THE CYBER TOOLCHAIN

NO. 020

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 20, 2026 · EVERY DAY

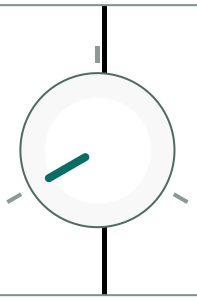
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 30 tools

OFFENSIVE SECURITY

vim ? help

◆ AI-Powered Offensive Security

Strix

Feedback ⓘ v1.1.0 NOTES ↗

Open-source AI penetration testing tool to find and fix your app's vulnerabilities.

Strix v1.1.0 adds 9 new security skills, SARIF 2.1.0 output, budget controls, and multi-target list scanning.

└──▷ GET THIS VERSION ▾

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

└──▷ TRY IT

Scan a list of targets (e.g., all apps in scope for a bug-bounty engagement) from a file rather than passing each on the command line.

```
$ strix --target-list ./targets.txt
```

Adds weak-password detection and LLM prompt-injection security skills.

Adds GCP and Auth0 reconnaissance skills.

Adds five new security skills covering OAuth, AWS, prototype pollution, deserialization, and Django.

Adds skill-directory registration for extending the skill set.

Emits SARIF 2.1.0 output for GitHub code-scanning and ASPM ingestion.

+10 more

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

bbot

Feedback ⓘ v3.0.0 NOTES ↗

The recursive internet scanner for hackers.

BBOT 3.0 brings new HTTP/DNS engines, unified findings, 10+ new modules, and pydantic-validated config.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.0.0
https://github.com/blacklanternsecurity/bbot.git
# already have the repo? check out this version:
$ git checkout v3.0.0
```

└──▷ TRY IT

Regenerate a stale BBOT config file after upgrading to 3.0 to pick up new default keys.

```
$ bbot --reset-config
```

New in-process, rate-limited HTTP engine (blasthttp) replaces the external httpx subprocess for all web traffic.

New DNS resolution engine (BlastDNS) with bounded LRU caches and noisy PTR flood filtering.

Pydantic-validated config and preset system catches typos and type errors before a scan starts.

Scope reworked: `-t/--targets` defines scope and `-s/--seeds` drives the scan; ASNs now usable as scan targets.

Findings and Vulnerabilities unified into a single event type carrying Severity and Confidence.

+14 more

└──▷ BREAKING ON UPGRADE

- ! `-s` now means `--seeds` (not `--silent`); scripts passing `-s` for quiet output will instead add a seed with no error. Silent output moved to `-S` / `--silent`.
- ! `--whitelist` was removed; use `-t/--targets` to define scope.
- ! `--allow-deadly` was removed.
- ! The `noisy` flag was renamed to `loud` (and other flags were similarly renamed).

- ! Module options are now pydantic-validated; the legacy per-module options dict hard-fails on upgrade.
- ! Some event types, event attributes, and config keys were renamed or removed.
- ! Modules wappalyzer, smuggler, digitorus, sitedossier, passivetotal, and wpscan were removed.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Portkey

Feedback ⓘ snapshot-20260720 NOTES ↗

Portkey is an AI gateway and observability platform that helps developers manage, monitor, and optimize LLM API calls with routing, caching, and reliability features.

└──▷ USE IT

Gate every agentic request through Zscaler AI Guard so your existing Zscaler Detection Policies apply to LLM inputs and outputs.

json

```
In your Portkey config, add the Zscaler AI Guard guardrail with your
policyId:
{
  "guardrails": {
    "beforeRequestHook": {
      "provider": "zscaler",
      "policyId": "<your-zscaler-policy-id>",
      "timeout": 10000
    },
    "afterRequestHook": {
      "provider": "zscaler",
      "policyId": "<your-zscaler-policy-id>",
      "timeout": 10000
    }
  }
}
```

Route GCP Cloud Storage log writes from an AWS-hosted gateway through Workload Identity Federation to avoid storing long-lived GCP service account keys in AWS.

```
$ Set these environment variables on your gateway deployment:
GCP_WIF_AUDIENCE=<your-wif-audience>
GCP_WIF_SERVICE_ACCOUNT_EMAIL=<your-
sa>@<project>.iam.gserviceaccount.com
```

Adds Secret References so Portkey fetches credentials at runtime from AWS Secrets Manager, Azure Key Vault, or HashiCorp Vault via `secret_mappings`, keeping keys out of Portkey config entirely.

Enables GCP Workload Identity Federation log delivery from AWS-hosted gateways to Google Cloud Storage via `GCP_WIF_AUDIENCE` and `GCP_WIF_SERVICE_ACCOUNT_EMAIL`, eliminating long-lived GCP keys in AWS.

Adds `customHost` support on the Bedrock Guardrails plugin so guardrail evaluation hits private or regional Bedrock-compatible endpoints instead of default public URLs.

Adds `tool_choice`, `stream_options`, and streaming usage reporting to DeepSeek `deepseek-chat`; maps `reasoning_effort` to thinking mode and surfaces `reasoning_content` in streams for `deepseek-reasoner`.

Maps Portkey metadata to Vertex AI resource labels and adds enterprise search grounding via `enterpriseWebSearch` / `enterprise_web_search` with separate cost attribution.

+2 more

WAS THIS USEFUL? 👍 👎

noir

Feedback ⓘ v1.2.0 NOTES ↗

Hunt every Endpoint in your code, expose Shadow APIs, map the Attack Surface.

Noir v1.2.0 adds CLI app surface detection, 9 stdlib HTTP servers, schema-driven APIs, enterprise legacy stacks, and major perf gains.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.2.0 https://github.com/owasp-noir/noir.git
# already have the repo? check out this version:
$ git checkout v1.2.0
```

Detects CLI application attack surfaces (argv, flags, env vars) across 17 languages and 40+ libraries including Cobra, Click, clap, Commander, Thor, and picocli.

Adds built-in HTTP server analyzers for Go (`net/http`), Python (`http.server`), Node (`node:http`), Java (`com.sun.net.httpserver`), C# (`HttpListener`), Ruby (WEBrick), Crystal, Dart, and Zig.

Adds schema-generated API analyzers for Strapi, Directus, Payload CMS, Supabase/PostgREST, Hasura, and Appwrite — extracting routes and params from schema files rather than source routes.

Adds enterprise legacy stack support: CFML/ColdFusion, Classic ASP, ASP.NET WebForms, and CFML frameworks Taffy, ColdBox, Wheels, and FW/1.

Adds BEAM language support: Erlang (Cowboy, Elli) and Gleam (Wisp).

+9 more

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

DataDome

Feedback ⓘ Fastly Compute module 1.4.0 (2026-07-20)

2 RELEASES · 2026-07-20 BLOG ↗ LATEST STABLE

DataDome is a bot protection and fraud prevention platform that detects and blocks malicious traffic, scraping, and cyberattacks in real time.

DataDome Fastly Compute 1.4.0 adds MCP protocol support with HTTP message signature collection per RFC 9421.

└──▷ USE IT

Enable MCP traffic inspection by pointing DataDome at your MCP endpoint in the Fastly Compute service config.

```ini
datadome_mcp_endpoint = "/mcp"
```

Adds Model Context Protocol (MCP) support with new `datadome_mcp_endpoint` configuration parameter (defaults to `/mcp`) to specify the MCP endpoint.

Collects `Signature`, `Signature-Agent`, and `Signature-Input` headers to support HTTP message signatures as defined by RFC 9421.

Collects `Skyfire-Pay-ID` header from incoming requests.

▾ 1 more release in this issue · 2026-07-20

> JavaScript Tag 5.9.0 (2026-07-20) BLOG ↗ STABLE
>
> DataDome JS Tag 5.9.0 adds challengeType to event payloads and the displayDataDomeResponsePage function for style adaptation.
>
> Adds `challengeType` parameter to event payloads, exposing the type of challenge triggered in downstream event handling.
>
> Accepts `challengeType` in the `displayDataDomeResponsePage` function, enabling style adaptation based on the challenge type being displayed.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

iLEAPP

Feedback ⓘ v2026.2.0 NOTES ↗ CODE ↗

iOS Logs, Events, And Plist Parser

iLEAPP v2026.2.0 adds Biome, C2PA AI provenance, and United Airlines artifacts with full iOS 26 and Python 3.13/3.14 support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2026.2.0 https://github.com/abrignoni/iLEAPP.git
# already have the repo? check out this version:
$ git checkout v2026.2.0
```

└──▷ TRY IT

Parse an iOS 26 extraction to capture new Biome and C2PA artifacts in a full HTML/TSV report.

```
$ ileapp -t zip -i /path/to/ios26_extraction.zip -o /path/to/output/
```

Limit a parse run to only the new Biome and C2PA modules using a saved profile, to reduce processing time on large extractions.

```
$ ileapp -t fs -i /path/to/extraction/ -o /path/to/output/ -m
/path/to/biome_c2pa.ilprofile
```

Adds 14+ new Biome database artifacts parsing SQLite stores under Library/Biome/databases/ (app openings, cable plug events, backlight events, CarPlay activity, intelligence platform entities, and more).

Adds Biome Set.db artifacts for installed apps and other syncable streams under Library/Biome/sets/.

Adds Biome stream artifacts for harvested mail/messages, reads, screen time, and keyboard data.

Adds C2PA Content Provenance artifact parsing AI-generated image provenance, including signer certificate, signing time, Media State column, and case-insensitive media path matching.

Adds new United Airlines app artifact module extracting travel data from the United Airlines iOS app.

+3 more

WAS THIS USEFUL?  

◆ Reverse Engineering

HyperDbg

Feedback ⓘ v0.22 2 RELEASES · 2026-07-05 → 2026-07-20

NOTES ↗ LATEST STABLE

State-of-the-art native debugging tools

HyperDbg v0.22 adds linked-list traversal commands, struct support in the script engine, and a pname filter for !pt.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.22 https://github.com/HyperDbg/HyperDbg.git
# already have the repo? check out this version:
$ git checkout v0.22
```

New `dl` command traverses linked lists in virtual address space directly from the debugger.

New `!dl` command traverses linked lists using physical addresses.

New SDK API exports for linked-list traversal over both virtual and physical memory.

Adds `pname` option to the `!pt` command to filter process trace by process name.

Script engine now supports structures, enabling struct-aware scripting and automation.

+1 more

▾ 1 more release in this issue · 2026-07-05 → 2026-07-20

v0.21 NOTES ↗ STABLE

HyperDbg v0.21 adds Intel PT parameter parsing, TID/PID helpers, hyperperf infrastructure, and an SDK load-state API.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.21 https://github.com/HyperDbg/HyperDbg.git
# already have the repo? check out this version:
$ git checkout v0.21
```

Adds user-mode and kernel-mode Intel PT parameter parser for fine-grained control of processor trace configuration.

> Adds thread (TID) and process (PID) helper functions for Intel PT, enabling per-thread and per-process tracing workflows.
>
> Introduces the hyperperf (Hardware Performance Counter) project structure as a new HyperDbg SDK module.
>
> Adds hypertrace, hyperevade, and hyperperf DLL files to the SDK, expanding the available module surface.
>
> Exports a new SDK API for checking whether any HyperDbg module is currently loaded.
>
> +2 more

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

malcontent

Feedback ⓘ v1.25.0 NOTES ↗

#supply #chain #attack #detection

malcontent v1.25.0 adds guarddog v3 rules and refreshes third-party YARA rule sets for improved supply-chain threat detection.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.25.0 https://github.com/chainguard-dev/malcontent.git
# already have the repo? check out this version:
$ git checkout v1.25.0
```

└──▷ TRY IT

Scan a suspicious PyPI package directory using the newly included guarddog v3 rules to catch malicious supply-chain behaviors.

```
$ mal scan --third-party <path-to-package-dir>
```

Adds guarddog v3 detection rules, expanding coverage of malicious package behaviors.

Updates bundled third-party YARA rules (Avast, Elastic, FireEye, Mandiant, Nextron, ReversingLabs) to June 2026 snapshots.

Upgrades yara-x engine to v1.19.0, improving underlying rule evaluation performance.

WAS THIS USEFUL? 👍 👎

YARA-X Feedback ⓘ v1.19.0 NOTES ↗

A rewrite of YARA in Rust.

YARA-X v1.19.0 adds missing PE machine types and new compiler warnings for single-byte and duplicate patterns.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.19.0 https://github.com/VirusTotal/yara-x.git
# already have the repo? check out this version:
$ git checkout v1.19.0
```

Adds missing machine architecture types to the `pe` module, expanding PE header analysis coverage.

New compiler warning when a rule defines a single-byte pattern, helping authors catch low-specificity signatures early.

New compiler warning when a rule contains duplicate patterns, surfacing redundant definitions at compile time.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Vectra AI Feedback ⓘ v9.13 - Observe (2026-07-21) BLOG ↗

Vectra AI is a cybersecurity platform that uses AI-driven threat detection and response to identify and prioritize attacks across hybrid cloud, network, and identity environments.

Adds behavioral signatures for HTTP-based attacker tooling, including malicious file downloads/uploads and certificate-related abuse via trusted Windows utilities and offensive frameworks.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Sophos Central

Feedback ⓘ snapshot-20260720 NOTES ↗

Sophos Central is a cloud-based management platform that unifies administration of Sophos security products, including endpoint, network, email, and cloud protection.

Sophos Central adds Network Map, Live Discover for AP6, DNS Protection APIs, Protected Browser extension, and a wave of new alerts and email capabilities.

└──▷ HOW TO FIND IT

View the Network Map to understand firewall-to-switch-to-AP6 topology and spot unexpected device relationships at a glance.

📍 In Sophos Central, go to My Products › Network › Network Map.

New Network Map provides a visual topology of all Sophos firewalls, switches, and AP6 access points with search and relationship views.

Live Discover queries now support AP6 Series access points, enabling on-demand SQL queries against wireless data in the Sophos data lake.

New Protected Browser extension lets admins apply security policies and monitor user activity in Chrome and Edge without requiring migration to the full Protected Browser.

Web control adds filtering for Generative AI sites, more granular per-site controls, and simplified policy management.

New alerts for Account Health Check (AHC) issues now surface unaddressed AHC findings directly in the alerts feed.

+29 more

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ

Extended inactivity and max session timeout to 12 hours (2026-06-30) BLOG ↗

Trend Vision One is a unified cybersecurity platform that integrates extended detection and response (XDR) across endpoints, networks, email, and cloud environments.

Trend Vision One extends max inactivity and session timeout limits from 8 hours to 12 hours.

└─▷ HOW TO FIND IT

Extend session limits for SOC analysts running multi-hour threat investigations without re-authentication prompts.

> 📍 In the console, go to Administration → Console Settings, open the Timeout section, and set Max session timeout and Inactivity timeout to up to 12 hours.

Raises the configurable ceiling for Max session timeout from 8 hours to 12 hours, allowing administrators to keep analysts signed in longer during extended investigations.

Raises the configurable ceiling for Inactivity timeout from 8 hours to 12 hours (capped at the Max session timeout value), reducing re-authentication interruptions during complex tasks.

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Laurel

Feedback ⓘ v0.8.0 NOTES ↗

Transform Linux Audit logs for SIEM usage

Laurel v0.8.0 adds SHA-256 exec hashing, UNIX socket logging, env-var prefix matching, and unified process spawning tracking.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.8.0 https://github.com/threathunters-
io/laurel.git
# already have the repo? check out this version:
$ git checkout v0.8.0
```

New `SYSCALL.SPAWNED_BY` field unifies process-lineage tracking across both single-exec and fork/exec scenarios.

Supports prefix matching for environment variables, enabling selective capture of env vars by prefix.

Adds logging output to a UNIX domain socket for integration with local log consumers.

Supports calculating SHA-256 hashes for executed binaries, enabling integrity checking of process execution.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 7.260715.0 7 RELEASES · 2026-06-24 → 2026-07-15

NOTES ↗ LATEST STABLE

Open Cyber Threat Intelligence Platform

OpenCTI 7.260715.0 adds connector log TTL and ingestion history in Redis, plus authentication validity policy improvements.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 7.260715.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260715.0
```

Adds connector logs TTL and ingestion history storage in Redis for improved connector observability.

Improves the authentication validity policy update workflow.

▾ 6 more releases in this issue · 2026-06-24 → 2026-07-15

7.260710.0

NOTES ↗ STABLE

OpenCTI 7.260710.0 adds admin-forced password resets, Custom Views date ranges & attribute widgets, OAuth SMTP dynamic tokens, and security coverage custom views.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 7.260710.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260710.0
```

└──▷ HOW TO FIND IT

Narrow a Custom View on an entity page to a specific investigation window without changing global filters.

📍 In the console, open any entity page, navigate to a Custom View tab, and use the new date range picker in the view configuration to set a start and end date.

Force a user to reset their password at next login for account hygiene or after a suspected compromise.

> 📍 In the console, go to Settings › Security › Users, select the target user, and use the new `Force password change` action.

Enables admins to force a password change for users via a new authentication capability.

Adds date range configuration directly within Custom Views on entity pages.

Adds an attribute widget for Custom Views.

Extends Custom Views to the security coverage section.

Supports dynamic token (OAuth) for SMTP email configuration.

+3 more

`7.260309.0-lts.6` `NOTES ↗` `STABLE`

OpenCTI 7.260309.0-lts.6 improves automatic enrichment on element creation/update and enhances JSON mapper regex extraction.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch 7.260309.0-lts.6 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260309.0-lts.6
```

Improves automatic enrichment triggered on creation or update of elements.

Enhances JSON mapper extraction with improved regexp support.

`7.260706.0` `NOTES ↗` `STABLE`

OpenCTI 7.260706.0 ships a native MCP server, STIX security-coverage export, custom newsfeed views, and richer telemetry.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch 7.260706.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260706.0
```

└──▷ **HOW TO FIND IT**

Find your MCP server connection details to integrate OpenCTI with AI assistants via XTM One.

> 📍 In the OpenCTI console, go to your User Profile and locate the XTM One MCP server section to retrieve the Streamable HTTP endpoint URL and credentials.

Adds a native Model Context Protocol (MCP) server with Streamable HTTP transport, exposed through XTM One and surfaced in the user profile.

Enables export of security coverage data to STIX format.

Supports custom views as a resource type in the newsfeed.

Adds `TELEMETRY_TAGS` deployment dimensions and tracks active connectors by catalog identity.

Extends telemetry to cover full AI usage and product adoption metrics, including workflow publish counts.

+5 more

7.260701.0 NOTES ↗ STABLE

OpenCTI 7.260701.0 adds XTM Hub custom view deployment, stream URI deny lists, ES health monitoring, and workflow transition comments.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 7.260701.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260701.0
```

Enables deploying custom views received from XTM Hub directly into OpenCTI.

Adds a configurable deny list of URIs to block during stream ingestion.

Adds an Elasticsearch engine health monitor to surface ES issues proactively.

Supports adding comments on workflow transitions and viewing them in the UI, with notifications sent to assignees and participants.

7.260626.0 NOTES ↗ STABLE

OpenCTI 7.260626.0 adds draft filtering in dashboards, a new text editor, and tightened AI playbook agent controls.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 7.260626.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260626.0
```

└──▷ **HOW TO FIND IT**

Track draft workspace activity across your team by adding a dashboard widget filtered to draft entities.

> 📍 In the console, go to Dashboards › (your dashboard) › Add Widget, then in the entity filter select `Draft Workspace` as the entity type to scope the visualization to draft data.

Limit which identities an AI playbook component can impersonate when running automated enrichment, reducing privilege-escalation risk.

> 📍 In the console, go to Automation › Playbooks › (your playbook) › AI Agent component › `Run as` field, and select from the restricted list of the current user or available service accounts.

Adds drafts as filterable entities in dashboard widgets, enabling visualizations and metrics over draft workspaces.

Restricts the `AI agent run as` selection in playbooks to the current user and service accounts only, tightening execution scope.

Replaces the legacy text editor with a new editor solution across the platform.

Enables entity lookup with on-the-fly creation in form intakes for users who have only create/update-in-draft permissions, without triggering a full creation process.

`7.260624.0` `NOTES ↗` `STABLE`

OpenCTI 7.260624.0 adds bulk playbook enrollment, has/not-changed stream filters, CTEM shortcut, and trust-proxy IP whitelisting config.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch 7.260624.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260624.0
```

└──▷ **HOW TO FIND IT**

Use the has/not-changed operator in a playbook filter to trigger actions only when a specific attribute on an entity has changed since the last event.

📍 In the Playbook or Stream filter editor, add a filter condition, select the desired attribute, and choose the `has changed` or `not changed` operator from the operator dropdown.

Bulk-enroll a large set of existing entities into a playbook using smart filtering instead of enrolling one at a time.

📍 In the Playbooks UI, open a playbook, click `Bulk enrollment`, apply the desired entity filters, and confirm to enroll all matching entities at once.

Adds trust proxy configuration for IP whitelisting, enabling accurate client IP resolution when OpenCTI sits behind a reverse proxy.

Adds a CTEM Command Center shortcut to the top bar when XTM One is connected.

Adds a feature flag to revert to the old rich text editor alongside the new @filigran/rich-text-editor component.

Adds dynamic user resolution in the playbook `send to notifier` component, allowing recipients to be resolved at runtime.

Adds has/not-changed attribute filtering operators for playbook and stream filtering components.

`+3 more`

WAS THIS USEFUL? 👍 👎

Yeti

Feedback ⓘ 2.5.1 NOTES ↗

Your Everyday Threat Intelligence

Yeti 2.5.1 adds structured JSON logging, agentic functionalities, and a new session-retrieval endpoint.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 2.5.1 https://github.com/yeti-platform/yeti.git
# already have the repo? check out this version:
$ git checkout 2.5.1
```

Adds structured JSON logging support for the console output.

Introduces agentic functionalities to Yeti.

New API endpoint to retrieve a session by user and session ID.

WAS THIS USEFUL?  

Watcher

Feedback ⓘ v3.5.0 NOTES ↗

Watcher - Open Source AI-powered Cyber Threat Intelligence & Hunting Platform. Developed with Django & React JS.

Watcher v3.5.0 adds a Timeline audit trail, UDRP approval workflow, a reworked ThreatsWatcher dashboard, and auto-fit table pagination.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.5.0 https://github.com/thalesgroup-cert/Watcher.git
# already have the repo? check out this version:
$ git checkout v3.5.0
```

└──▷ HOW TO FIND IT

Review and approve a queued UDRP domain transfer so no won case promotes a domain without human sign-off.

📍 In the Watcher UI, navigate to the PendingActions list, locate the entry with status `pending`, review the payload, then click Approve (or Reject) to trigger or cancel the domain transfer.

Adds a full Timeline audit trail module that records every create, update, and delete action across all Watcher modules, with per-object history panels, field-level diffs, and actor avatars visible in every table row.

New `/api/timeline/events/` endpoint lets practitioners query the full history of any tracked object by content type and object ID.

Introduces a PendingActions approval workflow for UDRP case transfers: when a case is marked `won`, a PendingAction is queued for admin review instead of automatically calling transfer_to_legitimate_domains().

Adds API endpoints for listing, approving, and rejecting pending actions, secured by existing permission guards.

New `avatar_color` field on `UserProfile` automatically assigns a deterministic color to each user on first login, surfaced as color-coded initials throughout the UI.

+2 more

└──▷ BREAKING ON UPGRADE

! The new `timeline` Django app requires a database migration (`python manage.py migrate`) and must be present in `INSTALLED_APPS` with its URLs registered — pull

the updated `settings.py` and `urls.py` or add them manually before starting the service.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Cortex

Feedback ⓘ 4.1.0 NOTES ↗

Cortex: a Powerful Observable Analysis and Active Response Engine

Cortex 4.1.0 adds audit logging for user actions in analyzer/responder catalogs.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 4.1.0 https://github.com/TheHive-
Project/Cortex.git
# already have the repo? check out this version:
$ git checkout 4.1.0
```

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

OpenCVE

Feedback ⓘ v3.0.0 NOTES ↗

Vulnerability Intelligence Platform

OpenCVE 3.0.0 ships Automations — trigger-driven alert and report workflows with conditions, actions, and webhooks.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.0.0 https://github.com/opencve/opencve.git
# already have the repo? check out this version:
$ git checkout v3.0.0
```

└──▷ USE IT

Disable public self-registration on a self-hosted instance to lock down new account creation.

```yaml
ENABLE_REGISTER: false
```

New Automations engine lets teams define CVE workflows with triggers, conditions, and actions (notifications, assignments, status changes, webhooks).

Alert automations react to matching CVEs and execute configurable actions including notifications, assignments, status changes, or webhooks.

Report automations deliver daily or weekly CVE digests with configurable conditions and timezone support.

Default report and alert automations are created automatically during project onboarding.

Reports now include CVE and CVSS summaries plus weekly AI-generated summarization.

+14 more

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta

Feedback ⓘ snapshot-20260720 NOTES ↗

Immuta is a data security platform that automates access control, privacy policy enforcement, and compliance monitoring across cloud data environments.

Immuta adds managed groups for agentic policies, Redshift viewless discovery, Delta Sharing enforcement, and OAuth 2.0 for Alation.

└──▷ HOW TO FIND IT

Restrict AI-sensitive data sources so only human (non-agent) accounts can subscribe, using the new built-in All Agents group.

> 📍 In the Immuta UI, go to Policies › Subscription Policies › New Policy. Under `Who is denied`, select the built-in group `All Agents`. Apply the policy to data sources tagged as AI-restricted.

Switch Alation integration to OAuth 2.0 to eliminate 90-day refresh-token rotation for the Immuta–Alation connection.

> 📍 In the Immuta UI, go to App Settings › Integrations › Alation, select `OAuth 2.0` as the authentication method, and supply your OAuth 2.0 client credentials.

Enable automated sensitive data discovery on a Redshift viewless data source to auto-tag PII columns.

> 📍 In the Immuta UI, open your Amazon Redshift viewless data source, navigate to the Sensitive Data Discovery settings, and enable automated discovery. Immuta will scan columns using regex, dictionary, and column-name-regex identifiers and apply matching tags automatically.

Adds three built-in system-managed groups (All Agents, All Users, All Identities) usable in subscription and data policies to enforce system-wide access rules without manual group management.

Supports row filtering and column masking on Unity Catalog managed, external, and foreign tables in Databricks Delta Sharing shares.

Adds OAuth 2.0 authentication option for Immuta's Alation integration, eliminating recurring API key rotation requirements.

Extends Databricks Unity Catalog tag ingestion to foreign catalogs (Lakehouse Federation) by falling back to system.information_schema tag views.

Embeds Immuta documentation and an AI-powered chatbot directly in the Immuta UI, accessible from the top menu Documentation icon.

+1 more

└─▷ **BREAKING ON UPGRADE**

- ! Snowflake policy role names will be renamed as part of the back-end architecture update rolling out beginning July 23, 2026; existing policy roles will be recreated under new names (user role names remain unchanged).
- ! Integration error banners for Snowflake integration validation failures will no longer be displayed in the existing location; errors are now only available through the new connections experience.
- ! When a data source is disabled in Immuta, data owners' Immuta-managed grants in the data platform will be revoked (behavior change release: on by default 8/1, enabled for all 9/1).

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

authentik

Feedback ⓘ version/2026.5.5

2 RELEASES · 2026-07-07 → 2026-07-15 NOTES ↗ LATEST STABLE

The authentication glue you need.

authentik 2026.5.5 adds number and boolean support in AKQL queries on JSON fields and licensing usage totals in the admin UI.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch version/2026.5.5
https://github.com/goauthentik/authentik.git
# already have the repo? check out this version:
$ git checkout version/2026.5.5
```

Adds number and boolean type support in AKQL queries on JSON fields via the API search interface.

Adds usage totals display to the licensing section in the admin web UI.

▾ 1 more release in this issue · 2026-07-07 → 2026-07-15

> version/2026.2.5 NOTES ↗ STABLE
>
> authentik 2026.2.5 adds SCIM discovery support and in-process per-IP rate throttling.
>
> └──▷ GET THIS VERSION ▾
>
> ```
> $ git clone --branch version/2026.2.5
> https://github.com/goauthentik/authentik.git
> # already have the repo? check out this version:
> $ git checkout version/2026.2.5
> ```
>
> Adds discovery support to the SCIM provider, enabling automatic capability negotiation with SCIM endpoints.
>
> Introduces in-process per-IP rate throttling to limit abusive request rates without an external component.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ RegScale 6.32.2.0 (2026-07-20) BLOG ↗

RegScale is a compliance operations platform that automates and continuously monitors regulatory compliance, security controls, and risk management for organizations.

RegScale 6.32.2.0 expands AI provider support to Azure OpenAI, AWS Bedrock, and OpenAI-compatible endpoints.

Supports Azure OpenAI, AWS Bedrock, and OpenAI-compatible endpoints as AI providers, broadening deployment options for enterprise environments.

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v3.19.2

4 RELEASES · 2026-06-22 → 2026-07-10 NOTES ↗ LATEST STABLE

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 150+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

CISO Assistant v3.19.2 adds asset sync for ServiceNow, reworked risk acceptance workflow, expanded documents management, and Slovak localization.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.19.2 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v3.19.2
```

Supports attaching evidences to data breach records.

Adds extra status options for findings tracking and exceptions.

Adds an autologin page capability.

Adds Slovak (sk) localization.

Adds more customization options for GDPR Processings.

+8 more

▾ 3 more releases in this issue · 2026-06-22 → 2026-07-10

v3.19.1 NOTES ↗ STABLE

CISO Assistant v3.19.1 adds applied control progress as a selectable column and multiple import improvements.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.19.1 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v3.19.1
```

Adds applied control progress as a selectable column in list views, surfacing remediation status at a glance.

Delivers multiple import improvements for loading controls and frameworks.

v3.19.0 NOTES ↗ STABLE

CISO Assistant v3.19.0 adds SCIM provisioning, offline AI embedding, managed portals, and admin MFA reset

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.19.0 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v3.19.0
```

Adds SCIM provisioning and IdP groups mapping for automated user lifecycle management.

Pre-bakes embedding and reranker models into the image for fully offline AI assistant operation (expect larger image size).

Introduces managed portals feature.

Enables admins to reset MFA for other users.

Adds typical evidence and annotation fields to audit/framework Excel exports.

+6 more

v3.18.3 NOTES ↗ STABLE

CISO Assistant v3.18.3 adds evidence visibility within domains and centralized domain field rendering in forms.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.18.3 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v3.18.3
```

Displays evidence items within their domain view, making it easier to review evidence by organizational scope.

Centralizes folder/domain field rendering across model forms for a consistent UI experience.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Windsurf

Feedback ⓘ snapshot-20260720 NOTES ↗

Windsurf is an AI-powered code editor that combines an intelligent coding assistant with agentic capabilities to help developers write, edit, and debug code faster.

Devin Desktop v3.5.17 adds inline network policy controls, Fast Context in Local sessions, and instant worktree-backed sessions.

└──▷ **USE IT**

Restore the status bar in the Agent Command Center after the default changes to hidden.

```json
In user settings (JSON), add:
{
  "workbench.statusBar.visible": true
}
```

Adds type-specific icons to @`-mention` pills in Cascade for clearer context at a glance.

New worktree-backed sessions open instantly and remain interactive while the worktree is created in the background.

Enables configuring a session's network policy and granting or denying network access requests inline from chat, without switching to the web app.

Adds Fast Context support in Devin Local sessions.

The Hooks tab in Devin Local Customizations now lists configured hooks with their source and trigger events.

+5 more

└──▷ **BREAKING ON UPGRADE**

! The status bar is now hidden by default in the Agent Command Center; set `"workbench.statusBar.visible": true` in user settings to restore it.

WAS THIS USEFUL? 👍 👎

Cline

Feedback ⓘ cli-v3.0.45 2 RELEASES · 2026-07-19 NOTES ↗

CODE ↗ LATEST STABLE

Autonomous coding agent as an SDK, IDE extension, or CLI assistant.

Cline CLI v3.0.45 adds Kimi K3 model, halves install size, and surfaces version numbers in hub/doctor status output.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch cli-v3.0.45 https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout cli-v3.0.45
```

Adds Kimi K3 as an available ClinePass model.

Reduces `npm i -g cline` install footprint from ~640 MB to ~285 MB by making Claude Code and Codex provider packages optional, loaded on demand.

Hub status output now includes CLI and core version numbers.

Supports automatic OAuth credential refresh and run retry, enabling unattended runs that survive token expiry.

▾ 1 more release in this issue · 2026-07-19

sdk/sdk/v0.0.65 NOTES ↗ STABLE

Cline SDK v0.0.65 adds Kimi K3 model support, on-demand provider loading, and new @cline/ui chat components.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch sdk/sdk/v0.0.65
https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout sdk/sdk/v0.0.65
```

Makes Claude Code and Codex provider SDKs optional peer dependencies loaded on demand, dramatically reducing install size.

Adds Kimi K3 to the bundled ClinePass model fallback list.

Enables automatic one-time retry after refreshing expired OAuth credentials on runs.

> Introduces new `@cline/ui` agent chat components with Storybook and npm packaging.
>
> Updates the bundled model catalog.
>
> +1 more

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ v2.1.215 NOTES ↗ CODE ↗

Claude Code is an agentic coding tool that lives in your terminal, understands your codebase, and helps you code faster by executing routine tasks, explaining complex code, and handling git workflows - all through natural language commands.

Claude Code v2.1.215 makes `/verify` and `/code-review` user-invoked only — no more automatic runs.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.1.215 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.215
```

└──▷ TRY IT

Manually trigger a code review after making changes, instead of waiting for Claude to decide to run it automatically.

```
$ /code-review
```

Run verification on demand after a refactor to confirm correctness on your own schedule.

```
$ /verify
```

The `/verify` and `/code-review` slash commands now run only on explicit user invocation, giving practitioners full control over when those checks execute.

└──▷ BREAKING ON UPGRADE

! Claude no longer automatically runs the `/verify` and `/code-review` skills; any workflow that relied on Claude triggering these autonomously must now invoke `/verify` or `/code-review` explicitly.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI

Feedback ⓘ rust-v0.145.0-alpha.24 NOTES ↗

CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI gains audio input support, inline visualization links in TUI, and remembered working directory for resumed sessions.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.145.0-alpha.24
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.24
```

Adds audio input variants (remote and local) to user input protocols, enabling audio to be forwarded to the Responses API.

Renders inline visualization links in the TUI — `::codex-inline-vis{file="..."}` directives in assistant Markdown become browser-openable links to generated artifacts.

Enables users to remember the working directory when resuming sessions (`tui_resume_cwd`).

Allows path-backed agents to be selected from the TUI agent picker.

Renders streamed Markdown incrementally in the TUI for smoother streaming output.

+2 more

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

ToolHive

Feedback ⓘ v0.40.1 9 RELEASES · 2026-06-30 → 2026-07-20

NOTES ↗ LATEST STABLE

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

ToolHive v0.40.1 adds per-session prompt serving in vMCP and audit logging for authorization denials on the proxy runner path.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.40.1 https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.40.1
```

Serves MCP prompts on a per-session basis in VirtualMCP (vMCP), enabling session-scoped prompt isolation.

Audits authorization denials on the proxy runner path, giving operators a log trail of rejected requests.

Respects DCR-negotiated token endpoint authentication methods, improving standards-compliant OAuth flows.

▾ 8 more releases in this issue · 2026-06-30 → 2026-07-20

v0.39.0 NOTES ↗ STABLE

ToolHive v0.39.0 adds AWS Bedrock compatibility flag to the LLM setup command.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.39.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.39.0
```

└──▷ TRY IT

Configure ToolHive to use an AWS Bedrock-compatible LLM endpoint instead of a default provider.

```
$ thv llm setup --bedrock
```

v0.38.0 NOTES ↗ STABLE

ToolHive v0.38.0 adds Codex App LLM setup support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.38.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.38.0
```

v0.37.0 NOTES ↗ STABLE

ToolHive v0.37.0 adds podTemplateSpec support for MCPRemoteProxy and exposes `--allow-docker-gateway` in the workload API.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.37.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.37.0
```

└──▷ USE IT

Customize the pod spec for an MCPRemoteProxy deployment — useful when you need node selectors, resource limits, or custom annotations on the proxy pod.

```yaml
spec:
  podTemplateSpec:
    metadata:
      annotations:
        my-org/team: security
    spec:
      nodeSelector:
        kubernetes.io/arch: amd64
      containers:
        - name: proxy
          resources:
            limits:
              memory: 256Mi
```

Adds `spec.podTemplateSpec` support to `MCPRemoteProxy`, enabling Kubernetes-native pod customization for remote MCP proxies.

Exposes `--allow-docker-gateway` option in the workload API, allowing workloads to reach the Docker host gateway network.

Adds runtime-stage environment variables to protocol Dockerfiles, enabling environment configuration at container runtime.

Adds MCP conformance CI validation for `thv run proxy`, improving protocol compliance guarantees.

Validates `aud` and `resource` claims in returned ID-JAG JWTs, strengthening authentication token verification.

v0.36.0 NOTES ↗ STABLE

ToolHive v0.36.0 adds OpenAI Codex CLI direct-mode support to the `thv llm setup` command.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.36.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.36.0
```

v0.35.0 NOTES ↗ STABLE

ToolHive v0.35.0 adds rate-limiting observability, HTTP OIDC issuer support, and new backend discovery APIs.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.35.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.35.0
```

Adds metrics and tracing for rate limiting, enabling observability into throttling behavior.

Supports HTTP upstream OIDC issuers for in-cluster development environments.

Adds `ListBackends` and `LookupBackend` methods to the vMCP core interface for programmatic backend discovery.

v0.34.0 NOTES ↗ STABLE

ToolHive v0.34.0 adds plugin lifecycle management, Claude Desktop gateway support, and Cedar nested-map claim authorization.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.34.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.34.0
```

Adds plugin install/list/info/uninstall commands plus a MaterializationAdapter for managing ToolHive plugins.

Adds Claude Desktop as a supported LLM gateway client.

Supports nested map claims in Cedar authorization for the `act` claim, enabling richer policy evaluation.

Allows custom HTTP headers on the OpenAI embedding client for flexible authentication and routing.

Filters the upstream auth chain via a configurable callback hook (`UpstreamFilter`) exposed through the `authserver.New` facade.

+2 more

└──▷ BREAKING ON UPGRADE

! The `--allow-docker-gateway` flag now actively grants host-level network access rather than merely unblocking it, which changes the effective network

permissions for containers using this flag.

! The Amp editor extension clients have been removed.

v0.33.0 NOTES ↗ STABLE

ToolHive v0.33.0 adds rate limiting on VirtualMCPServer and XAA auth strategy CRD support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.33.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.33.0
```

Supports configuring rate limits on VirtualMCPServer via the new rate-limit configuration.

Adds XAA authentication strategy support as a CRD for MCP server access control.

v0.32.0 NOTES ↗ STABLE

ToolHive v0.32.0 adds XAA cross-app auth, OBO strategy seam, upstream ID-token propagation, OpenAI-compatible embeddings, and Origin/DNS-rebind protection.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.32.0
https://github.com/stacklok/toolhive.git
# already have the repo? check out this version:
$ git checkout v0.32.0
```

└──▷ TRY IT

Allow a browser-based MCP client served from a custom origin to reach ToolHive's SSE transport after the wildcard CORS removal.

```
$ thv run --transport sse --allowed-origins https://my-web-
app.example.com some/mcp-server
```

Adds XAA (Cross-Application Access) outgoing auth strategy implementing the ID-JAG two-step token exchange (RFC 8693 → RFC 7523) for per-backend cross-application access tokens.

Surfaces upstream ID tokens through the auth middleware, consolidating access and ID token retrieval into a single bulk lookup.

Adds the vMCP OBO (on-behalf-of) strategy seam — a new optional `OBO` field on `BackendAuthStrategy` with strategy registration and override hook.

Enables the vMCP optimizer to use an OpenAI-compatible embedding client via optional `embeddingProvider` / `embeddingModel` config fields (defaults to TEI, so existing configs are unchanged).

Adds `insecureAllowHTTP` to `EmbeddedAuthServerConfig` to allow an `http://` issuer for in-cluster VirtualMCPServer deployments, with admission-time validation.

+2 more

└─▷ **BREAKING ON UPGRADE**

! ToolHive now validates the HTTP `Origin` header and removed `Access-Control-Allow-Origin: *` from the legacy SSE transport; browser clients on a non-`http://localhost` origin must pass the new `--allowed-origins` flag to `thv run` or `thv proxy`.

! `pkg/vmcp/server.New` dropped its `discovery.Manager` parameter (7 → 6 args) and the `pkg/vmcp/discovery` package was removed; out-of-tree Go code calling `server.New` must remove the `discovery.Manager` argument, ensure `Config.Aggregator` is non-nil, and pair `Config.AuthzMiddleware` with `Config.Authz` or receive `ErrInvalidConfig`.

WAS THIS USEFUL? 👍 👎

Giskard

Feedback ⓘ giskard-scan/v1.0.0b3 3 RELEASES · 2026-07-13

NOTES ↗ PRE-RELEASE

Open-Source Evaluation & Testing library for LLM Agents

Giskard Scan adds LIDAR and DeepTeam third-party integrations for broader AI vulnerability coverage.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch giskard-scan/v1.0.0b3 https://github.com/Giskard-AI/giskard-oss.git
# already have the repo? check out this version:
$ git checkout giskard-scan/v1.0.0b3
```

Integrates LIDAR scan as a third-party scan provider via `third_party_scan`.

Integrates DeepTeam as a third-party scan provider for additional red-teaming coverage.

▾ 2 more releases in this issue · 2026-07-13

giskard-checks/v1.0.2b5 NOTES ↗ PRE-RELEASE

Giskard Checks v1.0.2b5 adds LiDAR and DeepTeam third-party scan integrations to the vulnerability scanner.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch giskard-checks/v1.0.2b5
https://github.com/Giskard-AI/giskard-oss.git
# already have the repo? check out this version:
$ git checkout giskard-checks/v1.0.2b5
```

Integrates LiDAR scan into the third-party scan pipeline via `third_party_scan`.

Adds DeepTeam as a third-party scan integration for agent vulnerability scanning.

giskard-core/v1.0.1b5 NOTES ↗ PRE-RELEASE

Giskard adds crescendo/GCG/do-not-answer/sycophancy scan generators, HarmBench dataset, vulnerability scan API, singleturn/multiturn mode, and new checks.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch giskard-core/v1.0.1b5
https://github.com/Giskard-AI/giskard-oss.git
# already have the repo? check out this version:
$ git checkout giskard-core/v1.0.1b5
```

└──▷ **USE IT**

Run a full vulnerability scan against an agent with the new programmatic API, including multi-turn mode.

```python
import asyncio
from giskard.scan import vulnerability_scan

async def main():
    await vulnerability_scan(
        target=my_agent,
        description="A customer support chatbot for an e-commerce
platform.",
        languages=["en"],
    )

asyncio.run(main())
```

Export suite results to Hub-compatible format for upstream reporting or dashboards.

```python
result = await suite.run()
hub_payload = result.to_hub_format()
```

Adds crescendo-like multi-turn adversarial scenario generator for red-teaming LLM agents.

Adds GCG (Greedy Coordinate Gradient) injection scenario generator for adversarial suffix attacks.

Adds do-not-answer scan generator to probe refusal behavior.

Adds sycophantic scenario generator derived from a knowledge base.

Adds quality scan with knowledge base scenario generator for RAG and agent quality evaluation.

+13 more

WAS THIS USEFUL?  

NeMo Guardrails

Feedback ⓘ v0.23.0 NOTES ↗

NeMo Guardrails is an open-source toolkit for easily adding programmable guardrails to LLM-based conversational systems.

NeMo Guardrails v0.23.0 adds tool-call rails, HF classifier rails, PII masking, context bloat detection, and a `/v1` /checks endpoint.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.23.0 https://github.com/NVIDIA-NeMo/Guardrails.git
# already have the repo? check out this version:
$ git checkout v0.23.0
```

└──▷ **USE IT**

Apply Polygraf PII detection and masking to user input before it reaches the LLM.

yaml

```yaml
rails:
  input:
    flows:
      - polygraf pii detection
      - polygraf pii masking
```

Adds streaming and non-streaming tool calling to `IORails`, with local rails that validate model-emitted tool calls and application-returned tool results.

New `/v1/checks` endpoint on the OpenAI-compatible server runs input or output rails standalone, returning `passed`, `modified`, or `blocked` status without generating a model response.

Adds lightweight Hugging Face classifier rails for input, output, and retrieval, backed by local Transformers, vLLM, KServe, and FMS.

Adds context bloat detection rails to flag oversized, repetitive, low-entropy, or padded input and retrieved content.

Adds Polygraf integration for PII detection and masking across input, output, and retrieval rails.

+4 more

└──▷ **BREAKING ON UPGRADE**

! Pydantic `>=2.5,<3.0` is now required; environments pinned to Pydantic 1.x must upgrade before installing v0.23.0.

WAS THIS USEFUL? 👍 👎

AI-Infra-Guard

Feedback ⓘ v4.1.15 NOTES ↗

A full-stack AI Red Teaming platform securing AI ecosystems via OpenClaw Security Scan, Agent Scan, Skills Scan, MCP scan, AI Infra scan and LLM jailbreak evaluation.

AI-Infra-Guard v4.1.15 adds optional model token in scan APIs, 6 new llama.cpp CVE rules, and 3 new MCP threat detection rules.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v4.1.15 https://github.com/Tencent/AI-Infra-Guard.git
# already have the repo? check out this version:
$ git checkout v4.1.15
```

Supports omitting the model token in `mcp_scan` and `ai_infra_scan` API calls, falling back to the system default model.

Adds 6 new llama.cpp CVE detection rules to the llama-cpp vulnerability pack.

Adds 3 new MCP threat detection rules covering tool poisoning, credential exfiltration, and command injection.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 20, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue