ⓘ

Endor Labs + SpaceXSI: Securing every stage of agentic software delivery | Endor Labs (2

BLOG ↗ RANK

Endor Labs is an application security platform built for AI-generated and human-written code, helping teams prioritize and fix real risk across the SDLC.

> Endor Labs shipped a broad integration with SpaceXSI's Grok Build, bringing AURI remediation agents, deterministic coding-agent governance hooks, MCP trust scoring, a package firewall, and AI SAST/SCA scanning directly into agentic coding sessions.

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

01 AURI Agent Kit for Grok Build NEW

93

The Endor Labs Agent Kit for Grok Build is installable from the marketplace or via `/add-plugin endorlabs` in the agent chat, bringing AURI remediation agents into sessions. The AI SAST Remediation agent proposes fixes with test plans for critical code flaws and can open a pull request on developer approval; the SCA Remediation agent queries scan results, checks package usage, and applies upgrades, code changes, or Endor Patches, then rescans to confirm no new vulnerabilities were introduced. Agent Kit workflows can also run outside the IDE in CI or cloud runs through the Grok Build SDK.

Bootstrap Endor Labs security scanning inside a Grok Build agent session without leaving the chat.

```
$ /add-plugin endorlabs
```

— Exact install command and two-agent mechanism both given

Endor Labs + SpaceXSI: Securing every s…

02 AI SAST and SCA scanning in Grok Build sessions NEW

70

Endor Labs AI SAST scans first-party code for security flaws, traces data flow from untrusted input to vulnerable calls, and flags exposed secrets before commit inside Grok Build sessions. Endor Labs SCA checks direct and transitive dependencies, with reachability analysis reducing SCA findings by up to 92%.

— Mechanism and a concrete 92% figure given, no exact command shown

Endor Labs + SpaceXSI: Securing every s…

THINNER COVERAGE BELOW

03 **Coding Agent Governance hooks** NEW 53

New governance hooks record every MCP tool call, shell command, and file operation and enforce policies before each action runs, with deterministic blocking that does not depend on model judgment.

— Mechanism described but no policy syntax or config surface named

Endor Labs + SpaceXSI: Securing every s…

04 **Package Firewall for agentic pipelines** NEW 37

Endor Labs Package Firewall integration blocks confirmed malicious packages at the registry level for teams routing installs through it.

— Brief description, no setup steps or routing config detail

Endor Labs + SpaceXSI: Securing every s…

05 **Endor Scores for MCP servers and skills** NEW 35

Endor Scores now rate each MCP server and skill for trust, mapped to the OWASP Top 10 for LLM Applications.

— Names the scoring concept but no mechanism or access path given

Endor Labs + SpaceXSI: Securing every s…

WAS THIS USEFUL?  

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

ILSpy

Feedback ⓘ v11.2 NOTES ↗ RANK

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

> ILSpy v11.2 adds native AOT ilspycmd packages for linux-x64 and win-x64 with dramatically faster startup, extends IL disassembly filtering, and adds cross-platform PDB dumping plus small GUI editor improvements.

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

01 **Native AOT packages for ilspycmd** NEW 88

Adds opt-in native AOT packages for `ilspycmd` targeting `linux-x64` and `win-x64`, reducing per-call startup from ~1.8 s to ~315 ms (4–7× faster), aimed at use cases like agent loops that spawn `ilspycmd` per query.

Disassemble IL for a specific type using the AOT `ilspycmd` package on Linux to get fast, sub-100 ms output in an automated analysis loop.

```
$ ilspycmd -il --type MyNamespace.MyClass MyAssembly.dll
```

— Concrete platforms and measured startup-time improvement v11.2

02 **Type and member filtering for IL disassembly** IMPROVED 78

Adds `--type` and `--member` selection support to `ilspycmd -il`, so IL disassembly now respects the same filtering flags as decompilation.

Disassemble IL for a specific type using the AOT `ilspycmd` package on Linux to get fast, sub-100 ms output in an automated analysis loop.

```
$ ilspycmd -il --type MyNamespace.MyClass MyAssembly.dll
```

— Named flags with runnable command example v11.2

THINNER COVERAGE BELOW

03 **Highlight all usages at caret** NEW 40

Adds an option to highlight all usages of the symbol at the caret in the editor.

— Describes behaviour but no exact UI path v11.2

04 **Cross-platform PDB dumping in ilspycmd** NEW 37

Adds cross-platform PDB dumping to `ilspycmd`.

— Named surface but no further mechanism given v11.2

05 **Search box in Options tab** NEW 35

Adds a search box to the Options tab in the ILSpy GUI.

— Exact UI location but no further detail v11.2

WAS THIS USEFUL? 👍 👎

refinery

Feedback ⓘ 0.11.3 CODE ↗ RANK

High Octane Triage Analysis

> refinery 0.11.3 ships seven new extraction/deobfuscation units, renames several script deobfuscator commands and meta variables to fix naming collisions, and tightens output formats and struct parsing — with multiple breaking changes to existing pipelines.

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

01 **Bit field and peek support in struct format modifiers** 75

BREAKING

In `struct`, the format modifier `!0` now reads a bit field, and `:` was added as a prefix option for peeking a value; existing format strings using `!0` for other purposes will behave differently.

— Exact syntax and behaviour change named, directly usable. 0.11.3

02 **Seven new processing units added** NEW 73

Adds `dau3` for deobfuscating AutoIt3 scripts, `xtwim` for unpacking WIM archives, `xtiex` for unpacking IExpress setup files, `dns` for parsing DNS packets from PCAPs (used alongside `pcap` and `udp`), `huffman` as a low-level option to directly handle huffman-compressed data, `batev` to extract environment variables set during batch emulation, and `batfs` to extract file system artifacts from batch emulation.

Extract DNS hostnames queried in a PCAP to quickly triage C2 or data-exfil domains.

```
$ emit capture.pcap | pcap | udp | dns
```

Deobfuscate a suspicious AutoIt3 script dropped by malware.

```
$ emit suspicious.au3 | dau3
```

Unpack a WIM archive and list its contained files for triage.

```
$ emit payload.wim | xtwim
```

— Names and describes all seven units; only three have runnable examples. 0.11.3

03 **Script deobfuscator commands renamed to avoid shadowing**

70

BREAKING

Script deobfuscator commands were renamed to stop them from shadowing real script interpreters: `js` is now `djs`, `ps1` is now `dps1`, `vba` is now `dvba`, and `php` is now `dphp`. Deobfuscation quality and speed were also improved in `djs` and `dps1`.

— Exact old-to-new command mapping given, improvement left vague. 0.11.3

04 **Meta variable renames across multiple units** BREAKING

70

The meta variable `offset` was renamed to `start` in `carve`, `rex`, `xtp`, `xtw`, all `carve*` units, `ngrams`, `perc`, `vsect`, `vsnip`, and `xtasar`. The meta variable `kind` was renamed to `type` in `xtw` and `xtpdf`. The `xt7z` unit now reports the archive checksum as `checksum` rather than `crc32`. Pipelines referencing the old names will no longer find the values.

— Every affected unit and variable named; migration is mechanical but unstated. 0.11.3

THINNER COVERAGE BELOW

05 **Path handlers always split on both separators** BREAKING

50

The `pd` and `pb` path handlers lost their separator argument and now always split on both slash and backslash; any invocation passing a separator argument will break.

— Clear before/after but no example command given. 0.11.3

06 **Output format changes in csv and xtpdf** IMPROVED

40

The `csv` unit now outputs minified JSON, and `xtpdf` unit metadata paths now end in `.json`.

— Two named units with concrete format change, no example shown. 0.11.3

07 **Matching end offset for start meta variable** IMPROVED

30

Every unit that reports a `start` meta variable now also reports the matching `end` offset.

— States the behaviour but gives no usage detail. 0.11.3

08 Minimum Python version raised to 3.11

BREAKING

25

The minimum supported Python version was raised from 3.10 to 3.11.

— Bare version-bump fact with no further detail. `0.11.3`

└──▷ BREAKING ON UPGRADE

! The minimum supported Python version was raised from 3.10 to 3.11.

! Script deobfuscator commands were renamed: `js` is now `djs`, `ps1` is now `dps1`, `vba` is now `dvba`, and `php` is now `dphp` — any pipeline or script using the old names will break.

! The meta variable `offset` was renamed to `start` in `carve`, `rex`, `xtp`, `xtw`, all `carve*` units, `ngrams`, `perc`, `vsect`, `vsnip`, and `xtasar` — pipelines referencing `offset` will no longer find the value.

! The meta variable `kind` was renamed to `type` in `xtw` and `xtpdf` — pipelines referencing `kind` from those units will no longer find the value.

! The `xt7z` unit now reports the archive checksum as `checksum` rather than `crc32` — pipelines referencing `crc32` from `xt7z` will break.

! In `struct`, the format modifier `!0` now reads a bit field and `:` was added as a prefix option for peeking — existing `struct` format strings using `!0` for other purposes will behave differently.

! The `pd` and `pb` path handlers lost their separator argument and now always split on both slash and backslash — any invocation passing a separator argument will break.

WAS THIS USEFUL? 👍 👎

Finite State Platform

Feedback ⓘ

October 26, 2026 - Current version and version order (2026-10-26) BLOG ↗

RANK

Finite State Platform is a firmware and software supply chain security tool that analyzes binaries and embedded devices to identify vulnerabilities, SBOM components, and risks.

> Finite State Platform's biggest additions this window are manual version pinning/ordering with audit trail logging and a full project ownership transfer workflow across API and UI, alongside a new org-wide version browser and several targeted API enhancements for folders, licenses, and dependencies.

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

01 **Project ownership transfer via API and UI** NEW 95

Adds `PUT /public/v0/projects/{projectId}/owner` to transfer ownership of a single Project to another user, with optional `previousOwnerRole` to set the departing owner's role, plus `PUT /public/v0/projects/owner` for bulk transfer accepting the same parameter. Exposes `owner`, `ownerId`, and `ownerDeactivated` fields with filter and sort support on `GET /public/v0/projects`, `GET /public/v0/projects/{projectId}`, and `GET /public/v0/folders/{folderId}/projects`. Adds a Change owner control to the Projects list and Project settings UI for single and bulk transfer, shows a warning when deactivating or deleting a user who still owns Projects, and logs ownership changes (previous and new owner) on the Audit trail.

Transfer a single Project to a new owner and demote the previous owner to Editor in one API call.

```
$ curl -X PUT
'https://platform.finitestate.io/public/v0/projects/{projectId}/o
wner' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"newOwnerId": "<userId>", "previousOwnerRole": "Editor"}'
```

Bulk-transfer ownership of multiple Projects at once — useful when offboarding a departing team member who owns many Projects.

```
$ curl -X PUT
'https://platform.finitestate.io/public/v0/projects/owner' \
  -H 'Authorization: Bearer <token>' \
  -H 'Content-Type: application/json' \
  -d '{"projectIds": ["<projectId1>", "<projectId2>"],
"newOwnerId": "<userId>", "previousOwnerRole": "Editor"}'
```

Find all Projects owned by a specific user, including any whose owner has been deactivated, using the new owner fields and filter.

```
$ curl -G 'https://platform.finitestate.io/public/v0/projects' \
  -H 'Authorization: Bearer <token>' \
  --data-urlencode 'filter[owner]=<ownerName>' \
  --data-urlencode 'sort=owner'
```

— Named endpoints, fields and params with three runnable examples

October 26, 2026 - Project ownership tr…

02 **Version pinning and ordering on Versions page** NEW

90

Adds ability to pin a Version as current on the Versions page, preventing new scans from replacing it, with changes recorded in the Audit trail; adds ability to reorder Versions and return to the latest Version from a pinned one. The API now exposes `previousVersion` on `GET /public/v0/versions/{id}` and a `sort=order` query parameter on `GET /public/v0/versions` (e.g. `projectId=<id>&sort=order:desc`) to retrieve Versions in the custom order set on the Versions page.

Retrieve Versions for a project in the custom analyst-defined order, newest pinned release first, for downstream tooling or reporting.

```
$ curl -X GET
'https://platform.finitestate.io/public/v0/versions?projectId=
<id>&sort=order:desc' -H 'Authorization: Bearer <token>'
```

Pin a stable release Version as current so that ongoing scans for new firmware builds do not overwrite it during an active audit review.

📍 In the Finite State Platform, go to the Project › Versions page, select the target Version, and choose `Set as current`. The Version is marked Current and the change is recorded in the Audit trail.

— Named endpoint, field and param with runnable example and UI steps

October 26, 2026 - Current version and …

03 **Subfolder project listing and counts in Folders API** NEW

83

Adds `includeSubfolders=true` on `GET /public/v0/folders/{folderId}/projects` to return projects in a folder and all its subfolders in a single request, and `includeSubfolderProjectCount=true` on `GET /public/v0/folders` and `GET /public/v0/folders/{folderId}` to include each folder's project count across subfolders.

Retrieve all projects nested under a folder and its subfolders in a single API call — useful for inventorying an entire product line organized in a folder hierarchy.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/public/v0/folders/{folderId}/pro
jects?includeSubfolders=true'
```

— Named endpoints and params with runnable example

October 22, 2026 - Folder project lists…

04 **In-place dependency version change** IMPROVED

80

Supports changing a project dependency's version via `dependencyProjectVersionId` on `PUT /public/v0/project-dependencies/{dependencyId}` without removing and re-adding the dependency.

Retarget a dependency to a different version in-place instead of deleting and re-adding it — useful when a project upgrades a shared component.

```
$ curl -s -X PUT -H 'Authorization: Bearer <token>' -H 'Content-
Type: application/json' -d '{"dependencyProjectVersionId": "
<newVersionId>"}'
'https://platform.finitestate.io/public/v0/project-
dependencies/{dependencyId}'
```

— Named endpoint and field with runnable example October 22, 2026 - Folder project lists…

05 **'All versions' tab on Projects page** NEW

62

Adds an All versions tab on the Projects page, positioned between All projects and Folders, listing every visible version across the organization sorted newest first, with search to filter by project and row-click to open a version directly; archived project versions are excluded.

Find a specific firmware version anywhere in the organization without navigating into each project individually.

> 📍 1. Go to the Projects page. 2. Click the `All versions` tab (between `All projects` and `Folders`). 3. Type a project name in the search box to narrow the list. 4. Click the matching row to open that version directly.

— Clear UI location and behavior but no API surface

October 23, 2026 - All versions on the …

06 **Archived project filter for versions** NEW 60

Adds optional `archived` parameter on `GET /public/v0/versions` to filter versions by whether the owning project is archived.

Filter versions to only those belonging to active (non-archived) projects, reducing noise in version listings used for triage or reporting.

```
$ curl -s -H 'Authorization: Bearer <token>'
'https://platform.finitestate.io/public/v0/versions?
archived=false'
```

— Named endpoint and param with runnable example, thin mechanism

October 22, 2026 - Folder project lists…

THINNER COVERAGE BELOW

07 **Deprecated SPDX license visibility** NEW 59

Includes deprecated SPDX licenses in `GET /public/v0/licenses` by default, listed after current licenses, with `filter=false` to exclude them. Lists deprecated SPDX licenses in the Add Licenses UI with a `Deprecated` badge, including those that share a name with a current license.

— Named endpoint and param but no runnable example

October 22, 2026 - Folder project lists…

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-10-18 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds a disabled-by-default `html_scraper` auxiliary option to `auxiliary.conf.default` with setup instructions.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion Feedback ⓘ ECS cluster settings in your inventory: NOTES ↗

Plerion is a cloud security platform that detects misconfigurations, vulnerabilities, and threats across cloud environments and provides remediation guidance.

ECS cluster configuration fields appear on each cluster asset after the next scan, enabling custom checks to test KMS and ECS Exec settings.

NEW SUPPORTED PLATFORM Amazon ECS cluster settings inventory

WAS THIS USEFUL? 👍 👎

Rapid7 InsightCloudSec

Feedback ⓘ

changelog · Oct 26, 2026 BLOG ↗ RANK

Rapid7 InsightCloudSec is a cloud security platform that monitors multi-cloud environments to detect misconfigurations, enforce compliance policies, and manage risk across AWS, Azure, and Google Cloud.

> Rapid7 InsightCloudSec's October release added a Lifecycle Management page for deprecated components, tag export options for the Compliance Scorecard, and updates to the BotFactory API.

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

01 **Lifecycle Management page** NEW 40

A new Lifecycle Management page for tracking and managing deprecated components.

— Names the page but not how it works 2026-10-26-4082a7fd

02 **Compliance Scorecard tag export** IMPROVED 30

Adds tag export options to the Compliance Scorecard.

— No detail on export formats or scope 2026-10-26-4082a7fd

03 **BotFactory API updates** IMPROVED 15

Updates the BotFactory API with new capabilities.

— No specifics on endpoints or capabilities named 2026-10-26-4082a7fd

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Google distroless

Feedback ⓘ changes since 2026-10-20

CODE ↗

Language focused docker images, minus the operating system.

distroless diff tool now fails fast when no image targets are selected, and batches Bazel builds to cut elapsed time by up to 34%.

The `diff` tool exits with an explicit error when `--only <target>` matches no image targets in the query file, rather than silently succeeding.

WAS THIS USEFUL? 👍 👎

Chainguard

Feedback ⓘ Week of 2026-10-26 NOTES ↗

Chainguard provides trusted open source artifacts for every layer of your modern software stack—containers, language libraries, and VM images.

Chainguard offers a grace period of up to six months of continued rebuilds and security updates for eligible end-of-life images.

- NEW SUPPORTED PLATFORM commercial-gitlab-* EE images (with FIPS variants)
- NEW SUPPORTED PLATFORM sentry-* images (with FIPS variants)
- NEW SUPPORTED PLATFORM automq image with FIPS variant
- NEW SUPPORTED PLATFORM gke-deploy image with FIPS variant
- NEW SUPPORTED PLATFORM aws-elasticache-controller image with FIPS variant
- NEW SUPPORTED PLATFORM node-red image
- NEW SUPPORTED PLATFORM sloth image with FIPS variant
- NEW SUPPORTED PLATFORM vllm-omni image with FIPS variant
- NEW SUPPORTED PLATFORM aws-sns-controller image with FIPS variant
- NEW SUPPORTED PLATFORM gitleaks image with FIPS variant
- NEW SUPPORTED PLATFORM kubernetes-csi-driver-smb image with FIPS variant
- NEW SUPPORTED PLATFORM mpi-operator image with FIPS variant
- NEW SUPPORTED PLATFORM supabase-storage-api image
- NEW SUPPORTED PLATFORM snuba image with FIPS variant
- NEW SUPPORTED PLATFORM supabase-studio image
- NEW SUPPORTED PLATFORM symbolicator image with FIPS variant
- RETIRED SUPPORTED PLATFORM EOL images entering grace period

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Prowler

Feedback ⓘ 5.46.0 NOTES ↗ CODE ↗ RANK

This repo hosts Prowler OSS (CLI and Local Server). Prowler is the world's most widely used open-source cloud security platform that automates security and compliance across any cloud environment for Humans and AI Agents.

> Prowler 5.46.0 introduces an Issues page in Prowler Cloud, adds new checks for AWS Bedrock, Cognito and WAFv2 plus a GCP Vertex AI Workbench service, rolls out PHP runtime version checks across App Engine, Elastic Beanstalk and Lightsail, improves Azure runtime and Google Workspace edition-aware checks, and restricts findings/resources metadata APIs to the caller's provider groups.

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

01 New and improved AWS Bedrock security checks IMPROVED

90

Adds `bedrock_guardrail_sensitive_information_filter_enforced`, which flags Bedrock guardrails whose PII entities or custom regexes take no action on the output path, and `bedrock_model_invocation_logging_text_data_enabled`, which verifies that model invocation logging delivers prompt and response text instead of only naming a destination. Also improves `bedrock_api_key_no_long_term_credentials` to report long-term API keys that never expire, which were previously skipped and produced no finding.

Run the new WAFv2 and Bedrock checks together to audit whether Web ACLs and Bedrock guardrails meet baseline security requirements.

```
$ prowler aws --checks wafv2_webacl_managed_rule_groups_enabled
bedrock_guardrail_sensitive_information_filter_enforced
bedrock_model_invocation_logging_text_data_enabled
bedrock_api_key_no_long_term_credentials
```

— Named checks with mechanism and a runnable scan command 5.46.0

02 GCP Vertex AI Workbench checks NEW

90

Adds a new GCP `vertexai` service with five Workbench checks — `vertexai_workbench_instance_public_ip_disabled`, `vertexai_workbench_instance_root_access_disabled`, `vertexai_workbench_instance_terminal_access_disabled`, `vertexai_workbench_instance_file_download_disabled`, and

`vertexai_workbench_instance_auto_upgrade_enabled` — mapped to PCI DSS 4.0 and CCC MLDE requirements.

Run only the new Vertex AI Workbench checks against a GCP project to identify publicly reachable or over-privileged notebook instances.

```
$ prowler gcp --checks
vertexai_workbench_instance_public_ip_disabled
vertexai_workbench_instance_root_access_disabled
vertexai_workbench_instance_terminal_access_disabled
vertexai_workbench_instance_file_download_disabled
vertexai_workbench_instance_auto_upgrade_enabled
```

— Five named checks, compliance mapping, and a runnable command `5.46.0`

---

03 **Azure App Service runtime version checks improved** IMPROVED

85

Raises the default `php_latest_version` for `app_ensure_php_version_is_latest` from `8.2` to `8.5` in both the SDK and the Prowler App scan configuration template, and makes `app_ensure_php_version_is_latest`, `app_ensure_python_version_is_latest`, and `app_ensure_java_version_is_latest` pass any runtime at or above the configured minimum instead of failing every version newer than the configured one.

— Named checks and config key with exact before/after default value `5.46.0`

---

04 **Metadata APIs scoped to caller's provider groups** BREAKING

85

Restricts `GET /resources/metadata`, `GET /resources/metadata/latest`, `GET /findings/metadata`, and `GET /findings/metadata/latest` to honour the provider groups of the caller role instead of aggregating every provider in the tenant.

Scope findings metadata API calls to only the provider groups the caller role can see, confirming the new per-role scoping is in effect.

```
$ curl -s -H 'Authorization: Bearer <token>' 'https://<prowler-api>/api/v1/resources/metadata/latest'
```

— Four named endpoints with method and path, plus a runnable request `5.46.0`

---

05 **PHP runtime version checks across platforms** NEW

75

Adds a new GCP `appengine` service with the `appengine_version_php_runtime_supported` check flagging App Engine serving versions on PHP runtimes without security fixes, the `elasticbeanstalk_environment_php_platform_updated` check flagging Elastic Beanstalk PHP environments on branches without security fixes or outdated platform versions, and the `lightsail_instance_php_blueprint_updated` check flagging Lightsail instances created from PHP blueprints for manual verification of the PHP release.

— Three named checks across platforms, but no example command given 5.46.0

06 **Edition-aware Google Workspace findings** IMPROVED 75

Findings now state when a feature is not available on the tenant edition (e.g. Drive DLP rules, session control, shared drive sharing on Business Starter), requiring the optional `apps.licensing` scope, and checks now pass with the reason stated when an app is turned off for everyone and no organizational unit or group re-enables it.

— Named scope and specific edition examples but no runnable step 5.46.0

07 **New AWS security checks** NEW 60

Adds the `cognito_user_pool_account_recovery_sms_not_preferred` check, which verifies that Cognito user pools do not rank SMS first for account recovery, and the `wafv2_webacl_managed_rule_groups_enabled` check, which verifies that WAFv2 Web ACLs reference at least one managed rule group.

Run the new WAFv2 and Bedrock checks together to audit whether Web ACLs and Bedrock guardrails meet baseline security requirements.

```
$ prowler aws --checks wafv2_webacl_managed_rule_groups_enabled
bedrock_guardrail_sensitive_information_filter_enforced
bedrock_model_invocation_logging_text_data_enabled
bedrock_api_key_no_long_term_credentials
```

— Named check but no example or deeper mechanism Named check with a runnable scan command

5.46.0

THINNER COVERAGE BELOW

08 **Issues page and overview cards in Prowler Cloud** NEW 55

Adds an Issues page and overview cards that group scan findings into issues with verdicts, verdict filters, accessible issue details, and scoped finding evidence.

— UI feature named but no exact navigation path or config given 5.46.0

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Threat Intelligence

Feedly Threat Intelligence

Feedback ⓘ

Credential Leak Dark Web Agent: Monitor exposed accounts (2026-10-27) BLOG ↗

Feedly threat intelligence is a collect, analyze, share actionable intel faster - Feedly TI

New Credential Leak Dark Web Agent monitors up to five email domains per agent, with password-policy filtering, AI summaries, Insights Cards, Trends tab, and API/CSV export.

NEW INTEGRATION Adds the Credential Leak Dark Web Agent, monitoring up to five email domains per agent (including subsidiaries and acquisitions), checking for new matches every 24 hours.

Adds password-policy filtering to hide leaks where no password meets the organization's complexity requirements, reducing noise from non-compliant credentials.

Adds Credential Leaks Insights Cards showing affected email-and-password pairs, source excerpts, First seen and Status (New or Recycled) fields, and links to associated Threat Actor Insights Cards and Cyberattack Insights Cards.

Adds a Trends & Visualizations tab showing how findings change over time, their sources, and which accounts appear most frequently.

Supports exporting the Insights Card credential table as CSV (with passwords, if present) or retrieving findings and Insights Card content — including passwords — through the Feedly API.

+2 more

WAS THIS USEFUL? 👍 👎

◆ SIEM

Databricks

Feedback ⓘ

Editing files in Unity Catalog volumes is now generally available (2026-10-27)

BLOG ↗ RANK

Databricks is a cloud-based data lakehouse platform for data engineering, analytics, machine learning, and AI workloads, built on Apache Spark.

> Databricks shipped GA in-place editing of Unity Catalog volume files, added OpenAI's GPT-6.1 Terra model to Unity Gateway, and made Lakeflow Performance Observability generally available with an AI-assisted query optimization flow.

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

01 Lakeflow Jobs performance observability and Optimize button

NEW 77

Lakeflow Jobs now shows job-level query performance and data volume metrics for supported tasks, with query profiling for serverless compute, SQL Warehouses, and Lakeflow Pipelines. An Optimize button in the Lakeflow Jobs UI opens Genie Code, loaded with the performance insights, to suggest query improvements; this capability will soon be enabled by default for workspaces with the compliance security profile enabled.

Review query performance bottlenecks and open Genie Code for AI-assisted query tuning after a serverless job run.

> 📍 In the Lakeflow Jobs UI, open a completed job run, select a supported task, view the query performance metrics and data volume metrics panel, then click Optimize to open Genie Code with the performance insights loaded.

Review job-level data volume and query performance metrics for a serverless Lakeflow job to identify slow queries before they affect SLAs.

> 📍 In the Lakeflow Jobs UI, open a job run, navigate to the Query Performance tab to view metrics and insights, then click the Optimize button to open Genie Code for AI-assisted query improvements.

— Names metrics, supported compute types, and exact Optimize/Genie Code flow

Improved Lakeflow Performance Observabi… Improved Lakeflow Performance Observabi…

THINNER COVERAGE BELOW

02 **Editing files in Unity Catalog volumes** NEW 55

Files stored in Unity Catalog volumes can now be edited directly in the Databricks workspace file editor, accessible from Catalog Explorer or the schema browser in a notebook or SQL editor, without needing to download and re-upload them. This capability is now generally available.

Edit a config or data file in a Unity Catalog volume directly without leaving the workspace.

> 📍 In the workspace, open Catalog Explorer, navigate to your Unity Catalog volume, select a file, and choose Edit to modify it in the workspace file editor — no download or re-upload needed.

— Clear UI navigation path but no deeper mechanism detail

Editing files in Unity Catalog volumes …

03 **OpenAI GPT-6.1 Terra on Unity Gateway** NEW 35

OpenAI's GPT-6.1 Terra model is now available on Unity Gateway via the Foundation Model APIs.

— Names model and API surface but no usage detail

OpenAI GPT-6.1 Terra is now available o…

NEW INTEGRATION QuickBooks Online connector (Beta) in Databricks Lakeflow Connect

NEW INTEGRATION PayPal connector (Beta) in Databricks Lakeflow Connect

> QuickBooks connector is available in Beta (2026-10-26) BLOG ↗ STABLE

> PayPal connector is available in Beta (2026-10-24) BLOG ↗ STABLE

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

Feedback ⓘ

Tenzir Node v6.20.0: A new data model for faster pipelines (2026-10-26) BLOG ↗

RANK

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

> Tenzir v6.20.0 introduces the `ai_decide` operator for typed AI-driven decisions, broadens secrets support across functions, adds encryption and checksum functions, extends integrations with DuckDB, ClickHouse, SentinelOne and Windows Event Forwarding, and changes the default insert mode for `to_clickhouse`.

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

01 **ai_decide operator with typed decisions** NEW 75

Adds the `ai_decide` operator for making typed decisions in pipelines, backed by a new Decisions endpoint, with support for Cloudflare Workers AI and OpenAI backends.

— Names operator, endpoint, backends but no usage example

Tenzir Node v6.20.0: A new data model f…

02 **Secrets support expanded across functions** NEW 70

Secrets can now be passed as function arguments, secret keys are supported for `hmac` with secrets censored on assignment, and secret seeds are supported for Crypto-PAn anonymization.

Pass a secret value as a function argument to `hmac` without exposing the key in pipeline code.

```
$ yield hmac('sha256', secret('MY_HMAC_KEY'), data)
```

— Includes a runnable example using `hmac` and `secret`

Tenzir Node v6.20.0: A new data model f…

THINNER COVERAGE BELOW

03 **New data manipulation and checksum functions** NEW 55

Adds the `type_kind` function to inspect a value's type, plus Luhn check digit and CRC-32 checksum functions, OCSF functions, and functions for converting records to key-value entries and record values to a list.

– Names `type_kind` but other functions left unnamed

Tenzir Node v6.20.0: A new data model f…

04 **SentinelOne integration enhancements** IMPROVED 55

Adds SentinelOne Data Lake API keys, summaries computed directly in SentinelOne, and filter and time-window pushdown for SentinelOne via the console API.

– Names API keys and pushdown mechanism but no config example

Tenzir Node v6.20.0: A new data model f…

05 **`to_clickhouse` default insert mode changed** BREAKING 55

`append` is now the default mode for `to_clickhouse`, changing prior insert behavior on upgrade.

– Names config default change but no migration steps

Tenzir Node v6.20.0: A new data model f…

06 **Named in-memory DuckDB databases** NEW 50

Adds named in-memory DuckDB databases with the ability to read from and write to DuckDB databases.

– Names the integration but gives no config or command

Tenzir Node v6.20.0: A new data model f…

07 **Encryption functions and Parquet encryption** NEW 45

Adds encryption functions and Parquet encryption using a shared AES key.

– Mentions AES key but no function names or flags

Tenzir Node v6.20.0: A new data model f…

08 **New Suricata schema types** NEW 45

Adds Suricata schema types for NTP, LDAP, LLMNR, DNP3, drop, and frame events.

– Names all event types but no usage example Tenzir Node v6.20.0: A new data model f…

09 **Windows Event Forwarding without Windows collector** NEW 35

Adds support for Windows Event Forwarding without requiring a Windows collector.

– Brief description, no setup detail Tenzir Node v6.20.0: A new data model f…

10 **Read events from ClickHouse catch-all tables** NEW 35

Adds support for reading events from ClickHouse catch-all tables.

– Minimal description, no schema or query example

Tenzir Node v6.20.0: A new data model f…

11 **PRI-prefixed and Cisco syslog header parsing** NEW 35

Adds parsing support for PRI-prefixed text and Cisco syslog headers.

— States formats but no parser name or example Tenzir Node v6.20.0: A new data model f…

└──▷ **BREAKING ON UPGRADE**

! `append` is now the default mode for `to_clickhouse`, changing prior insert behavior on upgrade.

WAS THIS USEFUL?  

◆ Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

View scan progress while an AI scan is running (2026-10-26) BLOG ↗ RANK

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

> Trend Vision One added live progress streaming for in-progress AI Artifact Scanner scans and introduced granular, per-task approval controls for remote support access.

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

01 **Partial results streaming during AI scans** NEW 93

The `visionone` output destination gains `partial-upload: true` to enable uploading partial results while an AI scan is still running, plus `partial-frequency: <int>` to control how often results are sent as a completion-percentage threshold (defaults to 25%). The TrendAI Artifact Scanner interactive configuration wizard exposes the `partial-upload` option, and the AI Application Security dashboard shows a Scan in Progress banner with completed-attack count, a progress bar, a Cancel scan button, and auto-refreshing results tables and widgets as findings arrive.

Stream findings to Vision One every 10% of scan completion so your team can triage high-severity results before the full scan finishes.

yaml

```
outputs:
  - type: visionone
    partial-upload: true
    partial-frequency: 10
```

— Names exact config keys, default, and command example

View scan progress while an AI scan is …

02 **Granular remote support access controls** NEW 87

Under Administration → Support Settings, remote support is split into Proactive service console access (read-only access for customer success advisors, solution engineers, incident response, and threat researchers) and Troubleshooting and diagnostics (technical support console access, diagnostic package collection, and remediation utility deployment). Each access type can be set to Automatic or Manual approval mode, gated behind a master Enable TrendAI™ remote support toggle; diagnostic packages (logs, memory dumps, endpoint state) are stored in Azure in the US region for 21 days.

Grant the incident response team read-only console access and require manual approval for any diagnostic package collection.

> 📍 In Trend Vision One, go to Administration → Support Settings. Enable TrendAI™ remote support. Under `Proactive service console access`, enable access for the incident response team as part of an active service agreement. Under `Troubleshooting and diagnostics`, enable 'Collect diagnostic packages from endpoint agents without console access' and set its approval mode to `Manual`.

Allow automatic deployment of remediation utilities during an active support case — noting that diagnostic package collection must also be enabled as a prerequisite.

> 📍 In Trend Vision One, go to Administration → Support Settings. Enable TrendAI™ remote support. Under `Troubleshooting and diagnostics`, enable 'Collect diagnostic packages from endpoint agents without console access', then enable `Deploy remediation utilities without console access`, and set both to `Automatic` approval.

— Exact UI navigation and setting names, no API endpoint

`Updated remote support settings (2026-1…`

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

OXO

Feedback ⓘ v2.20.0 NOTES ↗ RANK

OXO is a security scanning orchestrator for the modern age.

> OXO v2.20.0 adds a logout subcommand for managing local authentication and improves OpenRouter model reliability with regional endpoint routing.

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

01 **Auth logout subcommand** NEW 70

A new `oxo auth logout` subcommand deletes the locally stored authorization token, useful for revoking a cached session on a shared or CI machine after a pipeline run.

Revoke the locally cached OXO session token, e.g. on a shared or CI machine after a pipeline run.

```
$ oxo auth logout
```

— Exact runnable command given, but mechanism beyond deletion unspecified v2.20.0

THINNER COVERAGE BELOW

02 **Regional routing for OpenRouter models** IMPROVED 30

OpenRouter models are now routed through regional endpoints for improved availability.

— No specifics on regions or configuration v2.20.0

WAS THIS USEFUL? 👍 👎

defectdojo

Feedback ⓘ 3.4.300 NOTES ↗ RANK

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo 3.4.300 adds a product-level export API, bulk Jira unlinking, per-test location attribution, and Siemens SSB VulnID support.

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

01 **Product export API with NDJSON stream** NEW 50

A new product export API returns a manifest plus an NDJSON stream for bulk export of product data.

— No exact endpoint path or usage example given 3.4.300

02 **Bulk unlink findings from Jira** NEW 45

Findings can now be unlinked from Jira in bulk via the bulk edit action in both the UI and the API.

— Names UI and API surface but no exact endpoint 3.4.300

03 **Siemens SSB VulnID support** NEW 45

Adds support for Siemens SSB VulnIDs, configurable via `settings.dist.py` / `local_settings.py`.

— Names exact config files but no further mechanism 3.4.300

04 **Location attribution to reporting test** IMPROVED 40

Each location reference is now attributed to the specific test that reported it, backed by a supporting database migration.

— Internal data model change with no user action noted 3.4.300

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Censys

Feedback ⓘ October 26, 2026 (2026-10-26) BLOG ↗

Censys is a internet-wide scanning and search platform that discovers and monitors hosts, services, and certificates to help security teams assess their exposure.

Adds a fingerprint observation API endpoint in Adversary Investigation and a PHP2Shell ARC dashboard for tracking worldwide PHP exposure over time.

Adds a new fingerprint observation API endpoint in the Adversary Investigation module to retrieve which TLS certificate chains an IP has presented, on which ports, and when — enabling infrastructure-change spotting and host pivoting.

NEW DETECTION UPDATED DETECTION Adds the Censys ARC PHP2Shell dashboard to track worldwide PHP exposure over time.

NEW DETECTION New Platform software fingerprints: changedetection.io, CLIProxyAPI, MarkLogic Server, ManageEngine OpManager

UPDATED INTEGRATION Palo Alto XSOAR integration v3.2.3

UPDATED INTEGRATION ServiceNow TISC integration v1.1.0

UPDATED INTEGRATION SentinelOne Singularity Operations Center integration

UPDATED INTEGRATION Microsoft Sentinel integration v3.0.1

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Auth0 Feedback ⓘ

Identity Conversion Suite is now in Early Access (2026-10-26) BLOG ↗ RANK

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

> Auth0 introduced the Identity Conversion Suite in Early Access, a new Enterprise paid add-on bundling Anonymous Sessions and Experiment Center for improving visitor-to-user conversion.

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

01 **Anonymous Sessions for pre-login visitors** NEW 30

As part of the new Identity Conversion Suite (Early Access, paid add-on for Enterprise plans), Auth0 adds Anonymous Sessions to maintain session context for visitors before they sign up or log in.

— Describes purpose but no config or API details Identity Conversion Suite is now in Ear…

02 **Experiment Center for login A/B testing** NEW 30

Also part of the Identity Conversion Suite, Experiment Center lets teams A/B test login experiences and review results in metrics dashboards.

— Mentions dashboards but no concrete metrics or setup steps

Identity Conversion Suite is now in Ear…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer

Feedback ⓘ 1.39.0 (2026-10-26) BLOG ↗ RANK

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

> Runlayer tightened MCP gateway security with cookie masking in approvals, stricter connector redirect and MCP request-size limits, locked-down server scoping for federated agent accounts, and CLI-only local stdio servers, while also shipping a new `runlayer_agent_trigger` Terraform resource and exposing MCP `resources/templates/list` through the gateway.

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

01 Terraform resource for agent triggers NEW

88

The new `runlayer_agent_trigger` Terraform resource manages an agent's event and poll triggers as code; changing `trigger_type`, `event_key`, or `account_id` replaces the trigger, and it requires agent triggers enabled for the organization.

Define an agent's poll trigger as Terraform code so it is version-controlled and repeatable across environments.

hcl

```
resource "runlayer_agent_trigger" "nightly_poll" {
  agent_id     = runlayer_agent.my_agent.id
  trigger_type = "poll"
  event_key    = "nightly"
  account_id   = "<account_id>"
}
```

— Named resource and fields with a runnable Terraform example 1.38.0 (2026-10-26)

02 Federated agent server scoping enforcement BREAKING

76

For federated agent accounts (signed in through an identity provider), tokens are now restricted to only the servers they were issued for; the setting shows as on and permanently locked, and existing API keys on those accounts are refused at the server level. Before unbinding such an account from its identity provider, any unneeded keys must be revoked.

— Clear mechanism and migration note, no runnable command 1.39.0 (2026-10-26)

03 MCP resources/templates/list exposed through gateway NEW

74

The Runlayer gateway now exposes the MCP `resources/templates/list` method so clients can list a connector's resource templates from upstream servers.

Verify that a connector's MCP resource templates are now visible to clients via the Runlayer gateway.

```
$ curl -X POST https://<your-runlayer-host>/mcp \
  -H 'Content-Type: application/json' \
  -d
'{"jsonrpc":"2.0","id":1,"method":"resources/templates/list","par
ams":{"connector_id":"<connector_id>"}}'
```

— Named method with a runnable curl example `1.38.0 (2026-10-26)`

---

04 **MCP request body size capped at 36 MiB** BREAKING

71

MCP endpoint requests are now limited to 36 MiB, returning HTTP 413 for anything larger. The limit is sized to allow a full 25 MB artifact upload via `create_artifact` / `update_artifact`.

— Exact limit and status code, but no direct remediation command `1.39.0 (2026-10-26)`

---

05 **Connector redirects restricted to same origin** BREAKING

67

Connector URL redirects are now restricted to within the server's own origin — redirects to another host, another port, or from `https://` down to `http://` are no longer followed, though a same-host `http://` to `https://` upgrade is still permitted. Connector configurations that relied on cross-origin redirects must be updated to use the final destination URL directly.

— Precise before/after behaviour with migration guidance `1.39.0 (2026-10-26)`

---

06 **Connector page filtering, badges, and deletion** NEW

61

The Connectors page gains a Hosting filter to narrow connectors by type (Local , Remote , or Runlayer Deploy), replaces the Official badge with a Third-party icon for connectors not built by the vendor while removing the Secure badge, and lets admins delete a connector directly from the ⋯ menu on its connector page instead of only from Manage connectors.

Filter the Connectors page to show only connectors you can self-host, to audit your local attack surface.

📍 In the Runlayer web app, go to Connectors, open the Hosting filter, and select `Local`.

— Names filter values and badges with a clear UI starting point `1.39.0 (2026-10-26)` `1.38.0 (2026-10-26)`

THINNER COVERAGE BELOW

07 **Local stdio servers run only via Runlayer CLI** BREAKING

54

Local stdio servers are now restricted to run only on the user's machine via the Runlayer CLI; Runlayer no longer starts a local stdio server on its own servers when called through the proxy URL.

— States the new restriction but no setup steps given `1.39.0 (2026-10-26)`

08 **Cookie masking in approval requests** NEW

47

Approval requests across the web app, Slack, and Microsoft Teams now mask every cookie field, including fields such as `auth_cookie` and `session_cookie`.

— Named fields and channels but no reader action path `1.39.0 (2026-10-26)`

09 **'Owned by me' tab rename** IMPROVED

25

The personal tab on the Agents, Skills, Plugins, and Artifacts pages has been renamed to `Owned by me`.

— Bare naming change with no further detail `1.38.0 (2026-10-26)`

└──▷ **BREAKING ON UPGRADE**

! Federated agent accounts (signed in through an identity provider) now have server scoping permanently on and locked — existing API keys on those accounts are refused at the server level; revoke any keys you no longer need before unbinding the account from its identity provider.

! Connector URLs that redirect to another host, another port, or from `https://` down to `http://` are no longer followed — update connector configurations to use the final destination URL directly.

! MCP endpoint requests with a body larger than 36 MiB are rejected with HTTP 413 — payloads exceeding this limit must be reduced or split.

! Runlayer no longer starts a local stdio server on its own servers when called through the proxy URL — local stdio servers must be run on the user's machine via the Runlayer CLI.

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

varmor

Feedback ⓘ v0.10.6 NOTES ↗ RANK

vArmor is a cloud-native container hardening system that leverages AppArmor/BPF/Seccomp and NetworkProxy technologies to enforce access control from system calls to application protocols — protecting workloads including AI Agents.

> vArmor v0.10.6 introduces AI-assisted tooling for building NetworkProxy egress allowlists, diagnosing policy and enforcement issues, and optimizing EnhanceProtect policies, alongside several breaking changes affecting enforcer validation, NetworkProxy policy changes, and MITM certificate handling.

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

01 **Enforcer value validation on policy create/update** BREAKING

80

Unknown, duplicate, or improperly separated `enforcer` values (e.g. `AppArmor-BPF`) are now rejected on policy creation and update; only supported names and concatenated combinations such as `AppArmorBPF` are accepted.

— Gives exact before/after values and the config field affected v0.10.6

02 **MITM CA bundle rotation after Manager upgrade** BREAKING

80

After upgrading the Manager, MITM policy Secrets and mounted CA bundles must be verified to no longer contain `emSign ECC Root CA - G3`; applications that cache trust stores must be reloaded, and application-owned CA stores must be updated separately.

— Names the exact certificate and the required remediation steps v0.10.6

03 **Automated Policy Optimization for EnhanceProtect** NEW

60

Refines `EnhanceProtect` policies by analyzing violation logs and proposing rule removals that require user approval before being applied.

— Names EnhanceProtect and the approval workflow, but no exact command shown v0.10.6

04 **NetworkProxy removal requires policy recreation** BREAKING

60

Removing NetworkProxy from a combined enforcer policy now requires full policy recreation; affected Pods must be cleaned up and recreated.

— States the required migration step but lacks exact commands v0.10.6

THINNER COVERAGE BELOW

05 Network Allowlist Advisor for NetworkProxy egress NEW

49

Builds and validates NetworkProxy egress allowlists from live application traffic, with a user-review step required before the allowlist is enforced.

— Describes mechanism and review step but no command or UI path given v0.10.6

06 Automated Diagnosis for policies and NetworkProxy NEW

43

Adds read-only troubleshooting that covers policy state, Pod injection, node enforcement, and NetworkProxy to help identify issues.

— Names the four areas covered but no interaction details v0.10.6

└──▷ BREAKING ON UPGRADE

! Unknown, duplicate, or improperly separated `enforcer` values (e.g. `AppArmor-BPF`) are now rejected on policy creation and update — only supported names and concatenated combinations such as `AppArmorBPF` are accepted.

! Removing NetworkProxy from a combined enforcer policy now requires full policy recreation; affected Pods must be cleaned up and recreated.

! After upgrading the Manager, MITM policy Secrets and mounted CA bundles must be verified to no longer contain `emSign ECC Root CA - G3`; applications that cache trust stores must be reloaded, and application-owned CA stores must be updated separately.

WAS THIS USEFUL?  

ToolHive

Feedback ⓘ v0.57.0 NOTES ↗ RANK

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive v0.57.0 hardens authentication and token flows across vMCP and `thv run`, adds backend-scoped vMCP authorization, and tightens workload secret mounting and transport-layer request handling.

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

01 **Authentication and token exchange hardening** NEW 73

Adds `dcrConfig.registrationProbe` with an `authorization-endpoint` strategy for dynamic client registration probing, supports `authorization_details` on token exchange requests to enable OAuth 2.0 Rich Authorization Requests (RAR) in vMCP token flows, and wires OIDC transport flags into `thv run` auth middleware config so OIDC settings can be passed directly at server launch. Also adds opt-in upstream subject identifiers to surface the upstream identity subject in token and auth flows, and lets embedders refuse refresh grants on dead upstreams to prevent silent token refresh against unreachable services.

— Names config key and command but mechanism details are brief v0.57.0

THINNER COVERAGE BELOW

02 **Workload auth secret mounting in inline mode** NEW 45

Mounts workload auth secrets in inline mode, allowing inline-mode deployments to consume workload identity credentials.

— Clear capability but no config key or flag given v0.57.0

03 **MCPRemoteProxy pod health classification** IMPROVED 45

Classifies `MCPRemoteProxy` pods by container readiness rather than phase alone, improving the accuracy of proxy pod health reporting.

— Names the component and the before/after mechanism v0.57.0

04 **Backend-scoped vMCP authorization and policy updates** NEW 42

Adds backend-scoped vMCP authorization, enabling per-backend access control policies on the Virtual MCP Server, and lets delegate clients change policy on restart so updated authorization policies take effect without full redeployment.

— Describes behavior but no named config surface v0.57.0

05 **Request body rejection on non-POST methods** IMPROVED

Rejects request bodies on non-POST methods, hardening the MCP HTTP transport against malformed or abusive requests.

— Named transport behavior but no endpoint or flag specifics v0.57.0

WAS THIS USEFUL? 👍 👎

Docker Sandboxes

Feedback ⓘ v0.50.0 NOTES ↗ RANK

Docker Sandboxes run AI coding agents in isolated environments on your machine or on Docker-managed cloud infrastructure

> Docker Sandboxes v0.50.0 adds cloud sandbox inspection, configurable host mounts and persistent kit volumes, tightens network rules with query-parameter filtering, and closes a symlink-escape gap in `sbx cp` alongside two notable Windows/local-build breaking changes.

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

01 **Cloud sandbox inspection command** NEW 85

Adds `sbx --cloud inspect <id|name>` to show details for a single cloud sandbox, including published port URLs; supports `--format json` and `--format yaml` for scripting.

Inspect a specific cloud sandbox and get its published port URLs in machine-readable form for scripting or CI pipelines.

```
$ sbx --cloud inspect my-sandbox --format json
```

— Named command with format flags and a runnable example v0.50.0

02 **Configurable kit host mounts** NEW 80

Adds `--mount HOST:SANDBOX_PATH` flag to `sbx create` and `sbx run` for selecting a host directory for a kit-declared host mount; environment files support the same bindings through `mounts`.

Create a sandbox and bind a host directory to a kit-declared mount target so the agent can read/write a specific host path.

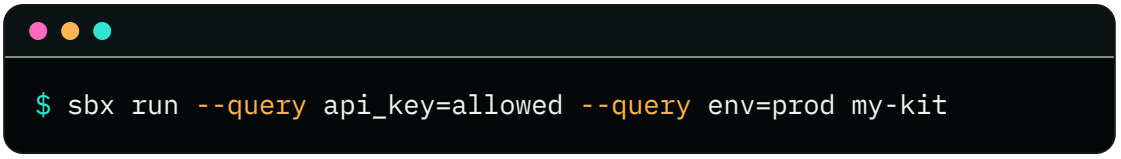
```
$ sbx create --mount /home/user/project:/workspace/project my-kit
```

— Flag and env key named with a runnable example v0.50.0

03 **HTTP query-parameter network rules** NEW 65

Adds a repeatable `--query name=value` flag to HTTP network rules to require specific query-parameter restrictions on outbound requests.

Restrict outbound HTTP access to requests that include specific query parameters, useful for scoping API calls to a known tenant or environment.



— Named flag with runnable example, limited explanation of scope v0.50.0

04 **sbx cp symlink escape protection** IMPROVED 65

`sbx cp` rejects symbolic links that escape a sandbox workspace when reached through an alternate host path, including a symlink, junction, macOS `/private/tmp`, or a Linux bind mount of a workspace subdirectory.

— Mechanism and attack paths named, no example command v0.50.0

THINNER COVERAGE BELOW

05 **Skip host commands flag for Environment MCP** NEW 55

Adds `--skip-host-commands` flag to disable Environment MCP host-execution commands for a sandbox.

— Flag named but no example or further mechanism given v0.50.0

06 **Clone workspace file preservation** IMPROVED 55

Sandboxes created with `--clone` now preserve kit-provided workspace files and report clone initialization failures.

— Flag named, behavior described, no example v0.50.0

07 **Windows daemon same-user pipe check** BREAKING 55

On Windows, the updated CLI verifies that the daemon's named pipes belong to the same Windows user; older elevated daemons must be stopped with the previous CLI or Task Manager before using the updated CLI.

— Mechanism and required migration step both given v0.50.0

08 **Persistent kit volumes across updates** NEW 50

Persistent kit volumes now retain data across kit updates and sandbox removal and recreation with the same name. Existing sandbox records keep their previous storage behavior — persistent kit volumes only apply to sandboxes created with v0.50.0 or later.

— Behavior and migration cutoff described, no command shown v0.50.0

09 **SSL trust store override warnings** IMPROVED 50

On macOS and Windows, `sbx diagnose` and daemon logs now warn when `SSL_CERT_FILE` or `SSL_CERT_DIR` replaces the system trust store.

— Named env vars and command, limited mechanism detail v0.50.0

10 **Local source kit rebuild on upgrade** BREAKING 45

Each local source kit rebuilds once after upgrading, because local source kits are now reused only when their digests match recorded local builds.

— Cause explained but no command or config to act on v0.50.0

11 **sbx exec exit status on failure** IMPROVED 40

`sbx exec` now returns exit status 125 when it cannot run the command, for example when the sandbox does not exist.

— Specific exit code named, no broader context v0.50.0

12 **Gordon clears tool approval grants on resume** IMPROVED 40

Gordon clears Approve all and Always allow grants when a conversation resumes, requiring fresh approval before running tools again.

— Behavior described but no named surface or config v0.50.0

└──▷ **BREAKING ON UPGRADE**

! Windows: older elevated daemons must be stopped with the previous CLI or Task Manager before using the updated CLI, because the updated CLI verifies that the daemon's named pipes belong to the same Windows user.

! Each local source kit rebuilds once after upgrading, because local source kits are now reused only when their digests match recorded local builds.

! Existing sandbox records retain their previous storage behavior — persistent kit volumes only apply to sandboxes created with v0.50.0 or later.

WAS THIS USEFUL?  

Backslash

Feedback ⓘ

AI Plugin Inventory: Uncover Hidden Risks | Backslash - Backslash (2026-10-26)

BLOG ↗ RANK

Backslash detects and prioritizes reachable vulnerabilities in cloud-native application code and dependencies.

> Backslash shipped a broad AI agent security layer in this window, anchored by a new Plugin Inventory with child-component risk scoring, plus hooks discovery, local model governance, AI rules vetting, inline host guardrails, and compliance-ready audit trails.

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

01 AI Plugin Inventory with child-component risk scoring NEW

85

Adds a Plugins tab that discovers every AI plugin installed on developer endpoints, surfacing publisher, source repository, unpacked contents, risk severity, associated agents, and last sync timestamp, fully filterable and sortable. Child-component risk scoring assigns a plugin's overall severity by the highest risk found among its nested `scripts`, MCP tools, `hooks`, and `rules`, preventing a low-risk wrapper from concealing a critical child component. MCP server and agent-skill mapping ties every MCP connection and skill to its parent plugin and host agent, surfacing over-permissive tool definitions and untrusted content sources.

— Names the tab and nested component types but no limits

AI Plugin Inventory: Uncover Hidden Ris…

02 Hooks Inventory across coding agents NEW

73

Automatically discovers hooks across Claude Code, Codex, Cursor, and Copilot, tracking whether a hook was added by Backslash, a developer, or a third-party plugin. Combines deep code parsing and digital-signature validation for hook files to verify component provenance and detect unauthorized script tampering before execution, while vetting scripts against backdoor and credential-theft threat models.

— Names agents and mechanism but no UI path given

AI Plugin Inventory: Uncover Hidden Ris…

03 AI Rules file vetting and quarantine NEW

73

Scans standing configuration files (`.cursorrules`, `AGENTS.md`, `CLAUDE.md`) for instruction poisoning and exposed credentials, with guardrails that automatically quarantine unsafe files at session start.

— Names exact config files and quarantine trigger point

AI Plugin Inventory: Uncover Hidden Ris…

04 **Local AI model discovery and governance** NEW 65

Adds local model discovery on macOS and Windows model directories for runtimes such as Ollama, LM Studio, and llama.cpp, governing locally-run models under the same allow/deny controls as remote foundation models.

— Names runtimes and platforms but no config detail

AI Plugin Inventory: Uncover Hidden Ris…

THINNER COVERAGE BELOW

05 **Automated compliance audit records and run tracing** NEW 50

Delivers automated audit records and host-level run tracing aligned to regulatory frameworks including EU AI Act, NIS2, DORA, and SOC, without requiring manual developer reporting.

— Names frameworks but no report format or access point

AI Plugin Inventory: Uncover Hidden Ris…

06 **Inline host-level guardrails for agent actions** NEW 45

Enforces inline host-level guardrails that block unapproved command executions, credential access, and data egress before actions execute, rather than alerting after the fact.

— Describes behavior but no config or trigger names

AI Plugin Inventory: Uncover Hidden Ris…

WAS THIS USEFUL?  

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 27, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue