THE CYBER TOOLCHAIN

NO. 042

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 11, 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 54 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

adscan

1 RELEASE · 2026-08-10 NOTES ↗

Free Active Directory pentesting tool for Linux. Automates AD and LDAP enumeration, Kerberoasting, ASREPRoast, password spraying, ADCS/ESC1, DCSync, RBCD, gMSA, shadow credentials, NTLM relay and credential dumping across 104 techniques, mapping BloodHound-compatible attack paths to Domain Admin.

> ADscan v11.1.0 adds cross-forest escalation chains for Kerberos delegation and MSSQL linked-server compromise, a new `ADSCAN_TRUST_SCOPE` environment variable to control trust-domain collection, and engagement-type-aware behavior for domain collection, password resets, and MSSQL account cleanup, alongside Tier 0 membership counts and an interactive SMB host-enrichment scope prompt.

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

01 **Trust-scope control via ADSCAN_TRUST_SCOPE** NEW 83

The `ADSCAN_TRUST_SCOPE` environment variable (values `all` or `origin`) controls which domains are collected across forest trusts during a scan, and can be overridden per scan via the web scan-setup domain-enumeration policy.

Run an audit scan locked to the origin domain only, preventing collection of trusted forests outside the agreed legal scope.

```
$ ADSCAN_TRUST_SCOPE=origin adscan ci auth --type audit --
interface eth0 --domain corp.local --dc-ip 10.0.0.1 -u alice -p
'S3cr3t!'
```

— Exact env var, values, and runnable example given. v11.1.0

02 **Interactive SMB host-enrichment scoping** NEW 66

On large directories, ADscan now prompts the operator to choose an SMB host-enrichment scope: the recommended host set, a custom count, all hosts, or skipping enrichment entirely; reports and platform state how many hosts were enriched out of the reachable total.

— Describes prompt options and reporting, no flag named. v11.1.0

03 **Cross-forest MSSQL linked-server-chain compromise** IMPROVED 64

The MSSQL linked-server-chain compromise attack path now extends across forest boundaries, coupling SYSTEM-on-DC compromise to directory replication in the target trusted domain and terminating the attack path there rather than at the controller.

— Clear mechanism and before/after, no config surface named. v11.1.0

04 **Engagement-type-aware Force Change Password handling** IMPROVED

62

Force Change Password steps now follow engagement type in unattended runs: CTF scans reset the target password and continue automatically, while audit scans skip the reset unless an operator explicitly approves it; machine account password resets are refused outright in every mode.

— Detailed behavior per mode, no exact flag shown. v11.1.0

05 **Cross-forest Kerberos delegation escalation path** NEW

60

ADscan now surfaces cross-forest Kerberos ticket delegation as an escalation route, drawing the two forests as one blast radius in the attack path map and adding remediation guidance to disable the delegation.

— Names the path and remediation but no exact command. v11.1.0

06 **Engagement-type-aware trust and domain collection** IMPROVED

60

Trust relationships are always mapped regardless of collection scope; domain collection now follows engagement type, with CTF scans collecting all reachable domains including newly unlocked trusted forests unattended, while audit scans default to the origin domain only.

— Clear before/after by engagement type, no flag named. v11.1.0

THINNER COVERAGE BELOW

07 **Workspace-exit cleanup of MSSQL follow-up accounts** NEW

58

At workspace exit, ADscan prompts whether to delete any account it created during an MSSQL follow-up (default: yes); declining records the exact removal command so nothing is silently left behind.

— Default behavior stated but exact removal command not shown. v11.1.0

08 **Tier 0 membership counts in reports** NEW

53

Reports and the platform now show how many accounts already hold Tier 0 privilege via group membership alongside the count reachable through attack paths, with already-privileged accounts excluded from the attack-path figure.

— Explains what's shown but no UI path or export named. v11.1.0

09 **Account choice for MSSQL SYSTEM-escalation and HasSession**

IMPROVED

49

MSSQL SYSTEM-escalation follow-up and HasSession exploitation steps now offer a choice between minting a fresh privileged account or elevating an account already under the operator's control.

— Describes the choice but no exact command shown. v11.1.0

10 **HasSession escalation targets local Administrators** IMPROVED

43

HasSession escalation now elevates to the host's own local Administrators group on member servers instead of always adding to Domain Admins.

— Clear before/after but brief, no command shown. v11.1.0

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Openlayer

2 RELEASES · SEEN 2026-08-11 CHANGELOG ↗

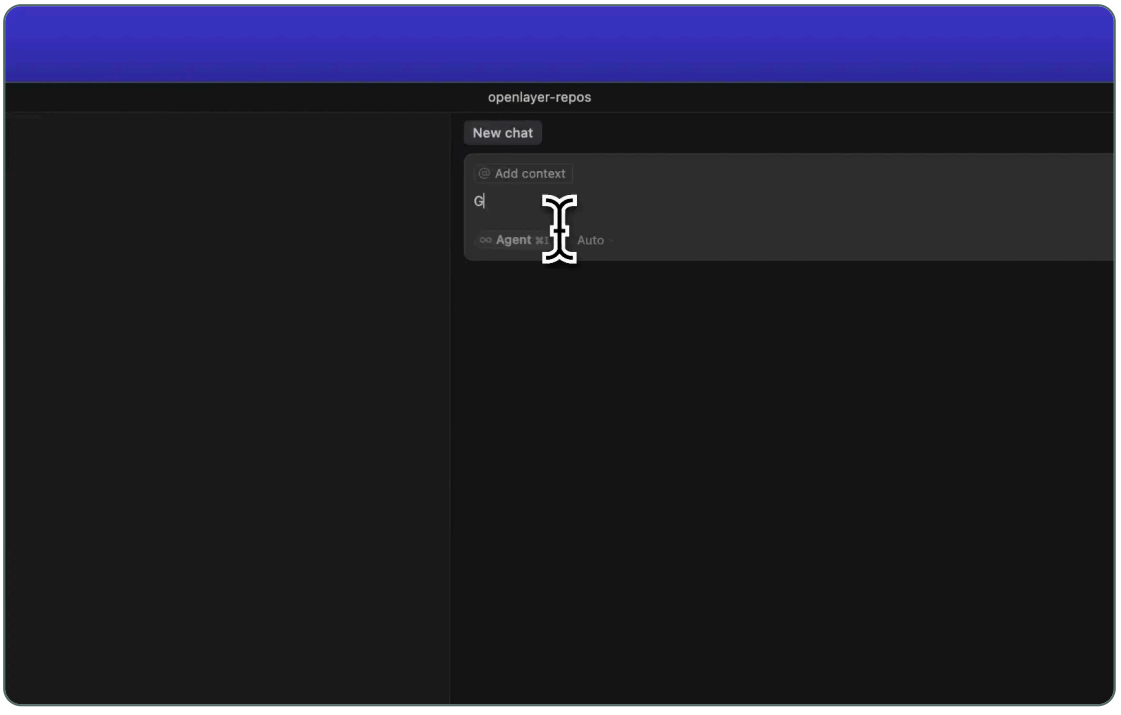
Openlayer is a platform for monitoring and evaluating large language model performance, quality, and safety in production.

> Openlayer's August 2026 release is headlined by a hosted OAuth-protected remote MCP connector, packaged Agent Skills for AI coding assistants, and AI-generated summaries across sessions, traces, and test results, alongside dozens of CLI, Gateway, tracing, governance, and on-prem additions.

└──▷ WHAT SHIPPED · 42 FEATURES most completely described first ? what's the number?

01 **Remote OAuth-protected MCP connector** NEW 95

The Openlayer MCP server is now hosted at `https://mcp.openlayer.com`, OAuth-protected so no API key or local install is needed, and supports switching workspace directly from chat for users in multiple workspaces. Integrates with Claude desktop/claude.ai via Settings → Connectors → Add custom connector using `https://mcp.openlayer.com`, Claude Code via `claude mcp add --transport http openlayer https://mcp.openlayer.com`, Cursor via `.cursor/mcp.json` (`"url": "https://mcp.openlayer.com"`), VS Code via the `MCP: Add Server` command or `.vscode/mcp.json` (`"type": "http"`, `"url": "https://mcp.openlayer.com"`), Codex via `codex mcp add openlayer --url https://mcp.openlayer.com`, and other clients via the `mcp-remote` bridge (`npx -y mcp-remote https://mcp.openlayer.com`).



Register the Openlayer remote MCP server in Claude Code so the assistant can access your Openlayer workspace without managing an API key.

```
$ claude mcp add --transport http openlayer
https://mcp.openlayer.com
```

Add the Openlayer MCP server to a VS Code project for use with GitHub Copilot or other MCP-aware extensions.

```json
{
  "servers": {
    "openlayer": {
      "type": "http",
      "url": "https://mcp.openlayer.com"
    }
  }
}
```

Bridge the remote connector through mcp-remote for clients (e.g. Windsurf, Zed) that lack native remote-MCP + OAuth support.

```json
{
  "command": "npx",
  "args": ["-y", "mcp-remote", "https://mcp.openlayer.com"]
}
```

— Full runnable command set per client, complete `changelog-20260811-b3d33883` `snapshot-20260811`

02 **Openlayer Agent Skills for coding assistants** NEW 95

Packaged Agent Skills teach AI coding assistants to integrate a codebase with Openlayer — tracing, offline evals, tests, guardrails, and CI gating — covering both LLM/agent apps and traditional ML models (scikit-learn, XGBoost, and similar). Install via `npx skills add openlayer-ai/openlayer-skills --skill "openlayer"` into Codex, GitHub Copilot, Gemini CLI, Windsurf, opencode and other agents, or manually by symlinking/copying the `skills/openlayer` folder from `github.com/openlayer-ai/openlayer-skills`; ships as a Claude Code plugin, a Cursor plugin, or via the skills CLI. The skill uses `openlayer.json` + `tests.json` for offline evals, integrates the Openlayer MCP server for workspace queries and doc search, and authenticates via `OPENLAYER_API_KEY` (and `OPENLAYER_BASE_URL` for self-hosted instances).

Install the Openlayer skill into all detected coding agents on your machine in one command.

```
$ npx skills add openlayer-ai/openlayer-skills --skill
"openlayer"
```

Authenticate your agent to Openlayer, or point it at a self-hosted instance.

```
$ export OPENLAYER_API_KEY=<your-api-key>
export OPENLAYER_BASE_URL=<your-instance-url>  # self-hosted
only
```

— Exact install command and config files, fully actionable changelog-20260811-b3d33883 snapshot-20260811

---

03 **Python auto-instrumentation via `openlayer.init()`** NEW

90

openlayer.init() automatically patches every installed supported LLM SDK — OpenAI, Anthropic, Mistral, Groq, Gemini, LiteLLM, Portkey, Google ADK, and Azure — in one call; opt out with `auto_instrument=False`.

— Exact call and flag named, fully actionable snapshot-20260811

---

04 **AI-generated summaries for sessions, traces, and test results**

NEW

80

Session detail views open with an AI summary explaining what the user was trying to do, what the agent did, and where it struggled (Projects > Sessions > open a session). Trace detail views get an AI summary replacing line-by-line span reading (open any trace from the Traces view). Test results include a summary of dominant failure modes with drill-down to the exact matching rows. Summaries are generated on demand and cached so reopening the same session or result is instant.

— Exact navigation paths and caching behaviour given changelog-20260811-b3d33883 snapshot-20260811

---

05 **CLI project bootstrap via `openlayer init`** NEW

75

`openlayer init` is a guided command that takes a fresh directory to a pushable, linked, monitored project with tracing instrumentation configured in one step.

Bootstrap a new project from scratch — tracing, links, and monitoring configured in one shot — without manual setup steps.

```
$ openlayer init
```

Bootstrap a new project with tracing fully wired in one guided step, suitable for kicking off a new LLM app or CI environment.


```
$ openlayer init
```


— Names the exact runnable command, mechanism kept high-level changelog-20260811-b3d33883 snapshot-20260811

06 **Claude Agent SDK tracing** NEW 75

A drop-in import auto-traces Claude Agent SDK runs in TypeScript. Python tracing instruments the full agent loop, capturing tool calls, MCP tools, subagent nesting, session continuity, cost, and tokens.

— Detailed mechanism, no import statement given changelog-20260811-b3d33883 snapshot-20260811

07 **Semantic search filters for test subpopulations** NEW 75

A `semantic` match type in test subpopulation filters uses embeddings to match rows by topic or intent rather than exact or keyword wording — useful for free-text columns like user messages, model outputs, and transcripts — and composes with existing metadata conditions (date ranges, user IDs, tags).

Narrow a test to rows about billing issues even when users phrase the problem differently (`charged twice`, `invoice looks wrong`).

📍 In your project, go to Tests and create or edit a test. Under subpopulation filters, select the free-text column (e.g. user message), set the match type to `semantic`, and enter the phrase `billing problems`. The test will evaluate only rows semantically similar to that phrase.

— Named match type and worked example given changelog-20260811-b3d33883 snapshot-20260811

08 **PATCH endpoint for project metadata** NEW 70

A new PATCH endpoint updates a project's name and description via API, surfaced in the Python, Go, Java, and Ruby SDKs and the MCP server.

– Method and resource named, exact path not given changelog-20260811-b3d33883 snapshot-20260811

09 **Expanded workspace API key scopes** IMPROVED 70

Workspace API keys can now be used for updating projects, cost-override endpoints, test retrieval/update/deletion, and seven previously session-only read routes.

– Named scopes, no exact endpoint paths listed changelog-20260811-b3d33883 snapshot-20260811

10 **Per-project data retention windows** NEW 70

Set a rolling retention window per project — 7, 30, 90, 180, or 365 days, or indefinite — after which monitoring data older than the window is deleted.

– Exact window options and deletion behaviour named changelog-20260811-b3d33883 snapshot-20260811

11 **Gateway spend visibility and admin controls** NEW 70

Adds a per-model breakdown on the usage endpoint, a per-user spend aggregation endpoint, projected savings shown when creating a limit, key rotation, and provider secrets settable from the admin UI.

– Names each endpoint/control, no exact paths given changelog-20260811-b3d33883 snapshot-20260811

12 **Unified trace config via `init()`, `configure()` deprecated** DEPRECATED 70

Trace configuration in the SDK is unified around init(); settings now merge incrementally instead of resetting on each call. configure() is deprecated in favor of init().

– Names both functions and the behaviour change snapshot-20260811

13 **CSV and XLSX export formats** IMPROVED 65

`openlayer export` now supports CSV and XLSX output formats in addition to JSON, and adds the ability to include scores in the export.

– Names formats and command, no flag syntax shown changelog-20260811-b3d33883 snapshot-20260811

14 **Multimodal Gateway `/responses` and `/messages` endpoints** IMPROVED 65

The Gateway's `/responses` and `/messages` endpoints now accept and return media, not just text, adding multimodal support to prompt and response handling.

– Named endpoints, method not specified changelog-20260811-b3d33883 snapshot-20260811

15 **MCP toolset expanded to 53 tools with governance access**

IMPROVED 65

MCP tooling grows from 44 to 53 tools, adding governance and lifecycle access (frameworks, rules, and rule results) plus 12 new session, trace, and test-lifecycle tools reachable through MCP.

– Concrete tool counts, no tool names listed changelog-20260811-b3d33883 snapshot-20260811

16 **Granular token usage and cost tracking** NEW 65

Cached, reasoning, and audio token categories are now priced and stored, with per-API-key cost visible in the UI; LLM-as-judge evaluator runs are cost-tracked per test result.

– Named token categories and cost-tracking scope changelog-20260811-b3d33883 snapshot-20260811

17 **Headless CLI mode for CI and agents** NEW 60

Non-interactive CLI mode lets `link`, `login`, and resource creation run without prompts, enabling the CLI to be driven from AI coding agents and CI pipelines.

– Names affected subcommands, no exact flag syntax given changelog-20260811-b3d33883 snapshot-20260811

18 **Semantic (model-based) Gateway routing** NEW 60

Routing rules can now score each prompt in-process with a model and route it to a strong or weak target by threshold, alongside existing static first-match rules.

– Mechanism described, no config keys shown changelog-20260811-b3d33883 snapshot-20260811

19 **Deployment admin area for on-prem operators** NEW 60

A new deployment admin area, gated by a dedicated deployment-admin role, gives operators background-job queue health, metrics, history, and license/feature status.

– Names role and data surfaced, no navigation path changelog-20260811-b3d33883 snapshot-20260811

20 **On-prem deployment options for storage and secrets** NEW 60

Adds CSI Secrets Store support, external/managed ClickHouse, native Oracle Object Storage as a storage backend, and a multi-arch Gateway image for ARM hosts.

— Names every storage/secrets surface added changelog-20260811-b3d33883 snapshot-20260811

THINNER COVERAGE BELOW

21 **Sorting and filtering on more list endpoints** IMPROVED 55

Adds sorting and filtering to the workspace rules, project members, and user sessions list endpoints, plus server-side group filtering on rules tables.

— Named endpoints, no query params shown changelog-20260811-b3d33883 snapshot-20260811

22 **Dedicated Gemini and Vertex AI tracing** NEW 55

Adds dedicated tracing for the unified Gemini SDK including Vertex AI in Python and TypeScript, where previously only the legacy SDK was auto-instrumented.

— Names SDKs and languages, no import/API shown changelog-20260811-b3d33883 snapshot-20260811

23 **Free public LLM cost explorer** NEW 55

A free public tool at `llm-costs.openlayer.com` lets anyone browse and compare LLM pricing across providers, including per-token-category pricing (cached, reasoning, audio).

— Names the URL, no usage steps given changelog-20260811-b3d33883 snapshot-20260811

24 **LangChain `ToolMessage.artifact` capture** NEW 55

The Python and TypeScript SDKs now capture LangChain `ToolMessage.artifact` data, including document lists, promoted to trace-level context.

— Names the exact field captured snapshot-20260811

25 **Per-user Gateway limits and guardrails** NEW 50

Usage limits and guardrails in the Gateway can now target an individual user, including per-user caps within a team.

— Names scope, no config surface shown changelog-20260811-b3d33883 snapshot-20260811

26 **Bulk session deletion with cascade** NEW 50

Multi-select sessions and delete them in bulk, with a full cascade to their traces, scores, and embeddings.

— Names cascade scope, no UI path given changelog-20260811-b3d33883 snapshot-20260811

27 **Bulk test actions and custom metric file management** NEW 50

Duplicate several tests at once and schedule runs for several tests at once; view, download, and delete uploaded custom metric files.

– Names actions, no UI path or endpoint given changelog-20260811-b3d33883 snapshot-20260811

28 **Timeline (Gantt-style) trace view** NEW 50

A new timeline view for traces shows every span's start time, duration, and nesting side by side.

– Describes layout, no navigation path given snapshot-20260811

29 **Webhooks for test-suite results** NEW 45

Subscribe a webhook endpoint to test-suite outcomes and get notified when tests pass or fail.

– No endpoint path or payload spec given changelog-20260811-b3d33883 snapshot-20260811

30 **Governance and compliance tracking additions** NEW 45

Projects can now have a lifecycle stage and evidence documents attached for governance tracking. Deactivating a rule can include a reason, which is displayed on the rule page. Results pages now show which tests are required for compliance, in context.

– Three thin governance additions, no mechanism depth snapshot-20260811

31 **Browser-based CLI login** IMPROVED 40

CLI login supports a device-code browser flow, letting users sign in without pasting an API key.

– Thin description, no command shown changelog-20260811-b3d33883 snapshot-20260811

32 **IBM watsonx Orchestrate integration** NEW 40

Adds a native integration to monitor IBM watsonx Orchestrate agents, modeled on the existing Agentforce and Copilot Studio integrations.

– Named integration only, no setup steps given changelog-20260811-b3d33883 snapshot-20260811

33 **NPI as a native PII check** NEW 40

US National Provider Identifier (NPI) is now a native PII type, detectable in Contains-PII tests and redactable in traces.

– Names the check and two use sites changelog-20260811-b3d33883 snapshot-20260811

34 **Helm chart for the Gateway** NEW 40

An official Helm chart is now available for deploying the Gateway on Kubernetes.

– Names the artifact, no chart repo or values shown changelog-20260811-b3d33883 snapshot-20260811

35 **Gateway API key invite links** NEW 40

Admins can generate links that let users self-provision their own Gateway API keys.

– Thin one-line description, no UI path given snapshot-20260811

36 **Missing-guardrails governance rule** NEW 35

A new integrity test flags traces that are missing an expected guardrail.

– Behaviour stated, no rule setup shown changelog-20260811-b3d33883 snapshot-20260811

37 **Pause and resume individual monitoring tests** NEW 35

Individual monitoring tests can be paused and resumed without disabling the whole project.

– Simple stated behaviour, no UI path given changelog-20260811-b3d33883 snapshot-20260811

38 **Workspace-level dashboards** NEW 35

At-a-glance cost, usage, and project health across the workspace, with drill-in widgets.

– Generic description, no navigation path given changelog-20260811-b3d33883 snapshot-20260811

39 **Suggested filter values** NEW 35

Filter fields suggest likely values as you type, served from the backend and respecting other active filters.

– Behaviour described, no field or endpoint named changelog-20260811-b3d33883

40 **Rename data sources from row menu** NEW 35

Data sources can now be renamed directly from the data source row menu.

– Names the UI control, thin description snapshot-20260811

41 **Content safety guardrail type** NEW 30

Adds a content safety guardrail type that can be applied to Gateway requests and responses.

– Bare name of the new guardrail type changelog-20260811-b3d33883 snapshot-20260811

42 **Row data download in development mode** NEW 25

Adds the ability to download row data while in development mode.

— Bare one-line description of the addition `snapshot-20260811`

└──▷ **BREAKING ON UPGRADE**

! configure() is deprecated in favor of init() for trace configuration in the SDK; settings now merge incrementally instead of resetting on each call.

WAS THIS USEFUL? 👍 👎

Semgrep ⓘ 2 RELEASES · SEEN 2026-08-11 CHANGELOG ↗

Semgrep is a static analysis tool that scans code for bugs, security vulnerabilities, and policy violations using customizable rules.

> Semgrep expanded its public policies API with new vocabulary for representing scan policy options, overhauled the Agentic Workflows log viewer for easier debugging, and published the unified policies migration timeline.

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

01 breaking_change and full_scan fields in policies API NEW

65

The public policies API vocabulary now includes `breaking_change` and `full_scan` fields, enabling scan policies that use those options to be represented through the API. The API remains in public beta for organizations using unified policies.

– Names exact API fields but no endpoint path given snapshot-20260811 changelog-20260811-dd95ddfd

THINNER COVERAGE BELOW

02 Structured step-log viewer for Agentic Workflows IMPROVED

55

Agentic Workflows step logs now use a structured viewer that parses run logs into individual searchable and filterable fields, replacing the previous single block of unstructured text.

– Describes mechanism but gives no UI path or command snapshot-20260811 changelog-20260811-dd95ddfd

03 Unified policies migration banner with schedule DEPRECATED

55

The unified policies migration banner now displays the migration schedule: migrations begin August 24, 2026, with the current Policies version sunsetting November 1, 2026.

– Gives concrete dates but no action steps beyond viewing banner snapshot-20260811

WAS THIS USEFUL? 👍 👎

Veracode

1 RELEASE · 2026-08-10 BLOG ↗

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

> Veracode's biggest integration update lets teams install, configure, uninstall, and schedule scans for the GitHub Workflow Integration directly from the Veracode Platform, while the Java API Wrapper gains new Findings, Sandbox, and paginated Applications endpoints, and Azure DevOps, GitLab, and Visual Studio integrations pick up smaller quality-of-life and deprecation changes.

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

01 Scan scheduling for GitHub and Azure DevOps integrations NEW

85

GitHub Workflow Integration v2.0.0 adds scan scheduling for Static Analysis, SCA, and IaC scans directly from the Veracode Platform for organizations using either the platform or GitHub Marketplace install path. Azure DevOps Workflow Integration v2.3.9 adds the ability to schedule scans for every organization where the integration is installed, driven by `veracode.yml` configuration, and v2.3.10 adds an `analysis_branch` key to `veracode.yml` so scheduled scans run against a specified branch instead of the default branch.

Target a specific branch for scheduled scans in an ADO workflow instead of the default branch.

```yaml
# veracode.yml
analysis_branch: main-release
```

— Names config key and scan types across two integrations with example.

Integration updates (2026-08-10)

02 Java API Wrapper endpoint coverage and reliability IMPROVED

80

Java API Wrapper v26.3.19.0 adds page size support to the Platform Applications API (`/appsec/v1/applications`) to improve efficiency when retrieving large application lists, and adds support for the paginated Sandbox REST API to retrieve sandbox lists using pagination. v26.2.18.0 adds support for the Platform Findings REST API to access open and mitigated application findings, and implements a configurable retry mechanism for the Platform Applications API (`appsec/v1/applications`) to handle transient failures and timeouts.

— Names four API endpoints and mechanisms, but no code sample.

Integration updates (2026-08-10)

03 **Platform-managed install, uninstall, and permissions for GitHub integration**

IMPROVED 75

GitHub Workflow Integration v2.0.0 enables installing, configuring, and updating credentials for the Veracode GitHub Integration directly from the Veracode Platform, removing the requirement for manual GitHub Marketplace steps. v2.1.1 adds a new uninstallation workflow accessible from the Veracode Platform to remove the app from GitHub organizations, and reduces the GitHub Veracode Workflow App's permissions by removing administration access, preventing the app from altering or removing installed repositories.

— Describes install/uninstall flow and permission scope reduction clearly.

Integration updates (2026-08-10)

04 **`debug` config key across workflow integrations** NEW

75

Adds a `debug` configuration option to `veracode.yml` across Azure DevOps Workflow Integration v2.3.6, GitHub Workflow Integration v1.0.6, and GitLab Workflow Integration v2.1.4, for debugging and troubleshooting scan pipelines.

Enable debug output in a veracode.yml-driven workflow (GitHub, GitLab, or ADO) to troubleshoot scan pipeline issues.

yaml

```yaml
# veracode.yml
debug: true
```

— Named config key with runnable example, but no output behaviour detail.

Integration updates (2026-08-10)

05 **VeracodeFlawImporter task replaces deprecated task** DEPRECATED

75

The `VeracodeFlawImporter` task replaces the deprecated Veracode Flaw Importer task (space-separated name) in Veracode Azure DevOps extension v3.35.0; both are currently supported, but the space-separated task is deprecated and scheduled for removal in a future release.

— Names both task identifiers and clear migration path. Integration updates (2026-08-10)

06 **Updated default runner image for GitHub Workflow Integration**

IMPROVED 70

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

— Exact image name and component versions given, no usage steps.

Integration updates (2026-08-10)

---

07 **SCA agent-based scan policy enforcement in Azure DevOps**

70

BREAKING

Azure DevOps Workflow Integration v2.3.3 enforces `Build Action: error` custom policy behavior for SCA agent-based scans, causing the scan to fail when vulnerabilities are found.

— Names policy behavior and effect; no migration guidance given.

Integration updates (2026-08-10)

---

08 **Visual Studio extension scan scope, analytics, and VS2026 support**

NEW

60

Veracode Scan for Visual Studio v0.8.0 adds a Scan Scope setting to choose between scanning the full project directory or only the active solution, and adds usage analytics providing insights into plugin adoption, feature usage, and scan activity. v0.7.2 adds support for Visual Studio 2026.

— Lists three additions with a UI setting but no exact navigation path.

Integration updates (2026-08-10)

---

09 **Jira Server/Data Center 9.x support ending** DEPRECATED

60

Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends November 30, 2026; customers must upgrade to Jira 10 or later to maintain a supported configuration.

— Exact end date and required upgrade path given, no further mechanism.

Integration updates (2026-08-10)

THINNER COVERAGE BELOW

---

10 **Autopackaging support for Apex, VB6, and Perl** NEW

40

GitHub Workflow Integration v2.0.0 adds autopackaging support for Apex, VB6, and Perl.

— Bare list of newly supported languages, no mechanism detail.

Integration updates (2026-08-10)

└──▷ BREAKING ON UPGRADE

! The `Veracode Flaw Importer` task (with spaces) in the Veracode Azure DevOps extension is deprecated and scheduled for removal in a future release; replace it with `VeracodeFlawImporter`.

! Support for the Veracode Integration for Jira on Jira Server and Jira Data Center 9.x ends November 30, 2026; upgrade to Jira 10 or later to maintain a supported configuration.

WAS THIS USEFUL?  

cve-lite-cli

1 RELEASE · 2026-08-11 NOTES ↗

Fast, developer-friendly JS/TS dependency vulnerability scanner with local lockfile scanning, OSV matching, direct vs transitive visibility, --fix, JSON output, and practical remediation guidance.

> cve-lite-cli v1.29.0 extends multi-folder usage filtering, surfaces cooldown warnings in SARIF for GitHub Code Scanning, deepens DM001 maintenance-risk detection with graduated severities, and adds scan-completeness reporting.

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

01 **Expanded DM001 maintenance-risk detection** IMPROVED 90

With `--check-maintenance`, deprecated npm packages now produce their own `medium -severity` finding, separate from constraint-drag findings which remain `high -severity`. Packages abandoned for 2+ years also show a `Last release` context line.

Fail CI on any high-severity constraint-drag finding (DM001) while also auditing override hygiene across all workspace packages.

```
$ cve-lite /path/to/monorepo --check-maintenance --check-overrides --fail-on high
```

— Names flag, severity split, context line, and a runnable command v1.29.0

02 **Multi-folder support for usage flags** IMPROVED 70

`--usage` / `--only-used` now work correctly in multi-folder scans: each subfolder gets its own source-usage annotation and filtering, isolated from other subfolders. Previously these flags were accepted but silently ignored outside single-folder mode.

— Names flags and before/after behavior, no example given v1.29.0

03 **Cooldown warnings surfaced in SARIF** IMPROVED 65

Release-cooldown warnings now appear in SARIF output as `properties.cooldownWarning` on the result, making them visible in GitHub Code Scanning.

— Names exact SARIF property but no direct command shown v1.29.0

THINNER COVERAGE BELOW

04 **Scan completeness reporting** NEW 45

Terminal output (compact and verbose modes) and HTML reports now render scan completeness, surfacing packument-fetch failures and other scan gaps.

— Describes what changed but no config or command to invoke it v1.29.0

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical ⓘ 1 RELEASE · 2026-08-10 NOTES ↗

Infisical is the open-source platform for secrets, certificates, and privileged access management.

> Infisical v0.162.19 extends its privileged access management (PAM) to non-human identities, allowing machine identities to be granted PAM permissions.

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

01 **Machine identity support in PAM** NEW 33

PAM access now supports machine identity (MI), enabling non-human identities to be granted privileged access management permissions rather than restricting PAM grants to human users.

— Describes capability but no config keys or steps given. v0.162.19

WAS THIS USEFUL? 👍 👎

GitGuardian SaaS

1 RELEASE · 2026-08-10 BLOG ↗

GitGuardian SaaS is a cloud-based security platform that detects and prevents secret leaks in Git repositories and source code.

> GitGuardian SaaS adds honeytoken-based endpoint intrusion detection and programmatic PAT/SAT token management, rolls out ten new secret detectors, and extends source health monitoring to more VCS providers.

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

01 **Honeytoken-based endpoint intrusion detection** NEW 90

The `ggshield honeytoken plant` command plants a decoy AWS credential on developer endpoints; any access to the decoy triggers an alert naming the compromised machine, its IP address, user agent, and timestamp. A new Honeytoken protection card in the Endpoints dashboard reports the share of the fleet with active honeytoken coverage and lists unprotected endpoints.

Plant a honeytoken on a developer workstation to detect credential-stealing malware the moment it reads the credential file.

```
$ ggshield honeytoken plant
```

Plant a honeytoken decoy credential on a developer machine — run once per endpoint via MDM job or manually to activate tripwire detection.

```
$ ggshield honeytoken plant
```

— Full mechanism and a runnable command are given.

Catch credential stealers on your endpo…

02 **Detector precision and recall improvements** IMPROVED 80

Improves precision for `github_oauth_token`, `generic_password`, `base64_generic_high_entropy_secret`, and `generic_high_entropy_secret` detectors so values that fail format validation are no longer reported. Expands `generic_password` to ignore passwords in public proxy-subscription configs (shadowsocks/vmess/vless/trojan/hysteria URIs, Telegram MTProto links, clash/mihomo node lists), and expands `generic_encryption_key` to ignore Talend encrypted passwords (`enc:system|routine.encryption.key.vN:...`). Improves recall for PostgreSQL credential detectors: usernames containing dashes (e.g. `workflow-prod-`

`user`) are now matched, and the username length bound is raised from 20 to 30 characters.

— Before/after values and exact detector names given. product docs

---

03 **Source health monitoring for more VCS providers** IMPROVED

65

Extends source health management to Bitbucket Cloud, Bitbucket Data Center, GitLab, and Azure DevOps: GitGuardian now pauses real-time ingestion and historical scans on unreachable sources, auto-resumes on recovery, and surfaces an actionable recovery step.

— Behaviour and scope named, but no UI path given. product docs

---

04 **Ten new secret detectors added** NEW

60

Adds the Lovable Access Token, Lovable Git Token, Slack Webhook Trigger URL, n8n API Key, Atlassian Auth Key, AWS Cognito Access Token, OpenAI Access Token, Vercel OIDC Token, and OpenCode API Key detectors, plus a Supabase Service Role JWT checker for Supabase `service_role` JWTs.

— Names all ten detectors but no usage detail. product docs

THINNER COVERAGE BELOW

---

05 **Programmatic PAT/SAT creation and audit logging** NEW

45

New API endpoints let you create Personal Access Tokens (PATs) and Service Account Tokens (SATs) programmatically, enabling automated token rotation workflows; tokens created via the API must have a finite TTL. Scans triggered via a Personal Access Token now emit an audit log event.

— No exact endpoint paths or methods given. Catch credential stealers on your endpo…

---

06 **Amadeus OAuth detector disabled** DEPRECATED

35

The Amadeus OAuth credentials detector and checker have been disabled as the Amadeus self-service API has been discontinued.

— States the removal and reason, no migration step. product docs

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The Amadeus OAuth credentials detector and checker have been disabled as the Amadeus self-service API has been discontinued.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

juice-shop

1 RELEASE · 2026-08-10 NOTES ↗

OWASP Juice Shop: Probably the most modern and sophisticated insecure web application

> Juice Shop v20.2.0 adds three new challenges, similarity-based cheat detection for challenge submissions, a user profile role field, and moves the frontend to Angular 22 while replacing Grunt with a custom build script.

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

01 **Build process switched from Grunt to custom script** BREAKING

60

Replaces Grunt with a custom packaging script (`scripts/package.mjs`) for the build process.

— Names exact replacement script for build pipeline. v20.2.0

THINNER COVERAGE BELOW

02 **Similarity-based cheat detection for challenge submissions** NEW

55

Adds cheat detection for source file overlap in challenge submissions using Dice Coefficient similarity scoring, and supports n-gram similarity checks for the System Prompt Extraction challenge.

— Names two concrete detection mechanisms but no user-facing action. v20.2.0

03 **`allowScripts` configuration key added** NEW

55

Adds `allowScripts` configuration key to `package.json` and `frontend/package.json` .

— Exact config key and files named, but purpose unexplained. v20.2.0

04 **Three new challenges added to challenge set** NEW

45

Adds new challenges Misplaced IaC Files , Login Cloud Admin (including a coding challenge), and Vulnerable Infrastructure (including a coding challenge) to the challenge set.

— Names three challenges but no mechanism or navigation detail. v20.2.0

05 **Role field and deluxe membership link on profile page** NEW

40

Adds a role field and deluxe membership link to the user profile page.

— Clear UI location but no further mechanism. v20.2.0

06 **Frontend upgraded to Angular 22** IMPROVED 40

Upgrades frontend to Angular 22.x and Angular Material 22.x.

— Names versions but no behavioural detail. v20.2.0

07 **'Deprecated Interface' challenge disabled by default on Docker/Heroku**

IMPROVED 40

Disables the dangerous Deprecated Interface challenge on Docker/Heroku by default.

— States the change but not how to re-enable it. v20.2.0

WAS THIS USEFUL? 👍 👎

Invicti ⓘ 1 RELEASE · 2026-08-06 BLOG ↗

Invicti is an automated web application security scanning platform that identifies vulnerabilities like SQL injection and XSS through dynamic and interactive testing.

> Invicti's platform release window centers on HashiCorp Vault-backed secrets management, a dedicated API scan profile, custom scan-speed throttling, gRPC multi-file proto upload, and LDAP user sync, alongside a long tail of smaller improvements to LSR reporting, vulnerability triage, reporting and installer options.

└──▷ **WHAT SHIPPED · 18 FEATURES** most completely described first ? what's the number?

01 HashiCorp Vault integration for Secrets Manager NEW

65

Invicti Platform's Secrets Manager now integrates with HashiCorp Vault, letting credentials be retrieved directly from Vault when configuring scan authentication instead of storing them directly in Invicti.

Pull scan authentication credentials from HashiCorp Vault instead of storing them directly in Invicti, reducing secret sprawl.

> 📍 In Invicti Platform, go to Secrets Manager, add a new secret, choose HashiCorp Vault as the provider, supply your Vault address and token, then reference that secret in the Authentication settings of your scan target.

— Names Vault and Secrets Manager, gives UI steps to wire it up product docs

02 Custom scan speed control IMPROVED

60

A new Custom scan speed option in scan configuration lets you set the number of concurrent requests and their send frequency for precise control over scan throughput; a fix ensures scans configured with a custom speed and throttling delay now correctly run at the configured request rate.

Throttle an API scan to a precise request rate to avoid overwhelming a rate-limited target during a scheduled scan.

> 📍 In Invicti Platform, open the scan configuration for your API target, select Custom scan speed, set the desired number of concurrent requests and the delay between request batches, then save and launch the scan.

— Names the exact option and mechanism, with a usage walkthrough product docs

THINNER COVERAGE BELOW

03 **Dedicated API scan profile and profile selection** NEW

55

A new dedicated API scan profile is available by default for all new API targets, providing comprehensive API vulnerability coverage out of the box; all scan profiles in an account are now selectable for API targets rather than being locked to OWASP Top 10 API.

– Explains before/after profile restriction but no exact config path product docs

04 **gRPC multi-file proto ZIP upload** NEW

55

Invicti now supports uploading multiple gRPC proto file definitions as a single ZIP archive for gRPC scanning, eliminating the need to import proto files one by one for services that span multiple files.

– Names the ZIP proto upload mechanism and its before/after benefit

Release 20260806 (2026-08-06)

05 **IDOR/BOLA authentication character limit increase** IMPROVED

55

The character limit for authentication values in IDOR/BOLA credentials was increased from `1,024` to `4,096` characters, resolving save errors for long JWT Bearer Tokens.

– Gives exact before/after limits and named credential type product docs

06 **LDAP user synchronization** NEW

50

Invicti Platform now supports LDAP user synchronization from OpenLDAP or Microsoft Active Directory for automated user provisioning and team membership management.

– Names the directory sources but no setup steps given product docs

07 **Login Sequence Recorder secrets access and HTTP reporting**

IMPROVED

50

Secrets configured for a target are now accessible inside the Login Sequence Recorder (LSR) during recording, enabling stored credentials instead of manual entry; LSR reports now also include the full HTTP request and response for each recorded step, with automatic obfuscation of sensitive tokens in headers and bodies.

– Names LSR and obfuscation behavior but no walkthrough product docs

08 **Website Discovery CSV export** NEW

50

Website Discovery results can now be exported to a CSV file for portable tracking and reporting outside Invicti Platform.

Export newly discovered websites to CSV after a Website Discovery run for offline triage and asset-tracking.

> 📍 In Invicti Platform, open Website Discovery results and click `Export to CSV` to download the full list of discovered sites as a portable file.

— Names the CSV export and gives exact UI steps to use it `product docs`

---

09 **Vulnerability triage: state persistence and severity filter** IMPROVED 45

Vulnerabilities marked as Ignored or False Positive are now carried forward to subsequent DAST scans, preserving assigned states across scan runs; the severity filter on the Vulnerabilities page is now also available in the expanded (individual vulnerability) view, not only the by-type view.

— Names both changes but no exact navigation steps `product docs`

---

10 **Azure Pipelines plugin-based CI/CD setup** IMPROVED 45

Azure Pipelines CI/CD integration setup now includes a direct link to the Invicti plugin in the Azure Marketplace, adding a plugin-based option alongside the existing script-based method.

— Names Azure Marketplace plugin as a clear starting point `product docs`

---

11 **'Last' position for ordinal monthly scan schedules** NEW 45

The `Last` position option is now available when configuring ordinal monthly scan schedules, e.g. last Monday or last Friday of the month.

— Names the exact scheduling option and its use case `product docs`

---

12 **Client certificate storage in Secrets Manager** NEW 40

Secrets Manager now supports storing client certificates by supplying the certificate in base64 format.

— Thin one-line description with no usage steps `product docs`

---

13 **HTTP response data in Comprehensive and SOX reports** IMPROVED 40

Comprehensive and Sarbanes-Oxley reports now include HTTP response data alongside request data.

— Names the two report types but stays brief `product docs`

---

14 **Offline activation for Windows on-premises installer** NEW 35

Offline activation for Windows on-premises installations is now supported directly from the Windows installer, requiring no internet access during setup.

— Names the installer capability but no configuration steps `product docs`

15 **Sensorless API Discovery out-of-scope visibility** IMPROVED

35

Sensorless API Discovery now surfaces APIs found on out-of-scope sites, broadening API attack-surface visibility beyond scanned targets.

— Names the feature and scope change but no usage path `product docs`

16 **AI Business Logic Recorder endpoint marking** IMPROVED

35

Endpoints identified by the AI Business Logic Recorder (BLR) are now marked in the site structure view for visibility into AI-assisted recording coverage.

— Names the feature but description stays brief `product docs`

17 **Advanced Target Settings visibility for non-owners** IMPROVED

35

The Advanced Target Settings section is now visible to all users with appropriate permissions, not only account Owners.

— States the permission change without further mechanism `product docs`

18 **Resizable list-view columns** IMPROVED

20

Columns across all list views in Invicti Platform are now resizable.

— Bare one-line UI change with no further detail `product docs`

WAS THIS USEFUL? 👍 👎

Bright Security

1 RELEASE · 2026-08-11 NOTES ↗

Bright Security is an application security platform that automates vulnerability detection and remediation through continuous scanning and testing throughout the development lifecycle.

> Bright Security's August 11, 2026 release adds an EXIF metadata leakage detection test, a repeater infrastructure dashboard widget, per-project authentication management, and new project classification and discovery tools.

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

01 **EXIF Metadata Leakage security test** NEW 78

A new EXIF Metadata Leakage test uploads canary images containing GPS coordinates, device identifiers, timestamps, and comment fields to file upload endpoints, then checks whether served copies leak this data across upload responses, redirects, API responses, rendered HTML, and linked image URLs. It reports Medium severity findings under `CWE-1230`, covers JPEG and PNG formats, and supports multipart and JSON-encoded uploads.

— Detailed mechanism and named CWE, formats, upload types. Release of August 11, 2026

02 **Repeater Insights widget on Organization Dashboard** NEW 60

A new Repeater Insights widget on the Organization Dashboard shows connected, active, idle, disconnected, and disabled repeaters alongside running, queued, and scheduled Scan and Discovery tasks across selected projects.

— Names widget and all status categories; UI path given. Release of August 11, 2026

03 **Project-scoped Authentications tab** NEW 60

A new Authentications tab on the Project page lets teams view, create, and edit project-scoped authentications, with automatic assignment to the current project on creation, without leaving the project context.

— Describes behaviour and exact UI location. Release of August 11, 2026

04 **Project classification and host-based search** NEW 60

New project classification attributes (Environment, Business Criticality, Exposure) are available for filtering on the Projects page and Organization Dashboard to prioritize security efforts. Host-based project search is now available in the Projects table and the project selector in the Organization Dashboard.

— Names all attributes and search locations, but no mechanism detail.

Release of August 11, 2026

THINNER COVERAGE BELOW

05 **Skip reason display in Entry Points table** IMPROVED 55

The Entry Points table now displays the skip reason when an entry point is omitted due to a missing Content-Type header, with additional tooltip guidance and the ability to sort entries by skip reason.

— Names table, cause, and new sort capability. Release of August 11, 2026

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

Velocidex Velociraptor ⓘ 2 RELEASES · 2026-08-10 NOTES ↗

CODE ↗

Velociraptor is an endpoint visibility and collection tool that gathers host-based state information using the Velociraptor Query Language (VQL).

> Velociraptor's v0.77.2 and v0.76.7 releases add VQL default parameter values, a new full-text index() plugin, programmatic hunt tagging, and a reworked OIDC role-assignment system with new config keys, while removing support for old client APIs and dropping the `--config_override` / `-o` CLI flag.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 OIDC role assignment rework and auth controls IMPROVED

90

Reworks automated OIDC role assignment, adds support for alphanumeric-only e-mail barrier authentication, and introduces `GUI.authenticator.claims.allow_unverified_email` and `GUI.authenticator.claims.override_acls` config keys for finer control over OIDC-based role assignment.

Enable OIDC login to accept users whose e-mail address is not yet verified by the identity provider.

yaml

```yaml
GUI:
  authenticator:
    claims:
      allow_unverified_email: true
```

— Named config keys with a runnable YAML example. v0.77.2 v0.76.7

02 Default parameter values for VQL functions NEW

85

VQL functions and stored queries can now define default parameter values, e.g. `LET F(X=2, Y=3) = X + Y`; the artifact verifier flags calls that omit required parameters.

Define a reusable VQL helper with default parameter values to simplify repeated queries.

```
$ LET Score(Severity='medium', Multiplier=2) = Severity + ' x' +
str(str=Multiplier)
SELECT Score(), Score(Severity='high'), Score(Multiplier=5) FROM
scope()
```

— Runnable VQL example and verifier behaviour both given. v0.77.2 v0.76.7

03 **dir_permissions parameter in copy()** NEW 65

Adds a `dir_permissions` parameter to the copy() VQL function for controlling directory permissions during file operations.

— Named flag and function but no example call shown. v0.77.2

04 **Tags argument in hunt() VQL function** NEW 65

Adds a Tags argument to the hunt() VQL function, enabling programmatic tagging of hunts at creation time.

— Named argument, no usage example provided. v0.77.2 v0.76.7

THINNER COVERAGE BELOW

05 **Full-text index() VQL plugin** NEW 55

Adds the index() VQL plugin to build local full-text indexes over collected data, including a batch mode for faster indexing.

— Named plugin and mode but no usage example given. v0.77.2 v0.76.7

06 **ZIP import via --definitions flag** NEW 55

Allows a ZIP file to be imported using the `--definitions` flag.

— Named flag gives a clear starting point. v0.77.2 v0.76.7

07 **Disk-based LRU for cache() and memoize()** IMPROVED 55

Updates the cache() and memoize() plugins to use disk-based LRU storage, enabling larger datasets to be cached without exhausting memory.

— Explains mechanism and benefit but no usage steps. v0.77.2 v0.76.7

08 **Shell and GUI session handling rework** IMPROVED 55

Reworks the shell interface and GUI to utilize sessions, with graceful timeout for shell sessions, and stores GUI state in `sessionStorage` for improved session continuity.

— Mechanism described but no config or path given. v0.77.2

09 **API support for artifacts collect command** IMPROVED 50

Adds API support to the `artifacts collect` command-line run mode.

– Named CLI mode but no detail on what API covers. v0.77.2 v0.76.7

10 **ADX flow-upload artifact** NEW 50

Adds a new flow-upload artifact for Azure Data Explorer (ADX).

– Named integration target, no config details. v0.77.2

11 **Filters on users list in management screen** NEW 35

Adds filters on the users list in the user management screen.

– Names UI location but no filter types listed. v0.77.2 v0.76.7

12 **Initial Velociraptor container implementation** NEW 35

Adds initial implementation of a Velociraptor container.

– No image name, registry, or usage instructions given. v0.77.2

13 **Removal of --config_override/-o flag** BREAKING 35

The `--config_override` / `-o` flag has been removed from the CLI.

– Named flag removal but no replacement noted. v0.76.7

14 **Elastic client library updated to 9.3** IMPROVED 30

Updates the Elastic client library to version 9.3.

– Version named but no behavioural change described. v0.77.2

15 **Message users directly from GUI** NEW 25

Adds the ability to message users directly from the GUI.

– Bare description with no UI path or mechanism. v0.77.2 v0.76.7

16 **Removal of old client API support** BREAKING 25

Support for old client APIs has been removed.

– No migration guidance or affected endpoints named. v0.77.2 v0.76.7

└──▷ BREAKING ON UPGRADE

! Support for old client APIs has been removed (Remove support for old client APIs).

! The `--config_override` / `-o` flag has been removed from the CLI.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

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

A fast, easy-to-use AF_XDP traffic generator for Linux

> Wireblast v0.4.0 adds a configurable memory budget for PCAP replay loads.

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

01 **Configurable memory budget for PCAP replay** NEW 60

Adds the `--pcap-memory` flag to make the memory budget for loading PCAP files configurable at runtime.

— Names exact flag but no default or limit values given v0.4.0

WAS THIS USEFUL? 👍 👎

◆ Firewalls & Network Defense

Sophos Central Firewall ⓘ

2026.35Released on August 11, 2026 NOTES ↗

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

Sophos Central Firewall Management gains recurring firmware upgrade schedules with weekly, monthly, and staged rollout options.

› Adds recurring firmware upgrade schedules in Sophos Central Firewall Management, supporting weekly or monthly cadences and staging firewalls to validate upgrades before broader rollout.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion ⓘ 1 RELEASE · SEEN 2026-08-11 NOTES ↗

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

> Plerion's AI agent Pleri can now remediate IaC issues directly in Azure DevOps by opening, reviewing, and abandoning pull requests.

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

01 **Pleri pull request remediation in Azure DevOps** NEW

50

Pleri can now raise pull requests in Azure DevOps repositories to remediate IaC issues, and can check pull request status, request a review, or abandon a pull request with a reason.

— Names actions and platform but no exact command or endpoint.

Pleri now raises pull requests in Azure…

WAS THIS USEFUL? 👍 👎

Rapid7 InsightCloudSec

1 RELEASE · 2026-08-11 DOCS ↗

BLOG ↗

Rapid7 InsightCloudSec is a cloud security platform that provides visibility, compliance monitoring, and threat detection across multi-cloud environments.

> Rapid7 InsightCloudSec's August 2026 release adds Azure Data Factory harvesting and query filters, four new CIS compliance packs, a batch of new misconfiguration Insights, and a paginated Insights Library for faster loading.

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

01 **Azure Data Factory harvesting and query filters** NEW

82

Adds harvesting support for Managed Identity configurations (system-assigned and user-assigned) and linked services credential storage method for Azure Data Factory resources. Adds the `Data Factory with/without Managed Identity` query filter (defaults to `without`) and the `Data Factory Linked Services Use Key Vault` query filter (defaults to identifying inline storage), plus the corresponding Insights `Data Factory without Customer Managed Key Encryption`, `Data Factory without Managed Identity`, and `Data Factory Linked Services Not Using Key Vault for Credentials`.

— Names exact filters, defaults and Insights but no CLI/API path 2026-08-11-18980834

02 **Four new CIS compliance packs** NEW

75

Adds the `CIS Amazon Elastic Kubernetes Service (EKS) Benchmark v2.0.0` for EKS clusters, `CIS Google Kubernetes Engine (GKE) Benchmark v1.9.0` for GKE clusters, `CIS Red Hat OpenShift Container Platform Benchmark v1.9.0` for OpenShift clusters, and `CIS AWS Compute Services Benchmark v1.1.0` covering EC2, EBS, ECS, Lambda, Lightsail, Elastic Beanstalk, Fargate, Batch, App Runner, and SimSpace Weaver, with 47 Insights mapped across AWS, AWS GovCloud, and AWS China regions.

— Named benchmarks and scope, but only a UI destination 2026-08-11-18980834

03 **New encryption and access-control Insights** NEW 70

New Insights: `PostgreSQL Database Instance without Customer Managed Key Encryption` identifies PostgreSQL instances not using Customer Managed Keys; `Aurora Database Cluster Encryption at Rest Disabled` identifies Aurora clusters without encryption at rest; `DocumentDB Database Cluster Encryption at Rest Disabled` identifies DocumentDB clusters without encryption at rest; and `Cache Instance SSL Access Not Enforced` identifies Azure Cache for Redis instances not enforcing SSL-only access, aligned with CIS Azure Database Services Benchmark Control 2.2.

— Each Insight named with target resource, no usage path given 2026-08-11-18980834

THINNER COVERAGE BELOW

04 **Paginated Cloud Security Insights Library** IMPROVED 50

Paginates the Cloud Security Insights Library by default, with a toggle to revert to the classic view, improving load times for large libraries.

— Describes behaviour change but no exact UI location given 2026-08-11-18980834

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Anchore Grype ⓘ 1 RELEASE · 2026-08-10 NOTES ↗ CODE ↗

A vulnerability scanner for container images and filesystems

> Grype v0.117.0 enriches CycloneDX SBOM output with vulnerable version ranges and restores CPE-based matching for Rust crates.

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

01 **Rust CPE-based vulnerability matching fix** IMPROVED 75

Fixes CPE matching for Rust crates, controlled via `match.rust.using-cpes` (or the `GRYPE_MATCH_RUST_USING_CPES` environment variable), so vulnerabilities reachable only via CPE, such as the openssl rust-crate, are no longer silently missed.

Enable CPE-based matching for Rust crates when scanning an image, so vulnerabilities reachable only via CPE (e.g. openssl rust-crate) are no longer silently missed.

```
$ GRYPE_MATCH_RUST_USING_CPES=true grype alpine:latest
```

— Names config key, env var and a runnable example command v0.117.0

THINNER COVERAGE BELOW

02 **Vulnerable version ranges in CycloneDX output** IMPROVED 40

CycloneDX SBOM output now includes vulnerable version ranges, giving richer remediation context in CycloneDX SBOMs.

— Names format and change but no mechanism detail v0.117.0

WAS THIS USEFUL? 👍 👎

Anchore syft ⓘ 1 RELEASE · 2026-08-10 NOTES ↗

CLI tool and library for generating a Software Bill of Materials from container images and filesystems

> Syft v1.51.0 adds detection of native Go FIPS 140 mode in scanned binaries and expands Go binary symbol capture with a new stdlib scope.

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

01 **Extended stdlib scope for Go symbol capture** NEW 52

Adds an `extended-stdlib` scope and expanded module patterns for symbol capture, improving coverage of Go binary analysis.

— Names the `extended-stdlib` scope but no usage command shown. v1.51.0

02 **Go binary FIPS 140 mode detection** NEW 42

Syft now detects native Go FIPS 140 mode in scanned binaries, surfacing FIPS compliance status directly in generated SBOMs.

— Describes what it detects but no flag or output field named. v1.51.0

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

NVIDIA OpenShell ⓘ 1 RELEASE · 2026-08-11 NOTES ↗

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

> OpenShell v0.0.103 adds the ability to bind static credentials to proxy provider endpoints, enabling credential injection for specific backend routes without dynamic resolution.

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

01 **Static credential binding for proxy provider endpoints** NEW

45

OpenShell's proxy can now bind static credentials to provider endpoints, enabling credential injection for specific backend routes without dynamic resolution.

— Explains mechanism but no config keys or commands given. v0.0.103

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

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 CLI Scanner v1.29.0 now evaluates base-image supply chain rules locally within vulnerability policies.

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

01 **Allow/Deny Base Image rules in CLI Scanner** NEW 62

Sysdig CLI Scanner v1.29.0 adds support for Allow Base Image and Deny Base Image supply chain rules in vulnerability policies, so policies that allow or deny specific base images are now evaluated locally by the CLI Scanner.

— Names exact rule types and scanner version but no usage steps.

Secure SaaS - August 11, 2026

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Trend Vision One

4 RELEASES · 2026-08-06 → 2026-08-12 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 added AskAI-driven quarantine management in the Outlook add-in, a new EXCEPT operator for DLP template exceptions, direct IoC search with bulk validation in the Threat Intelligence Hub, and agentless vulnerability/malware scanning for GCP Cloud Run Functions.

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

01 **IoC search in Threat Intelligence Hub** NEW 86

Analysts can enter a file hash (SHA-1 or SHA-256), IP address, domain, or URL at Threat Intelligence → Threat Intelligence Hub → Indicators to check whether Trend Micro threat intelligence has flagged it as malicious, linked it to a detection, or associated it with known threat campaigns or intelligence reports. Bulk IoC validation via list upload lets analysts validate multiple suspicious objects in a single operation instead of one at a time.

Quickly determine whether a suspicious file hash encountered during an investigation is tied to a known threat campaign.

> 📍 In the Trend Vision One console, go to Threat Intelligence › Threat Intelligence Hub › Indicators, enter the SHA-256 hash of the suspicious file, and review the malicious verdict, associated detections, and linked threat campaigns.

Bulk-validate a list of IoCs (hashes, IPs, domains, URLs) collected from a recent incident to triage which ones have known threat intelligence coverage.

> 📍 In the Trend Vision One console, go to Threat Intelligence › Threat Intelligence Hub › Indicators, upload your IoC list file, and review results showing malicious determinations, detection coverage, and associated threat campaigns for each entry.

— Exact navigation path, hash types, and bulk workflow with worked examples.

Search Threat Intelligence by IoC (2026…

THINNER COVERAGE BELOW

02 **GCP Cloud Run Functions support in Agentless Vulnerability and Threat Detection**

NEW 53

Agentless Vulnerability and Threat Detection extends vulnerability and malware scanning to GCP Cloud Run Functions, accessible via Cyber Risk Exposure Management → Continuous Risk Management → Threat and Exposure Management.

— Names the platform, scan type, and exact navigation path.

Agentless Vulnerability and Threat Dete…

03 **EXCEPT operator for DLP template exceptions** NEW 50

DLP templates now support an `EXCEPT` operator, allowing administrators to exclude specific expressions or keywords from a match condition alongside existing `AND` and `OR` logic, reducing false positives.

— Named operator and logic it combines with, but no config path shown.

DLP templates now support the EXCEPT op…

04 **AskAI quarantine search and management in Outlook add-in** NEW 47

AskAI in the Outlook add-in now lets end users search quarantined email by sender or subject directly from the chat interface, and enforces the permissions already assigned by administrators when performing quarantine actions, enabling end-user self-service while keeping security teams in control.

— Describes behaviour and permission model but no exact UI steps.

AskAI quarantine search and management …

WAS THIS USEFUL? 👍 👎

Sophos Central

2 RELEASES · SEEN 2026-07-20 → 2026-08-11 NOTES ↗

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

> Sophos Central's biggest addition this window is an AI Defense Early Access Program plus ML-powered Identity Risk Scores in ITDR, alongside a new Network Map visualization for firewalls, switches and AP6 access points and a lightweight Protected Browser extension for Chrome and Edge.

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

01 **Network Map topology visualization** NEW 62

A new Network Map provides a visual topology of all Sophos firewalls, switches, and AP6 access points in a Sophos Central account, including device relationship views and search capabilities.

— Names device types and views but only a UI path. July 20, 2026

02 **Identity Risk Scores in Sophos ITDR** NEW 60

Sophos ITDR now computes a single ML-powered Identity Risk Score per identity, combining incident likelihood with potential impact, and the score updates dynamically as findings and accounts change.

— Explains mechanism and behaviour but no UI path given. Latest news

THINNER COVERAGE BELOW

03 **Live Discover queries for AP6 access points** IMPROVED 55

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

— Names query tool and device series, gives a starting point. July 20, 2026

04 **Protected Browser extension for Chrome and Edge** NEW 55

A new Protected Browser extension lets administrators apply security policies and monitor user activity in Chrome and Edge without migrating users to the full Sophos Protected Browser.

— Names browsers and migration-free approach, but no install steps. July 20, 2026

05 **Generative AI category in DNS Protection** NEW 48

DNS Protection filtering policies gain a new Generative AI category, giving visibility into and control over AI services used on monitored networks.

— Names the exact filtering category and where to apply it. Latest news

06 **Sophos AI Defense Early Access Program** NEW 46

Sophos launched the Sophos AI Defense Early Access Program for Sophos EDR, XDR, and MDR customers on Sophos Fusion, enabling monitoring, control, and securing of AI usage across an environment.

— Names program and eligible products but no enrollment steps. Latest news

07 **Bulk license selection for partner add-ons** IMPROVED 45

Bulk license management for partners is improved: monthly-usage customer license changes now support explicit selection of add-ons and products when applied in bulk.

— Describes the change but not where to perform it. Latest news

08 **ITDR expanded to five new regions** IMPROVED 35

Sophos ITDR is now available to customers with tenants provisioned in Canada, Australia, Japan, India, or Brazil.

— Lists regions but no configuration or activation detail. Latest news

WAS THIS USEFUL? 👍 👎

Sophos Protection for Linux ⓘ 2 RELEASES · SEEN 2026-08-11

NOTES ↗

Sophos Protection for Linux is an endpoint security tool that provides threat detection, malware prevention, and vulnerability management for Linux systems.

> Sophos Protection for Linux expanded its Taegis data lake export with DNS query events and SNI connection details, and added kernel 7.1 support in this window.

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

01 **DNS query events and SNI details in Taegis data lake** NEW

68

Adds the `dns_query_events` event type to the Taegis data lake hydration pipeline, enabled via Helm charts configuration, and adds SNI (Server Name Indication) details to connection events in the data lake.

— Names event type and config mechanism but no exact keys. Version 5.18.2

THINNER COVERAGE BELOW

02 **Kernel 7.1 support** NEW

30

Adds support for kernel 7.1, extending the range of Linux environments where the sensor can run.

— Bare statement of new kernel support, no further detail. Version 5.19.0New

WAS THIS USEFUL? 👍 👎

Huntress

2 RELEASES · SEEN 2026-08-11 NOTES ↗

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

> Huntress shipped GA Monthly Executive Summary reports for stakeholder-friendly reporting and a new GA API for querying remote access connections across an environment.

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

01 **Remote access connections API (GA)** NEW 75

A new Remote Access section in the Huntress API lets you programmatically query all remote access connections observed across your environment, including ConnectWise ScreenConnect hosts, returning agent hostname, relay host, port, and first/last seen timestamps — data previously confined to the portal's ScreenConnect Hosts view.

— Names API area and returned fields, but no exact endpoint path.

Remote access connections API is now av…

02 **Monthly Executive Summary reports (GA)** NEW 65

A one-page, plain-language snapshot per organization highlighting active protection layers (EDR, ITDR, SIEM, etc.), signals investigated, and threats contained, aimed at non-technical stakeholders. Accessible under the Reports section of the Huntress portal and delivered via email link to partners for forwarding to end customers; it complements rather than replaces the existing full monthly/quarterly Threat Report and is described as the first in a planned series of reporting improvements.

— Names UI location and purpose but no config detail.

Monthly Executive Summary reports are n…

WAS THIS USEFUL? 👍 👎

◆ SIEM

Sumo Logic ⓘ 1 RELEASE · 2026-08-11 BLOG ↗

Sumo Logic is a cloud-based monitoring and analytics platform for logs, metrics, and security data.

> Sumo Logic added a Shift+click drill-down on dashboard map charts that jumps straight to a pre-filtered Log Search.

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

01 **Drill-down from map charts to Log Search** NEW 70

Shift+click on a point or line in a dashboard map chart drills down into that location's data by redirecting to the Log Search page with the query auto-populated with the corresponding latitude and longitude values.

Investigate suspicious geo-located traffic without manually hunting for coordinates — jump straight from a map chart to a filtered log search.

> 📍 In a dashboard map chart, Shift+click a point or line on the map. Sumo Logic opens the Log Search page with the query pre-populated with the latitude and longitude for that location.

— Names exact gesture and resulting behavior, but narrow scope

Drill Down from Map Charts (Dashboards)…

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Elastic Detection Rules ⓘ 1 RELEASE · 2026-08-10 NOTES ↗

Detection Rules are customizable security conditions that automatically identify and alert on suspicious activities and threats within your environment.

> Detection Rules dev-v2.1.0 ships over 40 new detection rules spanning AWS, Azure AKS, GCP GKE, NFS/SIP/AMQP protocols, and database systems, alongside nested KQL query support for rule authoring and an XDR rule tag taxonomy for the Copilot PR review workflow.

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

01 **AWS cloud detection rules** NEW 61

Seven new AWS-focused detection rules: `AWS Bedrock High Risk Filesystem or Execution Tool Invocation`, `AWS Batch Job Submitted with Container Override by Unusual Identity`, `AWS IAM Permission Boundary or Guardrail Policy Deleted by Unusual Identity`, `AWS IAM User Created Own Access Key`, `AWS SageMaker Execution Role Passed by Unusual Principal`, `AWS S3 Bucket ACL Modified to Allow Public Access by New Identity`, and `AWS EC2 NACL Entry Created or Replaced Allowing All Traffic by New Identity`.

— Names every rule but gives no detection logic. dev-v2.1.0

02 **Azure AKS detection rules** NEW 61

Nine new Azure AKS detection rules covering cluster and Kubernetes API abuse: `Azure AKS Service Account Token Created via TokenRequest API`, `Azure AKS Certificate Signing Request Created or Approved`, `Azure AKS CoreDNS or Kube-DNS Configuration Modified`, `Azure AKS Secret get or list with Suspicious User Agent`, `Azure AKS Kubernetes Events Deleted`, `Azure AKS Potential API Enumeration by User`, `Azure AKS Suspicious Self-Subject Review via Service Account`, `Azure AKS Ephemeral Container Added to Pod`, and `Azure AKS User Exec into Pod`.

— Enumerates all rules; no per-rule logic given. dev-v2.1.0

THINNER COVERAGE BELOW

03 **Network protocol abuse detection rules** NEW 55

New rules for NFS, SIP, and AMQP protocol abuse: `Potential NFS Destructive Operation Burst`, `NFS AUTH_SYS Root UID Access`, `Potential SIP REGISTER Brute Force`, `Potential SIP Extension Enumeration`, and `Successful AMQP Multi-Queue Purge Burst`.

— Names five distinct protocol rules, no detection thresholds. dev-v2.1.0

04 **Database protocol abuse detection rules** NEW 55

New rules targeting database and RPC protocols: `First Time Seen Memcached Writer`, `Destructive MongoDB Command`, `MySQL User-Defined Function Injection`, `Postgresql COPY PROGRAM Command Execution`, `Cassandra JavaScript UDF Creation`, and `Thrift RPC Method from an External Client`.

— Names six database rules, no execution logic detailed. `dev-v2.1.0`

---

05 **Endpoint and identity evasion detection rules** NEW 54

New rules covering host evasion, credential abuse, and web-facing tampering: `Newly Observed RC4 Kerberos Service Ticket Request`, `Repeated Stalled TLS Handshakes via ALPN acme-tls/1 Extension`, `First Time Seen RMM Signer Across the Environment`, `Potential EDR-Freeze via WerFaultSecure Abuse`, `Suspicious UID Change to Root via Python`, `Direct Process Execution via setsid, nohup, or disown`, `Mark-of-the-Web Removal by an Unusual Process`, `Multiple SonicWall Login Failures Followed by Successful Login`, `Unusual File Creation via Web Server`, `PHP File Creation in WordPress Plugin Directory`, and `Potential Tunneling via Tailscaled`.

— Lists eleven named rules across varied techniques, no thresholds. `dev-v2.1.0`

---

06 **GCP GKE detection rule expansion** NEW 47

Expands GKE detection coverage with `GCP GKE Rule Conversions Part 3, 4, 5, 6` and `GCP GKE CSR Abuse Rules`.

— Names rule sets but not individual conditions. `dev-v2.1.0`

---

07 **Nested KQL query support for rule authoring** NEW 43

Adds nested KQL query support in `lib/kql`, enabling more expressive detection rule authoring.

— Names the module but not the new syntax mechanics.' `dev-v2.1.0`

---

08 **XDR rule tag taxonomy in Copilot PR review** IMPROVED 38

Adds an XDR rule tag taxonomy to the Copilot PR review workflow, improving automated rule classification.

— States purpose but no tag list or workflow mechanics. `dev-v2.1.0`

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Feedly Threat Intelligence ⓘ 1 RELEASE · 2026-08-11

CHANGELOG ↗

Feedly Threat Intelligence aggregates security feeds and threat data to help organizations monitor emerging cybersecurity threats and vulnerabilities.

> Feedly Threat Intelligence brought AI-driven Analyze and Research actions to Custom Intel Agents, added Censys lookups to IP Insight Cards, and gave Report Builder a generation progress indicator.

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

01 **Censys lookup in IP Insight Cards** NEW 40

IP Insight Cards gain a Censys lookup, surfacing Censys data directly in the IP context panel.

— Names integration and location but no detail on data shown changelog-20260811-f3605320

02 **Ask AI Analyze and Research in Custom Intel Agents** NEW 35

Custom Intel Agents now support Ask AI Analyze and Research actions, enabling AI-driven analysis and research workflows within agent configurations.

— Names the actions but not underlying mechanism changelog-20260811-f3605320

03 **Progress indicator in Report Builder** IMPROVED 20

Report Builder adds a progress indicator during generation, making long-running report creation more transparent.

— Minor UI change described briefly changelog-20260811-f3605320

WAS THIS USEFUL? 👍 👎

Flashpoint

1 RELEASE · 2026-08-11 CHANGELOG ↗

Flashpoint is a threat intelligence platform that aggregates data from darknet markets, forums, and other sources to help organizations identify security risks.

> Flashpoint's Echosec geospatial OSINT tool gains Natural Language Search, letting analysts query geospatial data sources using plain-language input.

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

01 **Natural Language Search in Echosec** NEW 40

Echosec, Flashpoint's geospatial open-source intelligence tool, adds Natural Language Search, enabling query-by-plain-language across geospatial data sources for AI-powered discovery.

— Names the feature and product but no mechanism detail. changelog-20260811-f901e8ae

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

1 RELEASE · 2026-08-11 CHANGELOG ↗

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

> Tines 42.6.0 brings a self-hosted LDAP configuration UI, a new Plan mode for Apps, and several usability improvements to dashboards and Cases.

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

01 **Self-hosted LDAP configuration UI** NEW 45

Adds a self-hosted LDAP configuration UI, enabling on-premises deployments to configure LDAP directly from the Tines interface.

– Names the surface but no config fields or steps changelog-20260811-5bd77e0b

02 **Plan mode for Apps** NEW 35

Introduces Plan mode for Apps, giving practitioners a way to model or preview app behavior before execution.

– Describes purpose but not mechanism or UI path changelog-20260811-5bd77e0b

03 **Dashboard time granularity controls** IMPROVED 35

Adds time granularity controls to dashboard charts, allowing finer or coarser resolution when visualizing workflow data.

– General description, no specific settings named changelog-20260811-5bd77e0b

04 **Time zone support in Cases and dashboards** NEW 35

Enables setting a working time zone in Cases Records and dashboards for consistent time-aware case management.

– Names feature area but no config detail changelog-20260811-5bd77e0b

05 **Improved group change reviews** IMPROVED 20

Improves group change reviews, enhancing visibility and control over permission group modifications.

– Vague description, no specifics on mechanism changelog-20260811-5bd77e0b

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Censys

1 RELEASE · 2026-08-10 BLOG ↗

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

> Censys's biggest release this window is the beta launch of AI Investigations, an autonomous agent that pivots across indicators to map related infrastructure and produce ranked reports; alongside it Censys deepened HTTP body search, added bulk tag management to its CLI, and shipped nine new service fingerprints.

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

01 **AI Investigations for infrastructure pivoting** NEW 85

New AI Investigations (beta, requires Adversary Investigation module access) in the Censys Platform web console autonomously analyzes IP addresses, domain names, host/port pairs, URLs, and pasted intel reports to discover related infrastructure in the Censys Internet Map. Prompts accept historical time context — single dates, month/quarter/half-year/year ranges, or explicit closed ranges like '20 January 2025 through 5 February 2025' — to scope the search; the agent pivots on shared parsed data attributes, interactively prompts the analyst when cluster choices arise, checks entity history, profiles web resources in parallel, and delivers a CenQL query targeting the discovered asset cluster for immediate use in the Platform. The Evidence tab gives a full audit trail of every tool call, decision tree, and pivot with per-entry detail views (e.g. `web.hostname: bancaporinternetbbmx.online`) and downloadable artifacts, while the Reports tab can produce up to three files: `narrative_llm` (full narrative with pivot visualizations and recommended actions), `intelligence_report.md` (Mermaid-diagram report with cluster table, CenQL queries, Platform links, and invalid IOCs as a separate section), and `body_intelligence.md` (HTTP response body analysis covering page titles, status codes, server banners, TLS certificate SANs, lure-pattern matches, and external `script-src` URLs). Closed beta limits: three concurrent investigations per user, 150 per organization per week, and no Censys Credits consumed.

Kick off an infrastructure pivot investigation from a threat-intel report drop — paste the raw indicator list and a date anchor so the agent scopes its search to the relevant activity window.

📍 In the Platform web console, go to AI Investigations. In the prompt window, paste the IP addresses and domain names from your intel report, then add: 'activity was noted on July 30, 2024.' Press Enter to start the investigation.

— Mechanism and output files fully named; only a UI walkthrough, not a runnable command.

August 10, 2026 (2026-08-10)

02 **Expanded HTTP body search depth** IMPROVED 85

HTTP body search depth for `host.services.endpoints.http.body` and `web.endpoints.http.body` expanded from the first 2 KB to the first 64 KB, enabling more targeted Platform queries across host service and IP-based web property endpoints.

Search for hosts whose HTTP response body contains a specific string beyond the old 2 KB limit — useful for finding embedded tokens, framework signatures, or configuration leaks deeper in large pages.

```
$ host.services.endpoints.http.body: "api_secret_key"
```

— Named fields, exact before/after limits, and a runnable query example.

August 10, 2026 (2026-08-10)

THINNER COVERAGE BELOW

03 **Bulk tag management via cencli tags** NEW 35

Adds `tags` commands to the `cencli` CLI for managing individual and bulk tags in the Censys Platform.

— Names the CLI and command group but no exact subcommand syntax given.

August 10, 2026 (2026-08-10)

04 **Nine new service fingerprints** NEW 35

Adds nine new fingerprints to the Platform: Omron PLC (hardware), HashiCorp Vault, InfluxDB, Jupyter Server, Oracle Java RMI Registry, Proxmox VE, JBoss JNP Naming Service, Synology DiskStation Manager (DSM), and Tarantool.

— Lists all nine named fingerprints but no query example or mechanism.

August 10, 2026 (2026-08-10)

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

defectdojo ⓘ 1 RELEASE · 2026-08-10 NOTES ↗

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo's v3.2 release lands a large feature set — Rules Engine 2.0, general availability of Threat Intelligence Enrichment, first-class multi-CWE and Vulnerability modeling, new SPDX/CSAF/OpenVEX/Promptfoo parsers, and a unified Connectors hub — alongside a new multi-format report generation API, and the 3.2.100 patch adds location migration tooling, permission scoping, and vendor data preservation fixes.

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

01 Multi-format report generation across API scopes NEW

95

Adds `POST /api/v2/findings/generate_report/` for all viewable findings plus scoped endpoints `POST /api/v2/products/{id}/generate_report/`, `POST /api/v2/engagements/{id}/generate_report/`, `POST /api/v2/tests/{id}/generate_report/`, `POST /api/v2/product_types/{id}/generate_report/`, and `POST /api/v2/endpoints/{id}/generate_report/`, plus Pro-alias endpoints `POST /api/v2/assets/{id}/generate_report/`, `POST /api/v2/organizations/{id}/generate_report/`, and `POST /api/v2/location/{id}/generate_report/`. Requests accept `report_type` (default `JSON`; also `HTML`, `CSV`, `Excel`), `include_finding_notes`, `include_finding_images`, `include_executive_summary`, and `include_table_of_contents`; posting `{}` returns a JSON report, while CSV/Excel are returned as file attachments via `Content-Disposition` with filenames like `product_1_findings.csv` and `test_42_findings.xlsx` (`findings.csv` / `findings.xlsx` for the global endpoint).

Export all findings you can view as a CSV file — useful for offline triage or stakeholder reporting.

```
$ curl -X POST \
  -H "Authorization: Token <your-api-token>" \
  -H "Content-Type: application/json" \
  -d '{"report_type": "CSV"}' \
  https://<your-instance>/api/v2/findings/generate_report/ \
  -o findings.csv
```

Generate an HTML report with executive summary and notes for a single engagement, scoped to only what you are authorized to see.

```
$ curl -X POST \
  -H "Authorization: Token <your-api-token>" \
  -H "Content-Type: application/json" \
  -d '{"report_type": "HTML", "include_executive_summary": true,
"include_finding_notes": true}' \
  https://<your-instance>/api/v2/engagements/7/generate_report/
\
  -o engagement_7_report.html
```

— Full endpoint list, params, and runnable curl examples. product docs

02 **Background-import status fields and polling pattern** IMPROVED

80

Documents the response shape for `POST /api/v2/import-scan/` when `background_import=true` or the `api_async_import` system setting is enabled, including `background_import`, `status`, `findings_parsed`, `test_id`, `engagement_id`, `product_id`, `product_type_id`, and `message` fields (notably omitting `statistics` and `deduplication_complete`), plus a polling pattern via `GET /api/v2/tests/{test_id}/` tracking `status` values (`Processing`, `Post Processing - Deduplication`, `Post Processing - False Positive History`, `Processed`, `Failed`) and a `processing` field.

— Named response fields and statuses, no runnable example. product docs

03 **Threat Intelligence Enrichment reaches general availability**

NEW

80

Threat Intelligence Enrichment is GA with signed threat-intel bundles, downgrade hysteresis, and new list and dashboard surfaces. Adds a threat-intel risk floor based on EPSS, KEV, or other exploitability signals (requires Threat Intelligence Enrichment enabled), plus a two-stage KEV/EPSS pipeline that projects the worst score per vulnerability onto Findings, along with bulk cloud-enrichment reads and import-time enrichment.

— Several named mechanisms, but no exact config steps. product docs

04 **Unified Connectors hub for upstream and downstream integrations**

IMPROVED 75

Introduces a Connectors group in the DefectDojo Pro UI sidebar merging the old API Connectors and Integrations pages into `Connectors > Upstream Connectors` (pulling findings/asset inventory from scanners) and `Connectors > Downstream Connectors` (pushing findings to issue trackers; Beta on Cloud). Legacy URLs auto-redirect; each connector is a tile with logo/description/action and a real-time name search; Upstream adds an All / Asset / Product / Finding type filter and configured connectors show health status, last operation, total/mapped record counts, and a Manage Configuration menu with Manage Records & Operations, Edit Configuration, and Delete Configuration.

— Named UI paths and actions, but navigation only, no API. product docs

05 **Multi-CWE findings and normalized Vulnerability entity** IMPROVED

65

Findings can now carry multiple CWEs across the API, Vue UI, and universal parser; vulnerability IDs are normalized into a first-class Vulnerability entity with ordered references, per-vulnerability KEV/EPSS enrichment columns, and vulnerability aliases.

— Detailed model change described, no API example. product docs

06 **Rules Engine 2.0 visual automation (feature flag)** NEW

60

Rules Engine 2.0, gated behind a feature flag, lets teams build automation rules as visual node graphs reacting to Finding events, with per-run traces and a delivery outbox.

— Describes mechanism, no flag name or config given. product docs

07 **Location-to-endpoint migration and product-scoped location data**

IMPROVED 60

Adds a `migrate_locations_to_endpoints` operation — the reverse of the endpoint-to-location conversion — letting teams move location data back to endpoints, and scopes Location metadata and tags to the owning product via a database migration, enforcing product-level isolation on location data.

— Names the migration operation, no CLI usage shown. 3.2.100

THINNER COVERAGE BELOW

08 **Set-based and cross-product finding deduplication** NEW

55

Adds set-based deduplication matching Findings on their full set of vulnerability IDs and CWEs, including partial/subset matches, plus a new `global_locations` cross-product deduplication algorithm for shared locations.

— Names algorithm and matching logic, no config steps. product docs

09 **Jira migration to Downstream Connectors and `/assign_product` mapping**

IMPROVED 55

Adds one-button migration from classic Jira to Downstream Connectors; the public `/assign_product` endpoint can map Findings-type records again.

— Names endpoint and migration path, limited mechanism. product docs

10 **Deprecation of API-based pull parsers, Tool Type/Configuration, and dbbackup**

DEPRECATED 55

API-based pull parsers, Tool Type/Tool Configuration, and dbbackup are deprecated with end-of-life scheduled for v3.5.0 — integrations relying on these surfaces will break at that release.

— Names deprecated surfaces and exact EOL version. product docs

11 **SPDX, CSAF 2.0, OpenVEX, and Promptfoo parsers** NEW

50

Adds SPDX, CSAF 2.0, and OpenVEX interchange-format parsers, plus a Promptfoo (LLM eval and red-teaming) parser.

— Named parsers listed, no usage detail given. product docs

12 **Work Assignment for individuals (feature flag)** NEW

50

Work Assignment, gated behind a feature flag, assigns Findings and Risk Acceptances to individual people alongside existing group Owners, with a personal My Work queue.

— Names queue and scope, no config key given. product docs

13 **Federal Compliance pack (feature flag)** NEW 45

The Federal Compliance pack, gated behind a feature flag, adds a FedRAMP POA&M ledger and ConMon deliverables, CMMC Level 2 assessments, and control coverage.

— Names deliverables, no usage steps. product docs

14 **Review Claiming for Finding reviews (feature flag)** NEW

40

Review Claiming, gated behind a feature flag, lets a requested reviewer claim a Finding review so other eligible reviewers can see it is being handled.

– Explains mechanism, no config key named. product docs

15 **Raw CycloneDX VEX data preserved on findings** IMPROVED 40

Preserves raw CycloneDX VEX analysis data on parsed findings, retaining the original vendor-supplied VEX analysis fields rather than discarding them.

– Explains before/after, field names not given. 3.2.100

16 **Named POST permission on risk acceptance API** IMPROVED 40

Names the POST permission on the risk acceptance API, making the permission model explicit for that endpoint.

– Names endpoint, permission label not specified. 3.2.100

17 **Snyk raw reachability values on findings** IMPROVED 40

Renders Snyk's raw reachability values instead of a derived boolean, exposing the original Snyk reachability string on findings.

– Names vendor and change, field name unspecified. 3.2.100

18 **JFrog connector scoped to latest build** IMPROVED 35

JFrog connector now scopes artifact-mode Findings to each artifact's latest build.

– Single-sentence behavior change, no scope detail. product docs

19 **Location drift matching across reimports** NEW 30

Location drift matching keeps a Finding tracked as its locations change across reimports.

– One-line description, mechanism unspecified. product docs

20 **Menu 2.0 settings reorganization (feature flag)** IMPROVED 30

Menu 2.0, gated behind a feature flag, reorganizes the Settings menu with a new All Settings hub.

– Names new hub, minimal further detail. product docs

21 **Permission-scoped CI/CD Infrastructure details** IMPROVED 30

Scopes CI/CD Infrastructure detail fields to the matching configuration permission, restricting access to infrastructure metadata by permission level.

— One-line permission change, no specific field names. 3.2.100

22 **Markdown support in Notes** NEW 15

Notes now support Markdown.

— Bare statement, no formatting detail. product docs

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! API-based pull parsers, Tool Type/Tool Configuration, and dbbackup are deprecated with end-of-life scheduled for v3.5.0 — integrations relying on these surfaces will break at that release.

WAS THIS USEFUL? 👍 👎

Cloudmersive ⓘ 1 RELEASE · 2026-08-11 CHANGELOG ↗

Cloudmersive is an API platform providing cloud-based document processing, image manipulation, and data conversion services for developers.

> Cloudmersive shipped incremental updates across its Private Cloud CDR, DLP AI, and Virus APIs, adding content-cleaning options, batch job overrides, and expanded file format detection.

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

01 **Batch job API overrides in DLP AI API** IMPROVED 35

Adds additional batch job API overrides to the Cloud DLP AI API `v2-8-0` and Private Cloud DLP AI API `v2-8-0`.

— Vague on what overrides do or how to use them changelog-20260811-0787ff67

02 **Python model file format support in Virus API** IMPROVED 30

Extends the Cloud Virus API `v11-10-10` to recognize additional Python model file formats.

— No specifics on which formats or how detection changed changelog-20260811-0787ff67

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

OpenSSH

2 RELEASES · SEEN 2026-08-11 CHANGELOG ↗

Portable OpenSSH

> OpenSSH 10.5 adds a new `ssh -Z` mode for previewing which keys will be tried during authentication, FIDO key flag editing and smarter FIDO ordering, several sshd_config and hardening tweaks, and now requires ECC support in libcrypto.

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

01 **ssh -Z key preview mode** NEW 75

New `ssh -Z user@host` mode prints the public keys that will be tried for pubkey authentication, in the exact order they will be attempted — useful for diagnosing auth failures or unexpected key selection before actually connecting.

Preview exactly which keys SSH will offer — and in what order — before connecting to a host, useful for debugging pubkey auth failures.

```
$ ssh -Z user@host
```

Inspect which keys `ssh` will present for public key authentication and in what order, before actually connecting — useful for diagnosing auth failures or unexpected key selection.

```
$ ssh -Z user@host
```

— Names exact command; lacks deeper mechanism detail OpenSSH 10.5

changelog-20260811-292795c6

02 **sshd_config Match block directive support** IMPROVED 65

Re-allows `PAMServiceName` inside a `Match` block in `sshd_config` (had been incorrectly disabled in 10.4), and allows `ChannelTimeout` and `RekeyLimit` to be set inside `sshd_config` `Match` blocks.

— Names exact config keys; no worked config example changelog-20260811-292795c6

03 **ECC requirement in libcrypto for builds** BREAKING 60

Portable OpenSSH now requires ECC (Elliptic Curve Cryptography) support in libcrypto, including support for the NISTP521 curve; builds against a libcrypto without ECC will fail. The `--without-openssl` build configuration is not affected.

— Names exact requirement and unaffected build flag OpenSSH 10.5

changelog-20260811-292795c6

THINNER COVERAGE BELOW

04 **FIDO key flag editing on passphrase reset** NEW 55

ssh-keygen(1) can now set or clear the `touch-required` and `verify-required` flags on FIDO private keys when resetting a private key's passphrase.

— Names flags and tool but no example command OpenSSH 10.5 changelog-20260811-292795c6

05 **Low-friction FIDO key ordering in pubkey auth** IMPROVED 50

ssh(1) now reorders FIDO key trials during pubkey authentication to try low-friction keys (no user presence required) first, and keys requiring PIN or biometric verification last, reducing unnecessary authentication friction.

— Explains behaviour but no configurable surface given OpenSSH 10.5 changelog-20260811-292795c6

06 **Non-blocking ssh-keyscan banner reading** IMPROVED 35

Makes `ssh-keyscan` banner reading non-blocking so a stuck server no longer stalls a multi-host keyscan.

— Names tool and fix but minimal detail on mechanism changelog-20260811-292795c6

07 **Earlier public key type validation** IMPROVED 35

Moves the public key type check against allowed algorithms to before key parsing, reducing pre-authentication attack surface.

— States the change but no numbers or config to act on changelog-20260811-292795c6

08 **sshd-session process title labeling** IMPROVED 35

sshd(8) now uses setproctitle(3) to label the `sshd-session` process when acting as a post-authentication monitor, improving process visibility.

— Names mechanism but no visible reader action OpenSSH 10.5

└──▷ BREAKING ON UPGRADE

! Portable OpenSSH now requires ECC (Elliptic Curve Cryptography) support in libcrypto, including support for the NISTP521 curve; builds against a libcrypto without ECC will fail. The `--without-openssl` build configuration is not affected.

! Portable OpenSSH now requires ECC (Elliptic Curve Cryptography) support — including the NISTP521 curve — in libcrypto; builds against a libcrypto without ECC will fail. The `--without-openssl` build configuration is not affected.

WAS THIS USEFUL? 👍 👎

DNSFilter ⓘ 1 RELEASE · 2026-08-10 BLOG ↗

DNSFilter is a security tool that filters DNS queries to block malicious domains, ads, and unwanted content at the network level.

> DNSFilter's Windows Roaming Client v3.6.10 adds DNS PreCheck v2 to preserve policy enforcement against DNS-hijacking networks, plus automatic recovery for a crashed tray process.

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

01 **DNS PreCheck v2 for hijacked networks** NEW 65

The Windows Roaming Client adds `DNS PreCheck v2`, which maintains policy enforcement on networks that hijack or redirect DNS traffic. It rolls out gradually and is not automatically enabled on upgrade; the tray updates to display `DNS PreCheck v2` to confirm the active mode.

— Names mechanism and tray indicator but no config/flag to control rollout.

08/10 August 10 – Windows Roaming Clien…

THINNER COVERAGE BELOW

02 **Automatic tray process recovery** IMPROVED 40

The system tray process now restarts itself automatically if it stops unexpectedly, keeping tray status and controls available without requiring a manual client or machine restart.

— Describes behavior but no configuration or trigger detail given.

08/10 August 10 – Windows Roaming Clien…

WAS THIS USEFUL? 👍 👎

tailscale

1 RELEASE · 2026-08-11 BLOG ↗

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

> Tailscale's Kubernetes Operator v1.102.2 adds in-cluster PeerRelays, workload identity federation, and expanded IPv6/dual-stack support, alongside smaller Helm and logging improvements.

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

01 **Workload identity federation for Tailnet resource** NEW

45

Adds workload identity federation configuration for the `Tailnet` custom resource.

– Names the custom resource but no config fields shown.

Tailscale Kubernetes Operator v1.102.2 …

02 **4via6 support in connector and egress proxy resources** NEW

45

The operator now supports `4via6` in connector and egress proxy resources when egressing from a dual-stack cluster.

– Names the mechanism and resources but no setup steps.

Tailscale Kubernetes Operator v1.102.2 …

03 **In-cluster PeerRelays via custom resource** NEW

40

The Kubernetes Operator now supports deploying PeerRelays in-cluster through a new custom resource.

– Names a new custom resource but no schema or fields given.

Tailscale Kubernetes Operator v1.102.2 …

04 **Helm annotations for operator deployment resource** IMPROVED

40

Adds the ability to apply annotations to the operator's deployment resource via Helm.

– Clear Helm-based action point but no key name given.

Tailscale Kubernetes Operator v1.102.2 …

05 **IPv6 support in Egress ProxyGroups** NEW

30

Egress ProxyGroups now support IPv6.

– Bare statement of support, no mechanism or scope detail.

Tailscale Kubernetes Operator v1.102.2 …

06 **Reduced operator log noise** IMPROVED 30

Reduces operator log noise by excluding superfluous entries for resources that do not contain annotations.

— Explains behavior change but no config to act on.

Tailscale Kubernetes Operator v1.102.2 …

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

SailPoint

4 RELEASES · 2026-08-10 BLOG ↗

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

> SailPoint's headline addition this window is support for requesting entitlements on behalf of Machine Identities through the ISC Request Center, paired with a major overhaul of approval reminders/escalations and email routing, new MCP client governance in Agentic Fabric, Data Segmentation extended to Human Identities, and report management changes in Non-Employee Risk Management.

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

01 **Approval reminders, escalations, and CC/BCC notifications** NEW

90

Adds an `Admin > System Settings > Feature Settings > Approval Settings` tab for globally configuring approval reminders, escalations, and timeouts, including cron-string cadence, multi-level escalation chains, and a fallback approver set to an identity or governance group. Adds per-access-item reminder/escalation configuration under Access Requests for Entitlements, Access Profiles, and Roles that overrides the global config, and supports independent setups (reminders without escalations, escalations without reminders, custom escalation targets outside the management chain). Email templates now support configuring CC:/BCC: recipients from Manager, Identity, Governance Groups, Org Admins, and/or a static email address, restricted on the `cloud_user_password_expiration`, `cloud_user_password_reset_token`, `cloud_user_unlock_token`, `cloud_user_verification_token`, and `onboarding_password_reset` templates.

.

Configure a global escalation policy with a cron-based cadence and a governance-group fallback approver for all access-request approvals.

> 📍 In the console, go to Admin › System Settings › Feature Settings, select the Approval Settings tab, set the reminder and escalation cadence using a Cron string, define the number of escalation levels, assign a Governance Group as the fallback approver, and save.

Add CC: recipients to an approval notification template so managers are copied on access-request decisions.

> 📍 In the console, open the target email template (e.g., an access-request approval template), navigate to the CC:/BCC: configuration, select recipients from Manager, Governance Group, or enter a static email address, and save the template.

— Exact config paths and template names given, fully actionable.

Enhancement: Approvals Migration and No…

02 **Regional AWS SES routing for identity emails** BREAKING

80

New identities (email addresses and domains) are now created and verified in the tenant's primary AWS region instead of always routing through `us-east-1`, and a new `region` field on `EMAIL_SENT` audit events records which AWS region sent each email. Pending approvals open longer than 90 days began expiring starting August 10, 2026, with remaining pre-migration approvals closed; domains with a custom `MAIL FROM` attribute are NOT auto-migrated and must be manually re-added/re-verified in the primary region, and existing verified sender addresses likewise are not auto-migrated and keep sending from `us-east-1` until manually removed and re-added.

Migrate a verified sender domain to your primary AWS region so emails no longer route through us-east-1.

📍 In the console, go to the email identities section, locate the domain showing the `us-east-1` region label, remove it, and re-add the same domain — the new entry will replace the existing record and trigger verification in your primary region automatically.

— Names fields and behavior; gives manual migration steps but no exact command.

Enhancement: Approvals Migration and No…

03 **Machine Identity access requests via Request Center** NEW

65

Users can now request entitlements on behalf of Machine Identities through the Identity Security Cloud (ISC) Request Center for Machine Identity Security customers. Manager-assigned approvals for these requests route to the Machine Identity owner instead of a human manager, and requesters must select a target account when a Machine Identity has multiple accounts on a source.

— Clear behavior and routing logic, no exact command or endpoint.

New Capability: Machine Identity Access…

04 **Report deletion and visibility in Non-Employee Risk Management**

IMPROVED

60

Report creators can delete their own report definitions from the Reports page in Non-Employee Risk Management, which also removes all associated exports. The Reports All

tab now shows only reports the logged-in user created or that were shared with them, preventing exposure of sensitive metadata in report titles to unauthorized users.

— Clear UI behavior and access-scoping, no exact config key given.

New Capability: Non-Employee Risk Manag…

THINNER COVERAGE BELOW

05 MCP client discovery and governance in Agentic Fabric NEW

55

Adds an MCP Clients page to the non-human identity registry giving a central view of all discovered MCP clients connected to AI agents, discoverable from endpoint agents, browser extensions, and cloud connectors. SailPoint sources now aggregate and govern MCP clients alongside AI agents, credentials, and IAM roles objects using datasets and resources, treating MCP clients as governed non-human identity assets that can be reviewed for their access to tools, systems, data sources, and credentials.

— Names the registry page and object types but no config or command. product docs

06 Data Segmentation extended to Human Identities IMPROVED

45

Data Segmentation in Identity Security Cloud now covers Human Identities, not just Roles and Entitlements, letting administrators restrict which identities a given user can see. This enables delegated administration at the subsidiary or team level without exposing cross-organizational identity data to scoped admins.

— Describes scope of extension but no config key or navigation path.

Enhancement: Data Segmentation for Iden…

└──▷ BREAKING ON UPGRADE

- ! Pending approvals that have been open for longer than 90 days will begin expiring starting August 10, 2026; any remaining open pre-migration approvals will be closed.
- ! Domains with a custom `MAIL FROM` attribute configured will NOT be automatically migrated to the primary region — owners must manually re-add and re-verify the domain in their primary region or emails risk being blocked by SMTP servers.
- ! Existing verified sender email addresses are not auto-migrated and continue to send from `us-east-1`; they must be manually removed and re-added to trigger verification in the primary region.

WAS THIS USEFUL? 👍 👎

Auth0 ⓘ 1 RELEASE · 2026-08-10 BLOG ↗

Auth0 is a cloud-based identity and access management platform providing authentication, authorization, and user management solutions for applications.

> Auth0 closed a gap in Universal Login by extending Custom Prompts enforcement to Social and Enterprise connections, so signup fields and consent checkboxes can no longer be bypassed via those sign-in methods.

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

01 **Custom Prompts extended to Social and Enterprise connections**

IMPROVED 80

Custom signup/login fields and consent checkboxes configured in Universal Login now apply to Social and Enterprise connections, closing the gap where users could bypass those prompts by signing in via a social button or enterprise identity provider. No configuration changes are required — existing Custom Prompts setups automatically extend to all connection types, and can be reviewed or updated via Dashboard under Branding → Universal Login → Enhance screens with partials, or via the Management API.

— Names exact Dashboard path and Management API, but no example.

Custom Prompts now capture the same fie…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Clerk ⓘ 1 RELEASE · 2026-08-10 BLOG ↗

Clerk is a user authentication and identity management platform that provides sign-up, login, and user profile management for web and mobile applications.

> Clerk Billing now supports discounts and promo codes, letting subscriptions carry percentage or fixed-amount reductions redeemable at checkout.

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

01 **Discounts and promo codes in Clerk Billing** NEW 45

Clerk Billing adds support for applying percentage or fixed-amount discounts directly to a customer's subscription, plus promo codes that customers can redeem at checkout.

— Describes what was added but no UI path or API detail.

Discounts and promo codes for Billing (…

WAS THIS USEFUL? 👍 👎

Ory ⓘ 1 RELEASE · SEEN 2026-08-11 NOTES ↗

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

> Ory shipped a single configuration addition to Hydra letting operators explicitly disable specific OAuth 2.0 grant types.

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

01 **Disable specific OAuth 2.0 grant types in Hydra** NEW 65

Ory Hydra's configuration schema gains the `oauth2.grant_types_disabled` config key, which accepts a list of OAuth 2.0 grant types to explicitly disable.

— Names exact config key but no behavior details or example

Ory Hydra, Ory Kratos, Ory Keto v26.3.5…

WAS THIS USEFUL? 👍 👎

1Password

4 RELEASES · 2026-08-11 BLOG ↗

1Password is a password manager that securely stores, generates, and autofills login credentials and sensitive information across devices.

> 1Password's 8.12.32 releases across platforms add a Claude AI agentic sign-in integration and a cluster of browser security and usability refinements, alongside smaller platform-specific fixes for Android, Mac and Windows.

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

01 **Shell plugin credential save setting and flow in browser**

IMPROVED 70

Adds a setting in `Settings > Autofill & save` to choose whether 1Password offers to save API credentials for shell plugins; shell plugin credential saves now use the same autosave notification flow as standard logins.

— Names exact settings path and behaviour change 1Password in the browser 8.12.32 (2026-…

02 **Phishing prevention settings relocated in browser** IMPROVED

65

Moves phishing warning settings to `Settings > Security & privacy > Phishing prevention`, with a new separate toggle for the popup phishing warning.

— Exact settings path given, minimal further mechanism

1Password in the browser 8.12.32 (2026-…

THINNER COVERAGE BELOW

03 **Agentic sign-in via 1Password for Claude** NEW 50

1Password for Claude lets Claude sign in to websites on your behalf by filling logins or one-time codes after per-request user approval, without sharing credentials with Claude. Shipped for 1Password for Mac and 1Password in the browser.

— Explains mechanism but no exact setup path or command

1Password for Mac 8.12.32 (2026-08-11) 1Password in the browser 8.12.32 (2026-…

04 **Secret Key paste protection in browser** NEW 50

Pasting your 1Password Secret Key into any website is now blocked and opens the extension to warn you.

— Clear behaviour described, no config or flag named

1Password in the browser 8.12.32 (2026-…

05 **Auto-Type confirmation for untrusted apps** IMPROVED 50

When using Auto-Type to fill into an unsigned or untrusted app, 1Password now prompts for confirmation instead of silently failing.

— Clear before/after behaviour, no exact UI location

1Password for Windows 8.12.32 (2026-08-…

06 **Error message for custom icon moves with file storage disabled**

IMPROVED

50

Adds a clear error message when attempting to move or duplicate an item with a custom icon while file storage is disabled by an admin policy.

— Specific scenario named, no exact error text given

1Password for Mac 8.12.32 (2026-08-11)

07 **'Hide on this page' suppresses notifications** IMPROVED

45

Selecting Hide on this page now suppresses all 1Password notifications on that page, except Device Trust notifications.

— Specific exception noted but limited scope detail

1Password in the browser 8.12.32 (2026-…

08 **App icons in 'Type in window' menu** IMPROVED

45

The Type in window window selection menu now displays an app icon next to each open window when filling a username or password into another app.

— Names the menu and change but nothing to configure

1Password for Windows 8.12.32 (2026-08-…

09 **Date format follows system or in-app language** IMPROVED

40

Item date fields now respect the date format chosen in system settings, or the language set within 1Password when it differs from the system language. Applies on Mac and Windows.

— Describes behaviour, no UI path or setting named 1Password for Mac 8.12.32 (2026-08-11)

1Password for Windows 8.12.32 (2026-08-…

10 **Login save prompts for split username/password screens** NEW

30

1Password for Android now offers to save logins when a site presents the username and password on separate screens.

— Single-sentence description with no further mechanism

1Password for Android 8.12.32 (2026-08-…

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

2 RELEASES · 2026-08-10 → 2026-08-11 BLOG ↗

RegScale is a compliance and regulatory management platform that automates governance, risk, and compliance (GRC) workflows for organizations.

> RegScale's CLI gained a Splunk integration that pulls saved-search results into assets, findings, and control evidence, while the platform added application hierarchy and visibility controls to App Builder.

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

01 Splunk integration for asset, finding and evidence sync NEW

88

New `splunk sync` CLI command runs saved searches listed in the `splunkSavedSearches` config and imports assets, findings, and control evidence into RegScale. The underlying Splunk integration supports search, streaming export, saved searches, search job management, and asset/finding processing.

Pull Splunk saved search results into RegScale to automatically populate assets, findings, and control evidence without manual export/import.

```
$ regscale splunk sync
```

— Names command and config key, includes runnable example. CLI 6.39.0 (2026-08-11)

THINNER COVERAGE BELOW

02 Application hierarchy and visibility controls in App Builder

NEW

47

Adds the ability to define a hierarchy of applications on the App Management page, plus a global visibility toggle within the application's Advanced Security tab.

— Names UI locations but no mechanism or example given. RegScale 6.33.0.1 (2026-08-10)

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

Qualys Policy Compliance ⓘ 1 RELEASE · 2026-08-10 BLOG ↗

Qualys Policy Compliance is a cloud-based security tool that monitors IT infrastructure for regulatory compliance violations and policy adherence.

> Qualys Policy Compliance launches Audit Fix, an add-on that autonomously remediates failed compliance controls via Cloud Agent, alongside broader real-time coverage for auto-scaling cloud workloads and expanded framework support.

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

01 **Audit Fix autonomous remediation module** NEW 75

The new Audit Fix add-on deploys pre-built Custom Assessment and Remediation (CAR) scripts through the existing Qualys Cloud Agent to autonomously remediate failed compliance controls at scale. Each CAR script is version-controlled and mapped to its specific benchmark control (CIS, DISA STIG, or NIST requirement); after execution an automatic rescan updates control status from Failed to Passed, generating timestamped, audit-ready evidence across more than 90 compliance frameworks.

— Explains mechanism and evidence but no exact UI/API steps. 2026-08-10-1486c8f0

THINNER COVERAGE BELOW

02 **Expanded compliance framework and platform coverage** IMPROVED 50

Policy Compliance now covers more than 500 platforms and environments, continuously checking against CIS Benchmarks, DISA STIGs, PCI-DSS, and NIST frameworks for compliance drift detection.

— States coverage breadth but no usage steps. 2026-08-10-1486c8f0

03 **Real-time assessment for auto-scaling cloud workloads** IMPROVED 40

Continuous real-time assessment now extends to new and auto-scaling cloud workloads — assets are automatically activated and evaluated the moment they spin up, with no separate scan window required.

— Describes behavior but no named config or trigger. 2026-08-10-1486c8f0

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Kiro ⓘ 1 RELEASE · 2026-08-11 CHANGELOG ↗

> Kiro's 1.0.288 release adds support for the open Agent Plugin format, session pinning and CLI hand-off in Agent Focus Mode, an update checker, and direct-to-editor spec/task navigation.

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

01 **Agent Plugin format support for Powers** NEW 70

Kiro now supports Powers packaged in the open Agent Plugin format, installable from a local folder or a GitHub URL. Plugins bundle skills and MCP together for sharing across compatible agent tools.

– Names install sources and bundled contents but no exact command

changelog-20260811-431df658

THINNER COVERAGE BELOW

02 **Session rail controls in Agent Focus Mode** NEW 45

In Agent Focus Mode's session rail, sessions can be pinned with Pin Session to keep them at the top of their section, and an existing session can be resumed in Kiro CLI via Open with Kiro CLI directly from the rail.

– Names the exact UI actions but describes them only briefly changelog-20260811-431df658

03 **Check for Updates control** NEW 45

A new Check for Updates button appears in the settings pane, plus an update row shown in the rail footer when a new Kiro version is available.

– Exact UI location given, minimal further detail changelog-20260811-431df658

04 **Spec documents open directly in editor** IMPROVED 30

Spec documents and task rows in the context panel now open directly in the editor, rather than in a separate view.

– States the change but no mechanism or scope detail changelog-20260811-431df658

└──▷ **ALSO FROM THESE RELEASES**





WAS THIS USEFUL? 👍 👎

Devin Desktop

2 RELEASES · SEEN 2026-08-11 CHANGELOG ↗

Devin Desktop is an AI-powered IDE that automates coding tasks, debugging, and software development workflows for developers.

> Devin Desktop's window centers on new API authentication and extensibility groundwork — personal access tokens, a cross-surface plugin system, and an official Terraform provider — alongside a v3.7.16 release that reworks Settings/Keyboard Shortcuts into editor tabs, adds shareable local conversations, composable permission rules, and per-server MCP logging.

└──▷ WHAT SHIPPED · 18 FEATURES most completely described first ? what's the number?

01 Editor tabs replace modal Settings/Keyboard Shortcuts

BREAKING

85

Settings and Keyboard Shortcuts now open as regular editor tabs instead of modal overlays by default; set `"workbench.editor.useModal": "all"` to restore the previous modal behavior.

Keep Settings and Keyboard Shortcuts as modal overlays after upgrading, restoring the pre-v3.7.16 behavior.

```json
"workbench.editor.useModal": "all"
```

Restore the modal overlay for Settings and Keyboard Shortcuts if your workflow depends on the old behavior.

```json
{
  "workbench.editor.useModal": "all"
}
```

— Exact config key and default change given, no migration detail beyond that

changelog-20260811-8112423e v3.7.16

02 Shareable Devin Local conversation links

NEW

80

New Share Conversation action uploads a sanitized Devin Local transcript — system prompts, tool definitions, and secrets stripped, paths normalized — and copies a team-visible link, accessible from the actions menu below a completed turn.

Share a completed Devin Local conversation with your team — transcripts are sanitized before upload.

📍 Open the actions menu below a completed turn › Share Conversation

— Names the action and exact sanitization steps plus UI path changelog-20260811-8112423e v3.7.16

03 **Customizations panel: subagents, search, plugin scope** NEW

80

Adds a Subagents section and cross-section search across skills, subagents, rules, hooks, plugins, and MCP servers, plus a refresh action for plugins. Plugins installed from Customizations are added to personal plugins by default (available in Devin Cloud and other devices), with an Install locally option to limit the install to the current device.

— Names every named surface added to the panel changelog-20260811-8112423e v3.7.16

04 **Composable permission rules across levels** IMPROVED

70

Permission rules from enterprise, mode, user, project, and subagent levels now compose predictably: an explicit `deny:` always wins, always allow options only appear when the grant would actually take effect, and denials name the layer that denied them.

— Mechanism is clear but no config file or command shown changelog-20260811-8112423e v3.7.16

05 **Editable queued messages** NEW

60

Queued messages can now be edited — click the edit button or press ⬆ in an empty input to pull the last queued message back before it runs.

— Exact keystroke and button given, small scope changelog-20260811-8112423e v3.7.16

06 **Clickable diff summary opens multi-diff editor** IMPROVED

60

The `+X -Y` diff summary on a collapsed agent group is now clickable, opening a multi-diff editor scoped to that group's files.

— Exact UI element named with clear resulting behavior changelog-20260811-8112423e v3.7.16

07 **Personal access token API authentication** NEW

60

The API now supports Personal Access Tokens, letting users authenticate directly as themselves without a service user or separate account; sessions are automatically attributed to the authenticated user.

— Explains mechanism and attribution but no token creation steps product docs

THINNER COVERAGE BELOW

08 **Per-server MCP output log channels** IMPROVED 55

Each MCP server's logs are now separated into its own `MCP: <server>` output channel for isolated debugging.

– Names channel pattern but no further mechanism changelog-20260811-8112423e v3.7.16

09 **Immediate revert buttons with turn cancellation** IMPROVED 55

Revert buttons now appear as soon as a prompt is sent; reverting while an agent is working cancels the turn first, making mid-turn reverts reliable.

– Behavior described but no exact UI location given changelog-20260811-8112423e v3.7.16

10 **Worktree support: terminal command and error surfacing** IMPROVED 55

`Terminal: Create New Terminal in Editor Area` now works inside worktrees, and starting an agent in a new worktree that fails now surfaces the underlying git error instead of silently falling back to the workspace.

– Names exact command but limited mechanism beyond that changelog-20260811-8112423e v3.7.16

11 **Remote-aware Codemaps and MCP config files** IMPROVED 55

Codemaps and MCP configuration files now open from the remote machine when connected over WSL, SSH, or a dev container.

– Names remote transports but no further detail changelog-20260811-8112423e v3.7.16

12 **Enterprise ACU-based session size chips** IMPROVED 50

Session size chips now use the same enterprise ACU thresholds as the web app, with auto-continuation billing notices shown inline in the transcript.

– Names ACU thresholds concept but no exact values changelog-20260811-8112423e v3.7.16

13 **Fork conversations into new tabs** IMPROVED 45

Duplicate session and send as fork now open the fork in a new tab, keeping the original conversation intact.

– Simple UI behavior, thinly described changelog-20260811-8112423e v3.7.16

14 **Bounded, paginated large-file reads via ACP** IMPROVED 45

Large-file reads through ACP are now bounded and paginated instead of repeatedly rereading overflow files.

– States mechanism change but no limits or numbers `changelog-20260811-8112423e` `v3.7.16`

15 **Plugin system across Devin surfaces** NEW 45

A new plugin system supports installing, authoring, and governing plugins across Devin cloud sessions, the CLI, and Devin Desktop.

– High-level description with no commands or config named `product docs`

16 **Cascade options hidden when disabled** IMPROVED 40

Cascade-specific configuration and customization options are hidden when Cascade is disabled for the team.

– Simple conditional UI behavior, minimal detail `changelog-20260811-8112423e` `v3.7.16`

17 **Official Terraform provider for Devin resources** NEW 40

A new official Terraform provider enables managing Devin resources as infrastructure-as-code.

– Names the provider but no resources or usage detail `product docs`

18 **Plugins section renamed to Extensions** BREAKING 30

The `Plugins` section for IDE extension marketplaces has been renamed to `Extensions`.

– Rename noted with no further scope or migration detail `changelog-20260811-8112423e`

└──▷ ALSO FROM THESE RELEASES

.
.
.
.

└──▷ BREAKING ON UPGRADE

! Settings and Keyboard Shortcuts now open as regular editor tabs instead of modal overlays by default; set `"workbench.editor.useModal": "all"` to restore the previous modal behavior.

! The `Plugins` section for IDE extension marketplaces has been renamed to `Extensions`.

WAS THIS USEFUL? 👍 👎

Cline ⓘ 1 RELEASE · 2026-08-11 NOTES ↗ CODE ↗

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

> Cline v4.1.8 expands Vertex AI model support with custom model IDs, Fable 5, and full regional catalogs, while consolidating unattended-run approval into a single auto-approve menu.

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

01 Auto-approve menu consolidation for unattended runs BREAKING

75

The auto-approve menu is now the single source of truth for unattended runs; existing Yolo Mode or auto-approve-all setups are automatically migrated to auto-approving every action. The Yolo Mode toggle (`yoloModeToggled`) is removed from settings as part of this migration.

— Names removed setting key and migration behavior clearly v4.1.8

02 Expanded Vertex model support and catalogs IMPROVED

60

Adds manual entry of any Vertex model ID, including models not yet listed in the catalog, adds support for Fable 5 on Vertex, and now shows the full Vertex model catalog for every region instead of a hardcoded list of global-endpoint models. Unsupported region/model combinations now fail at request time with recovery guidance.

— Names Vertex, Fable 5, catalog behavior but no exact UI path v4.1.8

└──▷ **BREAKING ON UPGRADE**

! The Yolo Mode toggle (`yoloModeToggled`) is removed from settings; users who had it enabled are migrated to auto-approving every action via the auto-approve menu.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI

2 RELEASES · 2026-08-10 NOTES ↗ CODE ↗

Lightweight coding agent that runs in your terminal

> Codex CLI's two alpha releases add Amazon Bedrock support (hosted web search and remote compaction), raise the MCP namespace description limit from 1,000 bytes to 512 KiB, and introduce a `--psp` process-scoped routing flag alongside a long tail of smaller CLI, tool-registry, and security hardening changes.

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

01 **MCP namespace description limit raised** IMPROVED 75

Raises the MCP namespace tool-spec description limit from `1,000 bytes` to `512 KiB`, truncating only at a UTF-8 character boundary when the new limit is exceeded; a later release further preserves complete MCP namespace descriptions.

– Exact byte limits given, but no user-facing flag. rust-v0.148.0-alpha.6 rust-v0.147.0-alpha.6.6

02 **PSP routing flag for app-server** NEW 75

Adds `--psp` flag to `app-server` to enable process-scoped PSP routing for first-party ChatGPT requests.

Enable process-scoped PSP routing for first-party ChatGPT requests when starting the app server.

```
$ codex app-server --psp
```

– Exact flag and runnable command given. rust-v0.148.0-alpha.6

03 **Amazon Bedrock hosted web search** NEW 65

Enables hosted web search for Amazon Bedrock sessions, resolving unsupported live and indexed search modes to cached search and disabling the tool entirely when cached search is prohibited by managed requirements; a following release also enabled cached web search for Bedrock sessions.

– Mechanism described but no config flag given. rust-v0.148.0-alpha.6 rust-v0.147.0-alpha.6.6

THINNER COVERAGE BELOW

04 **externalAgentConfig/detect serialization queue** IMPROVED 55

Gives `externalAgentConfig/detect` its own global serialization queue so configuration reads, writes, and external agent imports remain responsive while detection runs.

– Named internal endpoint, but no user action described. `rust-v0.147.0-alpha.6.6`

05 **Cursor project path resolution bounded** IMPROVED 55

Bounds Cursor project path resolution to 128 probes using common filename separators instead of recursively walking the directory tree, rejecting ambiguous matches and unsafe encoded components.

– Detailed mechanism, but no user-triggerable action. `rust-v0.147.0-alpha.6.6`

06 **Thread submission queues and paginated turns** IMPROVED 50

Adds durable per-thread user submission queues for more reliable turn handling, and supports `includeTurns` reads for paginated threads.

– Named field `includeTurns`, limited mechanism detail. `rust-v0.148.0-alpha.6`

07 **Tool registry and namespace support** NEW 45

Adds tool registry collision policy configuration to control how name conflicts are resolved, supports deferred loading for freeform tools and deferred custom tools in tool search, and supports custom tools in namespaces.

– Grouped thin one-liners, no config keys named. `rust-v0.148.0-alpha.6`

08 **Secret redaction from command execution items** IMPROVED 45

Redacts secrets (including bearer tokens) from app-server command execution items.

– Named secret type, no exact log location given. `rust-v0.148.0-alpha.6`

09 **Host skill root loading via codex-skills crate** IMPROVED 45

Adds host skill root loading and moves skill policy resolution into the `codex-skills` crate.

– Named crate, but no usage instructions. `rust-v0.148.0-alpha.6`

10 **Plugin install skips non-regular entries** IMPROVED 40

Skips symbolic links and other non-regular entries when installing a plugin into the store instead of rejecting the installation.

– Behavior change described, no command shown. `rust-v0.147.0-alpha.6.6`

11 **Code-mode dual-WebSocket transport** NEW 40

Adds a dual-WebSocket transport for code mode, enabling a new host-transport path for code-mode sessions, and times out stalled code-mode host requests.

— New transport described briefly, no config surface named. `rust-v0.148.0-alpha.6`

12 **Local project directory trust prompts** IMPROVED 40

Prompts before trusting local project directories, and automatically trusts undecided local projects after confirmation.

— Behavior described, no exact command or flag. `rust-v0.148.0-alpha.6`

13 **Bare Git repository rejection** IMPROVED 40

Rejects implicitly discovered bare Git repositories to reduce unintended repo scope.

— Scope described, no exact trigger given. `rust-v0.148.0-alpha.6`

14 **MCP session robustness improvements** IMPROVED 35

Negotiates MCP extensions per app-server session and handles late MCP startup results that arrive after the lag timeout.

— Grouped one-line mentions, no mechanism detail. `rust-v0.148.0-alpha.6`

15 **Paginated TUI transcript history** NEW 35

Adds paginated transcript history loading in the TUI, unlocking browsing of long session histories.

— UI feature named, no exact navigation given. `rust-v0.148.0-alpha.6`

16 **Amazon Bedrock remote compaction** NEW 30

Enables remote compaction for Amazon Bedrock sessions.

— One-line description, no mechanism given. `rust-v0.148.0-alpha.6`

17 **Token budget context identity and naming** IMPROVED 30

Makes token budget context identity configurable and identifies agents by name in token budget context.

— Grouped thin mentions, no config key named. `rust-v0.148.0-alpha.6`

18 **Concurrent exec-server request dispatch** NEW 25

Adds opt-in concurrent exec-server request dispatch.

— Opt-in noted, no flag or mechanism given. `rust-v0.148.0-alpha.6`

19 **Disable built-in image viewer** NEW 25

Allows disabling the built-in image viewer.

— Setting mentioned but no config key given. `rust-v0.148.0-alpha.6`

20 **Leaf models in multi-agent v2** NEW 20

Supports leaf models in multi-agent v2.

— Bare mention, no explanation of leaf models. `rust-v0.148.0-alpha.6`

21 **ChatGPT cookies in HTTP clients** NEW 20

Supports configured ChatGPT cookies in HTTP clients.

— Bare mention, no config detail. `rust-v0.148.0-alpha.6`

22 **Codex attribution in pull request bodies** NEW 20

Links Codex attribution in pull request bodies.

— Bare mention with no mechanism. `rust-v0.148.0-alpha.6`

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

holmesgpt ⓘ 1 RELEASE · 2026-08-10 NOTES ↗

SRE Agent - CNCF Sandbox Project

> HolmesGPT 0.39.0 adds a new Atlassian Rovo MCP data source and extends the GitHub MCP integration to support multi-organization GitHub App installs.

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

01 **GitHub MCP addon multi-org and credential-based skill pulling**

IMPROVED 32

The GitHub MCP addon now supports GitHub App installations spanning multiple organizations, and can pull skills using GitHub App credentials.

– Describes two related changes but no config keys or commands given 0.39.0

02 **Atlassian Rovo MCP integration** NEW 30

HolmesGPT can now use Atlassian Rovo as a data source via the MCP addon.

– Names the integration but no config or usage detail 0.39.0

WAS THIS USEFUL? 👍 👎

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

Build resilient agents.

> LangGraph 1.2.11 adds per-node control over tracing behaviour.

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

01 **Per-node tracing control via trace_policy** NEW 60

`add_node` now accepts a `trace_policy` parameter that lets developers control tracing behaviour on a per-node basis.

– Names exact parameter and method but no behaviour options detailed. 1.2.11

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

Ollama ⓘ 3 RELEASES · 2026-08-10 → 2026-08-11 NOTES ↗ CODE ↗

Get up and running with Kimi-K2.6, GLM-5.2, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.

> Ollama's latest releases add two new local models — NVIDIA's Nemotron 3.5 Lightning MoE with built-in speculative decoding, and Muse Glimmer with MLX multimodal support and ready-made coding-agent/assistant integrations — alongside a streaming thinking-trace view in the TUI and a fix for namespaced tool calls in the OpenAI Responses API.

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

01 Muse Glimmer model support and agent integrations NEW

90

Adds the Muse Glimmer model runnable via `ollama run muse-glimmer` on NVIDIA, AMD, and additional platforms, plus a `muse-glimmer:30b-mlx` tag for the 30B multimodal variant via `ollama run muse-glimmer:30b-mlx`. Adds `ollama launch claude --model muse-glimmer`, `ollama launch pi --model muse-glimmer`, `ollama launch openclaw --model muse-glimmer`, and `ollama launch hermes --model muse-glimmer` to wire the model into local coding-agent and personal-assistant frameworks. On Apple Silicon, adds DFlash block-diffusion drafting support and image input support to the MLX engine for multimodal vision workloads with Muse Glimmer.

Run Muse Glimmer locally with Claude Code as a coding agent backend — no cloud API key required.

```
$ ollama launch claude --model muse-glimmer
```

Use Muse Glimmer with OpenClaw to self-host a personal AI assistant across messaging platforms.

```
$ ollama launch openclaw --model muse-glimmer
```

Run Muse Glimmer locally on Apple Silicon for interactive chat or to test multimodal (image) input.

```
$ ollama run muse-glimmer:30b-mlx
```

Wire Muse Glimmer into Claude Code as a local coding agent on Apple Silicon.

```
$ ollama launch claude --model muse-glimmer:30b-mlx
```

Launch Muse Glimmer with the OpenClaw personal assistant framework for WhatsApp, Telegram, Slack, or Discord.

```
$ ollama launch openclaw --model muse-glimmer:30b-mlx
```

— Lists every run and launch command and tag verbatim. v0.32.8 v0.32.7

---

02 **Nemotron 3.5 Lightning model support** NEW 88

Adds `nemotron-3.5-lightning`, an NVIDIA 30B mixture-of-experts model with 3B active parameters, runnable via `ollama run nemotron-3.5-lightning`. Adds the Nemotron 3 architecture to the MLX backend, including Mamba2/recurrent layers, MoE routing, and quantized NVFP4/MXFP8 expert paths with Metal-optimized block-mapped kernels, and serves the model's multi-token prediction head as a built-in self-draft speculator, enabling speculative decoding without a separate draft model.

Run the new Nemotron 3.5 Lightning 30B MoE model locally for always-on agent workloads.

```
$ ollama run nemotron-3.5-lightning
```

— Names architecture, quantization paths and decoding mechanism plus run command. v0.32.9

---

03 **Streaming thinking traces in agent TUI** IMPROVED 60

Streams live thinking deltas beneath a `Thinking ↓ N tokens` row in the agent TUI, collapsing to a persistent `Thought` marker once the assistant begins responding, with a `ctrl+o` toggle to show or hide the trace.

— Names UI elements and toggle but no deeper mechanism given. v0.32.7

---

04 **Namespaced tool declarations in Responses API** IMPROVED 60

Expands namespace tool declarations in the OpenAI Responses API, resolving previously undeclarable namespaced tool calls by unfolding nested `tools` arrays into namespace-qualified flat function names.

— Names API and mechanism but no endpoint or example given. v0.32.7

WAS THIS USEFUL? 👍 👎

vMLX

1 RELEASE · 2026-08-11 NOTES ↗

vMLX - JANGTQ Uber Compressed MLX Models - L2 Disk Cache (survives restart) + L1 Paged (super fast ttft) + Hybrid SSM Scheduler + Cont Batching + etc!

> vMLX 1.6.27 adds full support for the Muse Glimmer 30B vision/video model, wires up reasoning-effort control across the Anthropic and Ollama API surfaces, and quadruples the prefix cache to 262,144 tokens with automatic session migration.

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

01 **Reasoning effort control via Anthropic and Ollama APIs** IMPROVED

80

Wires up Anthropic `thinking.budget_tokens` so reasoning chains can be capped, and fixes Ollama `think` string levels (`'low'` , `'medium'` , `'high'`) which now actually engage thinking and select effort depth — previously they were silently ignored.

Cap Muse Glimmer 30B reasoning budget via the Anthropic API so long chains do not run unbounded.

```
$ curl http://localhost:8000/v1/messages \
  -H 'Content-Type: application/json' \
  -d '{
    "model": "local",
    "max_tokens": 4096,
    "thinking": {"type": "enabled", "budget_tokens": 2000},
    "messages": [{"role": "user", "content": "Explain zero-
knowledge proofs."}]
  }'
```

Drive reasoning effort from an Ollama client using the new `think` string level instead of a boolean.

```
$ curl http://localhost:8000/api/chat \
  -H 'Content-Type: application/json' \
  -d '{
    "model": "local",
    "think": "high",
    "messages": [{"role": "user", "content": "Analyze this
network capture for lateral movement."}]
  }'
```

— Runnable curl examples for both named API fields. v1.6.27

02 **Muse Glimmer 30B model support** NEW

65

Adds full support for Muse Glimmer 30B: text tower, windowed vision tower, recipient-routed reasoning rail, and ATEM tool dialect, with all three bundles (`JANG_2D` , `4M` , `6M`) supported.

— Names components and bundles but no run command. v1.6.27

03 **Prefix cache expanded to 262,144 tokens** IMPROVED

60

Prefix cache index expanded from roughly 64,000 to 262,144 tokens, with automatic migration of existing sessions, so long conversations no longer report zero reuse.

— Concrete before/after numbers but no config or command. v1.6.27

04 **Muse Glimmer 30B quality and UI improvements** IMPROVED

60

Muse Glimmer 30B single- and multi-image prompts now correctly identify distinct objects in the same turn, video clips are read in temporal order, Chat Settings exposes all four reasoning depths — Low / Medium / High / Extra High — replacing the previous Auto/On-only toggle, and each reasoning depth maintains its own cache lineage so switching back to a prior depth resumes that chain rather than starting over.

— UI path named (Chat Settings) but no exact flags. v1.6.27

THINNER COVERAGE BELOW

05 **Consistent paged RAM for vision models** IMPROVED

40

Paged RAM is now consistently enabled for vision models between the app and command line, eliminating a silent cache-reuse penalty that previously only affected the app.

— Explains fix but gives no toggle or command. v1.6.27

06 **Control markup scrubbing on malformed streams** IMPROVED

40

Control markup (recipient markers) is now scrubbed from visible Muse answers on malformed or truncated streams, while reasoning text and tool-call bodies are left intact.

— Describes behavior precisely but no user-facing control. v1.6.27

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NoLabs nono ⓘ 1 RELEASE · 2026-08-10 NOTES ↗ CODE ↗

Sandbox any AI agent in seconds - zero setup, zero latency.

> nono v0.73.0 adds glob pattern support to filesystem path fields, allowing more flexible wildcard-based sandbox policies.

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

01 **Glob patterns in filesystem path fields** IMPROVED 60

Sandbox policy profiles now support glob patterns in filesystem path fields such as `fs_read` and `fs_write`, enabling wildcard-based filesystem access rules instead of exact paths only.

— Names exact fields but no examples or syntax details. v0.73.0

WAS THIS USEFUL? 👍 👎

ToolHive ⓘ 1 RELEASE · 2026-08-10 NOTES ↗

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

> ToolHive v0.42.1 closes a fail-open gap in vMCP tool exposure with deny-by-default visibility, tightens MCPOIDCConfig validation of issuer/JWKS URLs, and adds MCP tool annotations for composite tools alongside external OIDC token exchange via trusted_issuers — while two breaking changes enforce stricter Cedar authorization Content-Type checks and block direct calls to excluded vMCP tools.

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

01 Stricter validation and plain-HTTP opt-in for MCPOIDCConfig

95

BREAKING

Adds `insecureAllowHTTP: true` under `spec.inline` of `MCPOIDCConfig` to explicitly opt in to plain-HTTP issuer and JWKS URLs in dev/test environments. Inline configs without it now validate both `issuer` and `jwksUrl` at reconcile time, and any inline `MCPOIDCConfig` with a plain-HTTP, malformed, or scheme-less `issuer` or `jwksUrl` flips to `Valid=False` on its next reconcile, blocking reconciliation of every `MCPServer`, `MCPRemoteProxy`, and `VirtualMCPServer` referencing it; switch to HTTPS or add `insecureAllowHTTP: true` to fix.

Allow a dev/test MCPOIDCConfig to use a plain-HTTP in-cluster Keycloak without breaking reconciliation after upgrading.

yaml

```yaml
apiVersion: toolhive.stacklok.dev/v1beta1
kind: MCPOIDCConfig
metadata:
  name: keycloak-auth
spec:
  type: inline
  inline:
    issuer: http://keycloak:8080/realms/toolhive
    jwksUrl:
http://keycloak:8080/realms/toolhive/protocol/openid-
connect/certs
    insecureAllowHTTP: true
```

— Exact field, failure state and remediation with example v0.42.1

02 Deny-by-default tool visibility in vMCP NEW

85

Adds `aggregation.defaultToolVisibility: deny` to vMCP config so only workloads explicitly listed in `aggregation.tools` have their tools advertised, closing the fail-open gap where adding a workload to a group silently exposed it.

Lock down a vMCP group so only explicitly listed workloads expose tools — prevents newly added workloads from silently advertising all their tools.

```yaml
aggregation:
  defaultToolVisibility: deny
  tools:
    - workload: github
    - workload: jira
```

— Named config keys plus a runnable YAML example `v0.42.1`

---

03 **Excluded vMCP tools no longer directly callable** BREAKING

80

vMCP tools excluded via `aggregation.tools` `filter`, per-workload `excludeAll`, or global `excludeAllTools` are no longer directly callable on the Modern (2026-07-28) path — direct `tools/call` now returns JSON-RPC error `-32602` at HTTP 400 instead of executing; un-filter the tool or wrap it in a composite tool to keep it reachable.

— Exact error code, path name and workaround given `v0.42.1`

---

04 **Content-Type enforcement under Cedar authorization** BREAKING

75

With Cedar authorization enabled (`--authz-config` or `authzConfig` in the CRD), `POST` requests without `Content-Type: application/json` (including a missing header) now return `400` instead of being forwarded; all MCP POST requests must send `Content-Type: application/json`.

— Exact flags and required header, no code sample `v0.42.1`

---

05 **TOOLHIVE_API_TIMEOUT environment variable** NEW

70

Adds `TOOLHIVE_API_TIMEOUT` environment variable to override the CLI's API client timeout for `thv skill` and `thv ai-plugin` commands, with a new default of 10 minutes.

— Exact env var and commands but no usage example `v0.42.1`

---

06 **External OIDC token exchange via trusted_issuers** NEW

70

The embedded auth server now accepts `trusted_issuers` configuration, enabling agents to exchange subject tokens minted by external OIDC providers (Entra, Okta, Keycloak) for ToolHive-scoped delegated tokens under an RFC 8693 fail-closed consent policy.

— Names the config key and protocol but no setup example `v0.42.1`

---

07 **MCP tool annotations for composite tools** NEW 65

Composite tools now support MCP tool annotations — `readOnlyHint`, `destructiveHint`, `idempotentHint`, and `openWorldHint` — with a fail-closed safety floor derived from step tools when none are set explicitly.

— Names all four annotation fields, no config example `v0.42.1`

└──▷ BREAKING ON UPGRADE

! With Cedar authorization enabled (`--authz-config` or `authzConfig` in the CRD), `POST` requests without `Content-Type: application/json` (including a missing header) now return `400` instead of being forwarded; all MCP POST requests must send `Content-Type: application/json`.

! vMCP tools excluded via `aggregation.tools` `filter`, per-workload `excludeAll`, or global `excludeAllTools` are no longer directly callable on the Modern (2026-07-28) path — direct `tools/call` now returns JSON-RPC error `-32602` at HTTP 400 instead of executing; un-filter the tool or wrap it in a composite tool to keep it reachable.

! `MCPOIDCConfig` resources of `spec.type: inline` with a plain-HTTP, malformed, or scheme-less `issuer` or `jwksUrl` now flip to `Valid=False` on their next reconcile and block reconciliation of every `MCPServer`, `MCPRemoteProxy`, and `VirtualMCPServer` referencing them; switch to HTTPS or add `insecureAllowHTTP: true` under `spec.inline`.

WAS THIS USEFUL? 👍 👎

SecureAgentics Adrian ⓘ 1 RELEASE · 2026-08-11 CODE ↗

Open-source runtime AI agent security tool - monitors and controls AI agents, catching malicious tool use, prompt injection, and policy drift in real time, before the agent acts.

> Adrian's Python SDK v1.1.0 ships a Claude Code plugin with server-driven execution modes and human-in-the-loop approval, alongside PII filtering, an Anthropic SDK integration, and dashboard improvements to the Events view.

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

01 Claude Code plugin with server-driven execution modes NEW

80

A new Claude Code plugin (`integrations/claude-code/adrian_cc/agent.py`) connects Claude Code hooks to the Adrian backend via WebSocket, supporting three server-driven execution modes: `MODE_ALERT` (fire-and-forget), `MODE_BLOCK` (wait for verdict, block per policy), and `MODE_HITL` (wait for human approval on the dashboard). HITL verdicts from the Adrian backend are forwarded to Claude Code for native inline approval workflows.

— Names file path and all three modes, no usage command. python-sdk-v1.1.0

02 PII filtering for Claude Code plugin NEW

70

Adds PII filtering to the Claude Code plugin (`adrian_cc/pii/_patterns.py`), detecting 12 PII categories: `EMAIL` , `PHONE` , `SSN` , `CREDIT_CARD` , `IP_ADDRESS` , `DATE_OF_BIRTH` , `IBAN` , `PASSPORT` , `STREET_ADDRESS` , `POSTAL_CODE` , `DRIVER_LICENSE` , and `AWS_KEY` .

— Lists all 12 categories and file path, but no setup steps. python-sdk-v1.1.0

THINNER COVERAGE BELOW

03 Severity badges and verdict metadata in Events view IMPROVED

50

Severity badges based on MAD codes were added to the Events table and mobile cards, and the event list API response now includes latest verdict metadata.

.

— Names MAD codes and API response field, exact UI path unclear. python-sdk-v1.1.0

04 **LangChain/LangGraph instrumentation consolidated** IMPROVED

40

LangChain/LangGraph instrumentation was extracted into `langchain_handler.py`, consolidating that integration surface.

— Names the new file but not what changed behaviourally. python-sdk-v1.1.0

05 **Anthropic SDK integration** NEW

30

Adds Anthropic SDK integration, enabling Adrian instrumentation for agents built with the Anthropic Python SDK.

— No mechanism, file, or setup details given. python-sdk-v1.1.0

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian ⓘ 1 RELEASE · 2026-08-10 BLOG ↗

Obsidian is a personal knowledge management app that lets users create, link, and organize markdown notes locally with powerful search and graph visualization.

> Obsidian Mobile 1.13.6 focuses on small tab-switcher and settings UI refinements rather than major new capabilities.

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

01 **Swipe-to-close gesture in tab switcher** IMPROVED 40

Tabs in the mobile tab switcher can now be dismissed with a swipe gesture, and remaining tabs animate to their new position after a tab is closed.

— Names the gesture and behavior but no mechanism detail

Obsidian 1.13.6 Mobile (Public) (2026-0…

02 **Icon labels in mobile toolbar settings** IMPROVED 35

The mobile toolbar settings page now displays icons alongside each command to make commands easier to identify.

— Names UI location but lacks further mechanism Obsidian 1.13.6 Mobile (Public) (2026-0…

03 **Checkbox scaling with font size** IMPROVED 20

Checkboxes now scale proportionally to the configured font size.

— Bare description with no configuration path given

Obsidian 1.13.6 Mobile (Public) (2026-0…

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare ⓘ 1 RELEASE · 2026-08-10 BLOG ↗

Cloudflare is a global network platform providing CDN, DDoS protection, DNS, security, and developer services for websites and applications.

> Cloudflare's biggest window items are a managed Apache Iceberg catalog built directly into R2, general availability for MySQL in Hyperdrive and for Turnstile Spin, and unified billing across Workers AI and AI Gateway, alongside new outbound traffic controls for Sandbox, new Radar BGP API endpoints, and Docker/Kubernetes support for Mesh.

└──▷ **WHAT SHIPPED · 16 FEATURES** most completely described first ? what's the number?

01 **Turnstile Spin general availability** NEW 94

Turnstile Spin is now GA: `wrangler turnstile widget create` creates a widget and returns the sitekey and secret; the Set up with Spin flow in the Turnstile dashboard lets you enter domains and generate a widget with sitekey, secret, and agent prompt in one step; pasting the Spin prompt into Claude Code, Cursor, Codex, OpenCode, or GitHub Copilot Chat lets the agent fetch the skill, create the widget, embed it, and wire `siteverify` automatically; it includes server-side `siteverify` insertion snippets for Next.js (App Router and Pages Router), Astro, SvelteKit, Hugo, and vanilla HTML; and an automated validation step runs a real Turnstile token through the protected endpoint to confirm it passes, then replays it to confirm the endpoint rejects it, stopping with diagnostics on failure.

Quickly scaffold a Turnstile widget from the terminal and get the credentials needed to wire siteverify by hand.

```
$ wrangler turnstile widget create
```

— Runnable CLI command plus full setup flow and frameworks.

Turnstile - Turnstile Spin is now gener…

02 **MySQL support in Hyperdrive GA** NEW 91

MySQL and MySQL-compatible databases are now generally available in Hyperdrive — connect from Workers using `mysql2/promise` via `env.HYPERDRIVE.host`, `.user`, `.password`, `.database`, and `.port` binding properties, with no code changes required for existing drivers, ORMs, or query builders.

Connect a Workers function to a MySQL database via Hyperdrive using the GA mysql2 integration.

```javascript
import { createConnection } from 'mysql2/promise';

export default {
  async fetch(request, env, ctx) {
    const connection = await createConnection({
      host: env.HYPERDRIVE.host,
      user: env.HYPERDRIVE.user,
      password: env.HYPERDRIVE.password,
      database: env.HYPERDRIVE.database,
      port: env.HYPERDRIVE.port,
      disableEval: true, // Required for Workers compatibility
    });
    const [results, fields] = await connection.query('SHOW
tables;');
    ctx.waitUntil(connection.end());
    return new Response(JSON.stringify({ results, fields }), {
      headers: { 'Content-Type': 'application/json' },
    });
  },
};
```

— Full runnable code example with binding properties. product docs

03 **Unified billing across Workers AI and AI Gateway** NEW

90

Workers AI and AI Gateway now support unified billing via a Workers AI billing setting on the gateway set to Unified Billing, allowing prepaid AI Gateway credits to pay for Workers AI inference with a shared credit balance. Routing requests through AI Gateway now supports a `gateway: { id: 'default' }` option in env.AI.run() calls, which auto-creates a gateway on the first authenticated request. Three frontier Workers AI models — `@cf/moonshotai/kimi-k2.6`, `@cf/moonshotai/kimi-k2.7-code`, and `@cf/zai-org/glm-5.2` — are now accessible without the Workers Paid plan when billed via AI Gateway credits, at 50 requests per minute per account per model.

Route a Workers AI inference request through AI Gateway for observability, caching, and unified billing using the `gateway` option.

```javascript
const response = await env.AI.run(
  '@cf/zai-org/glm-5.2',
  { messages: [{ role: 'user', content: 'What is the capital of France?' }] },
  { gateway: { id: 'default' } }
);
```

— Names setting, code option, and models, with runnable example. product docs

04 **Radar BGP path and upstream API endpoints** NEW 88

New `/bgp/routes/paths/{asn}` Radar BGP API endpoint returns ordered AS path segments an AS uses to reach Tier-1 networks, with a `collector` parameter to scope to a single RouteViews collector. New `/bgp/routes/upstreams/{asn}/timeseries` endpoint returns the share of an AS's observed paths per upstream over time, with `limit` and `ipVersion` parameters.

Retrieve AS-level BGP path segments toward Tier-1 networks for a specific ASN, scoped to a single RouteViews collector.

```
$ curl 'https://api.cloudflare.com/client/v4/radar/bgp/routes/paths/13335?collector=route-views.oregon' \
  -H 'Authorization: Bearer <API_TOKEN>'
```

— Runnable curl example with named endpoints and parameters. product docs

05 **R2 Data Catalog for Apache Iceberg** NEW 78

R2 Data Catalog is a managed Apache Iceberg REST catalog built directly into R2 buckets, letting Spark, Snowflake, PyIceberg, DuckDB, Trino, ClickHouse, and R2 SQL query R2 as a data warehouse or lakehouse with no separate catalog infrastructure required. It supports ACID transactions, optimized metadata indexing, and full schema evolution (add, rename, delete columns without rewriting data), and enables cross-cloud and cross-region analytics — log analytics, BI, data pipelines — with no egress fees under R2's zero-egress model.

— Mechanism and integrations named, no setup command shown. product docs

06 **Docker and Kubernetes support for Cloudflare Mesh** NEW 78

Cloudflare Mesh nodes can now run as Docker containers via the `cloudflare/mesh` image on Docker Hub, supporting `amd64` and `arm64` architectures, Docker Compose, Kubernetes StatefulSet, Kubernetes sidecar, and CI/CD pipeline deployment patterns.

— Image and deployment patterns named, no exact command. product docs

07 **Cloudflare Skills and sandbox@next preview SDK** NEW

73

Adds Cloudflare Skills for coding agents with three install targets: `sandbox-stable` for the current stable package, `sandbox-next` for `@cloudflare/sandbox@next` (recommended for new projects), and `sandbox-migrate-to-next` to port existing apps. Also introduces `@cloudflare/sandbox@next` as a 1.0 preview SDK for early adoption of the upcoming Sandbox SDK 1.0 release.

— Named install targets given but no full install command. product docs

08 **Queue message peek endpoint** NEW

73

New `POST /accounts/{account_id}/queues/{queue_id}/messages/peek` endpoint inspects queued messages without consuming them.

Inspect pending messages in a Cloudflare Queue without consuming them, useful for debugging or monitoring queue depth.

```
$ curl -X POST
'https://api.cloudflare.com/client/v4/accounts/{account_id}/queu
es/{queue_id}/messages/peek' \
  -H 'Authorization: Bearer <API_TOKEN>'
```

— Runnable curl example against the named endpoint. product docs

09 **Live log streaming for Tunnel** NEW

70

A new Live logs tab under Networking > Tunnels > (tunnel detail view) in the Cloudflare dashboard streams real-time logs from single or multiple `cloudflared` connectors, with filtering by log level, event type (HTTP, TCP, UDP, or `cloudflared` internal), and HTTP method.

— Exact dashboard navigation given, no CLI equivalent. product docs

10 **AS connectivity graph and upstream chart in Radar** NEW

68

Cloudflare Radar adds an AS-level connectivity graph and an Upstream providers stacked-area chart to AS pages (e.g. AS13335), with a Show full paths toggle and an `ipVersion` selector for IPv4/IPv6.

— Exact UI page and controls named, no API surface. product docs

11 **Outbound traffic control in Sandbox** NEW

65

Cloudflare Sandbox now supports blocking, allowing, and intercepting outbound HTTP requests from sandboxes, with the ability to inject authorization headers in outbound handlers so credentials stay inside the Worker rather than the sandbox.

— Mechanism described but no code example given. product docs

12 **ByteDance Seedance 2.5 video model** NEW 60

Adds `seedance-2.5`, ByteDance's audio-video generation model supporting 30-second video creation, extended storytelling, audio and visual editing, white-model control, and green-screen editing.

— Model and capabilities named but no invocation example. product docs

THINNER COVERAGE BELOW

13 **URL intelligence lookup endpoint** NEW 50

New `GET /accounts/{account_id}/intel/url` endpoint retrieves URL intelligence data for a given URL.

— Endpoint named but no request example given. product docs

14 **Radar Researcher AI assistant and WebMCP support** NEW 49

Cloudflare Radar adds Radar Researcher, a beta AI-powered assistant for exploring Internet trends in plain language via voice or text, accessible from the header on any Radar page. Radar also adds WebMCP support for browser-based AI agents.

— WebMCP support mentioned with no further detail. product docs

15 **Cloudflare One Client for Windows update** IMPROVED 28

New Cloudflare One Client for Windows (version 2026.6.905.0) released.

— Only a version number is given, no changelog detail. product docs

16 **Dataset deletion endpoint** NEW 21

New endpoint to delete an account or zone dataset.

— No method, path, or example provided. product docs

WAS THIS USEFUL?  

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 11, 2026

Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe