THE CYBER TOOLCHAIN NO. 106

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 13, 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 42 changes

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 18 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.

GOVERN

Answer who — or what — actually wrote the code you shipped

vim ? help

JFrog's AppTrust traceability links a release back through to the originating prompt, and its open Agentic Process Evidence spec gives auditors a format for agent activity rather than a vendor log dump. GitHub's durable push attribution, enforceable as a ruleset condition, keeps the pusher identity attached to a commit, so 'an agent did this' stops being an untestable claim at review time.

JFrog · GitHub

PROTECT

Scan the native code you have been shipping unscanned

Full C coverage plus rapid-response remediation means C dependencies and first-party native code go through the same pipeline as managed-language code. Teams carrying C in the build have generally had to run a separate, noisier toolchain and reconcile findings by hand.

Endor Labs

PROTECT

Gate secret changes and stop scanning the whole org at once

Secret scanning is now scoped per project instead of org-wide, so one noisy repo's findings stop drowning everyone else's queue, and global approvals put a human in front of secret mutations before they reach a running workload.

infisical

IDENTIFY

Keep a takeover scan from wandering into internal space

Subjack adds SRV record takeover checks and swaps its domain-availability lookups to RDAP, which stops the memory blowups that killed long runs mid-scan. OpenOSINT's exposure guard restricts investigations to public-only targets — the cheap guardrail that keeps an OSINT run from accidentally probing your own internal ranges.

subjack · OpenOSINT

OFFENSIVE SECURITY

◆ Reconnaissance & OSINT

OpenOSINT ⓘ v2.35.0 NOTES ↗ RANK

Openosint is a join the world's most widely adopted, AI-powered developer platform where millions of developers, businesses, and the largest open source community build software that advances humanity.

> OpenOSINT v2.35.0 adds a network exposure guard to restrict investigations to public-only targets.

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

01 **Public-only network exposure guard** NEW 60

A new `--public-only` flag restricts investigations to public-only targets, guarding against unintended network exposure.

— Names exact flag but no mechanism detail beyond restriction. v2.35.0

WAS THIS USEFUL? 👍 👎

subjack ⓘ v3.1.0 NOTES ↗ RANK

Subjack is a Go-based tool that checks subdomains for subdomain takeover vulnerabilities by detecting dangling CNAME records pointing to unclaimed third-party services.

> Subjack v3.1.0 adds opt-in SRV record takeover scanning, replaces its domain availability backend with RDAP to fix out-of-memory crashes, and improves resolver configuration and AXFR detection.

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

01 **Opt-in SRV record takeover detection** NEW 75

A new `-srv` flag lets users opt in to SRV record takeover checks, issuing roughly 28 DNS queries per domain to cover services like SIP, XMPP, LDAP, Kerberos, IMAP, and CalDAV, in addition to standard CNAME-based detection.

Scan for SRV record takeovers alongside standard CNAME checks — useful when targets run SIP, XMPP, or LDAP services with potentially expired backend domains

```
$ subjack -w subdomains.txt -srv -t 50 -o results.json
```

— Names flag, query count, and protocols covered v3.1.0

02 **RDAP-based domain availability checks** IMPROVED 60

Domain availability checks now use RDAP (RFC 9083) with a whois fallback instead of the heavyweight zonedb dependency, eliminating out-of-memory crashes on small VPSes.

— Explains mechanism and bug fixed, no direct command v3.1.0

03 **Resolver list host:port support** IMPROVED 60

The `-r` flag for specifying a resolver list now accepts entries in `host:port` format in addition to bare IP addresses.

Use a custom resolver list with non-standard ports (e.g. an internal DNS forwarder) to route all subjack queries through specific resolvers

```
$ subjack -w subdomains.txt -r resolvers.txt -ssl -o
results.json
```

— Names flag and format change with example command v3.1.0

THINNER COVERAGE BELOW

04 **AXFR detection fix** IMPROVED 20

Fixes a bug in AXFR (zone transfer) detection, though the specific prior behavior or cause is not detailed.

— Only a bare mention of a fix, no detail v3.1.0

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

merlin ⓘ v2.2.1 NOTES ↗ RANK

> Merlin v2.2.1 expands its release bundle with additional server binaries for macOS and Linux ARM platforms.

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

01 **Additional macOS and Linux ARM server binaries** NEW 50

The release bundle now includes macOS (Darwin) `arm64` and Linux `armv7` and `arm64` server binaries, alongside the existing Windows x64, Linux x64, and macOS x64 builds.

— Names exact platforms/architectures but no further detail v2.2.1

WAS THIS USEFUL? 👍 👎

merlin-agent ⓘ v2.5.1 NOTES ↗

└──▷ GET THIS VERSION

```
$ git clone --branch v2.5.1 https://github.com/Ne0nd0g/merlin-agent.git
# already have the repo? check out this version:
$ git checkout v2.5.1
```

COVERAGE

NEW supported platform

macOS (Darwin) arm64 and Linux armv7 / arm64 agent binaries

All of merlin-agent's coverage changes →

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

Dreadnode ⓘ File and folder uploads in Web Chat NOTES ↗ RANK

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

> Dreadnode ships file, folder, and pasted-image uploads in Web Chat alongside a custom `severity_policy` parameter for AI red-teaming and structured binary-analysis findings, plus smaller improvements to session memories, capability dependencies, sharing, and eval tracking.

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

01 **Custom severity policy for AI red-teaming** NEW 95

Adds a `severity_policy` parameter to `generate_attack` so you can define custom risk taxonomies — thresholds, matrices, and aliases — that the agent applies when classifying findings, e.g. generate_attack(severity_policy={"thresholds": {"critical": 9.0, "high": 7.0, "medium": 4.0}, "aliases": {"rce": "critical"}}).

Apply a custom severity taxonomy during an AI red-team engagement so findings are classified against your org's own risk thresholds rather than defaults.

python

```
generate_attack(severity_policy={"thresholds": {"critical": 9.0,
"high": 7.0, "medium": 4.0}, "aliases": {"rce": "critical"}})
```

— Named param with mechanism and runnable example File and folder uploads in Web Chat

02 **Capability tool directories via dependencies.path** NEW 85

Adds a `dependencies.path` field in `capability.yaml` so capabilities can declare additional tool directories for agents running in E2B sandboxes, e.g. `dependencies:\n path:\n - /home/user/.cargo/bin\n - /usr/local/go/bin`.

Declare extra tool directories in a capability so agents can find binaries installed outside the default PATH in E2B sandboxes.

yaml

```
dependencies:
  path:
    - /home/user/.cargo/bin
    - /usr/local/go/bin
```

— Named config file, field, and runnable example File and folder uploads in Web Chat

03 **Filter project memories by session ID** NEW 75

Supports filtering project memories by session ID via the API, SDK tool, and `ApiClient.list_project_memories`, e.g. ApiClient.list_project_memories(project_id="<project_id>", session_id="<session_id>").

Trace all memories a specific session created or modified when investigating an agent run.

```python
ApiClient.list_project_memories(project_id="<project_id>",
session_id="<session_id>")
```

— Named API method with runnable example File and folder uploads in Web Chat

04 **Structured binary-analysis output** IMPROVED 70

The binary-analysis capability now emits structured findings and assets — recovered keys, vulnerability root causes, malicious capabilities, extracted payloads, and embedded endpoints — trackable through the review-and-submit workflow.

— Lists concrete output types but no API or command shown

File and folder uploads in Web Chat

THINNER COVERAGE BELOW

05 **File and folder uploads in Web Chat** NEW 55

Adds file, folder, and pasted-image uploads directly in Web Chat; uploads land in the runtime workspace so the agent can reference them immediately.

— Describes behavior but no exact UI path or flag File and folder uploads in Web Chat

06 **Jump to tool call from Media tab** IMPROVED 55

Adds a `?tool_call=<id>` linkable URL parameter and a View in transcript button in the Media tab to jump directly to the tool call that logged an artifact.

— Named URL parameter but only a UI path to use it File and folder uploads in Web Chat

07 **Memories tab in Sessions** NEW 50

Adds a Memories tab to Sessions showing every memory the session created, updated, closed, or reopened, with links to full memory detail.

— Describes UI addition without further mechanism File and folder uploads in Web Chat

08 **Version visibility and updates in capabilities picker** IMPROVED 45

Web Chat capabilities picker now shows the installed version of each capability and offers an inline update when a newer version is available.

— Clear UI behavior but no further technical detail File and folder uploads in Web Chat

09 **Share action for Chat sessions** NEW 40

Adds a Share action to Chat session menus so a private session can be made visible to the workspace team without leaving Chat.

— Thin description of a UI-only action File and folder uploads in Web Chat

10 **Worker sessions on eval sample detail page** IMPROVED 40

Eval sample detail page now lists worker sessions the agent started for that attempt, with links to each session in Agent Sessions.

— Describes UI listing with no further mechanism File and folder uploads in Web Chat

COVERAGE

NEW compliance framework

AIRT dataset library — OWASP LLM and Agentic coverage mapping

All of Dreadnode's coverage changes →

Web Chat Interface Redesign NOTES ↗ STABLE

Pre-launch sandbox config panel, bulk finding triage, Assessment.record_attack_result() SDK method, dataset attribution metadata in SDK push, bulk memory close/reopen, safe clickable links in agent content.

- Adds Assessment.record_attack_result() SDK method so custom-target attacks (multistep tool attacks, indirect injection flows) can emit traces and materialize findings in the platform.
- Adds bulk triage on the Agent Output page — select multiple findings and set their triage status in a single action, with per-row results and a confirmation toast.
- Adds pre-launch sandbox configuration panel in the web chat session view: teams can choose capabilities, compute, secrets, and idle timeout before a sandbox starts, then inspect its runtime status, installed capabilities, health issues, and logs in a panel beside the chat.
- Adds dataset attribution metadata to SDK push: datasets now carry license, tags, language, task categories, and source provenance, populating the license badge, tag filters, and Source section in the platform.
- Adds bulk close and reopen for project memories: select multiple memories on the Memories page and close or reopen them via the bulk-action bar.

+3 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

graudit ⓘ v4.1 NOTES ↗

graudit is a grep-based source code auditing tool that uses regex signature databases to flag potential security vulnerabilities across many programming languages.

Adds a rule linter script for validating graudit signature databases.

└──▷ GET THIS VERSION

```
$ git clone --branch v4.1 https://github.com/wireghoul/graudit.git
# already have the repo? check out this version:
$ git checkout v4.1
```

COVERAGE

NEW detections Regex DoS (ReDoS) rules, Secrets detection rules, Default rules

NEW supported platforms iOS rules, Ruby rules

UPDATED detection

Extended regex correctness updates across multiple rule databases

All of graudit's coverage changes →

WAS THIS USEFUL? 👍 👎

GitHub ⓘ

Client telemetry user attribution and configurable capture (2026-10-12)

BLOG ↗ RANK

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

> GitHub shipped granular, opt-in telemetry controls for Copilot enterprise deployments, durable push attribution for commits with a new ruleset condition, and an Anonymous authentication option for Dependabot's organization-level private registries.

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

01 User attribution and granular capture for Copilot telemetry

NEW 95

A new `telemetry.capture.identity` setting attributes Copilot activity to a user by appending `user.name`, `process.user.name`, and `host.name` to OpenTelemetry data; it is off by default and kept separate from existing pseudonymous identifiers. Five independently toggleable capture classes — `prompts`, `responses`, `toolArguments`, `toolOutput`, and `policyDetail` — are each off by default and available in Copilot CLI, the GitHub Copilot app, and VS Code sessions using Agent Host, while the legacy `captureContent` shorthand still covers all five content classes but never enables identity capture. Explicit managed values take precedence over local settings, SDK configuration, and environment variables.

Enable user attribution plus tool-argument capture for an audit pipeline — without collecting prompts or responses — in your enterprise managed settings.

yaml

```yaml
telemetry:
  capture:
    identity: true
    toolArguments: true
    prompts: false
    responses: false
    toolOutput: false
    policyDetail: false
```

— Exact config keys and a runnable managed-settings example

Client telemetry user attribution and c…

02 Anonymous authentication for Dependabot org registries NEW

75

Dependabot now supports `Anonymous` as a selectable authentication method for organization-level private registries, letting registries that allow unauthenticated reads — such as internal package mirrors — work without a placeholder token workaround. Anonymous-authenticated registries apply to every granted repository, including those using only Dependabot security updates with no `dependabot.yml` file; note that an existing registry's authentication method cannot be changed, so switching a placeholder-token registry to `Anonymous` requires deleting and re-creating it via Settings > Secrets and variables > Private registries > New private registry.

Add a credential-free internal package mirror as an organization-level Dependabot registry so all granted repos can reach it without a token.

> 📍 In your organization, go to Settings › Secrets and variables › Private registries › New private registry, enter the registry URL, and set the authentication method to `Anonymous`.

— Exact UI path and auth value, plus noted re-creation limitation

`Anonymous organization-level registries…`

---

03 Durable push attribution and committer/pusher match rule NEW

60

Every commit pushed to a repository now carries a durable, authenticated pusher identity, independent of the author and committer metadata, and this identity is surfaced in the UI and API alongside the existing author and committer fields for spot-checks and investigations. A new repository rule condition, `Committer must match pusher`, blocks pushes where the committer metadata does not match any verified email address of the authenticated pusher.

— Names the new rule but no concrete API/UI steps given

`Push attribution for commits in limited…`

┗—▷ BREAKING ON UPGRADE

! You cannot change the authentication method of an existing registry — to switch a placeholder-token registry to `Anonymous`, you must delete it and re-create it via Settings > Secrets and variables > Private registries > New private registry.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical ⓘ v0.166.5 NOTES ↗ RANK

Infisical is an open-source platform for managing secrets, API keys, and environment variables, with encryption, access controls, and audit logs for developer teams.

> Infisical v0.166.5 rescoped secret scanning to individual projects, added global approvals for secret changes, and introduced nested JSON import as folder hierarchies.

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

01 **Nested JSON import as folder hierarchies** NEW 41

Secrets import now supports nested JSON objects, converting them into folder hierarchies when imported.

— Names the JSON import behavior but no exact command or path v0.166.5

02 **Project-scoped secret scanning** IMPROVED 37

Secret scanning moved to a single-project product scope, so all scanning operations and results are now scoped to individual projects rather than an organization-wide scope.

— Describes the scope change but no config or UI path v0.166.5

03 **Global secret-change approvals** NEW 10

Adds global approvals for secret changes, per the release summary, though no further detail on configuration or workflow is given.

— Only named in summary with no supporting detail v0.166.5

WAS THIS USEFUL? 👍 👎

◆ Software Supply Chain Security

JFrog ⓘ

Governing the Agentic SDLC Through Continuous Policy Enforcement (2026-10-12)

BLOG ↗ RANK

JFrog software supply chain solutions provide an end-to-end pipeline to control your binaries from build to production. Power your software updates to the edge.

> JFrog shipped a new framework for governing AI coding agents across the SDLC, adding prompt-to-release traceability in AppTrust and an open Agentic Process Evidence specification for auditing and policy enforcement of agent activity.

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

01 Agentic Process Evidence Protocol for AI agent governance NEW

81

The Agentic Process Evidence Protocol is a four-step framework that captures agent session telemetry (user prompts, thinking tokens, file edits, tool calls), cryptographically signs and attaches evidence to SDLC entities, evaluates that evidence at policy-as-code release gates before promotion, and persists session logs for post-release auditing. Gates can enforce checks such as approved LLM usage, approved MCP tool usage, human accountability, human review presence, and absence of unapproved lateral operations. Each agentic process is captured as five signed categories of facts: agent identity (agent version, harness version, model ID), execution traceability (full session logs as structured JSON covering every tool call, API call, input, and context source), provenance, integrity, and policy compliance. Session logs are bound to git commits and software releases, enabling retroactive searches — for example, identifying all releases produced using a later-identified malicious MCP tool.

— Mechanism and scope fully described but no config key or command given.

Governing the Agentic SDLC Through Cont…

02 Open Agentic Process Evidence specification NEW

76

JFrog publishes an open Agentic Process Evidence specification on GitHub (`https://github.com/jfrog/agentic-process-evidence`) for integration by agent frameworks (Cursor, OpenCode, Claude Code, or custom AI workers), regulatory bodies, and SDLC platform providers.

— Exact repository URL and named target frameworks make it directly actionable.

Governing the Agentic SDLC Through Cont…

THINNER COVERAGE BELOW

03 **Prompt-to-Release Traceability in AppTrust** NEW 59

Prompt to Release Traceability in JFrog AppTrust connects AI agent session logs to signed SLSA-compatible provenance attestations, binding them to git commits, artifacts, and released versions.

— Names the feature and SLSA linkage but no API or config path given.

Governing the Agentic SDLC Through Cont…

WAS THIS USEFUL? 👍 👎

Endor Labs ⓘ

Open source carries the world. Patching it at Mythos-scale can't fall to maintainers a

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 expanded into rapid-response vulnerability remediation and added full C language coverage to its AI-driven scanning stack.

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

01 **AI SAST and buildless scanning for C** NEW 70

Endor Labs AI SAST now has full C language support with buildless scanning available in the IDE and on PRs. It pairs with C SCA in the same platform so first-party C code and open-source C dependencies are covered under a single analysis run, and integrates into AURI by Endor Labs to provide function-level call paths, reproducible evidence, and policy enforcement across any agent, model, or CI stage.

— Names IDE/PR surfaces and AURI integration, no exact command

Why Endor Labs AI SAST for C finds what…

02 **Zero-Day Patches via Project Akrites** NEW 63

Endor Labs launched Zero-Day Patches, delivering AI-generated, verified fixes for novel open source vulnerabilities in under 24 hours. It joined the Linux Foundation's Project Akrites as a founding member alongside Anthropic, OpenAI, Amazon, Chainguard, Cisco, Citi, IBM, JPMorgan Chase, and Microsoft, routing coordinated vulnerability disclosure and remediation through a shared channel, and applies function-level reachability analysis across direct and transitive dependencies to rank which zero-day findings are actually reachable from a customer's application code.

— Names program, members, timeline, but no self-serve config or flag

Open source carries the world. Patching…

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion ⓘ Agent Registry check follows Agent Registry to its new home:

NOTES ↗ RANK

> Plerion migrated its Agent Registry check to AWS's new standalone Agent Registry service ahead of the Bedrock AgentCore API shutdown, automatically closing findings from legacy registries.

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

01 Agent Registry check moves to new AWS service BREAKING

77

Check `PLERION-AWS-1110` ('Agent registry publishes new records without review') now runs against registries in the new AWS Agent Registry service instead of the deprecated Bedrock AgentCore registry API, ahead of the API's 30 October 2026 shutdown. The Plerion stack must be updated to the latest version to gain `agent-registry` read permissions needed to read the new registries; as a breaking consequence, registries created in AgentCore before 6 August are no longer checked and any open findings from those legacy registries are automatically closed on upgrade.

— Names exact check ID, permission, and cutoff dates clearly

Agent Registry check follows Agent Regi…

└──▷ **BREAKING ON UPGRADE**

! Registries created in AgentCore before 6 August are no longer checked and their open findings resolve — any findings from those legacy AgentCore registries will be automatically closed on upgrade.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Databricks ⓘ

Workspace-level package repositories for Databricks Apps are generally available (2026

BLOG ↗ RANK

Databricks is a unified data analytics platform built on Apache Spark that provides data engineering, machine learning, and collaborative notebook-based analytics in the cloud.

> Databricks brought Feature Views to general availability for batch, streaming, and request-time data, while adding Beta previews of agent memory pipelines, Unity Catalog governed tag propagation, and a Unity Gateway Skills REST API; it also shipped GA Lakeflow pipeline parameters, GA workspace-level package repositories for Databricks Apps, and default JDBC Unity Catalog connections for compliance security profile workspaces.

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

01 **Governed tag propagation in Unity Catalog** NEW 85

Tag propagation in Unity Catalog (Beta) uses lineage to automatically apply a governed tag to derived tables and views whenever `CREATE TABLE AS SELECT`, `CREATE VIEW`, `INSERT INTO ... SELECT`, or `MERGE` runs, keeping classification and ABAC enforcement consistent at table and column levels as data moves downstream.

— Names exact triggering SQL statements and enforcement mechanism

Governed tag propagation is in Beta (20…

02 **Parameters for Lakeflow pipelines** NEW 82

Pipeline parameters for Lakeflow pipelines are now GA: define key-value pair defaults in pipeline settings, override them at update start, and reference them in SQL with named parameter syntax, so a single pipeline can run against different catalogs, schemas, or tables without editing source code.

— Explains override and reference mechanism, lacks exact syntax example

Parameters for Lakeflow pipelines are n…

03 **Unity Gateway Skills API and developer tools** NEW 80

The Unity Gateway Skills REST API (Beta) supports create, read, update, list, delete, and finalize operations, and is available across the Databricks CLI, Python SDK, Go SDK, Go Modular SDK, Java SDK, and JavaScript SDK.

— Names all operations and SDKs but no endpoint path given

Unity Gateway Skills API and developer …

04 **Agent memory pipelines** NEW 65

Agent memory pipelines (Beta) connect a session store to a memory store and use an LLM to distill durable facts, preferences, and decisions from stored conversations into memory entries that are recalled in later sessions. A memory pipeline can run on demand or on an approximately 24-hour schedule.

— Mechanism and scheduling described but no API or command named

Agent memory pipelines are in Beta (202…

THINNER COVERAGE BELOW

05 **Default JDBC Unity Catalog connections for compliance workspaces**

IMPROVED 56

JDBC Unity Catalog connections — including SSL certificate support — will become available by default for workspaces with the compliance security profile enabled, allowing external database access via custom JDBC drivers governed by Unity Catalog.

— Names SSL support and driver governance but no config path

JDBC Unity Catalog connections will soo…

06 **Feature Views reach general availability** NEW 53

Feature Views are now generally available for batch, streaming, and request-time data workloads, enabling declarative feature definition and computation from batch or streaming data for model training and online serving. Workspaces with the compliance security profile enabled will gain access in mid-November 2026.

— Describes scope and rollout timing but no concrete API or syntax

Feature Views are now generally availab… Feature Views will soon be available by…

07 **Workspace package repositories for Databricks Apps** NEW 50

Workspace admins can configure default Python and npm package repositories for Databricks Apps, with app-level overrides still supported. This capability is now generally available.

— Names Python/npm repos but no config key or UI path given

Workspace-level package repositories fo…

COVERAGE

NEW integration Amazon Bedrock Mantle model provider support in Unity Gateway

All of Databricks's coverage changes →

Unity Gateway supports Amazon Bedrock Mantle (2026-10-09) BLOG ↗ STABLE

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir ⓘ changelog · Oct 13, 2026 CHANGELOG ↗ RANK

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

> Tenzir shipped a new internal data model aimed at faster pipeline execution, alongside Explorer usability improvements in the Platform UI.

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

01 **Explorer scrolling and error details** IMPROVED 40

Tenzir Platform v1.40.1 adds scrolling and error details to the Explorer.

— Names UI area and two concrete additions changelog-20261013-7f4b66b9

02 **New data model for faster pipelines** IMPROVED 25

Tenzir Node v6.20.0-rc.2 introduces a new data model designed for faster pipeline execution.

— Names the change but no mechanism or benchmarks given changelog-20261013-7f4b66b9

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

defectdojo ⓘ 3.4.100 NOTES ↗

DefectDojo is an open-source application security vulnerability management platform that consolidates, deduplicates, and tracks findings from many security scanning tools.

Adds system checks for unmodified default settings, limits account-recovery emails, applies API v3 token auth settings, and scopes endpoint choices on import forms.

└──▷ GET THIS VERSION

```
$ git clone --branch 3.4.100 https://github.com/DefectDojo/django-DefectDojo.git
# already have the repo? check out this version:
$ git checkout 3.4.100
```

- Adds system checks for settings left at their shipped values in `settings.dist.py` / `local_settings.py`.
- Adds `locations` as a prefetch option on the findings and products API endpoints.
- Limits account-recovery emails per address (new `settings.dist.py` / `local_settings.py` setting).
- Applies the token endpoint's account rules to HTTP Basic auth on the API (new `settings.dist.py` / `local_settings.py` setting).
- Requires a login for the report cover page (new `settings.dist.py` / `local_settings.py` setting).

+7 more

COVERAGE

UPDATED integration SonarQube REST API import — unique_id_from_tool

All of defectdojo's coverage changes →

WAS THIS USEFUL? 👍 👎

◆ Endpoint Management (UEM/RMM)

Microsoft Intune ⓘ Week of October 5, 2026 NOTES ↗

Microsoft Intune is a cloud-based unified endpoint management solution built on Zero Trust. Empower IT to manage, secure and govern devices, apps and endpoints.

Existing baseline profiles must be manually updated to the new version; the Configure NetBIOS settings policy is excluded pending broader Windows support.

› Adds the Windows 11, version 26H2 security baseline in Intune, reflecting current Microsoft security recommendations with new settings, updated default values, and revised security guidance.

└──▷ **BREAKING ON UPGRADE**

! Existing security baseline profiles do not automatically update to the Windows 11, version 26H2 baseline — admins must manually create a new profile or update an existing profile to the latest version, and must review settings before migrating, especially if existing profiles include customizations.

COVERAGE

NEW compliance framework Windows 11, version 26H2 security baseline

All of Microsoft Intune's coverage changes →

WAS THIS USEFUL?  

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Auth0 ⓘ

Tenant ACL Dashboard UX (2026-10-12) BLOG ↗ RANK

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

> Auth0 shipped a guided dashboard UI for managing Tenant ACL rules, replacing manual policy syntax with a visual builder, sequencer, and one-click toggles.

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

01 **Guided Tenant ACL rule management UI** NEW 69

Under `Security > Access Control`, Auth0 added a Guided Rule Builder with step-by-step prompts for scope, entity, and action selection to construct Tenant ACL policies without syntax errors; a Visual Rule Sequencing canvas to drag-reorder policy priorities and inspect exact rule execution order; and Instant Status Toggles to enable or disable individual Tenant ACL rules in a single click while Auth0 handles the underlying API calls.

Create and activate a new Tenant ACL rule without writing raw policy syntax — useful for teams who need to restrict tenant access quickly during an incident.

📍 1. Navigate to Security › Access Control in the Auth0 dashboard. 2. Click New Rule to open the Guided Rule Builder. 3. Step through the scope, entity, and action prompts and save. 4. Use the Instant Status Toggle to enable the rule immediately.

Adjust enforcement order when a newly added deny rule must take precedence over existing allow rules — prevents unintended access due to rule sequencing.

📍 1. Navigate to Security › Access Control in the Auth0 dashboard. 2. Open the Visual Rule Sequencing canvas. 3. Drag the deny rule above the relevant allow rules to set correct execution order. 4. Save the updated sequence.

— Names UI paths and components but no API/config surface

Tenant ACL Dashboard UX (2026-10-12)

WAS THIS USEFUL? 👍 👎

AGENT

◆ MCP Servers & Tooling

Runlayer ⓘ 1.28.1233 (2026-10-12) BLOG ↗

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

New Terraform resources for plugins, schedules, teams, org settings, policy attributes, and AI Watch; new single-resource API endpoints across gateway, MDM, OAuth, and AI providers; pre-save policy impact checks.

- Adds `runlayer_plugin` resource to the Terraform provider to define plugins with their connectors and skills as code.
- Adds `runlayer_agent_schedule` resource to the Terraform provider to deploy an agent's cron schedules alongside the agent itself.
- Adds `runlayer_team` and `runlayer_organization` resources to the Terraform provider to manage teams and organization settings as code.
- Adds `runlayer_policy_attribute_setting` resource to the Terraform provider to define policy-referenceable attributes as code.
- Adds Terraform provider resources to manage AI Watch enforcement settings, per-user AI Watch policies, and per-tool scanner overrides on a connector.

+18 more

COVERAGE

NEW detection 'Claw' family personal agent

NEW integrations

Muse Code hook, Bring-your-own Claude subscription via LLM Gateway

NEW supported platform GPT-6.1 Sol model support for agents

All of Runlayer's coverage changes →

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ToolHive ⓘ v0.52.0 NOTES ↗ RANK

ToolHive is an open-source tool from Stacklok that runs and manages MCP (Model Context Protocol) servers in isolated containers, with security controls and easy client integration.

> ToolHive v0.52.0 adds TLS encryption for Redis session storage across the Kubernetes operator and proxy runner, changes `thv llm teardown` to purge cached tokens by default, and enables one-hour Claude Code prompt caching.

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

01 **TLS for Redis session storage** NEW 80

Adds `spec.sessionStorage.tls` (including `caCertSecretRef` for a private CA) to the `MCPServer`, `MCPRemoteProxy`, and `VirtualMCPServer` Kubernetes resources, and enables TLS-encrypted Redis session storage in the proxy runner, vMCP, and their rate limiters, using system roots or a private CA bundle. A startup WARN is logged when Redis runs without TLS.

Encrypt Redis session storage in the operator with a private CA — required for production deployments where the Redis TLS cert is signed by an internal CA.

yaml

```yaml
spec:
  sessionStorage:
    tls:
      caCertSecretRef:
        name: redis-ca-secret
        key: ca.crt
```

— Names the config key, CA field and affected resources with a runnable example v0.52.0

02 **thv llm teardown token purge behavior** BREAKING 80

`thv llm teardown` now purges cached OIDC tokens by default when tearing down all configured LLM tools, exiting non-zero if the secrets provider is unavailable or token operations fail. Pass `--keep-tokens` to preserve cached tokens; `--keep-tokens` and `--purge-tokens` are mutually exclusive.

Tear down all LLM tools in CI without invalidating cached tokens needed by a subsequent `thv llm setup` step.

```
$ thv llm teardown --keep-tokens
```

— States exact flags, default behavior change and failure condition v0.52.0

03 DCRCredentialStore interface change BREAKING

70

Out-of-tree Go implementations of `pkg/authserver/storage.DCRCredentialStore` no longer compile until they add the new `UpdateDCRCredentialsIfUnchanged(ctx context.Context, creds *DCRCredentials, expected *DCRCredentials) error` method.

— Exact method signature given but no usage example v0.52.0

04 Claude Code one-hour prompt caching default IMPROVED

65

`thv llm setup` now configures Claude Code for one-hour prompt caching by default; pass `--short-prompt-cache` to opt out.

Set up Claude Code without one-hour prompt caching when shorter cache windows are preferred for cost or freshness reasons.

```
$ thv llm setup --short-prompt-cache
```

— Clear flag and default change but no mechanism detail v0.52.0

└──▷ BREAKING ON UPGRADE

! `thv llm teardown` now purges cached OIDC tokens by default on a full teardown (all tools removed) and exits non-zero if the secrets provider is unavailable or token operations fail. Scripts that relied on tokens surviving teardown must add `--keep-tokens`. `--keep-tokens` and `--purge-tokens` are mutually exclusive.

! Out-of-tree Go implementations of `pkg/authserver/storage.DCRCredentialStore` no longer compile until they add the new `UpdateDCRCredentialsIfUnchanged(ctx context.Context, creds *DCRCredentials, expected *DCRCredentials) error` method.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 13, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe