THE CYBER TOOLCHAIN

NO. 043

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 12, 2026 · EVERY WEEKDAY

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 41 tools

OFFENSIVE SECURITY

◆ Wireless & RF

rfcat

1 RELEASE · 2026-08-11 NOTES ↗

RfCat - swiss-army knife of ISM band radio

> rfcat v3.0.0 brings the tool into the Python 3 era, adding PySide6 support and fixing widespread Python 3 compatibility issues, including a now-working spectrum analyzer button.

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

01 **Python 3 compatibility fixes across codebase** IMPROVED 35

Extends Python 3 compatibility across the codebase, resolving str/bytes and print-statement issues that previously broke functionality under Python 3.

— Describes fix scope but no concrete modules or commands v3.0.0

02 **PySide6 support with working spectrum analyzer** NEW 30

Adds PySide6 support to the GUI, with the spectrum analyzer button now working correctly under Python 3.

— Names PySide6 and spectrum analyzer but no usage detail v3.0.0

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

Maigret

1 RELEASE · 2026-08-11 | NOTES ↗ | CODE ↗

Collect a dossier on a person by username from 3000+ sites

> Maigret's v0.6.4 release adds QQ and Bilibili account lookup, tags over 2,200 sites by country and topic for targeted scanning, grows the site database, enriches recursive search via socid_extractor, and revamps the web interface with live scan progress and session history; documentation separately adds a new legacy XMind mindmap report format.

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

01 **XMind mindmap report format** NEW — 80

Adds a `-X` / `--xmind` flag to generate a legacy XML XMind mindmap report, one file per username, that includes a manifest for compatibility with modern XMind readers.

Generate an XMind mindmap report for a username to visualize OSINT findings in XMind-compatible tools.

```
$ maigret --xmind <username>
```

— Exact flag and runnable example command given. product docs

02 **Site tagging by country and topic** NEW — 70

Tags 2,210 sites in the database by country and topic, enabling more targeted scans with `--tags`.

— Names flag and exact site count, no example command. v0.6.4

THINNER COVERAGE BELOW

03 **Profile enrichment via socid_extractor URL mutations** IMPROVED — 55

Adds enrichment via `socid_extractor` URL mutations, extracting more linked accounts and IDs from profile pages during recursive search.

— Names mechanism but gives no command or flag. v0.6.4

04 **QQ and Bilibili account lookup support** NEW 50

Adds QQ account search via the `qq_id` identifier type, enabling lookup of QQ users alongside usernames, and adds support for Bilibili account lookup.

— Names identifier type but no usage example given. v0.6.4

05 **Web interface live scan and session improvements** IMPROVED 45

Updates the web interface with live scan progress, settings persistence across sessions, scan history, and improved navigation.

— Lists UI features but no navigation path given. v0.6.4

06 **32 new sites added to site database** NEW 40

Adds 32 new sites to the database (21 in one batch, 11 in another), plus CSDN for Chinese-language username lookups.

— Names counts and one site, no further detail. v0.6.4

07 **Per-site encoding configuration in site database** NEW 30

Adds per-site encoding configuration to the site database, allowing correct handling of non-ASCII identifiers.

— Describes purpose only, no config key named. v0.6.4

08 **AI report prompt and default model update** IMPROVED 30

Updates the AI report prompt and default model, improving the quality of AI-generated investigation summaries.

— No model name or config key given. v0.6.4

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

T3MP3ST ⓘ changes since 2026-07-13 CODE ↗

autonomous red teaming platform; multi-agent offensive-security meta-harness

T3MP3ST gains LiteLLM gateway, Hugging Face + NanoGPT providers, SOCKS5 proxy, six OSINT adapters, subdomain-takeover detection, and multi-language white-box ingest.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-13 https://github.com/elder-plinius/T3MP3ST.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-13
```

└──▷ TRY IT

Preview a safe upstream sync before committing any changes, to verify what would be updated without touching git state.

```
$ pwsh scripts/update.ps1 -DryRun
```

- Adds `SOCKS5` egress proxy support plus hardened local-model path, enabling outbound traffic routing through a SOCKS5 proxy.
- Adds `LiteLLM` as an AI gateway provider, broadening the range of LLM backends T3MP3ST can reach.
- Adds Hugging Face provider and selectable local Claude Code models as new LLM backends.
- Adds native NanoGPT provider support.
- Adds six keyless recon/OSINT adapters with argument templates to the arsenal.

+16 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Veracode

1 RELEASE · 2026-08-12 CHANGELOG ↗

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

> Veracode added business unit selection for DAST targets, real-time in-app scan status notifications, and improved SSL certificate checks for a broken access control weakness.

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

01 **CAA record check detection for CWE-862** IMPROVED 50

Adds improved detection for CAA record checks on SSL certificates for CWE-862 in both Dynamic Analysis and DAST.

– Names CWE and scan types but no user action changelog-20260812-fefad1a0

02 **Business unit selection for DAST targets** NEW 45

Adds business unit selection when creating a web application or API specification target, or on the DAST configuration page.

– Names UI entry points but no mechanism detail changelog-20260812-fefad1a0

03 **Real-time in-app scan notifications** NEW 30

Adds real-time in-app notifications when scans complete, fail, or are cancelled.

– States the trigger states, no config or path given changelog-20260812-fefad1a0

WAS THIS USEFUL? 👍 👎

Openlayer

1 RELEASE · seen 2026-08-12 NOTES ↗

Openlayer is a platform for monitoring and evaluating large language model performance, quality, and safety in production.

> Openlayer's update centers on new APIs for drilling into test-result rows and deleting projects, plus MCP custom functions and downloadable metric bundles, alongside a breaking rename of system tags to categories.

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

01 **System tags renamed to categories** BREAKING 55

Openlayer's system tags are renamed to `categories`; user-defined labels remain called `tags`.

– Clear rename with before/after but no migration steps snapshot-20260812

02 **Row-level test result API** NEW 45

A new API endpoint lets users pull individual rows from specific test results, enabling programmatic drill-down into evaluation data.

– Names an API endpoint but no exact path or method given snapshot-20260812

03 **Project deletion API** NEW 45

A new API endpoint allows deleting projects, and this is available via the TypeScript SDK.

– SDK access named but no exact endpoint path or method snapshot-20260812

04 **Downloadable custom metrics bundles** NEW 40

Custom metrics bundles are now viewable and downloadable directly from the UI.

– Clear UI location but no further detail on format or scope snapshot-20260812

05 **Custom functions in MCP integration** NEW 30

The MCP integration now supports custom functions that can be called from Openlayer.

– States capability only, no mechanism or example given snapshot-20260812

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! Openlayer system tags are renamed to 'categories'; user-defined labels remain 'tags'.

WAS THIS USEFUL? 👍 👎

StepSecurity ⓘ 1 RELEASE · 2026-08-11 BLOG ↗

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

> StepSecurity's Self-Hosted Runners page now merges GitHub-registered and agent-reported runners into one fleet inventory, requiring admins to approve a new GitHub App permission to see organization-level runners.

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

01 **Unified self-hosted runner fleet inventory** NEW 85

A new `Self-hosted runners: Read-only` organization permission on the StepSecurity GitHub App lets it list runners registered at the organization level via the GitHub API, not just those reporting through the Harden-Runner agent; admins must approve this pending request under `Organization Settings → GitHub Apps` for org-level runners to appear. The Self-Hosted Runners page then merges org-level, repo-level, and agent-reported runners into one inventory, marking agent-less runners as 'Not installed' in the Harden Runner column, and each entry shows a discovery source, a registration scope badge (Org or Repo with repository name), and a lifecycle classification (Persistent or Ephemeral). A per-runner detail panel adds agent version and install time for agent-reported runners, and level, OS, runner group, labels, and observed job counts for GitHub-registered runners, alongside fleet search by name or label, status filtering, last-sync timestamps, and an on-demand refresh.

Approve the new GitHub App permission so your full runner fleet — including organization-level runners never seen by the agent — appears in the StepSecurity inventory.

> 📍 In GitHub, go to Organization Settings → GitHub Apps, locate the StepSecurity app, and approve the pending 'Self-hosted runners: Read-only' permission request.

— Names permission, exact UI path, badges and panel fields, no CLI command.

Self-Hosted Runners: Full Fleet Visibili…

WAS THIS USEFUL? 👍 👎

Cloudsmith

3 RELEASES · 2026-08-12 BLOG ↗

Cloudsmith is a cloud-based software supply chain platform for securely managing and distributing software packages and artifacts.

> Cloudsmith took its policy-as-code system and vulnerability scanning to general availability, and made the Rego policy editor catch mistakes as you type.

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

01 **Continuous risk detection reaches general availability** NEW

70

Continuous risk detection is now generally available, providing automatic vulnerability and malicious-package checking for every package in a workspace with no scheduled or manual rescans required. It replaces Trivy with OSV.dev as the data source powering vulnerability results in the web app and API, unifying it with the source that has driven policy evaluation since early access.

— Names data source swap and scope but no config detail

Continuous risk detection is now general…

THINNER COVERAGE BELOW

02 **Policy management reaches general availability** NEW 55

Cloudsmith's policy-as-code system for enforcing rules across the software supply chain is now generally available across the web app, API, and Terraform provider.

— Names three surfaces but no mechanism beyond GA status

Policy management is now generally avail…

03 **Inline Rego error checking in policy editor** IMPROVED

55

The policy editor now highlights Rego syntax and logic errors inline and in a dedicated error panel as you write, catching mistakes before saving. Autocomplete in the editor also suggests schema fields as you type.

— Describes mechanism but no config or API specifics

Catch policy errors before you save (202…

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

ILSpy

1 RELEASE · 2026-08-12 NOTES ↗

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

> ILSpy v11.0 ports its GUI from WPF to Avalonia for native Linux and macOS support, drops legacy Visual Studio add-in support, and adds several ilspycmd flags aimed at agentic/LLM decompilation workflows.

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

01 ilspycmd flags for agentic decompilation workflows NEW

85

`ilspycmd` gains a `--member` flag for single-member decompilation to reduce token spend in agentic/LLM workflows, a `--dump-table` metadata dump mode, and fuzzy matching for `-t` type name resolution.

Decompile a single method from an assembly to minimize output and reduce token usage when calling from an AI agent.

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

Dump a specific metadata table from an assembly for low-level inspection without full decompilation.

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

— Three named CLI flags with example commands for two v11.0

02 Avalonia port replaces WPF frontend BREAKING

60

ILSpy is ported from WPF to Avalonia, enabling native GUI support on Linux and macOS in addition to Windows; the WPF frontend has been entirely removed.

— Clear before/after but no migration steps given v11.0

THINNER COVERAGE BELOW

03 VS2017/2019 add-in removed BREAKING

50

The Visual Studio 2017/2019 add-in has been removed; only VS2022–2026 add-ins are supported going forward.

— States the change but no upgrade path detailed v11.0

04 **Load assemblies from NuGet feed or running process** NEW

50

Adds 'File / Open from NuGet feed' to load assemblies directly from a NuGet feed, and 'File / Open from running process' available on Windows, Linux and macOS.

— Named UI menu paths but no config or command v11.0

05 **Bookmarks for decompiled code view** NEW

45

Adds 'Window / Bookmarks' — a full-featured bookmarks implementation for the decompiled C# view.

— Named menu path, no detail on bookmark mechanics v11.0

06 **Omnibar breadcrumb/search bar** NEW

45

Adds 'View / Options / Display / Show Omnibar' — a breadcrumb/search bar shown at the top of decompiled code.

— Named UI toggle path but behavior undescribed v11.0

└──▷ BREAKING ON UPGRADE

! The WPF frontend has been entirely removed; ILSpy now uses the Avalonia frontend exclusively.

! The VS2017/2019 add-in has been removed; only VS2022–2026 add-ins are supported going forward.

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

Datadog GuardDog

1 RELEASE · 2026-08-12 | NOTES ↗ | CODE ↗

GuardDog is a CLI tool to Identify malicious PyPI and npm packages

> GuardDog v3.2.0 extends its malicious-package detection to a new ecosystem with Rust crates support, adds a new npm supply-chain tampering check, and changes default behavior for npm dependency verification.

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

01 Rust crates scanning on crates.io NEW

85

Adds `guarddog crates scan <package>` and `guarddog crates verify <Cargo.lock>` to scan Rust crates hosted on crates.io using the same YARA + metadata risk engine used for other ecosystems.

Scan a Rust crate from crates.io for supply-chain risks, getting JSON output for downstream processing.

```
$ guarddog crates scan serde --output-format=json
```

Verify all production Rust dependencies in a project's lockfile, ignoring dev-only crates.

```
$ guarddog crates verify /tmp/repo/Cargo.lock
```

— Names exact new subcommands and their scope, with runnable examples. v3.2.0

02 devDependencies excluded by default in npm verify BREAKING

60

`guarddog npm verify` now excludes npm `devDependencies` by default, reducing noise in CI pipelines that only care about production dependencies; setups that previously relied on devDependencies being scanned will no longer flag them unless explicitly included.

— Names exact command and behavior change, with migration note. v3.2.0

THINNER COVERAGE BELOW

03 npm provenance-disappearance detection NEW

45

A new npm metadata check flags packages when provenance attestations disappear between releases, detecting a potential supply-chain tampering signal.

— Explains what it detects but no command shown. v3.2.0

└──▷ BREAKING ON UPGRADE

! `guarddog npm verify` now excludes `devDependencies` by default; setups that previously relied on devDependencies being scanned will no longer flag them unless explicitly included.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Sophos DNS Protection

1 RELEASE · 2026-08-12 CHANGELOG ↗

Sophos DNS Protection is a security service that blocks malicious websites and threats by filtering DNS queries at the network level.

> Sophos DNS Protection introduces a dedicated Generative AI web category for policy-level control over AI site access, and renames its business category group to reflect the change.

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

01 Generative AI web category for filtering policies NEW

75

Adds the Generative AI category to the Business and generative AI category group, letting filtering policies block or allow generative AI websites independently of other categories. Supports an allowlist workflow: block the Generative AI category in a policy, then create a domain list of approved AI URLs and allow that domain list in the same policy to permit only sanctioned AI services.

— Names category, group, and exact allowlist workflow steps. changelog-20260812-0727b8a4

THINNER COVERAGE BELOW

02 Category group renamed to Business and generative AI BREAKING

40

The Business-relevant site categories category group is renamed to Business and generative AI to reflect the addition of the Generative AI category; documentation, runbooks, or policy references to the old name need updating.

— Names old and new group names but no migration mechanism. changelog-20260812-0727b8a4

└──▷ BREAKING ON UPGRADE

! The Business-relevant site categories category group is renamed to Business and generative AI ; any documentation, runbooks, or references to the old name will need to be updated.

WAS THIS USEFUL? 👍 👎

Sophos Central Firewall

1 RELEASE · 2026-08-12 CHANGELOG ↗

Sophos Central Firewall is a cloud-managed network security platform providing centralized firewall protection and threat defense for organizations.

> Sophos Central Firewall added the ability to stage firmware upgrades on select firewalls before wider rollout.

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

01 **Staged firmware upgrade validation** NEW 35

Supports staging firewalls to validate firmware upgrades before rolling them out to the broader firewall fleet.

— Describes capability but no config keys or steps given. changelog-20260812-9164359c

WAS THIS USEFUL? 👍 👎

Fastly

2 RELEASES · 2026-08-11 BLOG ↗

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

> Fastly's Terraform provider gained management support for Audit Log Event Mappings, while the Historical Stats and Real-Time Analytics APIs picked up new Private Access Token challenge metrics for troubleshooting.

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

01 Audit Log Event Mapping resource in Terraform provider NEW

70

Fastly Terraform Provider 9.6.0 adds `fastly_audit_log_event_mapping` as a new Terraform resource and data source for managing Audit Log Event Mappings.

— Names exact resource/data source but no config example

Fastly Terraform Provider 9.6.0 (2026-08…

THINNER COVERAGE BELOW

02 PAT challenge metrics in Historical Stats and Real-Time Analytics

NEW

50

The Historical Stats API and Real-Time Analytics API now include two new Private Access Token challenge metrics: a count of issued PAT challenges and a count of successful PAT challenges, enabling troubleshooting of Private Access Token flows.

— Names the two metrics but no field names or endpoint paths

Add Private Access Token (PAT) challenge…

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

kubescape

1 RELEASE · 2026-08-12 NOTES ↗

Kubescape is an open-source Kubernetes security platform for your IDE, CI/CD pipelines, and clusters. It includes risk analysis, security, compliance, and misconfiguration scanning, saving Kubernetes users and administrators precious time, effort, and resources.

> Kubescape v4.0.12 adds a scan-cancellation API endpoint, new scan flags, expanded output formats, MCP server scanning tools, Terraform manifest scanning, cloud registry vulnerability adaptors, and a switch to RegoV1-only policy evaluation, alongside registry-auth, streaming, and host-sensor performance work.

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

01 **New scan command flags** NEW 80

The workload `scan` command gains `--api-version` to specify the API version of the scanned resource, `--label-selector` to filter collected resources during a scan, and `--only-fixable` to gate the scan image severity threshold on fix availability.

— Three flags named with purpose but no usage examples. v4.0.12

02 **Expanded output formats across scan, list, and config commands** NEW 75

Adds `markdown` output format to the scan printer alongside existing JSON, SARIF, HTML, CSV, etc.; adds `csv` and `yaml` output formats to the `list` command; adds `--output json` to the `version` command for machine-readable version info; adds configurable output formats to `config view`; adds a CSV output format with remediation path evidence columns to scan results; and extends image scans to support all output formats, aligning `patch` and `scan image` output options.

— Formats and commands named but no sample output shown. v4.0.12

03 **MCP server scanning tools** NEW 65

Adds a `scan_container_image` MCP server tool for on-demand container image vulnerability scanning, a headless framework security scanner MCP tool, and a local IaC scanner MCP tool, all exposed via `kubescape mcpserver`.

— Tools named but no invocation details given. v4.0.12

04 **API endpoint to cancel in-flight scans** NEW 60

Adds `DELETE /v1/scan` API endpoint to cancel an in-flight scan.

— Endpoint named but no request or response details. v4.0.12

THINNER COVERAGE BELOW

05 **Cloud registry vulnerability adaptors and AKS mapping** NEW — 55

Introduces a native GCP Artifact Registry vulnerability adaptor (Container Image Vulnerability, Phase 2) and a Container Image Vulnerability (CIV) adaptor for Azure Container Registry, plus adds Azure AKS support to `ListEntitiesForPolicies` resource mapping.

— Adaptors named but no configuration steps given. v4.0.12

06 **Offline tarball image scanning** NEW — 50

Supports scanning local offline tarball images with `kubescape scan image`.

— Command named but no flag or example shown. v4.0.12

07 **Registry authentication improvements for image scans** IMPROVED — 50

Adds an automated in-cluster private registry credential resolver for image scans, fallback support for OpenShift external registry routes, and support for scoped registry auth for image scans.

— Mechanisms named but no config keys or flags given. v4.0.12

08 **Terraform manifest scanning** NEW — 45

Adds scanning of `kubernetes_manifest` and typed `kubernetes_*` resources within Terraform files.

— Resource types named but no scan command usage shown. v4.0.12

09 **RegoV1-only policy evaluation** BREAKING — 45

Switches Rego evaluation to RegoV1, dropping the v0 compatibility shim; policies relying on v0 Rego syntax will break.

— Impact stated but no migration steps given. v4.0.12

10 **Resource streaming for large clusters** IMPROVED — 40

Adds resource streaming for large clusters (feat(streaming)) to reduce memory usage during scans.

— Benefit stated but no configuration or limits given. v4.0.12

11 **Host-sensor query optimizations** IMPROVED — 35

Optimizes host-sensor queries with paging, rate limits, and caching.

— Mechanisms listed but no numbers or config given. v4.0.12

12 **Current field values shown alongside failed paths** IMPROVED

35

Surfaces current field values alongside failed paths in scan output (HTML and printer).

— Output change described but no example shown. v4.0.12

13 **Concurrent deduplicated image scan pipeline** IMPROVED

30

Implements a concurrent and deduplicated image scan pipeline (KS-ADV-02) to improve throughput.

— Improvement stated but no metrics or config given. v4.0.12

14 **Configurable workload kinds in resource prioritization** IMPROVED

30

Supports configurable workload kinds and dynamic pod template spec detection in resource prioritization.

— Capability named but no config key or example given. v4.0.12

15 **Duration-based environment variable configuration** NEW

30

Adds `ParseDurationEnvVar` for duration-based environment variable configuration.

— Function named but no usage example given. v4.0.12

16 **Runtime profile drift detection and remediation** NEW

25

Implements Runtime Profile Drift Detection and Remediation (KS-ADV-01).

— Named only by internal code, no mechanism described. v4.0.12

17 **CEL findings gain remediation paths** IMPROVED

20

CEL findings now carry the same remediation paths as their Rego equivalents.

— Change stated with no mechanism or example. v4.0.12

18 **Cloud exceptions preferred over CRD exceptions** IMPROVED

20

Cloud exceptions are now preferred over CRD exceptions (feat(exceptions)).

— Precedence change stated without configuration detail. v4.0.12

19 **Unified scan stdin input handling** IMPROVED

15

Unifies scan stdin input handling for consistent behavior across scan commands.

— Change described generically without specifics. v4.0.12

└──▷ BREAKING ON UPGRADE

! Rego policies are now evaluated with RegoV1 only — the v0 compatibility shim has been removed. Policies relying on v0 Rego syntax will break.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Wiz ⓘ 1 RELEASE · 2026-07-27 BLOG ↗

Wiz is a cloud security platform that identifies and remediates risks across cloud infrastructure and applications.

> Wiz launches Atlas, an autonomous multi-agent AI system for vulnerability research that ranked #1 on the CyberGym benchmark and uncovered 200+ previously unknown vulnerabilities in major open-source projects, backed by an internal model-routing benchmark and a strict evaluation-driven development process.

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

01 **Atlas autonomous AI vulnerability researcher** NEW 70

Wiz introduces Atlas, a multi-agent AI system for AI SAST that maps attack surfaces via a code property graph (CPG), hunts vulnerabilities in parallel using independent agents, and validates each finding through an adversarial debate — one agent argues exploitability, one argues against, a third adjudicates — before confirming real security impact by dynamically executing a generated exploit in a dedicated execution environment. Atlas ranked #1 on CyberGym, the public benchmark for AI-driven vulnerability research, with a 90.9% success rate, and uncovered 200+ previously unknown vulnerabilities in heavily audited open-source projects including grpc, dnsmasq, Kubernetes, gVisor, the Linux kernel, and containerd.

— Explains mechanism and results but no user-facing config or command. 2026-07-27-44780eb3

THINNER COVERAGE BELOW

02 **Cyber Model Arena for AI model routing** NEW 50

Atlas uses an internal benchmark called Cyber Model Arena to continuously evaluate newly released models across threat modeling, hunting, validation, and proof-generation tasks, routing each pipeline stage to whichever model currently wins on that task.

— Describes internal routing mechanism with no external interface. 2026-07-27-44780eb3

03 **Evaluation-Driven Development for Atlas pipeline** NEW 40

Atlas development follows Evaluation-Driven Development (EDD), requiring strict quality metrics for every AI component and stage, including a dedicated validation eval that scores candidate ratings against known false-positive and true-positive ground-truth sets, before any pipeline change is accepted.

— Internal QA process description, nothing for readers to act on. 2026-07-27-44780eb3

WAS THIS USEFUL? 👍 👎

Plexicus

3 RELEASES · seen 2026-08-12 NOTES ↗

> Plexicus expanded SCM onboarding across three releases—multi-account connections, nested GitLab groups and broader roles, and zero-touch first-boot setup—while also upgrading CVSS scoring and hardening its AI remediation path.

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

01 **Zero-touch first-boot onboarding** IMPROVED 50

Frictionless first-boot onboarding now completes end to end with zero manual intervention, covering initial boot, sample-repo scan, and the in-product GitHub App flow.

– Names steps covered but no exact commands or endpoints.

Plexicus v1.2.36: Frictionless First-Boo…

02 **GitLab onboarding for nested groups and broader roles** IMPROVED 45

GitLab repository connection now supports nested groups during onboarding, and onboarding was extended to broader access roles, reducing friction for teams with non-admin members.

– Describes scope but no exact UI path or config.

Plexicus v1.2.31: Smoother GitLab Onboar…

03 **Multiple SCM accounts per provider** NEW 30

Plexicus now supports connecting multiple SCM accounts per provider, enabling flexible repository onboarding across different accounts.

– States capability but no mechanism or limits given.

Plexicus v1.2.25: Multiple SCM Accounts …

04 **NVD-backed CVSS scoring** IMPROVED 30

CVSS scoring is upgraded to use NVD-backed data, improving accuracy of vulnerability severity scores across findings.

– Names data source but no scoring mechanics or example.

Plexicus v1.2.25: Multiple SCM Accounts …

05 **Canonical naming for BYOAI models** IMPROVED 20

Unified canonical naming for BYOAI (Bring Your Own AI) models across the platform for consistent model identification.

– Names feature but no naming scheme or config surface.

Plexicus v1.2.36: Frictionless First-Boo…

06 **Hardened Codex Remedium remediation path** IMPROVED 20

The Codex Remedium AI remediation path was hardened for more reliable remediation runs.

— No detail on what was hardened or how. Plexicus v1.2.36: Frictionless First-Boo…

WAS THIS USEFUL? 👍 👎

Aikido Security

2 RELEASES · seen 2026-08-12 NOTES ↗

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

> Aikido's biggest addition this window is code-based DSPM via a new Data Exposure Audit that traces sensitive data through source code without connecting to production data stores, alongside agentic PR analysis (Deep Review), Android pentesting, AI-built custom reports, and a grab-bag of integration and platform updates.

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

01 **Data Exposure Audit (code-based DSPM)** NEW 80

New Data Exposure Audit feature uses LLM agents to trace sensitive data from sources (request bodies, uploads, DB reads, config, application logic) to sinks (logs, third-party APIs, LLM prompts, caches, queues, storage) without connecting to any production data store. It detects credentials in URLs, API over-exposure (endpoints returning excess fields, missing tenant ownership checks), unsafe logging of unredacted PII and secrets, over-broad third-party and AI data sharing, weak storage (unencrypted regulated data, base64/MD5/SHA-1 misused as encryption), and incomplete deletion flows leaving downstream records and caches intact. It runs against already-connected repositories with no new data source to onboard, and each audit consumes credits from the Aikido Wallet.

— Mechanism, detection list and billing model all specified

New - DSPM: Find data exposure risks wit…

02 **Free monthly Aikido Credits for Deep Review and Security Audit** NEW 60

Pro and Advanced plans now receive 100–200 free Aikido Credits per month for use with Deep Review and Security Audit.

— Concrete plan names and credit amounts, no redemption steps

#53 - Deep Review, Android Pentests, Cus…

THINNER COVERAGE BELOW

03 **AI-built Custom Reports** NEW 55

New Custom Reports capability lets users describe a report in plain language; the Aikido Agent builds it from findings and makes it shareable across the workspace.

— Describes input method and output but no exact UI path

#53 - Deep Review, Android Pentests, Cus…

04 **Deep Review agentic PR analysis** NEW 50

New Deep Review feature performs agentic PR analysis that goes beyond the diff, examining surrounding code and linked repositories to catch business logic mistakes and cross-service security issues.

— Explains scope and purpose but no usage path #53 - Deep Review, Android Pentests, Cus…

05 **Per-repository PR check bypass in Training mode** NEW 50

Adds per-repository bypass of PR checks in Training mode.

— Names the mode and scope, no config steps #53 - Deep Review, Android Pentests, Cus…

06 **API for AutoFix settings** NEW 45

New API surface to read and update AutoFix settings.

— Names an API surface but no endpoint or schema given

#53 - Deep Review, Android Pentests, Cus…

07 **IDE extensions for VSCode and Visual Studio** NEW 45

Adds new IDE extensions for VSCode and Visual Studio.

— Names both IDEs, installable but no further detail

#53 - Deep Review, Android Pentests, Cus…

08 **Kubernetes agent updates: FIPS variants and container label annotations**

IMPROVED 45

Kubernetes team annotations are now surfaced as container labels in Cloud, and Kubernetes agents now have FIPS variants.

— Two thin Kubernetes changes grouped, no config keys given

#53 - Deep Review, Android Pentests, Cus…

09 **Android pentesting coverage** NEW 40

Aikido Pentest now covers native and hybrid Android apps, including those with biometric login.

— States scope only, no configuration detail #53 - Deep Review, Android Pentests, Cus…

10 **Zen Firewall agents for PHP and Ruby** NEW 40

Adds Zen Firewall agents for PHP and Ruby.

— Names two new agent languages, no further mechanism

#53 - Deep Review, Android Pentests, Cus…

11 **Repository filter for trends-over-time in Reports** IMPROVED

40

Adds a repository filter for trends-over-time in Reports.

— Names the exact report view affected, no further detail

#53 - Deep Review, Android Pentests, Cus…

12 **Aikido Terraform provider integration** NEW

35

Adds an Aikido Terraform provider integration.

— Bare mention, no scope or setup detail #53 - Deep Review, Android Pentests, Cus…

13 **GitLab issue resolution in integration** IMPROVED

35

GitLab integration now supports resolving issues.

— One-line capability, no workflow detail #53 - Deep Review, Android Pentests, Cus…

WAS THIS USEFUL? 👍 👎

Sysdig Secure

1 RELEASE BLOG ↗

Sysdig Secure is a container security platform providing runtime monitoring, vulnerability scanning, and compliance enforcement for cloud-native applications.

> Sysdig Secure's KSPM Analyzer (v1.49.3) now automatically detects CNI plugins on Kubernetes nodes, enabling automated evaluation of a previously manual CIS benchmark check.

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

01 **Automatic CNI detection for CIS benchmark 5.3.1** IMPROVED

60

KSPM Analyzer v1.49.3 adds automatic detection of Container Network Interface (CNI) plugins in use on each Kubernetes node, enabling CIS Kubernetes Benchmark v2.0.0 recommendation 5.3.1 to be evaluated automatically instead of requiring manual attestation.

— Names exact benchmark and version but no UI/API path given.

Secure SaaS - August 12, 2026

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

NVIDIA OpenShell ⓘ 1 RELEASE · 2026-08-12 NOTES ↗

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

> OpenShell v0.0.104 focuses on developer experience and platform reach, adding a credential-leak warning for CLI environment flags, Windows build support, and a statically-linked supervisor option.

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

01 **Credential-leak warning on --env flag** NEW 50

The CLI now warns when values passed to `--env` look like credentials, helping catch accidental secret exposure at the command line.

— Names the flag but not detection logic. v0.0.104

02 **glibc-static supervisor variant** NEW 30

Adds a glibc-static supervisor libc variant for environments that require static linking.

— Names the variant but no usage instructions. v0.0.104

03 **Windows compilation support** NEW 25

Adds Windows compilation support for the gateway and CLI components.

— States the addition with no build details. v0.0.104

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Sophos Central ⓘ 1 RELEASE · 2026-08-12 CHANGELOG ↗

Sophos Central is a cloud-based unified security management platform for protecting endpoints, servers, networks, and email across organizations.

> Sophos Central adds an ML-derived Risk Score for identities in ITDR, a Generative AI category in DNS Protection, and finer-grained bulk licensing controls for partners.

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

01 **Bulk add-on selection in partner licensing** IMPROVED 35

The partner portal now supports explicit selection of add-ons and products when applying bulk license changes to multiple monthly-usage customers simultaneously.

— Describes capability but no exact menu or field names. changelog-20260812-54f33ac7

WAS THIS USEFUL? 👍 👎

Huntress

1 RELEASE · 2026-08-12 CHANGELOG ↗

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

> Huntress shipped a new API for querying remote access connections, a monthly non-technical executive summary report, and an Early Access redesign of the ITDR dashboard packed with new triage, login-classification, and SIEM-search panels.

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

01 **Rapid Identity Triage panel in ITDR dashboard** NEW 75

The new Rapid Identity Triage panel lets you enter a user's email to get an immediate overview of all activity over the past 24 hours, with the ability to export a timeline, revoke sessions, or disable the account.

– Clear input and actions listed, no exact panel navigation path

changelog-20260812-4e8c96b9

02 **Quick SIEM Search panel for Entra logs** NEW 75

The new Quick SIEM Search panel inside the ITDR dashboard exposes pre-populated queries against Huntress-ingested Entra Unified Audit Log data, retained free for up to one year for all ITDR customers.

– Names retention window and log source but no query examples changelog-20260812-4e8c96b9

03 **Redesigned ITDR dashboard in Early Access** NEW 65

A redesigned ITDR dashboard is available in Early Access via the ITDR icon in left-hand navigation, running alongside the existing dashboard.

– Exact navigation path given, no rollout timeline detail changelog-20260812-4e8c96b9

04 **Failed Login Characterization view** NEW 60

A new Failed Login Characterization view surfaces failed login attempts globally with classification by residential proxies, VPNs, and datacenters.

– Describes classification but gives no usage steps changelog-20260812-4e8c96b9

WAS THIS USEFUL? 👍 👎

Automox

1 RELEASE · 2026-08-11 BLOG ↗

Automox is a cloud-based platform that automates patch management and security remediation across endpoints and servers.

> Automox added a new device software status that flags when an installed update is waiting on an application restart to take effect.

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

01 **Pending Application Restart software status** NEW 62

A new Pending Application Restart status appears in the Device Details > Software Table Status column for Windows devices where a package was installed but the application has not yet restarted to finalize the update. Currently supported for Google Chrome, Microsoft Edge, and OBS Studio.

·

— Names exact UI path and status but limited to three apps.

New Software Status for a Device: "Pendi…

WAS THIS USEFUL? 👍 👎

Trend Vision One

1 RELEASE · 2026-08-06 BLOG ↗

Trend Vision One is Trend Micro's unified cybersecurity platform providing threat intelligence, vulnerability management, and incident response across endpoints and networks.

> Trend Vision One's Helm Chart 3.5.1 adds a field for explicitly setting a cluster's resource ID during automated registration.

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

01 **Explicit cluster resource ID in Helm chart** NEW 70

The Helm chart values now include a `visionOne.resourceId` field to explicitly set the cluster resource ID, such as an AWS cluster ARN, during automated registration.

— Names exact config field and use case but lacks broader mechanism detail.

Helm Chart 3.5.1 (2026-08-06)

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Cotool

1 RELEASE · seen 2026-08-12 NOTES ↗

Cotool is an AI-powered code analysis and security tool that identifies vulnerabilities and provides automated fixes for software development.

> Cotool v0.63.0 adds three new integrations (Barracuda, OX Security, Supabase), a Slack emoji-reaction agent trigger, and several administrative controls for GitHub App scoping, session length, and alert escalation.

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

01 **New integrations for Barracuda, OX Security, Supabase** NEW

55

Adds a Barracuda email security integration for investigating and remediating email threats, an OX Security integration to search applications, artifacts, and SBOMs and investigate issues and attack paths, and a Supabase integration for organization and project security posture, configuration review, and project logs.

— Names three integrations and their scope, no setup steps given v0.63.0

02 **Per-organization session length setting** NEW 40

Adds a per-organization session length setting in Settings.

— Names Settings location but no default or range v0.63.0

03 **Slack emoji-reaction agent triggers** NEW 35

Adds Slack emoji-reaction triggers so reacting to a message can run an agent.

— Describes trigger mechanism but no configuration detail v0.63.0

04 **Minimum-severity filter for alert escalations** NEW 35

Adds a minimum-severity filter for escalated alert notifications.

— Names filter but no threshold values or config path v0.63.0

05 **Multi-organization GitHub App installation scoping** NEW

25

Adds multi-organization GitHub App installation scoping.

— Bare description with no mechanism or scope v0.63.0

WAS THIS USEFUL? 👍 👎

Splunk Security Content

1 RELEASE · 2026-08-12 NOTES ↗

Splunk Security Content

> Splunk Security Content v6.4.0 adds 27 new detection analytics for Linux and Windows endpoint threats, refreshes threat-intel lookup files, and permanently removes eleven legacy detections while retiring the Onboarding Assistant beta.

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

01 Permanent removal of eleven legacy detections

BREAKING

75

The following detections are permanently removed: `PowerShell - Connect To Internet With Hidden Window`, Regsvr32 with Known Silent Switch Cmdline, `Rundll32 CreateRemoteThread In Browser`, Splunk App for Lookup File Editing RCE via User XSLT, Splunk Code Injection via custom dashboard leading to RCE, Splunk Enterprise KV Store Incorrect Authorization, Splunk Information Disclosure on Account Login, Splunk Path Traversal In Splunk App For Lookup File Edit, Splunk RCE PDFgen Render, Windows Process Injection Of Wermgr to Known Browser, and Windows Process Injection With Public Source Path. Users relying on any of these should review the deprecated analytics documentation and migrate to recommended replacements before upgrading.

— Names all removed detections and gives migration guidance v6.4.0

02 Linux privilege-escalation, persistence and container abuse analytics

NEW

67

Adds 23 new Linux detections: Linux Binary Executed from Shared Memory Directory, Linux EFI Bootloader File Deletion, Linux File Creation In System Generator Directory, Linux Ghostscript Exploitation, Linux MOTD Script Added, Linux Netcat Outbound Connection, Linux Possible Bootloader Modification, Linux Possible GSM Privilege Escalation, Linux Possible Nimbuspwn Privilege Escalation, Linux Possible Privilege Escalation via PYTHONPATH, Linux Possible System Binary Backdoor, Linux Root Execution of id, Linux Shell History Access Via Command Line Utility, Linux Shell Pseudo Device Reverse Shell, `Linux Suspicious Child Process of PostgreSQL`, Linux Suspicious Docker Build Command Execution, Linux Suspicious GCC Invocation Building Init Shared Object, Linux Suspicious Privileged Container Execution, Linux Suspicious Redis Activity, Linux Suspicious Staging of Alternate System Files, Linux Suspicious XDG Autostart, Linux UDEV Rule Created, and Linux Usermod Root UID Set, covering bootloader tampering, container and database abuse, reverse shells, and multiple privilege-escalation techniques.

— Names all 23 rules but no mechanism beyond brief purpose each v6.4.0

03 **Windows recon and stealth-execution analytics** NEW 62

Adds three Windows detections: Windows Dir Piped to Findstr Activity for directory-listing recon piped into Findstr, Windows Network Sniffing Tool Executed for known sniffing utility execution, and Windows Powershell Commands from DNS TXT for PowerShell commands retrieved via DNS TXT record lookup.

— Names all three rules with brief detection purpose, no mechanism detail v6.4.0

THINNER COVERAGE BELOW

04 **Refreshed attacker_tools and malware_user_agents lookups**

IMPROVED 35

Refreshes the `attacker_tools` and `malware_user_agents` lookup CSV files with updated content to improve detection and investigation context.

— Names the two lookup files but no detail on what changed v6.4.0

05 **Onboarding Assistant beta removed from ESCU** BREAKING 33

The Onboarding Assistant beta is removed from ESCU and is no longer available; it will be reintroduced in Detection Studio in a future release.

— States removal and future reintroduction but no migration steps v6.4.0

└──▷ BREAKING ON UPGRADE

! The Onboarding Assistant beta is removed from ESCU and is no longer available; it will be reintroduced in Detection Studio in a future release.

! The following detections are permanently removed: `PowerShell - Connect To Internet With Hidden Window`, Regsvr32 with Known Silent Switch Cmdline, `Rundll32 CreateRemoteThread In Browser`, Splunk App for Lookup File Editing RCE via User XSLT, Splunk Code Injection via custom dashboard leading to RCE, Splunk Enterprise KV Store Incorrect Authorization, Splunk Information Disclosure on Account Login, Splunk Path Traversal In Splunk App For Lookup File Edit, Splunk RCE PDFgen Render, Windows Process Injection Of Wermgr to Known Browser, and Windows Process Injection With Public Source Path. Users relying on any of these detections should review the deprecated analytics documentation and migrate to recommended replacements before upgrading.

WAS THIS USEFUL? 👍 👎

Fibratus

1 RELEASE · 2026-08-12 NOTES ↗ CODE ↗

Security sensor for realtime threat detection and protection

> Fibratus v3.1.0 adds kernel callstack-based detection, a smart event-filtering pipeline with approvers, 11 new detection rules, and an overhauled signature verification pipeline — while removing the malicious/vulnerable driver filter fields and rules.

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

01 **11 new detection rules across attack techniques** NEW

70

Adds 11 new detection rules covering lateral movement, defense evasion, credential access, and suspicious DLL-loading techniques: 'Untrusted DLL loaded from masqueraded Windows directory', 'Suspicious LSASS process access', 'Potential privilege elevation via arbitrary section creation', 'CldApi DLL loaded by an unusual process', 'BindFlt DLL loaded by an unusual process', 'Direct disk device access', 'Suspicious registry symbolic link creation', 'Activation Context memory section hijacking', 'Untrusted DLL loaded from unusual directory', 'Sensitive registry hive dump', and 'Potential operation evasion via direct syscall'.

— Names all 11 rules verbatim but no config path given. v3.1.0

02 **Kernel thread callstack summary filter field** NEW

67

A new kernel thread callstack summary filter field resolves kernel return addresses to their corresponding driver modules and stitches callstack data to events via timestamps, enabling more reliable callstack-based filtering and detection.

— Names mechanism and purpose but no exact field name given. v3.1.0

THINNER COVERAGE BELOW

03 **Smart event filtering with approvers** NEW

50

Smart event filtering with approvers dynamically evaluates and discards events before they propagate through the full processing pipeline, substantially reducing processing volume while preserving detection-relevant events.

— Explains behaviour but no config key or flag named. v3.1.0

04 **Digital signature verification pipeline overhaul** IMPROVED

48

Overhauled the digital signature verification pipeline with improved async verification, better cache keying, eager metadata-store initialization, and more robust handling of freshly added signatures.

— Describes mechanism changes but no configurable surface named. v3.1.0

05 **Filtering and approver pipeline performance improvements**

IMPROVED 40

Delivers filtering and approver pipeline performance improvements including pattern precompilation, path string interning, object pooling, cached filter values, and more efficient wildcard matching.

— Lists mechanisms but no measurable numbers or flags. v3.1.0

06 **Removal of malicious/vulnerable driver filter fields** BREAKING 30

Malicious/vulnerable driver filter fields have been removed, along with their associated detection rules, as a breaking change in this release.

— States removal but no migration guidance or field names. v3.1.0

└──▷ **BREAKING ON UPGRADE**

! Malicious/vulnerable driver filter fields have been removed, along with their associated detection rules.

WAS THIS USEFUL? 👍 👎

◆ **Threat Intelligence**

OpenCTI ⓘ 1 RELEASE · 2026-08-11 NOTES ↗

Open Cyber Threat Intelligence Platform

> OpenCTI 7.260811.0 adds support for bulk-updating relationship timestamps through background tasks.

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

01 **Mass update of relation start/stop times** NEW 40

Background mass-operation tasks can now bulk update the `start_time` and `stop_time` fields on relations, rather than requiring per-relation edits.

— Names fields and mechanism but no usage example or limits. 7.260811.0

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

tailscale

1 RELEASE · 2026-08-12 CHANGELOG ↗

The easiest, most secure way to use WireGuard and 2FA.

> Tailscale's Kubernetes operator gains in-cluster PeerRelays, Helm annotation support and IPv6 egress on dual-stack clusters, alongside proxy timeout, MTU and metrics fixes.

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

01 **Cert issuance timeout raised in Kubernetes proxies** IMPROVED 55

The per-attempt cert issuance timeout in Kubernetes proxies is increased to 30 minutes, giving ACME challenges more time to complete under load.

– Concrete number and cause given, but no config key to set it.

changelog-20260812-0ad0de49

02 **New Prometheus metrics for DNS and users** NEW 55

Tailscale now exposes `tailscale_dns_split_nameservers` and `tailscale_users` metrics.

– Named metrics are directly queryable, but no context on use. changelog-20260812-0ad0de49

03 **MTU clamping on input and output interfaces** IMPROVED 45

MTU values are now clamped on both input and output interfaces, whereas previously only the output interface was clamped.

– Clear before/after but no user-facing setting to act on. changelog-20260812-0ad0de49

04 **Helm annotation support in Tailscale operator** NEW 25

The Tailscale operator now supports Helm annotations, per the release summary; no further detail on which annotations or how they're applied is given.

– Only named in the summary, no mechanism or example. changelog-20260812-0ad0de49

05 **IPv6 egress on dual-stack clusters** NEW 25

The Tailscale operator now supports IPv6 egress on dual-stack clusters, per the release summary; no further mechanism or configuration detail is provided.

– Only named in the summary, no mechanism or example. changelog-20260812-0ad0de49

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Ory ⓘ 1 RELEASE · seen 2026-08-12 NOTES ↗

Ory is an open-source identity and access management platform providing authentication, authorization, and user management services.

> Ory Network now lets workspace owners verify organization SSO email domains via DNS-TXT record, replacing manual verification.

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

01 **DNS-TXT verification for SSO email domains** IMPROVED 50

Workspace owners can verify their organization's SSO email domains using a DNS-TXT record, replacing the previous manual verification workflow.

– Names mechanism but no exact record format or UI path.

Ory Network, Ory Oathkeeper, Ory Polis, …

WAS THIS USEFUL? 👍 👎

JumpCloud

1 RELEASE · 2026-08-11 BLOG ↗

JumpCloud is a cloud-based directory platform that manages user identities, devices, and access across IT infrastructure.

> JumpCloud Agent 2.153.5 adds policy-based control over the Self Service App Catalog and reworks the macOS menu bar app's password reset and enrollment experience.

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

01 **macOS menu bar password reset and enrollment UX** IMPROVED

50

The macOS menu bar app now simultaneously displays password reset and MDM enrollment actions when both are pending so users complete all required device setup steps, keeps password status constantly visible with a more accessible Reset Password action, and lets users edit TOTP MFA digits in any order during a macOS password reset for more reliable code entry.

— Describes UI behavior changes without exact settings or flags

JumpCloud Agent Release Notes for August…

02 **Policy-controlled Self Service App Catalog** NEW

45

Administrators can now control Self Service App Catalog availability via policy, with the tray app reflecting the policy state to determine whether users see targeted software access.

— Names the feature and mechanism but no exact policy path

JumpCloud Agent Release Notes for August…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

DigiCert CertCentral ⓘ 1 RELEASE · 2026-08-11 NOTES ↗

DigiCert CertCentral is a platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across organizations.

> CertCentral adds a 'Use CN/SANs from CSR' option that overwrites Common name and SAN fields with values parsed from an uploaded CSR, with an undo option.

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

01 **Use CN/SANs from CSR on request forms** NEW 75

A new 'Use CN/SANs from CSR' option on certificate request forms overwrites the existing Common name and SANs fields with values parsed directly from the uploaded CSR, and includes an 'Undo' option to restore the previous values. It is available on DV, OV, and EV TLS; QWAC and QWAC PSD2; X9 PKI for TLS; and Private TLS request forms.

— Names exact option, undo behavior, and applicable form types. August 11, 2026

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Third-Party & Vendor Risk (TPRM)

Conveyor ⓘ 3 RELEASES · seen 2026-08-12 NOTES ↗

Conveyor is a tool that packages Java applications into native installers for Windows, macOS, and Linux with automatic updates.

> Conveyor sharpened its questionnaire workflow with AI-driven SME routing for all questionnaire types, added a SharePoint knowledge base integration, and made portals auto-export answers once auto-complete rules pass.

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

01 AI SME group auto-assignment for all questionnaire types

IMPROVED 73

The 'Assign the suggested group automatically' setting in `Settings > Workflow Automation > Team Routing` now covers all questionnaire types, not just intake questionnaires — unanswered or flagged questions are automatically routed to the matched SME group (e.g. Legal, Product) and that group is notified. This automatic routing now also works for questionnaires arriving via Salesforce, Slack, or direct team submission.

Enable automatic SME routing for all incoming questionnaires so no question stalls waiting for manual triage.

> 📍 In the Conveyor console, go to Settings › Workflow Automation › Team Routing, configure your SME groups, then enable 'Assign the suggested group automatically'.

— Names exact setting path and channels; UI steps given, no numbers.



02 Portal auto-export on auto-complete rule match IMPROVED

70

Portal questionnaires that pass auto-complete criteria now automatically export answers back into the portal, eliminating the manual Export step. This is enabled via Check auto-complete rules under Post answering in `Settings > Sales Self-Serve > Automation Rules`.

— Exact config path named; no worked example given.

🤖 Portals Now Auto-Export When Answers…

03 **SharePoint sync for knowledge base** NEW 60

Microsoft SharePoint is now a knowledge base integration, syncing documents and folders into Conveyor the same way Google Drive already does — including automatic ingestion of new files and nested subfolders for folder-level syncs. ConveyorAI answers draw from the live SharePoint version of a policy document rather than a stale exported copy, keeping compliance responses current without manual re-uploads.

Connect SharePoint so updated compliance policies are automatically reflected in ConveyorAI answers.

> 📍 In Conveyor, go to Integration Settings › Add Integration › SharePoint, then choose a folder to sync — new files added to that folder (including nested subfolders) will be pulled in automatically.

— Clear UI path given but no config keys or limits.

🔗 New Integration: Sync SharePoint Doc…

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Kiro ⓘ 2 RELEASES · seen 2026-08-12 NOTES ↗

> Kiro's headline addition is Cloud Sessions (preview), letting both the CLI and the IDE's Agent Focus Mode run agent sessions in a managed cloud sandbox with repo attachment and cross-machine resume. Alongside that, Agent Focus Mode gained a collapsible rail, in-place attention cards and titlebar navigation shortcuts, while the CLI picked up substring command matching and smoother keystroke handling.

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

01 Cloud Sessions across CLI and IDE NEW

95

Cloud Sessions (preview) run agent sessions in a managed cloud sandbox instead of locally. In the CLI, `kiro-cli chat --cloud` starts one, `--repo` (or the `/repo` picker) attaches a repository, and `--resume-id` reconnects to a running cloud session from any machine after disconnecting; cloud and local sessions now appear together in the session picker. In the IDE, Agent Focus Mode gains the same capability: start a cloud-hosted session alongside local ones, choose the repositories it runs against, and keep those repositories visible below the chat using the same agent and model selectors as local sessions.

Start a cloud session attached to a repo, then resume it later from a different machine without losing agent progress.

```
$ kiro-cli chat --cloud --repo <your-repo>
# later, from any machine:
kiro-cli chat --resume-id <session-id>
```

— Names every flag and gives a runnable command. Cloud Sessions Preview and Smarter Comma…

Cloud Sessions and Agent Focus Mode Upgr…

02 Agent Focus Mode session rail and navigation IMPROVED

70

The Agent Focus Mode sessions rail can be collapsed to an icon-only strip to reclaim chat space while every session stays one click away. New in-place attention cards surface next to the chat when a session needs input, showing the agent's question and options as buttons without forcing a session switch. Back/forward navigation across sessions and settings from the titlebar is added via `Cmd/Ctrl+[` and `Cmd/Ctrl+]`.

— Describes mechanism and gives exact shortcuts. Cloud Sessions and Agent Focus Mode Upgr…

THINNER COVERAGE BELOW

03 **Substring matching in slash command menu** IMPROVED 35

The slash command menu now supports substring matching, so partial typing finds commands mid-word rather than requiring a prefix match.

— Explains behaviour but no command or limit. Cloud Sessions Preview and Smarter Comma…

04 **Windows agent terminal pinned to PowerShell** IMPROVED 23

The agent's Windows terminal is now pinned to PowerShell.

— Names the change but no further detail. Cloud Sessions and Agent Focus Mode Upgr…

05 **Smoother TUI input during rapid keystrokes** IMPROVED 14

TUI input handling is smoother during rapid keystrokes.

— Bare one-line claim with no mechanism. Cloud Sessions Preview and Smarter Comma…

WAS THIS USEFUL? 👍 👎

gemini-cli ⓘ 1 RELEASE · 2026-08-11 NOTES ↗

An open-source AI agent that brings the power of Gemini directly into your terminal.

> Gemini CLI's headline addition is experimental model steering, letting users inject real-time guidance into a running agent without restarting it; alongside that, v0.55.1 quietly shipped a Cloud Run-based caretaker triage pipeline, a separate PR-generator agent, an eval coverage report command, and automatic tool registry discovery.

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

01 Experimental model steering for running agents NEW 93

Adds an `experimental.modelSteering` key to `settings.json` (disabled by default) that lets practitioners type free-text hints while the agent spinner is active — correcting paths, skipping steps, adding context, or redirecting effort — without stopping the session. The same toggle is reachable via the `/settings` command's 'Model Steering' entry instead of editing the file directly. Steering hints trigger a two-phase response: a fast small-model acknowledgment confirming receipt, then context injection at the agent's next turn with plan re-evaluation and minimal-diff task updates.

Enable model steering once in your settings file so all future long-running agent sessions accept mid-task hints.

```json
{
  "experimental": {
    "modelSteering": true
  }
}
```

While the agent is refactoring code and heads down the wrong path, steer it without restarting the task.

📍 `/settings` # Search for 'Model Steering', set to true # Then start your task: # › Refactor the database service # (spinner visible) type: # Actually, the utilities are in src/common/utils. # Press Enter — CLI acknowledges and injects the hint into the next turn.

— Names exact config key, command, and injection mechanism with runnable examples.

product docs

02 Caretaker agent triage pipeline NEW 66

Ships a Cloud Run webhook ingestion service for automated GitHub issue intake, a caretaker triage Cloud Run job workflow with an LLM-driven orchestrator, prompt hill-climbing, and automatic pre-close commenting, plus a caretaker evaluation framework with a judge runner, local golden issue collection, and Firestore sync tools for offline triage quality measurement. Also adds issue comment handling and re-triage in the ingestion pipeline, publishing of workable spec events to a ready-for-code Pub/Sub topic, a GCP deployment script for the caretaker services, and a Firestore schema update adding `error` and `pr_number` fields for triage tracking.

— Names every pipeline surface but is internal infra, not user-actionable. v0.55.1

THINNER COVERAGE BELOW

03 **PR-generator agent** NEW 52

Implements a PR-generator agent with an Antigravity runner, an iterative bug-fixing state machine, a Cloud Run job, Firestore dual-locking, and Workflows infrastructure.

— Named components listed but no usage path for end users. v0.55.1

04 **Eval coverage report command** NEW 35

Adds an eval coverage report command for measuring evaluation coverage across the codebase.

— Names the command's purpose but no exact invocation. v0.55.1

05 **Tool registry discovery** NEW 35

Adds tool registry discovery, enabling the CLI to find and enumerate registered tools automatically.

— States capability but no mechanism or usage detail given. v0.55.1

WAS THIS USEFUL? 👍 👎

Claude Code

1 RELEASE · 2026-08-11 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 v2.1.228 focuses on hardening synced skills against local privilege escalation, aligning file-write permissions across tools, and polishing session messaging, Vertex AI auth, and compaction feedback.

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

01 **Security hardening for claude.ai-synced skills** BREAKING

65

Skills synced from claude.ai no longer shadow local commands or MCP prompts, their descriptions are sanitized and labeled, and skill bodies can no longer run `!` commands or expand `@` files on your machine.

— Names mechanism and restrictions but no migration steps given. v2.1.228

THINNER COVERAGE BELOW

02 **Write tool overwrite rule aligned with Edit tool** IMPROVED

50

Newer models can now overwrite an existing file they have not read in the current session, matching the Edit tool's existing rules; older models still require a prior read before overwriting.

— Clear behavior change but no config or flag to invoke it. v2.1.228

03 **Cross-session message display improvements** IMPROVED 45

Sender and body of cross-session messages now display inline instead of as a collapsed line, and messages to Remote Control sessions on other machines show your Remote Control session name as the sender.

— Describes UI change but no exact navigation path given. v2.1.228

04 **Faster Vertex AI credential failure detection** IMPROVED

40

Expired or missing Google Cloud credentials for Vertex AI now fail within seconds instead of retrying for minutes.

— States before/after timing but no reproduction steps. v2.1.228

05 **Compaction progress feedback improvements** IMPROVED 30

The retry countdown and stall hint now appear during compaction alongside the existing progress bar.

— Minor UI addition described briefly without further detail. v2.1.228

WAS THIS USEFUL? 👍 👎

Cline

2 RELEASES · 2026-08-11 NOTES ↗ CODE ↗

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

> Cline's CLI and SDK releases both add Fable 5 to the Vertex model catalog and fix custom Vertex model ID routing, while the SDK also removes a hardcoded token cap from its summarizer.

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

01 Configurable max output tokens in summarizer IMPROVED

73

The summarizer now honors an explicit `max-output-tokens` setting, defaulting to 4096 and lowered when the model reports less capacity, replacing a previously hardcoded 1024-token cap. This enables compaction to work correctly on reasoning models.

— Names config key and default/prior values, no set command shown sdk/sdk/v0.0.73

THINNER COVERAGE BELOW

02 Fable 5 added to Vertex model catalog NEW

52

Adds `claude-fable-5` (Fable 5) to the Vertex model catalog. Because Vertex billing is region-dependent, its cost displays as unknown rather than a wrong fixed price.

— Names the model but no usage steps given sdk/sdk/v0.0.73 cli-v3.0.53

03 Custom Vertex model ID routing fixed IMPROVED

50

Custom Vertex model IDs are now passed through unchanged, routing Claude-style IDs to the Anthropic-on-Vertex path automatically instead of being mishandled.

— Mechanism described but no config example shown sdk/sdk/v0.0.73 cli-v3.0.53

WAS THIS USEFUL? 👍 👎

Goose

1 RELEASE · 2026-08-12 | NOTES ↗ | CODE ↗

an open source, extensible AI agent that goes beyond code suggestions - install, execute, edit, and test with any LLM

> Goose v1.46.0 ships a Hooks framework with PreToolUse denial, a new `goose review` command, 25+ new LLM providers and dozens of new models, plus broad environment-variable configurability, OAuth improvements, and a wave of CLI, desktop, and TUI enhancements.

└──▷ **WHAT SHIPPED · 21 FEATURES** most completely described first ? what's the number?

01 **New environment variables for configuration** NEW 75

Adds `LOCAL_WHISPER_LANGUAGE` to set the language for multilingual local voice transcription, `GOOSE_MAX_TOOL_RESPONSE_SIZE` to cap tool output size returned to the agent, `MAX_CODE_BLOCK_LINES` to limit lines rendered in code blocks, `GOOSE_FAST_MODEL` (consumed by `ModelConfig::with_fast`) to designate a fast model for lightweight tasks, `GOOSE_DOCS_ROOT` to redirect documentation lookups for air-gapped deployments, `GOOSE_OAUTH_CALLBACK_PORT` to pin a stable OAuth `redirect_uri` port across sessions, and `AZURE_OPENAI_AD_TOKEN` to support Azure AD (Entra ID) bearer-token authentication for the Azure OpenAI provider.

Pin a stable OAuth callback port so OAuth redirect URIs stay consistent across restarts, useful in environments with strict firewall rules.

```
$ GOOSE_OAUTH_CALLBACK_PORT=9876 goose session
```

Set the transcription language for local Whisper so voice input is recognized correctly without auto-detection overhead.

```
$ LOCAL_WHISPER_LANGUAGE=fr goose session
```

— Lists env vars with two runnable examples, not all covered. v1.46.0

02 **New session CLI commands and flags** NEW 75

Adds `--edit` flag to the session command to edit conversation history before forking a session, `/goal` slash command for agent self-evaluation before a session concludes, `/status` slash command to report current session state, `/model` slash command to switch the active model mid-session, tab completion and provider switching in the `model` command, a CLI command to list available skills with their token counts, and

automatic loading of global agent hints from `~/.agents/AGENTS.md` each session. Also adds Markdown as a format option for session export.

— Names exact flags and commands; no full syntax examples. v1.46.0

---

03 **`goose review` command for local code review** NEW 70

Adds `goose review`, a CLI command for running local code review workflows against the working directory, delegating checks via the default orchestrator.

Run a local code review against the current working directory, delegating checks via the default orchestrator.

```
$ goose review
```

— Runnable command shown; deeper mechanism unexplained. v1.46.0

---

04 **Desktop and TUI interface improvements** NEW 70

Surfaces per-message token count, cost, time-to-first-token (TTFT), and tokens-per-second stats in the UI, and tracks cache tokens separately for accurate cost reporting. Streams shell command output to the UI in real time as commands execute. Adds a resizable sidebar with a drag handle and persisted width, chat history search in the navigation panel, a search filter in the provider selection grid, and a TUI diff viewer. Adds desktop localizations for French, German, Italian, Portuguese, Indonesian, Malay, Vietnamese, zh-TW, Korean, Spanish, Japanese, Hindi, Russian, Turkish, and Simplified Chinese (zh-CN), plus language selection in the UI.

— Many UI features named but only navigation paths, no commands. v1.46.0

---

05 **25+ new LLM provider integrations** NEW 65

Adds support for Celeris, Friendli, Azure AI Foundry (multi-LLM), OllamaCloud (dynamic model discovery), Sakana AI (Fugu API), iFlytek Spark, Astron MaaS, Fireworks AI, Together AI, OrcaRouter, Perplexity, Alibaba/Qwen (DashScope), Databricks AI Gateway, Scaleway, NEAR AI Cloud, EmpirioLabs, xAI SuperGrok (OAuth subscription), Vercel AI Gateway, Routstr, FuturMix, oMLX, Atomic Chat, and Muse Spark 1.1 (Meta Models API), plus AWS Bedrock support for OpenAI GPT models via the mantle endpoint and support for OpenRouter request parameters.

— Names 25+ providers but no setup steps given. v1.46.0

---

06 **OAuth authentication improvements** NEW 65

Adds Hugging Face OAuth support with a dedicated auth tab in settings, adds proactive OAuth token refresh to avoid re-authentication on every session, and adds a `logo_uri` field to OAuth client ID metadata documents (CIMD) so authorization servers can display the goose logo on consent screens.

— Names logo_uri field and settings tab but no full flow. v1.46.0

07 **ACP protocol session and transport enhancements** NEW

65

Adds `_meta.hidden` field to the `ACP session/new` request to create hidden (background) sessions, and adds TLS support for `acp serve`.

— Names exact field and command; no further usage shown. v1.46.0

08 **goose:// deep links for sessions** NEW

65

Adds `goose://new-session` and `goose://resume` deep links, including support for an initial prompt via `goose://new-session`.

— Named URI schemes; no host app integration detail. v1.46.0

09 **Hooks framework with PreToolUse denial** NEW

60

Adds a Hooks framework with `PreToolUse` denial, letting operators intercept and block tool calls before execution to enforce policy on sensitive operations.

Block a specific tool from executing during a session by wiring a PreToolUse hook that denies it — useful for enforcing policy on sensitive operations.

```
$ GOOSE_MAX_TOOL_RESPONSE_SIZE=65536 goose session
```

— Names hook type but no config syntax for wiring hooks. v1.46.0

10 **Summon extension peek mode and delegate parameters** NEW

60

Adds Summon extension peek mode for async background tasks, plus `context` and `working_dir` parameters for delegate tasks.

— Names parameters but no usage example. v1.46.0

11 **`goose project` command removed** BREAKING

60

The `goose project` command (and its `p` alias) has been removed; CLI project support is no longer available.

— States removal clearly with alias, no migration path given. v1.46.0

THINNER COVERAGE BELOW

12 **New model support across providers** NEW

55

Adds support for GPT-5.6, GPT-5.5, GLM-5.2, Opus5 (with adaptive thinking), claude-sonnet-5, claude-fable-5, Cursor composer-2.5, Cursor-agent (fetched from CLI), latest

Gemini models, Gemini 3.x, MiniMax-M3/M2.7, and OVHcloud Qwen3.6-27B.

— Lists model names only, no selection instructions. v1.46.0

13 **Cross-language SDK bindings via UniFFI** NEW 55

Adds a UniFFI SDK for cross-language bindings, Python wheel publishing, and Kotlin FFI with Maven publish.

— Names packaging targets, no install instructions. v1.46.0

14 **Image support in sessions** NEW 50

Adds an image read tool for ingesting images directly in sessions, and forwards images and MCP embedded-resource blobs in Anthropic and Google provider formats.

— Describes capability, no tool name syntax shown. v1.46.0

15 **Local inference provider expansion** NEW 45

Supports MLX models in the local inference provider and adds Linux Vulkan support for local inference.

— States support added, no configuration details. v1.46.0

16 **Trusted ownership metadata for MCP apps** NEW 45

Binds MCP apps to trusted ownership metadata to restrict which extensions can be loaded.

— Mechanism named but no config surface given. v1.46.0

17 **GenAI semantic conventions in OpenTelemetry spans** NEW 45

Emits GenAI semantic convention attributes in OpenTelemetry (OTEL) spans.

— Names standard, no setup or exporter detail. v1.46.0

18 **Termux detection in installer** NEW 40

Adds Termux detection in the installer to automatically select the musl portable build.

— Explains behavior but no manual override option given. v1.46.0

19 **Encrypted Nostr session sharing and import** NEW 40

Adds encrypted Nostr session sharing and session import.

— Names protocol, no encryption or import mechanics detail. v1.46.0

20 **Quarterly recurrence in session scheduler** NEW 30

Adds a quarterly recurrence option to the session scheduler.

— One-line addition with no further detail. v1.46.0

21 **Structured summary output for context compaction** NEW

30

Adds structured summary output with template rendering for context compaction.

— Thin description, no template syntax or config shown. v1.46.0

└──▷ **BREAKING ON UPGRADE**

! The `goose project` command (and its `p` alias) has been removed; CLI project support is no longer available.

WAS THIS USEFUL? 👍 👎

AGENT

◆ **AI Agent Frameworks**

CrewAI ⓘ 1 RELEASE · 2026-08-12 NOTES ↗

Framework for orchestrating role-playing, autonomous AI agents. By fostering collaborative intelligence, CrewAI empowers agents to work together seamlessly, tackling complex tasks.

> CrewAI 1.15.15 extends flow telemetry to capture outcome, duration and human-in-the-loop signals.

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

01 **Flow outcome and HITL telemetry reporting** IMPROVED

30

Flow telemetry now reports flow outcome, duration, and human-in-the-loop signals.

— Names data reported but no mechanism or config surface. 1.15.15

WAS THIS USEFUL? 👍 👎

Object-Oriented Agents

1 RELEASE · 2026-07-30 NOTES ↗

NVIDIA Object Oriented Agents: the Pythonic way to build AI Agents.

> OO Agents v0.0.7 is a security-hardening release across the playground, viewer API, ShellTools and CodeAct, paired with a new NOOA cybergym agent example and official PyPI distribution.

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

01 **PyPI distribution via Trusted Publishing** NEW 75

The package is now published to PyPI via Trusted Publishing on GitHub Releases, enabling `pip install nooa` or `uv add nooa` installs from official releases.

— Gives exact install commands, fully actionable. v0.0.7

02 **AST decoder replaces eval() in CodeAct** IMPROVED 60

The `CodeAct` constructor-string coercion path no longer uses eval(); it now uses an AST decoder, removing arbitrary code execution from that code path.

— Names the exact fix and mechanism, no migration steps given. v0.0.7

THINNER COVERAGE BELOW

03 **Resource limits and auth on trace-viewer API** IMPROVED 55

The trace-viewer's viewer ingest path now has resource limits to cap memory and request size, and the viewer API requires an auth token and restricts CORS, narrowing the attack surface of the local trace viewer.

— Describes mechanism but no exact endpoint or flag names. v0.0.7

04 **Server-declared endpoint allowlist for playground models** IMPROVED 45

The playground custom-model endpoint and `api_key_env` are now constrained to server-declared pairs, preventing clients from supplying arbitrary LLM endpoints.

— Names the config field but no exact endpoint path. v0.0.7

05 **ShellTools confined to working directory** IMPROVED 45

`ShellTools` file operations are now confined within the current working directory, limiting the filesystem surface available to agent-executed shell commands.

— States the restriction, no configuration detail given. v0.0.7

06 **NOOA cybergym agent example** NEW 45

Adds a NOOA cybergym agent example under `examples/` demonstrating a cybersecurity-oriented agent built on the OO Agents framework.

— Names the example path but no further detail on its behaviour. v0.0.7

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ AI Model & Data Infrastructure

Hugging Face ⓘ 1 RELEASE · seen 2026-08-12 NOTES ↗

Hugging Face is a platform providing pre-trained machine learning models, datasets, and tools for natural language processing and computer vision.

> Hugging Face introduced granular, per-resource-group feature access controls, moving beyond a single org-wide role setting.

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

01 **Per-resource-group feature access controls** NEW 55

Organizations can now scope feature access controls to individual resource groups rather than applying a single org-wide role. For example, Jobs can be left open to all members while Inference Endpoints are restricted to admins and blog publishing is limited to a specific resource group.

— Explains mechanism and gives examples but no UI path or config key

Granular Feature Access

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Daytona ⓘ 2 RELEASES · seen 2026-08-12 NOTES ↗

Daytona is a development environment platform that enables standardized, reproducible coding spaces for teams and AI agents.

> Daytona 0.204.0 and 0.200.0 add snapshot management by name, outbound proxy configuration for sandboxes, typed git transport errors, and a round of security hardening covering TLS enforcement, CLI config file permissions, and configurable SDK timeouts.

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

01 **Snapshot management by name and filtering** IMPROVED 55

Adds support for snapshot `delete` and `activate` operations by ID or name across the SDK and CLI, and adds a `sourceSandboxId` filter parameter to the snapshot list SDK call.

– Names operations and filter param but no usage example

Snapshot operations by name and outbound…

02 **Outbound proxy configuration on sandbox create** NEW 55

Adds `outboundProxyUrl` parameter to the sandbox create SDK call for outbound proxy configuration.

– Names exact new parameter but no example of use

Snapshot operations by name and outbound…

03 **Configurable Go SDK HTTP timeout** IMPROVED 50

Go SDK HTTP timeout is now configurable, enabling reliable long-running executions that previously timed out.

– Explains benefit but no config key or flag named snapshot-20260812

04 **CLI config file permissions restricted to 0600** IMPROVED 50

CLI config file is now created with `0600` permissions, restricting access to the owning user only.

– Exact permission value named, no file path given snapshot-20260812

05 **Typed git transport error codes** NEW 45

Introduces `GIT_TRANSPORT_FAILED` and `GIT_REMOTE_REJECTED` typed error codes across all SDKs.

— Names error codes but no handling detail `Snapshot operations by name and outbound…`

06 **TLS verification enforced in Ruby Socket.IO client** IMPROVED

35

Enforces TLS verification in the Ruby Socket.IO client for hardened client security.

— States change but no mechanism or migration note `snapshot-20260812`

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

bpf-linker

1 RELEASE · 2026-08-12 NOTES ↗

BPF static linker

> bpf-linker v0.11.0 adds support for targeting BPF CPU v4 and disabling specific CPU features, alongside an upgrade to the LLVM 23 backend.

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

01 **CPU targeting flags: v4 and negative features** NEW 70

`--cpu` now accepts `v4` as a valid value to target BPF CPU v4 processors, and `--cpu-features` has been extended to accept disabled features using the `-feature` prefix (e.g. `--cpu-features=-dwarfris`), in addition to the existing `+feature` enable syntax.

Target a modern kernel with BPF CPU v4 and selectively disable a CPU feature in a single build.

```
$ bpf-linker --cpu=v4 --cpu-features=+alu32,-dwarfris -o
output.o input.o
```

— Names exact flags and syntax but limited mechanism detail. v0.11.0

THINNER COVERAGE BELOW

02 **LLVM 23 backend upgrade** IMPROVED 25

Updates the underlying LLVM backend to LLVM 23.

— Bare version bump with no further detail. v0.11.0

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 12, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe