THE CYBER TOOLCHAIN

NO. 049

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 18, 2026 · EVERY DAY

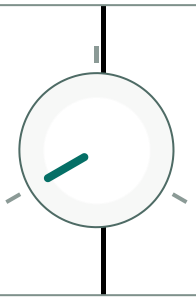
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 38 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

Vulners getsploit

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

Search and download public exploits from the Vulners database — online, or fully offline from a local SQLite FTS5 index.

> getsploit 3.0.0 introduces a fully offline SQLite FTS5 exploit index built and managed via `--update`, `--local`, and `--status`, moves online search onto the async Vulners Python SDK 4.1 (dropping direct HTTP access), and rounds out the CLI with multi-format output, `--mirror`, `--color`, and `--set-key` API-key handling.

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

01 Offline exploit search via `--local` FTS5 index NEW

95

Adds `--local` flag for fully offline exploit search against a local SQLite FTS5 index of the full exploit corpus (~1.7 GB), stored at `~/.getsploit/getsploit.db` (overridable via `$GETSPLOIT_HOME`). Supports full FTS5 expression syntax — including NEAR(...), boolean operators, column filters, and wildcards — across searchable columns `id`, `title`, `published`, `description`, and `source_data` (legacy `sourceData:` spelling still accepted), e.g. `title:eternalblue AND NOT description:metasploit`.

Build the local exploit index then search entirely offline — useful in air-gapped or no-network environments.

```
$ getsploit --update && getsploit --local 'title:eternalblue AND NOT description:metasploit'
```

Build the local offline index once, then search without a network connection or API key.

```
$ getsploit --update
getsploit --local 'title:eternalblue AND NOT description:metasploit'
```

Build a local offline index once, then search for exploits without any network access.

```
$ getsploit --update
getsploit --local 'title:eternalblue AND NOT
description:metasploit'
```

— Names flag, path, columns and runnable query syntax `product docs`

02

`--update` builds and refreshes the local exploit index NEW

88

Adds `--update` to download the full exploit corpus archive (~1.7 GB) over eight parallel range connections and build or atomically replace the local FTS5 index at `~/.getsploit/getsploit.db` (or `$GETSPLOIT_HOME`); a search running against the old database finishes undisturbed. The index is versioned with integrity checks, and existing FTS4 archives from Getsploit 2.x are automatically converted to FTS5. `--update` reports detailed progress — bytes and throughput for the download, and a bar with totals and time remaining for unpacking and document conversion — keeping the archive on disk rather than buffering it in memory.

— Names flag, size, connection count, path and progress mechanics `v3.0.0`

03 **Multi-format and adaptive terminal output** NEW 83

Adds `--format` flag supporting `json`, `jsonl`, and tab-separated (TSV) output, with `--json` retained as an alias for `--format json`; output automatically switches to stable tab-separated values with no color or decoration when redirected. Terminal display is rendered through Rich with adaptive `table`, `compact`, `JSON`, and `JSONL` modes.

Pipe JSON Lines results into jq to extract exploit IDs for downstream processing in a CI or reporting pipeline.

```
$ getsploit --format jsonl wordpress | jq -r .id
```

Stream exploit IDs matching a query into jq for downstream pipeline processing.

```
$ getsploit --format jsonl wordpress | jq -r .id
```

Pipe offline exploit IDs into jq or other tools using JSON Lines output.

```
$ getsploit --local --format jsonl wordpress 4.7 | jq -r .id
```

— Names every format flag and mode with default behaviour v3.0.0

04 **`--mirror` saves exploit bodies to local files** NEW 83

Adds `--mirror` flag to save each exploit body to its own file in a per-query directory, usable with both online (`--mirror`) and offline (`--local --mirror`) searches; files are created without following symlinks and never overwrite existing files.

Mirror exploit source files from an offline search for manual review before any execution.

```
$ getsploit --local --mirror eternalblue
```

Mirror all exploit source files for a CVE locally for offline review before use.

```
$ getsploit --local --mirror CVE-2024-3094
```

Download and save exploit source files locally for manual review before any analysis.

```
$ getsploit --local --mirror eternalblue
```

— Exact flag with concrete safety behaviour and runnable examples `product docs`

05 **API key storage via `--set-key` replaces `--api-key`** BREAKING

83

Adds `--set-key` to store a Vulners API key via a hidden confirmation prompt into a private file under the Getsploit home directory, with the `VULNERS_API_KEY` environment variable taking precedence over the stored file. The `--api-key VALUE` option is removed in v3.0.0, so credentials must now be supplied via `--set-key` or `VULNERS_API_KEY`.

— Names old and new flags plus the env var to use instead `v3.0.0`

06 **`--color` controls terminal rendering** NEW

75

Adds `--color` flag with values `always`, `never`, and `auto` (default) to control terminal color and decoration; wide terminals render a table, narrow terminals render stacked records, and redirected streams receive tab-separated output.

— Names flag values and exact rendering behaviour per mode `product docs`

07 **FTS4 databases migrate to FTS5 on update** BREAKING

70

Databases built by Getsploit 2.x (FTS4) are automatically migrated to FTS5 the next time `--update` runs; existing FTS4 proximity expressions such as `one NEAR/5 two` must be rewritten in the FTS5 NEAR(...) form to keep working.

— Exact before/after query syntax given for migration `v3.0.0`

08 **Online search pagination beyond 100 results** IMPROVED

70

Rebuilds `--count` above 100 to page through results rather than silently returning only the first 100, advancing by rows received; Vulners Lucene syntax (e.g. `title:wordpress AND description:"code execution"`) is supported for online queries with transparent pagination above 100 results, since no single API response exceeds 100 documents.

— Names flag and query syntax with the 100-doc limit explained `v3.0.0`

09 **Online search moves to Vulners SDK 4.1** BREAKING

60

Online search now runs asynchronously through the official Vulners Python SDK 4.1, replacing direct HTTP access to Vulners endpoints; runtime dependencies on `httpx`, `orjson`, and `texttable` are removed as a result.

— Explains mechanism and removed dependencies, nothing to run v3.0.0

THINNER COVERAGE BELOW

10 **`--status` reports local database state** NEW 50

Adds `--status` to report the local database location, size, and document count without requiring an API key.

Check local database size and path after upgrading to the new FTS5 index.

```
$ getsploit --status
```

— Exact command given, but only a one-line description v3.0.0

11 **Broader Python version and platform support** IMPROVED 35

Supports Python 3.11 through 3.14 on Linux, macOS, and Windows, built against the Vulners SDK 4.1+ with a reproducible `uv` lockfile.

— Plain compatibility statement with no further mechanism product docs

└──▷ **BREAKING ON UPGRADE**

! FTS4 proximity expressions such as `one NEAR/5 two` must be rewritten in the FTS5 NEAR(...) form after the database is migrated from Getsploit 2.x.

! The `--api-key VALUE` option is removed; use the API-key file instead.

! Direct HTTP access to Vulners endpoints is removed; all online search goes through the Vulners Python SDK.

! Runtime dependencies on `httpx`, `orjson`, and `texttable` are removed.

! Local FTS4 archives are converted to an FTS5 index during `--update`; rebuilt indexes use FTS5 query syntax, so FTS4 `NEAR/n` expressions must be rewritten as NEAR(...) in FTS5 form.

! FTS4 proximity expressions such as `one NEAR/5 two` in local queries must be rewritten in the FTS5 NEAR(...) form after migrating a Getsploit 2.x database.

! Databases built by Getsploit 2.x are migrated from FTS4 to FTS5 on the next `--update`; FTS4 proximity expressions such as `one NEAR/5 two` must be rewritten in the FTS5 NEAR(...) form.

! The `--api-key VALUE` option is removed; replace it with the `VULNERS_API_KEY` environment variable or `--set-key`.

! Local indexes rebuilt under v3.0.0 use FTS5 query syntax; FTS4 `NEAR/n` expressions must be rewritten as the equivalent FTS5 NEAR(...) form.

! Direct HTTP access to Vulners endpoints is removed; all online search now goes through the Vulners Python SDK.

! Runtime dependencies on `httpx`, `orjson`, and `texttable` are removed; environments that installed those transitively through getsploit may need adjustment.

WAS THIS USEFUL? 👍 👎

Bugcrowd

Feedback ⓘ 1 RELEASE · 2026-08-17 BLOG ↗

Bugcrowd is a crowdsourced security testing platform connecting organizations with ethical hackers to identify and report vulnerabilities.

> Bugcrowd launched Savant Pathseeker, a new autonomous outside-in penetration testing product (Early Access 1.0), with authenticated testing, WAF/bot evasion, inline asset management, and one-click Security Inbox integration for validated findings.

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

01 Authenticated testing behind login and external IdPs NEW

80

Pathseeker can log in with stored credentials and test content behind login pages, including apps that redirect through external identity providers such as Okta, Entra ID, or Google. To configure it, open the credential step in Autonomous Tests > New Autonomous Test, enter the test account username and password, and save — Pathseeker logs in automatically when it detects a login page and tests everything behind the authentication wall.

Set up an authenticated Pathseeker test so it can reach and assess content behind a login page, including an Okta or Entra ID redirect.

> 📍 In Autonomous Tests › New Autonomous Test, open the credential step, enter the username and password for the test account, and save. Pathseeker will log in automatically when the test detects a login page and test everything behind the authentication wall.

— Names IdPs and gives exact UI steps to enable it

Savant Pathseeker Early Access (1.0) (2…

02 Push to Inbox findings integration NEW

75

A Push to Inbox action promotes validated Pathseeker findings into the Security Inbox in one click; findings are labeled Autonomous Test Vulnerabilities and link back to the originating test run, and duplicate pushes are blocked. From the test Results tab, select one or more findings, click Push to Inbox, and confirm — the finding appears in Security Inbox within seconds, integrated alongside other detections and submissions without manual re-entry.

Promote a validated finding into the vulnerability-management workflow without manual re-entry.

📍 From the test Results tab, select one or more findings, then click `Push to Inbox` and confirm in the dialog. The finding appears in Security Inbox within seconds, labeled as `Autonomous Test Vulnerabilities`.

— Exact UI steps, label, and dedup behaviour all named

Savant Pathseeker Early Access (1.0) (2…

03 **Inline asset search and creation for test targeting** NEW

70

Practitioners can search an approved asset inventory or add and approve a new asset — name, URL, type, and business criticality — directly from Pathseeker test setup without switching tabs. In the target field, typing an asset name that has no match offers 'Create and approve a new asset?', which creates, approves, and adds the asset as a target in one action.

Add a net-new domain to scope without leaving the Pathseeker test setup wizard.

📍 In the target field, type the asset name. When no match is found, select 'Create and approve a new asset?', complete the name, URL, type, and business criticality fields, then save. Pathseeker creates, approves, and adds the asset as a target in one action.

— Names exact fields and the create-asset UI flow

Savant Pathseeker Early Access (1.0) (2…

04 **Savant Pathseeker autonomous pentest launch** NEW 60

Savant Pathseeker is a new product that runs autonomous, controlled outside-in penetration tests against approved internet-facing domains, web apps, and APIs, producing validated findings with evidence. It is now available as Early Access (1.0); launching a test requires organization owner permissions and targets must be pre-approved in the asset inventory before testing can begin.

— Names product, version and permission requirement but no exact launch command

Savant Pathseeker Early Access (1.0) (2…

THINNER COVERAGE BELOW

05 **Bot detection and WAF evasion** NEW 40

Pathseeker includes a bot detection and WAF evasion capability that adapts to perimeter controls within authorized boundaries, so test coverage reflects what a real attacker could reach.

— Describes mechanism but gives no configurable surface product docs

WAS THIS USEFUL?  

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Semgrep

Feedback ⓘ 1 RELEASE · SEEN 2026-08-18 NOTES ↗

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

> Semgrep's latest snapshot headlines with multi-finding Autofix pull requests that batch fixes for shared-rule findings, alongside an experimental v2 API for Code findings, several remediation and workflow UI refinements, and breaking changes to the Autofix toggle and rule deletion flow.

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

01 **Code Autofix toggle removed from Settings** BREAKING

60

The Code Autofix toggle has been removed from Settings; Supply Chain retains its own separate toggles for Upgrade Guidance and Supply Chain Autofix.

— Names exact toggles and settings location affected. snapshot-20260818

THINNER COVERAGE BELOW

02 **Branch filter on Workflow jobs tab** NEW

55

Adds a branch filter to the Workflow jobs tab on a project's Details page, with the primary branch pinned to the top of the list.

— Exact UI location given, no further mechanism. snapshot-20260818

03 **Confirmation required for rule deletion in Editor** BREAKING

55

Deleting a rule in Semgrep Editor now requires confirmation and automatically removes the rule from all policies; policies with no remaining conditions are also deleted.

— Describes cascading policy behavior and UI location. snapshot-20260818

04 **Multi-finding Autofix pull requests** NEW

50

When several findings in the same file share a rule, Semgrep now offers a single pull request or merge request to fix them all at once instead of one PR per finding.

— Explains scope and mechanism but no exact command or UI path. snapshot-20260818

05 **Name-based search in remediation policy workflows** NEW

45

Adds name-based search when configuring the `Run a workflow` action in a remediation policy.

— Names the exact action but limited further detail. snapshot-20260818

06 **Memories sorted by finding count** IMPROVED 45

Sorts suggested memories on Rules & Policies > Memories by the number of findings they affect.

— Clear UI path, minimal mechanism described. snapshot-20260818

07 **Affected packages on Advisories Details page** IMPROVED 45

Supply Chain's Advisories Details page now shows affected packages directly below the description.

— Exact UI location named, thin on mechanism. snapshot-20260818

08 **Free-tier AI scan limit message** IMPROVED 45

When a scan hits the free-tier AI scan limit, Semgrep now shows a `Free-tier AI scan limit reached` message instead of a generic internal error, and does not offer a Retry scan option.

— Describes before/after UX but no user action. snapshot-20260818

09 **Code findings in v2 API** NEW 35

Code findings are now available through the experimental v2 API.

— Names the API version but no endpoint path given. snapshot-20260818

└──▷ **BREAKING ON UPGRADE**

! The Code Autofix toggle has been removed from Settings; Supply Chain retains its own separate toggles for Upgrade Guidance and Supply Chain Autofix.

! Deleting a rule in Semgrep Editor now requires confirmation and automatically removes the rule from all policies; policies with no remaining conditions are also deleted.

WAS THIS USEFUL?  

StepSecurity

Feedback ⓘ 2 RELEASES · 2026-08-17 → 2026-08-18

BLOG ↗

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

> StepSecurity's Threat Center now tells you which incidents actually affect your environment, while Harden Runner adds a deny-list egress mode and AWS CodeBuild custom image support.

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

01 **Deny list egress policy in Harden Runner** NEW 95

Harden-Runner GitHub Action v2.21.0 adds a `denied-endpoints` input, used with `egress-policy: block`, to block specific destinations while leaving all other traffic reachable, avoiding the need to build a full allow list. Self-Hosted VM Harden Runner v1.9.0 adds the equivalent `denied_endpoints` config key with `egress_policy: block` for the same capability on self-hosted VMs. Deny lists for both the GitHub Action and Self-Hosted VM runner can now be managed centrally through the Policy Store instead of editing workflow or config files per job.

Block traffic to a known-bad domain and a wildcard subdomain range in a GitHub Actions workflow without allowlisting every other destination.

yaml

```yaml
- uses: step-security/harden-runner@v2.21.0
  with:
    egress-policy: block
    denied-endpoints: |
      malicious.example.com
      *.ads.example.org
```

— Exact flags, config keys and a runnable workflow example provided

Harden Runner Releases: August 17, 2026…

02 **Incident exposure scoping in Threat Center** NEW 81

Threat Center adds an Affecting You filter that toggles between All incidents and a scoped view of only incidents matching your environment, each with a live count, plus a You are affected badge shown on affected incidents in both views. Affected incidents get a matched-components breakdown listing compromised packages found in your repositories, split by open pull request vs. default branch, down to affected files, head commit, and matched version, and a matched-IOCs correlation view mapping incident

indicators of compromise (e.g. malicious network endpoints) against Harden-Runner runtime data with a per-run breakdown of repository, workflow, job, and run. New Affected packages and Exact version only notification granularity options let alerts fire only for incidents touching your own dependencies.

— Detailed mechanism and named UI controls, but no API/CLI surface

See Which Threat Center Incidents Affec…

THINNER COVERAGE BELOW

03 **AWS CodeBuild custom image support** NEW 38

Harden-Runner extends support to AWS CodeBuild runners using Linux containers on EC2 compute, covering the custom image scenario.

— Single-line mention with no configuration detail given

Harden Runner Releases: August 17, 2026…

WAS THIS USEFUL?  

◆ Web Application Security

DataDome

Feedback ⓘ 3 RELEASES · 2026-08-17 → 2026-08-18 BLOG ↗

DataDome is a bot management and API protection platform that defends applications against automated attacks and fraud using AI-driven detection.

> DataDome's biggest change this window is a breaking overhaul of Apache module URL exclusion handling, alongside a wave of GraphQL/MCP endpoint configuration additions across the Cloudflare Worker, OpenResty, and Node.js modules, plus enriched request headers for Cloudflare Worker logging.

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

01 Apache URL exclusion split into dedicated directives

BREAKING

95

Apache module 3.0.0 splits `DomeURIRegexExclusion` (now holding only custom URL exclusions) from a new `DomeStaticAssetExclusion` directive (enabled by default) that covers static-asset extensions including `avif`, `gz`, `json`, `xml`, and `zip`. It also adds `DomeRemoveMatrixParams` (default `On`) to strip matrix parameters (`;...`) per path segment per RFC 3986 before exclusion regex evaluation, and `DomeEncodedCharInclusion` (default `"%23|%25|%2f|%3b|%3f"`) to force inspection of static-looking paths containing encoded delimiters `%23`, `%25`, `%2f`, `%3b`, or `%3f`. Custom static-asset rules previously in `DomeURIRegexExclusion` must be migrated to `DomeStaticAssetExclusion`.

Migrate a 2.x Apache config to 3.0 by separating custom URL exclusions from static-asset extensions across the two new directives.

```
$ # Before (2.x)
DomeURIRegexExclusion="/health|/heartbeat|\.
(css|js|png|gif|jpg)$"

# After (3.0)
DomeURIRegexExclusion="/health|/heartbeat"
DomeStaticAssetExclusion="\.(css|js|png|gif|jpg)$"
```

— Names every directive and default, with a runnable migration config example

product docs

02 GraphQL/MCP endpoint config and enriched headers in Cloudflare Worker

NEW

90

Cloudflare Worker module 2.7.0 adds `DATADOME_GRAPHQL_ENDPOINT` (Terraform/Wrangler/dashboard config, defaults to `/graphql`) and its npm equivalent `graphQLEndpoint`, letting one or more comma-separated GraphQL endpoints be specified so operation name and type are extracted for analysis. It also adds `DATADOME_MCP_ENDPOINT` (defaults to `/mcp`) and npm `mcpEndpoint` to designate one or more MCP endpoints for MCP-property extraction, and exposes enriched request metadata via `ctx.dataDomeEnrichedHeaders` after calling dataDomeHandler.fetch() for downstream logging in Worker code.

Access enriched DataDome headers in your Cloudflare Worker to log or forward request classification metadata downstream.

typescript

```
async fetch(
  request: Request,
  env: Env,
  ctx: ExecutionContext,
): Promise<Response> {
  const response = await dataDomeHandler.fetch(request as any,
env, ctx);
  console.log(
    Object.entries(ctx.dataDomeEnrichedHeaders)
      .map(([name, value]) => `${name}: ${value}`)
      .join(' - '),
  );
  return response;
}
```

Override the default GraphQL and MCP endpoints when your app exposes them at non-standard paths.

```
$ DATADOME_GRAPHQL_ENDPOINT=/api/graphql
DATADOME_MCP_ENDPOINT=/api/mcp
```

— Names every config key and default, includes runnable code and env examples

Cloudflare Worker module 2.7.0 (2026-08…

03 **GraphQL endpoint config and logger validation in Node.js module**

NEW

70

Node.js module 2.6.0 adds a `graphQLEndpoint` parameter to specify the path(s) for the GraphQL endpoint, defaulting to `/graphql`, and adds validation of the `logger` parameter to ensure it contains every required method with the correct signature.

Override the default GraphQL path when your API mounts GraphQL at a non-standard route.

```json
{
  graphQLEndpoint: '/api/graphql'
}
```

— Names the config parameter and default, with a config override example

Node.js module 2.6.0 (2026-08-17)

04 **GraphQL endpoint and body-size config in OpenResty** NEW

65

OpenResty module 1.7.0 adds `datadome_graphql_endpoint` to specify the GraphQL endpoint (defaults to `graphql`) and `datadome_graphql_max_body_size` to limit the request body size parsed for GraphQL requests.

— Names both config parameters with a default value, no usage example

OpenResty module 1.7.0 (2026-08-18)

THINNER COVERAGE BELOW

05 **DebugMode parameter deprecated in Apache module** DEPRECATED

35

The `DebugMode` parameter in the Apache module is deprecated as of version 2.26, replaced by a new directive used to set the API request timeout (in ms) for both new and reused connections.

— Replacement directive's name isn't given, limiting actionability product docs

└──▷ **BREAKING ON UPGRADE**

! `DomeURIRegexExclusion` now holds only custom URL exclusions; static-asset extension exclusions previously in `DomeURIRegexExclusion` move to the new `DomeStaticAssetExclusion` directive — any custom static-asset rules in `DomeURIRegexExclusion` must be migrated.

! In 3.0.0, static-asset extension patterns must be moved out of `DomeURIRegexExclusion` into the new `DomeStaticAssetExclusion` directive; leaving them in `DomeURIRegexExclusion` duplicates or conflicts with the new default list.

! `DomeRemoveMatrixParams` is enabled (`On`) by default in 3.0.0, meaning matrix parameters are stripped before exclusion regex matching — setups relying on matrix params being passed through unchanged may behave differently.

! `DomeEncodedCharInclusion` is enabled by default in 3.0.0, causing static-looking paths containing `%23`, `%25`, `%2f`, `%3b`, or `%3f` to be inspected rather than excluded.

! The `DebugMode` parameter is deprecated since version 2.26 and replaced by a new directive used as 'API request timeout for both new and reused connections in ms'.

WAS THIS USEFUL? 👍 👎

◆ **Secrets Detection & Management**

infisical

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

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

> Infisical v0.162.20 adds automatic certificate renewal for KMIP and a standardized provider form for dynamic secrets, alongside a small UI usability tweak.

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

01 **Auto certificate renewal for KMIP** NEW 45

Infisical now automatically renews KMIP certificates, reducing the manual overhead of certificate lifecycle management.

— No mechanism, trigger, or config detail described. v0.162.20

02 **Escape key dismisses selected actions** NEW 45

Users can now press the Escape key to dismiss selected actions in the UI.

— Simple, self-contained UI shortcut. v0.162.20

03 **Shared provider form for dynamic secrets** IMPROVED 30

Adds a shared provider form contract for dynamic secrets, standardizing how dynamic secret providers are configured.

— No named providers, fields, or UI path given. v0.162.20

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

CAPEv2 Feedback ⓘ changes since 2026-08-17 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds a ColoredFormatter class to utils/dist.py for color-coded, ANSI-highlighted log output.

Adds `ColoredFormatter` class in `utils/dist.py` that color-codes log level labels by severity: cyan for DEBUG, green for INFO, yellow for WARNING, red for ERROR, and bold-red for CRITICAL, using ANSI escape codes.

Dynamically highlights key entities in log messages — worker names, task IDs, main/analysis IDs, and performance timings — using pre-compiled regex patterns (`RE_WORKER`, `RE_SETSTAT`, `RE_MAIN_TASK`, `RE_WORKER_TASK`, `RE_PERF`).

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Docker Compose

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

Define and run multi-container applications with Docker

> Docker Compose v5.5.0 improves image digest handling in `compose up` and teaches `compose pull` to respect `pull_policy` refresh windows to avoid redundant pulls.

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

01 **Pull policy refresh windows for compose pull** NEW 80

`compose pull` now honors `pull_policy` refresh windows (`daily`, `weekly`, `every_N`), skipping redundant pulls within the configured window. Set via the `pull_policy` key in `compose.yaml`, e.g. `pull_policy: weekly` under a service to skip re-pulling images pulled within the last 7 days.

Use a weekly pull policy so `compose pull` skips re-pulling images that were already pulled within the last 7 days.

yaml

```yaml
services:
  web:
    image: myapp:latest
    pull_policy: weekly
```

— Names config key, values, and runnable example v5.5.0

THINNER COVERAGE BELOW

02 **Image digest reconciliation in compose up** BREAKING 50

New image digest reconciliation logic in `compose up` re-evaluates digests more accurately to prevent unnecessary container recreation. As a side effect, existing containers may be recreated the first time `compose up` is run after upgrading to v5.5.0, since digests are re-evaluated under the new logic.

— Explains mechanism and upgrade impact but no config surface v5.5.0

└──▷ **BREAKING ON UPGRADE**

! Existing containers may be recreated the first time you run `compose up` after upgrading to v5.5.0, as image digests are re-evaluated using the new reconciliation logic.

WAS THIS USEFUL?  

Chainguard

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

Chainguard provides container image security and supply chain management solutions, offering hardened, minimal Linux images and vulnerability scanning.

> Chainguard launched a Public Skills MCP server for AI agents to discover and run hardened images, extended its Guardener GitHub App to migrate GitHub Actions workflows, and rolled out build-pinning for Libraries cache alongside sharper npm error reporting and 17 new hardened container images.

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

01 Public Skills MCP server for hardened agent skills NEW

96

A new Public Skills MCP server is reachable at `https://skills.cgr.dev/mcp` over Streamable HTTP with OAuth 2.0, letting any MCP-capable agent discover and load hardened agent skills without a separate install. It exposes five tools — `search_skills`, `list_skills`, `read_skill`, `read_skill_reference`, and `run_skill` — for ranked keyword search, paginated catalog browsing, manifest inspection, single-file fetching, and full bundle loading. `search_skills` and `list_skills` take `query` / `page_size` (default 50, max 200)/ `page_token`; `run_skill` returns `{ name, files, instructions }` in one call; `read_skill_reference` rejects path-traversal attempts like `../another-skill/...`. Skills are sourced from `skills.cgr.dev/public` and `skills.cgr.dev/chainguard`, pinned to the upstream commit they were hardened from. It can be added to Claude Code with `claude mcp add --transport http --scope user cgr-skills https://skills.cgr.dev/mcp`, with `local`, `project`, and `user` scopes, or shared project-wide via a committed `.mcp.json`.

Add the hardened skill catalog to Claude Code so it is available across all your projects and authenticate via OAuth.

```
$ claude mcp add --transport http --scope user cgr-skills
https://skills.cgr.dev/mcp
claude
/mcp
```

Share the Public Skills MCP server with everyone working in a repo by committing a project-scoped config file.

```json
{
  "mcpServers": {
    "cgr-skills": {
      "type": "http",
      "url": "https://skills.cgr.dev/mcp"
    }
  }
}
```

— Full mechanism, named endpoint, tools, flags, and runnable commands `product docs`

02 **Build pinning for Chainguard Libraries cache** BREAKING

95

A new `chainctl libraries cache` build-pinning feature records which copy (upstream mirror or Chainguard rebuild) an org first received for a package version and keeps resolving that version to the same copy, preventing `EINTEGRITY` checksum mismatches in npm. New commands `chainctl libraries cache opt-out` and `chainctl libraries cache opt-in`, both accepting `--ecosystem`, control per-organization, per-ecosystem pinning behavior, and `chainctl libraries cache list` shows which copy each held version was sourced from. As of August 24, 2026 this pinning is enabled by default for orgs using upstream fallback (JavaScript, Python, and Java ecosystems), so resolutions no longer silently switch between upstream and a Chainguard rebuild unless an org runs `chainctl libraries cache opt-out`; orgs without upstream fallback are unaffected.

Check which copy (upstream mirror vs. Chainguard rebuild) your org has pinned for each cached package version, to audit lockfile consistency before the August 24 build-pinning cutover.

```
$ chainctl libraries cache list
```

Opt out of build pinning for the npm ecosystem if your org relies on receiving the latest Chainguard rebuild of a version rather than staying on the first copy resolved.

```
$ chainctl libraries cache opt-out --ecosystem javascript
```

— Exact commands, flags, default-behavior change and opt-out path `Week of 2026-08-17`

03 **Guardener GitHub Actions migration (beta)** NEW

60

The Chainguard Guardener GitHub App (beta) now extends automated migration to GitHub Actions: it inventories Actions usage across org repositories, maps them to hardened Chainguard equivalents, and opens pull requests pinned to a specific SHA rather than a mutable tag. It runs in two modes — an upfront pass that surfaces existing Actions usage and opens migration PRs, and an ongoing standardization mode that watches workflow files and suggests equivalents as new upstream Actions appear.

— Describes mechanism and modes but no setup command shown `Week of 2026-08-17`

THINNER COVERAGE BELOW

04 Named 403 error codes for blocked npm packages IMPROVED

50

Chainguard Libraries now returns a `403` HTTP status with a named reason code (e.g., `MALWARE_DETECTED`) when npm requests a blocked package or version, replacing the previously unexplained `404`.

— Names the status and reason code but no follow-up action given `Week of 2026-08-17`

05 17 new hardened container images added NEW

50

Adds 17 new container images including `azurite`, `k0s-cni-node`, `peerdb-ui`, `scc`, `zitadel-login`, `aws-lambda-nodejs`, `localstack`, `kserve-router-fips`, `jdk-openssl-fips`, `jre-openssl-fips`, `kube-router`, and `commercial-nginx-ingress-plus`, with FIPS variants available for several.

— Names specific images but no pull command or catalog link `Week of 2026-08-17`

06 Removal of prometheus:3.9 image BREAKING

30

The `prometheus:3.9` container image is no longer available as of 2026-08-17, after its grace period ended.

— States removal but no migration guidance provided `Week of 2026-08-17`

└──▷ **BREAKING ON UPGRADE**

! As of August 24, 2026, Chainguard Libraries build pinning is enabled by default for organizations using Chainguard Libraries with upstream fallback (JavaScript, Python, and Java ecosystems): the first copy of a package version an org downloads is locked in, and subsequent resolutions will no longer silently switch between the upstream mirror and a Chainguard rebuild. To preserve the previous behavior (no pinning), run `chainctl libraries cache opt-out`. Organizations without upstream fallback are unaffected.

! The `prometheus:3.9` container image is no longer available as of 2026-08-17 (grace period ended).

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

XFA

Feedback ⓘ 2 RELEASES · SEEN 2026-08-18 NOTES ↗

XFA is a framework for building secure, auditable financial applications with cryptographic verification and transaction integrity.

> XFA expanded its device compliance checks with eight new browser-privacy policies for Chrome and Firefox, and added an organization-level default language setting for communications.

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

01 **Browser privacy device checks for Chrome and Firefox** NEW 76

XFA adds eight configurable per-organization device checks covering browser privacy: Chrome Safe Browsing (standard and enhanced modes), Chrome Do-Not-Track, Chrome DNS-over-HTTPS, Firefox HTTPS-Only, Firefox DNS-over-HTTPS, and Firefox Enhanced Tracking Protection (standard and strict modes). Each check can be set to warn or block, and violations surface in the desktop tray menu and the device detail view in the dashboard with one-click remediation links built into the agent.

– Names all eight checks and enforcement modes, but no exact config keys

Browser privacy checks (Safe Browsing &…

THINNER COVERAGE BELOW

02 **Organization default language setting** NEW 53

Adds an organization-level default language setting under `Settings > Organization`, used for emails and messages when a user's own language cannot be detected or is not set on their profile.

– Gives exact UI path but no further mechanism detail

Set a default language for your organiz…

WAS THIS USEFUL? 👍 👎

Huntress

Feedback ⓘ 1 RELEASE · SEEN 2026-08-18 NOTES ↗

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

> Huntress launched a GA'd Investigations dashboard giving a unified, exportable timeline of every investigation across Managed EDR, ITDR, and SIEM.

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

01 Investigations dashboard with chronological timeline NEW

70

A new Investigations dashboard, accessible via the top navigation in the Huntress portal, shows a chronological timeline of signals reviewed, actions taken, remediations applied, and final outcomes for both reported and closed-benign investigations. It unifies coverage across Managed EDR, ITDR, and SIEM, includes a summary section with key stats for quick situational awareness, and supports export of investigations to CSV and Excel for reporting and recordkeeping.

Access the full investigation timeline to review SOC actions and remediations on a closed or reported incident.

> 📍 In the Huntress portal, click Investigations in the top navigation, then select any reported or closed-benign investigation to view its chronological timeline of signals, actions, and outcomes.

— Names UI path, export formats, and covered products clearly.

View Chronological Timeline of Every Hu…

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ 4 RELEASES · 2026-08-17 → 2026-08-18

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 header address rewriting to Cloud Email Gateway Protection, reduced false positives in Entra ID MFA-disabled detection, expanded Correlated Intelligence custom detection signal conditions with new URL and domain-age matching, and let security playbooks target or exclude endpoints by operating system.

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

01 **Final URL domain registration age condition** NEW 90

A new Final URL domain registration age condition for Correlated Intelligence custom detection signals triggers when the domain of the final destination URL — after all redirects and short-link expansions — was registered within a configurable window of 0 to 366 days.

Flag emails whose final destination URL resolves to a freshly registered domain — a common indicator in phishing campaigns — by creating a Correlated Intelligence signal using the new age condition.

📍 In the Trend Vision One console, go to Email and Collaboration Security → Cloud Email Gateway Protection → Correlated Intelligence, create or edit a custom detection signal, add a condition of type Final URL domain registration age, and set the threshold to 30 days to catch domains registered within the last month.

— Mechanism, exact window, and a runnable console walkthrough with threshold

New detection signal customization cond…

02 **Operating system filter for security playbooks** NEW 70

Endpoint Response Playbooks and Automated Response Playbooks now support operating system as a condition parameter, with include or exclude modes to precisely target automated actions by OS.

Exclude unsupported operating systems from an automated response action to prevent failed executions against incompatible endpoints.

📍 In Trend Vision One, go to Workflow and Automation → Security Playbooks, open or create a playbook, add a condition, select Operating System, and choose

Exclude with the target OS.

— Clear UI navigation path but not a config/API reference

Filter security playbooks by operating …

03 **URL and URL domain conditions for Correlated Intelligence signals**

NEW 60

Cloud Email Gateway Protection's Correlated Intelligence custom detection signals gain a URL condition matching email URLs against a specified list of up to 1,024 URLs, and a URL domain condition matching the domain of any URL in an email against a specified list of up to 1,024 domains.

— Named conditions and limits, but no walkthrough given

New detection signal customization cond…

THINNER COVERAGE BELOW

04 **Header address rewriting in Cloud Email Gateway Protection**

NEW 48

Content Filtering policies in Cloud Email Gateway Protection now support header address rewriting for the From , Reply-To , To , and Cc fields, applying to both inbound and outbound email traffic.

— Names the fields and scope but no config steps Header address rewriting for inbound an…

05 **MFA Disabled detection excludes Entra ID security-defaults tenants**

IMPROVED 45

Accounts in Microsoft Entra ID tenants that have security defaults enabled no longer surface under the Microsoft Entra ID Account with Weak Sign-In Security Policy - MFA Disabled risk event, eliminating false positives for those organizations.

— Explains the fix precisely but requires no user action

More accurate MFA Disabled detection fo…

WAS THIS USEFUL?  

Perplexity numbat

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

Visibility into AI agent activity on endpoints, with on-device detection, optional pre-action blocking, and forensic reconstruction.

> numbat v0.2.0 adds controls for inspecting and limiting message content captured from AI agent activity.

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

01 **Message content inspection controls** NEW 28

Adds explicit message output controls for governing what content is emitted from inspected message bodies, and supports bounded message body inspection that limits content analysis to defined byte ranges.

— Describes behavior but names no flags, keys, or exact limits. v0.2.0

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Datadog Stratus Red Team

Feedback ⓘ

1 RELEASE · 2026-08-18 NOTES ↗ CODE ↗

Granular, Actionable Adversary Emulation for the Cloud

> Stratus Red Team's biggest addition this window is correlation-ID based isolation for running multiple attack-technique executions concurrently without state collisions, alongside a new EC2 AMI deregistration technique and configurable tagging/prefixing of AWS prerequisite resources.

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

01 Concurrent execution isolation via correlation IDs NEW

98

The `STRATUS_RED_TEAM_CORRELATION_ID` environment variable isolates each execution's Terraform state under `$HOME/.stratus-red-team/<technique-id>/<correlation-id>`, preventing concurrent runs of the same technique from sharing or destroying each other's resources. The Go library exposes runner.WithCorrelationID(id) for programmatic use, and technique authors can embed correlation IDs in cloud resource names via the Terraform variables `var.correlation.short` (8-character form) and `var.correlation.id` (full form), or via stratus.CorrelationShortID(providers.AWS().UniqueCorrelationId) in Go. A shared Terraform provider plugin cache at `$HOME/.stratus-red-team/plugin-cache` avoids re-downloading providers per run. `stratus status` now scopes its report to the specific execution when the correlation ID env var is set, falling back to the flat per-technique directory when unset, and existing flat state is automatically migrated into the per-execution directory on upgrade.

Run two isolated detonate-and-cleanup cycles of the same technique in parallel without them sharing Terraform state.

```
$ export STRATUS_RED_TEAM_CORRELATION_ID=$(uuidgen)
stratus detonate aws.persistence.iam-backdoor-user
stratus cleanup aws.persistence.iam-backdoor-user
```

Embed the short correlation ID in Terraform resource names so two concurrent executions request distinct cloud resources and never collide.

```hcl
locals {
  resource_prefix = "stratus-red-team-my-
technique-${var.correlation.short}"
}
```

Check the state of an isolated execution — without the correlation ID the technique appears COLD even though it is WARM under the execution's directory.

```
$ # Reports COLD (flat layout, no isolated state visible)
stratus status aws.persistence.iam-backdoor-user

# Reports WARM (scoped to this execution's directory)
STRATUS_RED_TEAM_CORRELATION_ID=$ID stratus status
aws.persistence.iam-backdoor-user
```

Run two parallel red-team sessions in the same AWS environment without state collisions by assigning each a unique correlation ID.

```
$ STRATUS_RED_TEAM_CORRELATION_ID=session-alpha stratus detonate
aws.impact.ec2-deregister-ami
```

— Full mechanism, exact paths, env var, Go API and Terraform vars named `v2.36.0`

02 Custom tags and name prefix for AWS prerequisites NEW

65

A `tags` config setting applies custom tags (via the AWS provider's `default_tags`) to Terraform-managed prerequisites for AWS techniques, enabling correlation IDs, ownership labels, and environment markers per technique. A `prefix` config setting prepends a custom, template-capable string to the names of Terraform-managed prerequisite resources. Both are controlled via the Stratus Red Team configuration file.

— Config keys named but exact file path and example missing `v2.36.0`

THINNER COVERAGE BELOW

03 **New EC2 AMI deregistration attack technique** NEW 50

Adds a new AWS attack technique, `aws.impact.ec2-deregister-ami` (Deregister an Amazon EC2 AMI), which deregisters an Amazon EC2 AMI to simulate disruption/impact. It was added to the attack technique library and to the MITRE ATT&CK coverage matrix.

— Technique id given but no usage example shown v2.36.0

WAS THIS USEFUL? 👍 👎

Tenzir

Feedback ⓘ 1 RELEASE · 2026-08-18 BLOG ↗

Tenzir is a security data platform that ingests, transforms, and analyzes security telemetry and threat intelligence data.

> Tenzir v6.12.0 adds native PCAPNG capture support and a new `accept_relp` operator for reliable Syslog ingestion over RELP.

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

01 **accept_relp operator for reliable Syslog** NEW 85

New `accept_relp` operator receives Syslog over the Reliable Event Logging Protocol (RELP), supporting plaintext, TLS, and mutual TLS connections. It preserves message boundaries and provides at-least-once delivery via bounded acknowledgements and backpressure.

Accept Syslog from RELP-capable log shippers (e.g. rsyslog with omrelp) and parse each message in-pipeline.

```
$ accept_relp "0.0.0.0:2514"
syslog = data.parse_syslog()
```

— Names operator, protocol modes and delivery guarantee with example.

Tenzir Node v6.12.0: Reliable Syslog in…

02 **PCAPNG capture support** NEW 80

The `write_pcap` operator gains a `format = "pcapng"` option to explicitly generate PCAPNG output, and `read_pcap` now auto-detects PCAPNG input, including multi-interface captures, variable timestamp resolutions, byte orders, and multiple sections.

Write captured packets as PCAPNG instead of legacy PCAP, useful when preserving multi-interface or high-resolution timestamp metadata.

```
$ write_pcap format = "pcapng"
```

— Names exact operator, flag and format with example command.

Tenzir Node v6.12.0: Reliable Syslog in…

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

Open Cyber Threat Intelligence Platform

> OpenCTI 7.260817.0 extends its data model by adding a scoring field to key threat intelligence entity types.

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

01 **Score field added to core entity types** NEW 45

The data model now includes a `score` field on the Threat Actor Group, Intrusion Set, and Malware entity types, allowing these objects to carry a numeric score attribute.

— Names exact fields and entity types but no usage or range details. 7.260817.0

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ 1 RELEASE · 2026-08-18 CHANGELOG ↗

runZero is an asset discovery and inventory platform that identifies and maps all connected devices across network infrastructure.

> runZero's latest release adds AWS Inspector vulnerability findings for EC2 and Lambda, broadens IPMI vulnerability detection, and ships a paginated, searchable reports page alongside several backend accuracy and performance improvements.

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

01 **AWS Inspector findings for EC2 and Lambda** NEW 50

The AWS integration now ingests AWS Inspector findings for EC2 and Lambda assets, surfacing vulnerability data for those resource types within runZero's asset inventory.

— Names AWS Inspector and asset types but no viewing mechanism.

changelog-20260818-ebfd60f1

02 **Paginated reports page with search and delete** IMPROVED 50

The reports page now displays a paginated list of all analysis task-based reports, with search and delete capabilities for managing them.

— Names the UI page but not exact navigation steps. changelog-20260818-ebfd60f1

03 **Refingerprint action re-evaluates service fingerprints** IMPROVED 50

The Refingerprint asset action now also re-evaluates service fingerprints for improved accuracy.

— Names exact action; no accuracy metrics given. changelog-20260818-ebfd60f1

04 **Connector tasks honor expiration thresholds** IMPROVED 40

Connector tasks now honor the configured attribute and vulnerability expiration thresholds during import.

— Names threshold types but not the config keys involved. changelog-20260818-ebfd60f1

05 **Custom integrations tolerate per-device errors** IMPROVED 35

Custom integrations can now continue processing results even if a single device causes errors.

— No detail on error types or configuration options. changelog-20260818-ebfd60f1

06 **Concurrent directory and device data processing** IMPROVED

30

Directory and device data are now processed concurrently, improving task performance.

— No numbers quantifying the performance gain. changelog-20260818-ebfd60f1

07 **Geolocation from additional integration attributes** IMPROVED

25

Asset geolocation now uses location data from additional integration-sourced attributes.

— Doesn't name which integrations or attributes. changelog-20260818-ebfd60f1

08 **Broader IPMI vulnerability detection** IMPROVED 20

The scanner now detects a broader range of IPMI vulnerabilities.

— No specifics on which vulnerabilities were added. changelog-20260818-ebfd60f1

WAS THIS USEFUL? 👍 👎

Censys

Feedback ⓘ 1 RELEASE · 2026-08-17 BLOG ↗

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

> Censys shipped an AI-generated ASM reporting feature, a new Kibana/Elasticsearch connector for enriching and pivoting on Censys data, two new ASM risk detections, and nine new asset fingerprints.

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

01 New ASM risk fingerprints for exposed devices NEW

78

Adds two new ASM risk fingerprints queryable via `risks.name: "Exposed Copeland XWEB Pro Refrigeration Controller"` and `risks.name: "Exposed Payara Server Default Page"`. New risk types may be disabled by default, so review risk type configuration in the ASM web console.

Hunt for newly exposed refrigeration controller management interfaces across your monitored attack surface.

```
$ risks.name: "Exposed Copeland XWEB Pro Refrigeration Controller"
```

Find any Payara Server installations on your attack surface that are still showing their default, unconfigured landing page — indicating they may be running with default credentials.

```
$ risks.name: "Exposed Payara Server Default Page"
```

— Exact queries given, plus config caveat, runnable today August 17, 2026 (2026-08-17)

02 Kibana/Elasticsearch connector for Censys Platform data NEW

67

A new Elasticsearch Kibana connector lets practitioners enrich IPs, web properties, and certificates with Censys Platform data; initiate rescans of hosts or web properties; retrieve event history for an IP address; and launch CensEye automated pivoting jobs to find assets related to an IP, web property, or certificate.

— Lists concrete capabilities but no setup/config steps August 17, 2026 (2026-08-17)

03 **AI-generated attack surface reports** NEW 62

Censys Assistant in ASM can now generate custom PDF-ready reports about your attack surface. This requires the organization to be migrated to the Platform for team management and the AI Features toggle enabled under Platform Organization Settings > Privacy.

— Names exact settings path but no report contents detail August 17, 2026 (2026-08-17)

THINNER COVERAGE BELOW

04 **Nine new device and software fingerprints** NEW 53

Adds nine new fingerprints to the Censys Platform covering: Copeland (Dixell) XWEB Refrigeration Monitoring System (hardware), TP-Link Omada Controller ECSP adopt/transfer ports (software), D-Link DI Series Router (hardware), D-Link DWL Series Access Point (hardware), Gradio (software), JHipster (software), LimeSurvey (software), Meilisearch (software), and Payara Server (software).

— Names all nine fingerprints but no query examples given August 17, 2026 (2026-08-17)

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

Rapid7 InsightVM

Feedback ⓘ 1 RELEASE · 2026-08-18 DOCS ↗

BLOG ↗

Rapid7 InsightVM is a vulnerability management platform that discovers, prioritizes, and remediates security vulnerabilities across IT assets.

> InsightVM 8.55.0 adds authenticated Oracle Database fingerprinting for more accurate vulnerability assessment and improves reliability of the Security Console scheduler.

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

01 **Authenticated fingerprinting for Oracle Database** NEW

45

InsightVM now performs authenticated fingerprinting of Oracle Database, identifying the installed Oracle Database version and associated patches to enable more accurate vulnerability assessment.

— Names the target and benefit but no config or command detail. 2026-08-18-2cea5e2b

02 **Security Console scheduler reliability improvements** IMPROVED

30

Improves the Security Console scheduler for more reliable handling of scheduled tasks, including jobs responsible for integrating data into the Console.

— Describes improvement generically without specifics on jobs affected.

2026-08-18-2cea5e2b

WAS THIS USEFUL? 👍 👎

defectdojo

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

Open-Source Unified Vulnerability Management, DevSecOps & ASPM

> DefectDojo Pro's biggest window items are a self-service Locations rollout with idempotent data-migration commands and RBAC-scoped API access, alongside SCIM 2.0 provisioning, CycloneDX VEX round-tripping, and a beta Messaging Connectors system routed through Rules Engine 2.0; DefectDojo also locked down private notes to true author/superuser-only visibility.

└──▷ **WHAT SHIPPED · 15 FEATURES** most completely described first ? what's the number?

01

Locations feature: toggle, backfills, and RBAC-scoped API

NEW

95

A self-service `Locations` toggle on the `Settings > Feature Flags` page in DefectDojo Pro replaces the previous environment-variable-only deployment change; once enabled it locks (`Cannot Be Disabled`) because reversing the location data migration is not yet supported. The `DD_V3_FEATURE_LOCATIONS` environment variable still governs Classic UI and `/api/v2` route wiring at boot, so a restart is recommended when enabling the toggle if either surface is in use. A new data-migration suite offers four idempotent, on-demand items — `migrate_endpoints_to_locations` , `migrate_components_to_dependencies` , `migrate_findings_to_code_locations` , and an identity rehash (locked until the first three complete) — each also runnable as a management command for scripted or air-gapped use. Separately, `GET /api/v2/location/` is now scoped by RBAC instead of requiring superuser access.

Run the endpoint-to-locations backfill as a management command in an air-gapped or scripted environment after enabling the Locations toggle.

```
$ python manage.py migrate_endpoints_to_locations
```

Enable Locations for the first time via the Pro UI, then run all three backfills and the identity rehash in sequence.

📍 1. Go to Settings › Feature Flags and enable the Locations toggle. 2. Under the Locations row, run migrate_endpoints_to_locations, then migrate_components_to_dependencies, then migrate_findings_to_code_locations. 3. Once all three complete, run the identity rehash. 4. If using the Classic UI or `/api` /v2, set `DD_V3_FEATURE_LOCATIONS=True` in your deployment config and restart.

— Names the toggle, env var, all four commands, and the exact endpoint `3.2.200`

02

Private notes now enforced as truly private

BREAKING

65

A note marked private is returned only to its author and to superusers, with consistent enforcement across the API, the UI notes tab, note popovers, finding/test/engagement note counts, generated reports, and issue-tracker sync. Private notes written by other users no longer appear anywhere for other viewers, so teams that relied on private notes remaining readable by all users must move that content into public notes before upgrading.

— Clear before/after and migration note, but no config surface named `product docs`

03 **SCIM 2.0 provisioning for users and groups** NEW 65

Adds SCIM 2.0 provisioning under `Connect > Authorization` (tagged Provisioning): identity providers can create, update, and deactivate DefectDojo users and manage groups directly; deactivating a user over SCIM also deletes that user's API tokens.

— UI path and behavior given, no API endpoint or payload shown product docs

04 **Qualys connector Host Tags filter** NEW 65

The Qualys connector now accepts a Host Tags filter to scope discovery to hosts carrying named Qualys asset tags; the filter applies to both detection downloads and host listing, reducing sync scope without downloading out-of-scope hosts.

Narrow a Qualys connector sync to only hosts tagged in Qualys, reducing download and sync time to in-scope assets only.

> 📍 In the DefectDojo UI, go to your Qualys connector connection settings, enter the exact tag name(s) in the Host Tags filter field, and save. Leave the field blank to discover all hosts.

— Exact field name plus UI steps to use it product docs

05 **Messaging Connectors (beta) via Rules Engine 2.0** NEW 60

New Messaging Connectors (beta) send alerts to Slack, Microsoft Teams, email, or an Amazon SNS topic, routed by Rules Engine 2.0 rules; requires the Messaging Connectors and `Rules Engine 2.0` feature flags to be enabled.

— Names channels and required flags, no setup steps product docs

THINNER COVERAGE BELOW

06 **CycloneDX SBOM/VEX/VDR round-trip** IMPROVED 50

CycloneDX SBOM, VEX, and VDR export and import now support a full round trip: a document exported from DefectDojo can be imported back in, and the raw CycloneDX VEX analysis is now preserved on parsed Findings.

— Describes the round trip but no format/version specifics product docs

07 **Classic report engine deprecated ahead of removal** DEPRECATED 50

The classic report engine (Report Builder, Report Templates, and Generated Reports) will be removed in v3.3.0 on September 8, 2026.

— Exact removal date and affected components given product docs

08 **Multi-connection providers and repository setup shortcut**

IMPROVED 45

A provider can now hold several connections rather than one, with setup scoped to the active connection; Add Repositories now opens directly on the repository step.

— Names the UI shortcut but not underlying data model product docs

09 **Reporting reaches general availability** IMPROVED 45

Reporting is now generally available, removing the BETA label; both the classic and new Report Builder offer one-click migration of existing report templates, and migration works even with the Reporting feature flag off.

— Describes migration option but no exact UI path or command product docs

10 **Rules Engine 2.0 rule scheduling** IMPROVED 40

A Rules Engine 2.0 rule can now be given its own schedule; enabling the feature flag now warns that a worker restart is required before the change takes effect.

— Names the flag but not the scheduling mechanism or format product docs

11 **Expire and Reinstate risk acceptance actions** NEW 35

Adds Expire and Reinstate to the risk acceptance menu and as API actions.

— Names the two actions but gives no endpoint or workflow detail product docs

12 **Customizable widget grids on View pages** NEW 35

The five View pages now use customizable widget grids, allowing per-page widget arrangement.

— Scope stated but no page names or config surface product docs

13 **Resizable table columns with saved preferences** IMPROVED 35

Table columns can now be resized with widths saved to table preferences; list tables now render per-column loading skeletons while data loads.

— UI behavior described, no settings key or scope named product docs

14 **test_type_name field on finding serializer** NEW 30

Adds a flat `test_type_name` field on the finding serializer, surfacing the test type name directly on finding API responses without a nested lookup.

— Names the field but gives no example response 3.2.200

15 **New connectors: Tenable WAS, Rapid7 InsightVM, and six others** NEW 30

Registers the Tenable Web App Scanning and Rapid7 InsightVM connectors, along with six previously unregistered connectors (unnamed in the source).

— Two connectors named, six others left unnamed by source `product docs`

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! Private notes written by other users no longer appear in the notes tab, note popovers, or note counts for findings, tests, and engagements. Teams that used the private flag as a label while relying on all users still reading the text must move that content into public notes before upgrading.

! The classic report engine (Report Builder, Report Templates, and Generated Reports) will be removed in v3.3.0 on September 8, 2026.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Clerk

Feedback ⓘ 1 RELEASE · 2026-08-18 CHANGELOG ↗

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 teams apply percentage or fixed-amount deals to subscriptions or let customers self-redeem codes at checkout.

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

01 **Discounts and promo codes in Clerk Billing** NEW 70

Clerk Billing adds reusable percentage or fixed-amount discounts that can be applied manually to any active subscriber's subscription from the Dashboard and revoked at any time, plus promo codes (e.g. `LAUNCH20`) customers can enter at checkout to self-redeem, with optional redemption caps, new-subscriber restrictions, and per-code usage tracking. Discounts support different amounts per billing period (monthly vs. annual) and can last a single billing cycle, a set number of cycles, or indefinitely.

— Names Dashboard flow and promo code example but no API/config surface.

changelog-20260818-8dcd3ce9

WAS THIS USEFUL? 👍 👎

ConductorOne

Feedback ⓘ 1 RELEASE · 2026-08-18 BLOG ↗

ConductorOne is an identity governance platform that automates access management, provisioning, and compliance workflows across enterprise systems.

> ConductorOne shipped an early-access Findings module that continuously detects identity risk on every connector sync, complete with per-finding actions, bulk triage, CEL-based transformation and routing rules, and a new Decoys feature that plants tripwire credentials to catch compromise attempts.

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

01 Findings module for identity risk detection and triage NEW

98

A new `Findings` module (early access) under `Identity security > Findings` continuously monitors for identity risk conditions — unowned service accounts, exposed credentials, misclassified identities, decoy credential use, and more — on every connector sync, replacing manual periodic reviews. Each finding supports per-finding actions from its detail view: `Accept risk` (requires justification and expiration date), `Suppress` (permanent, requires reason), `Snooze` (1, 7, 30, or 90 days or a custom date, requires reason), and `Remediate` (where supported). From `All findings`, bulk actions let you select up to all findings matching the current filter and apply `Snooze`, `Suppress`, `Accept risk`, or `Reopen` in one step, with a confirmation dialog showing how many selected findings each action affects. `Transformation rules` (at `Findings > Transformation rules`) match on `finding.severity`, `finding.state`, `finding.risk_score`, or `finding.annotations` via CEL expression, then set severity or add/remove annotations, running in ascending `Order`. `Routing rules` (at `Findings > Routing rules`) match on app, finding type, or a CEL expression, then suppress, snooze, accept, remediate, or dispatch (webhook, function, or automation), with only the first matching rule running in ascending `Order`. The finding detail view also adds an `Evidence` pane with a confidence score and an `Activity` audit trail covering creation, routing-rule evaluation, and evidence updates. Findings surface on the Identities & NHI dashboard, with every inventory table now filterable to identities with open findings.

Automatically escalate every decoy-credential finding to Critical and notify the security team the moment it is created, so compromise signals never wait in a queue.

📍 1. Go to Identity security › Findings, click the gear icon, toggle on `Decoy credential used`, and Save changes. 2. Go to Findings › Transformation rules › Create rule. Set Finding type to `Decoy credential used`. Add a Set severity transform, choose Critical, and save. 3. Go to Findings › Routing rules › Create rule.

Set Finding type to `Decoy credential used`. Set Action to `No action`. Add a Trigger automation dispatcher pointing to your team notification automation. Set Order lower than any broad catch-all rule, then save.

Suppress low-severity noise automatically so analysts only see findings that require attention, using a broad catch-all routing rule placed last in the order.

📍 1. Go to Findings › Routing rules › Create rule. 2. Leave App and Finding type unset (Match all). 3. Add a CEL match condition: finding.severity == `FINDING_SEVERITY_LOW` 4. Set Action to Suppress. 5. Assign an Order value (e.g. 9000) higher than all specific rules so it only fires when nothing else matches. 6. Save the rule.

— Names every action, rule field, order behavior, and UI path with worked examples

Findings and Decoys (2026-08-18)

THINNER COVERAGE BELOW

02 **Decoy tripwire credentials** NEW 50

A new Decoys feature lets teams plant a tripwire credential that looks real but grants no access; any use attempt raises a Critical finding while the actor sees only a normal authentication failure.

— Mechanism described but no setup command or UI path given

Findings and Decoys (2026-08-18)

WAS THIS USEFUL? 👍 👎

SailPoint

Feedback ⓘ 2 RELEASES · 2026-08-17 BLOG ↗

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

> SailPoint added nine new compliance connectors, introduced structured audit events for Non-Employee Risk Management workflows, and shipped a breaking IQService update that disables privileged services unless explicitly enabled and TLS-only.

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

01 IQService privileged services require explicit enable and TLS-only

95

BREAKING

IQService-August-2026 requires administrators to run `IQService.exe -g --enable <Service Name>` to activate the `ScriptExecutor` (native before/after script execution) and `gMSAPasswordData` (gMSA password retrieval) privileged services, e.g. `IQService.exe -g --enable ScriptExecutor` and `IQService.exe -g --enable gMSAPasswordData`; both services are silently disabled if TLS is not exclusively configured on the IQService instance, and running non-TLS ports alongside a TLS port is no longer supported. Auto-update push via UpdateService begins October 19, 2026, with a 60-day manual configuration window starting August 18, 2026 for administrators to configure and validate before then.

Before upgrading, explicitly enable the ScriptExecutor privileged service on a TLS-configured IQService instance to prevent script execution from silently stopping after the update.

```
$ IQService.exe -g --enable ScriptExecutor
```

Explicitly enable gMSA password retrieval so managed accounts that rely on gMSAPasswordData continue to function after the IQService-August-2026 upgrade.

```
$ IQService.exe -g --enable gMSAPasswordData
```

— Exact commands, service names, and migration deadline given

New Capability: Integration Service (IQ…

02 Audit events for Non-Employee Risk Management workflows NEW

83

The Audit Events API now emits dedicated events for workflow definition lifecycle actions (create, fork, deprecate, delete) and for workflow step changes (add, edit, delete, reorder), each carrying structured payload fields including workflow identity, definition label, actor, and timestamp. Existing workflow session audit events are enriched with `workflow_version_label` and `workflow_version_id` fields identifying the definition version active when a session started, and the new events can be filtered by workflow subject and definition identifiers using the existing Audit Events query API.

— Names event types and payload fields but not full API path

New Capability: Workflow Audit and Visi…

THINNER COVERAGE BELOW

03 **Nine new compliance connectors added** NEW 45

SailPoint's compliance integration catalog gains connectors for Cobot, Losant, Prismatic, Productboard, ServiceM8, StoryChief, SupportBee, UpCloud, and Ziflow.

— Names all nine connectors but no setup or usage detail product docs

└──▷ **BREAKING ON UPGRADE**

! `ScriptExecutor` and `gMSAPasswordData` privileged services will not run after upgrading to IQService-August-2026 unless TLS is exclusively configured on the IQService instance AND each service is explicitly enabled via `IQService.exe -g --enable <Service Name>`. Environments using either service without completing these steps will experience service interruptions on upgrade.

! Non-TLS ports must be disabled when switching an IQService instance to a TLS port; running both simultaneously is no longer a supported configuration.

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ 1 RELEASE · 2026-08-17 BLOG ↗

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

> JumpCloud's agent update makes Device Admin controls read-only when only a single option is available, tightening how managed users interact with eligible device actions on macOS and Windows.

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

01 **Read-only Device Admin controls for single options** IMPROVED

45

Device Admin controls are now read-only when only one selection is available on macOS and Windows, giving managed users clearer and safer control over eligible device actions.

– Describes behavior change but no config or command detail.

JumpCloud Agent Release Notes for Augus…

WAS THIS USEFUL? 👍 👎

1Password

Feedback ⓘ 4 RELEASES · 2026-08-17 BLOG ↗

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

> 1Password's beta channel (8.12.34-29) added passwordless account onboarding across Linux, Mac, and Windows, promoted the MCP Server out of Labs, introduced an Offline Items sidebar view, and expanded 1Password for Claude with sign-in provider support and access-reason transparency.

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

01 **Auto-lock support in KDE Plasma Wayland** NEW 60

Enables auto-lock support in KDE Plasma on Wayland and other Wayland environments that implement the `ext_idle_notifier_v1` protocol.

— Names the exact protocol but no user action required

1Password for Linux beta 8.12.34-29 (20…

THINNER COVERAGE BELOW

02 **MCP Server moved to Settings > Developer** IMPROVED 55

The 1Password MCP Server has graduated from experimental status and moved from `Settings > Labs` to `Settings > Developer` across the Linux, Mac, and Windows apps.

— Exact settings paths named, but only a relocation

1Password for Linux beta 8.12.34-29 (20… 1Password for Mac beta 8.12.34-29 (2026… 1Password for Windows beta 8.12.34-29 (…

03 **Offline Items sidebar entry and prompts** NEW 52

Adds an Offline Items entry to the sidebar on Linux, Mac, and Windows for items that cannot be synced with a 1Password account, including deleted items, and improves offline-item prompts to show more detail when items can't be synced.

— Names a UI location but no deeper mechanism given

1Password for Linux beta 8.12.34-29 (20… 1Password for Mac beta 8.12.34-29 (2026… 1Password for Windows beta 8.12.34-29 (…

04 **Claude access-reason transparency** NEW 52

1Password for Claude shows the reason Claude is requesting access to specific items when the user selects the button to the right of each item in the access prompt.

— Describes behavior and exact UI interaction point

1Password for Mac beta 8.12.34-29 (2026…

05 **Passwordless account onboarding** NEW 50

Users can add a 1Password account to the Linux, Mac, and Windows apps without re-entering their password, Secret Key, or SSO verification code.

— Explains what changed but no config or command to try it

1Password for Linux beta 8.12.34-29 (20… 1Password for Mac beta 8.12.34-29 (2026… 1Password for Windows beta 8.12.34-29 (…

06 **Clipboard history suppression for copied account details**

IMPROVED 43

Copying an account detail from the Manage Accounts pop-up no longer writes the value to the device clipboard history, on Linux and Windows.

— Clear behavior change but nothing for reader to configure

1Password for Linux beta 8.12.34-29 (20… 1Password for Windows beta 8.12.34-29 (…

07 **Accounts-and-vaults list display in Autofill passkey save**

IMPROVED 40

Improves the display of the accounts-and-vaults list when saving a passkey as a new item via Autofill on Mac.

— Minor UI fix described only briefly 1Password for Mac beta 8.12.34-29 (2026…

08 **Claude sign-in with supported providers** NEW 35

1Password for Claude now supports signing in with logins configured to use supported sign-in providers.

— Thin description, no named providers or mechanism

1Password for Mac beta 8.12.34-29 (2026… 1Password in the browser 8.12.34-29 (20…

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ 1 RELEASE · 2026-08-17 BLOG ↗

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

> Auth0 introduces Custom Rate Limits in Early Access, letting tenants cap Authentication API traffic on a per-client basis via API-driven policy configuration.

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

01 Custom Rate Limit Policies for Authentication API NEW

62

Adds Custom Rate Limit Policies (Early Access), configured via API-based self-service policy configuration, to cap the number of Authentication API / OAuth requests an individual client or group (all third-party apps, or CIMD apps) can consume, preventing one client from exhausting a tenant's overall rate limit entitlement. Each policy supports two enforcement modes: blocking (reject requests over the limit) or non-blocking (notifications only), enabling staged rollouts.

. . . .

— Describes mechanism and modes but no exact API endpoint or config key named.

Custom Rate Limits available in Early A…

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Cyera

Feedback ⓘ 6 RELEASES · SEEN 2026-08-18 NOTES ↗

Cyera is a data security platform that discovers, classifies, and protects sensitive data across cloud and on-premises environments.

> Cyera's latest releases expand Access Trail audit coverage to Outlook and NetApp, introduce Agent Guardian for AI agent discovery and runtime enforcement, add on-prem database support and partner integrations, launch an Incident Materiality Assessment service, deepen Cy's AI-assisted investigation, and introduce Projects for scoped multi-team access.

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

01 Cy AI investigation gains Saved Prompts and suggested next steps

IMPROVED 65

Cy adds Saved Prompts so teams can store and reuse proven natural-language investigation workflows across DSPM and DLP, adds suggested next steps that surface recommended follow-on actions after each query, and expands natural-language investigation to cover risky alerts, sensitive data exposure, access activity, and data movement patterns in one unified path.

— Names three concrete additions and their scope, offers a clear starting point New in Cy

02 Projects for scoped data, alert, and API access NEW

65

Introduces Projects, a construct that lets security teams define scoped workspaces limiting user visibility to specific data, alerts, and identities within a single Cyera tenant, enables scoped API access per project so integrations only surface data relevant to their assigned scope, and preserves organization-wide visibility for global users while restricting project-scoped users to their defined perimeter.

— Describes construct and scoping mechanism, a clear configuration starting point

Projects

03 Access Trail extended to Outlook and NetApp IMPROVED

60

Access Trail audit coverage now extends to Outlook, tracking message access, attachment access, send activity, deletions, inbox rules, folder permissions, and delegated access, and to NetApp on-prem file shares, tracking sensitive file activity across on-premises storage.

.

— Names specific audit actions and two systems, no config or API path

Access Trail for Outlook and NetApp

04 **Incident Materiality Assessment service** NEW 60

A new service that uses targeted data scanning and enriched classification to identify exactly what sensitive data was exposed, who accessed it, and when during a security incident, producing a defensible incident impact view suitable for legal, board, regulator, and disclosure decisions in place of broad blast-radius assumptions.

— Explains mechanism and use cases but delivered as a service, not self-serve

Incident Materiality Assessment

THINNER COVERAGE BELOW

05 **Agent Guardian for AI agent discovery and enforcement** NEW 50

Agent Guardian builds a live inventory of sanctioned and unsanctioned AI agents, mapping their models, tools, identities, and data connections, analyzes which agents introduce real risk based on their resources and data connections, and monitors risky tool calls at runtime to block unsafe actions before sensitive data is exposed.

.

— Explains discovery and enforcement mechanism but no concrete setup path Agent Guardian

06 **On-prem database support for MySQL, MariaDB, Oracle RAC, MSSQL, IBM Db2 for i**

NEW 45

Adds support for MySQL, MariaDB, Oracle RAC, MSSQL, and IBM Db2 for i, enabling discovery, classification, and governance of sensitive data in these on-prem enterprise databases.

— Names five databases and three capabilities but no config detail

Extending the Ecosystem

07 **Partner ecosystem data context integrations** IMPROVED 35

Extends Cyera's data context into partner workflows to support exposure prioritization, guardrail enforcement, automated remediation, and audit-ready evidence, though no specific partners are named.

— Describes purpose but names no specific partner or integration Extending the Ecosystem

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

DNSFilter

Feedback ⓘ 1 RELEASE · 2026-08-17 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 beta introduces SecureTransit, an encrypted VPN-layer privacy add-on, alongside a redesigned DNSFilter One Agent UI for monitoring protection and connection status.

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

01 **SecureTransit encrypted privacy layer** NEW 91

Adds SecureTransit, a policy-enforced encrypted privacy layer supporting `WireGuard` or `IKEv2` protocols, routing traffic through a DNSFilter-managed secure gateway to prevent network-level visibility of destinations on untrusted networks. Admins set policy mode to `Always On`, `Manual`, or `Disabled` per organization or per device from the existing DNS filtering dashboard, and SecureTransit blocks location trackers, mail trackers, data brokers, and ad networks at the network layer before data leaves the device. It is available as a per-device add-on with a 14-day free trial for existing customers, installed via the dashboard under `Deployments → Roaming Clients → Beta Channel`.

— Names protocols, policy modes, and exact dashboard install path.

08/17 August 17 – Windows Roaming Clien…

02 **DNSFilter One Agent UI** NEW 67

Introduces the DNSFilter One Agent, replacing the right-click tray menu with a full UI window supporting system light/dark mode, and showing protection state, filtering mode, last sync, and an activity view of blocked and allowed queries over the past hour. A SecureTransit panel within the Agent displays connection status, connected server on a map, region, protocol, session uptime, data transferred, and Connect/Disconnect controls.

— Describes UI panels and data shown but no config or command.

08/17 August 17 – Windows Roaming Clien…

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Daytona

Feedback ⓘ 2 RELEASES · SEEN 2026-08-18 NOTES ↗

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

> Daytona 0.205.0 adds warm pool management, spot GPU support and per-sandbox OpenTelemetry configuration, while the snapshot-20260818 release stabilizes fork and snapshot APIs and improves TypeScript SDK error handling.

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

01 **Per-sandbox OpenTelemetry endpoint override** NEW 60

Adds `otelEndpointOverride` to sandbox creation, enabling per-sandbox OpenTelemetry endpoint configuration.

– Names exact config key but no further mechanism detail

Warm pool management and spot GPUs

02 **Stable fork and snapshot creation API** IMPROVED 60

Promotes fork and create-snapshot operations from experimental to stable API, retiring the experimental aliases.

– States before/after status change but no exact endpoint names snapshot-20260818

THINNER COVERAGE BELOW

03 **Client-side error distinction in TypeScript SDK** IMPROVED 45

Distinguishes client-side argument validation errors from server-side 400 responses in the TypeScript SDK for clearer error handling.

– Explains behavior change but lacks specific error types or API snapshot-20260818

04 **Warm pool management APIs across SDKs** NEW 30

Adds warm pool management APIs across SDKs.

– Names capability but no endpoints or SDK specifics given

Warm pool management and spot GPUs

05 **Spot GPU support for sandboxes** NEW 20

Adds spot GPU support for sandboxes.

– Bare feature name with no mechanism or scope Warm pool management and spot GPUs

WAS THIS USEFUL? 👍 👎

Cursor

Feedback ⓘ 1 RELEASE · SEEN 2026-08-18 NOTES ↗

Cursor is an AI-powered code editor built on VS Code that helps developers write, edit, and debug code faster using artificial intelligence.

> Cursor launched Origin, a new git forge (early beta) at cursor.com/codebase offering repo hosting, GitHub mirroring, pull requests, per-repo settings/permissions, and third-party CI/deploy integrations, and separately added a private-worker-pool management API to the Cloud Agent for autoscaling self-hosted runners.

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

01 Private worker pool management API for Cloud Agent NEW

98

The Cloud Agent API adds private-worker-pool management endpoints: `GET /v0/private-workers` lists connected self-hosted pool workers with filterable `status` and `scope` (`team_pool` / `personal`) params, returning `workerId` , `labels` , `isInUse` , `repoOwner` , `repoName` , `workspaceRootPath` , `connectedAtMs` , `teamId` , `serviceAccountId` , `displayName` , plus `totalCount` / `nextPageToken` pagination; `GET /v0/private-workers/pools` lists durable pools filterable by `scope` and `includeStale` , exposing `poolName` , `connectedWorkerCount` , `inUseWorkerCount` , `firstSeenAtMs` , `lastSeenAtMs` , `isStale` , and optional repo metadata for scale-to-zero autoscaling dashboards; `DELETE /v0/private-workers/pools` deregisters a durable pool by `scope` / `poolName` (optionally `repoOwner` / `repoName`), returning `{ "deregistered": true }` ; `GET /v0/private-workers/pending-requests` lists unassigned pool requests with queue depth via `totalCount` for queue-depth-based autoscaling, filterable by repository URL; and `POST /v0/private-workers/claim` atomically reserves a pending request (`bcId`) for a specific `workerId` before it starts, with the claimed id passed via the `CURSOR_AGENT_WORKER_ID` environment variable (or hidden `--worker-id` flag) to `agent worker --pool <name> --worker-dir <path> start` . Omitting `repoOwner` / `repoName` when targeting a pool with `env.type: "pool"` routes to a repo-less pool, and an unknown pool name returns `400` immediately instead of queueing.

Poll pending request queue depth to decide when to spin up additional pool workers for autoscaling.

```
$ curl --request GET \
  --url "https://api.cursor.com/v0/private-workers/pending-
requests" \
  -u "$CURSOR_API_KEY:"
```

Atomically claim a pending request for a specific worker id, then start that worker so the bridge registers the claimed identity.

```
$ # 1. Claim the pending request
curl --request POST \
  --url "https://api.cursor.com/v0/private-workers/claim" \
  -u "$CURSOR_API_KEY:" \
  -d '{"bcId": "bc-00000000-0000-0000-0000-000000000002",
"workerId": "pw_123"}'

# 2. Start the worker with the reserved id
export CURSOR_AGENT_WORKER_ID="pw_123"
agent worker --pool gpu --worker-dir /workspace start
```

List idle team pool workers to monitor utilisation before scaling down capacity.

```
$ curl --request GET \
  --url "https://api.cursor.com/v0/private-workers?
status=idle&scope=team_pool&limit=50" \
  -u "$CURSOR_API_KEY:"
```

— Five named endpoints with fields, params, and runnable curl examples. product docs

02 **Origin git hosting platform** NEW 88

Origin, Cursor's git forge (early beta) at `cursor.com/codebase`, supports creating repositories, pushing and pulling with standard git, browsing and searching code, and sharing with Cursor teams. Repos can be cloned via HTTPS at `https://origin.cursor.com/{owner}/{repo}.git` from the Code dropdown, or evaluated by adding Origin as a second git remote alongside GitHub with `git remote set-url --add --push`. A new `origin` CLI installs via `curl -fsSL https://downloads.cursor.com/origin/install.sh | sh` and authenticates with `origin auth login` before git operations. Origin follows the Privacy Mode of the namespace owner (teams on legacy privacy mode must switch to Privacy Mode to enable it) and is available in early beta on Pro, Teams, and Enterprise plans (all paid plans, excluding opted-out enterprise orgs), not on free plans.



Clone a repo hosted on Cursor Origin using HTTPS after authenticating with the CLI.

```
$ origin auth login
git clone https://origin.cursor.com/acme/checkout.git
```

Mirror an existing GitHub repo to Origin while keeping both remotes in sync during evaluation.

```
$ git remote set-url --add --push origin
git@github.com:acme/checkout.git
git remote set-url --add --push origin
https://origin.cursor.com/acme/checkout.git
git push origin main
```

— Names CLI install command, clone URL, and remote setup verbatim. `Origin Code Hosting`

---

03 **GitHub mirroring and bi-directional sync** NEW 70

A `Sync from GitHub` workflow at `cursor.com/codebase` mirrors a GitHub repository into Origin, copying git history, branches, tags, and pull requests, and keeps Origin continuously updated as the GitHub source changes, with sync status visible under `Settings → General`. Pushes to a mirrored repo pass through to GitHub, which remains the source of truth and continues to run its own CI; access mirrors existing GitHub read/write permissions. Pull requests on a mirrored repo sync bi-directionally with GitHub, and cloud agents can be attached directly to mirrored repos to work against the full commit history.



Mirror an existing GitHub repo into Origin so your team can browse, search, and run agents on it without leaving Cursor.

> 📍 1. Go to cursor.com/codebase 2. Select `Sync from GitHub` 3. Choose your GitHub organization and repository 4. Confirm the sync 5. When complete, open the Origin copy to browse files and commit history, or clone from the green `Code` button

— Describes the mirror workflow, sync status location, and source-of-truth behavior.

`Origin Code Hosting`

---

04 **Pull requests in Origin** NEW 70

Origin's `Pull Requests` tab (at `cursor.com/codebase`) lists open pull requests and opens a review page with `Activity`, `Commits`, `Checks`, and `Files Changed` tabs supporting line-level comments, reviewer requests, and reviews from the diff. Pull requests can be merged once reviews and CI checks are satisfied, with conflicts surfaced before merge; mirrored GitHub repositories expose GitHub pull requests inside Origin, syncing comments and reactions bidirectionally within seconds, and GitHub-assigned reviews can be completed and merged from Cursor. Cloud agents can open pull requests on Origin automatically as part of a task, and per-repo agent integration lets you ask questions about browsed code, request changes, update PRs, or push branches from the same interface.

— Names PR review tabs and merge/sync mechanics but only UI navigation.

`Origin Code Hosting`

---

05 **Per-repo settings, permissions and branch protections** NEW 66

Each Origin repository has a Settings tab with General, Permissions, Rules and Protections, and Apps sections: General shows a Sync Status indicator for GitHub-mirrored repos; Permissions lets you review and manage access with Internal (visible to all Cursor team members) and Private (explicit grant required) visibility modes; Rules and Protections configures branch rules and merge protections per repository. At the team level, a Permissions UI lets admins claim a codebase name, enable or disable Origin at any time, and grant repo-creation rights to members.

— Names every settings section and permission mode, but only UI paths.

Origin Code Hosting

06 **Third-party app integrations (Vercel, Depot, Buildkite)** NEW

61

A per-repo Apps tab (with codebase-level Manage Apps) supports connecting Vercel (auto-provisions a preview deployment for every PR), Depot (runs CI, including existing GitHub Actions workflows, on Origin-hosted repos), and Buildkite (runs GitHub Actions workflows and native Buildkite pipelines on Origin-hosted repos); Depot and Buildkite are restricted to Origin-hosted repositories, while GitHub-mirrored repos continue CI on GitHub. Admins install and manage these apps across all Origin repos from codebase settings at `cursor.com/codebase/settings/apps`, with per-repo activation visible from the repository's Apps tab.

— Names three integrations and their restriction scope, UI path only. Origin Code Hosting

THINNER COVERAGE BELOW

07 **Origin API access via internal apps** NEW

35

Internal apps can be created under codebase settings to generate app credentials for calling the Origin API, authenticating via app JWTs and installation access tokens.

— Only names the auth mechanism, no endpoint or setup steps given. product docs

08 **Multi-repo Cloud Agent tasks** IMPROVED

23

Cloud Agent tasks can now target up to 20 repositories in a single task.

— Single bare fact with no mechanism or usage path given. product docs

WAS THIS USEFUL?

Cline

Feedback ⓘ 2 RELEASES · 2026-08-18 NOTES ↗ CODE ↗

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

> Cline's desktop beta added cross-device cloud agent sessions with GitHub-linked repo/branch handoff, plus voice input, model-driven image generation, media playback, and a new beta release channel.

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

01 Cloud agent sessions with cross-device sync NEW 85

Adds cloud agent sessions (remote provisioning, transcript synchronization, approval recovery, model switching, branch picker, and onboarding panel), gated behind an explicit Cloud sessions toggle in Settings. Users connect GitHub, pick a repository and branch, and hand sessions off between devices with transcripts, approvals, and queued prompts kept in sync, switching models mid-session; the onboarding flow now includes a GitHub integration step. Also includes early proof-of-concept support for running these sessions in SSH remote environments.

Hand a long-running coding session off to another device without losing context — enable cloud sessions in Settings, then pick a repo and branch to start a synced cloud session.

> 📍 In the desktop app, go to Settings and enable the Cloud sessions toggle, then select a GitHub repository and branch during or after onboarding to start a session that syncs transcripts, approvals, and queued prompts across devices.

— Names the Settings toggle and onboarding flow with a worked example.

desktop-v0.0.14-beta.1 desktop-beta

02 Generated media forwarding through CLI ACP layer NEW 85

Propagates generated media (images, audio, video) through the CLI ACP layer — images are forwarded as `agent_message_chunk` content with `type: image`, `data`, and `mimeType` fields; non-image modalities are forwarded as text placeholders `[Generated <modality>: <mediaType>]`.

— Exact field and placeholder format named, directly usable by integrators. desktop-beta

THINNER COVERAGE BELOW

03 Beta release channel for desktop app NEW 50

Adds a beta release channel surfaced in-app, backed by the `desktop-beta` feed for the auto-updater. Beta channel installs side-by-side with the stable app and updates independently, leaving stable installs unaffected.

– Explains install behaviour but no config or command given. desktop-beta desktop-v0.0.14-beta.1

04 **Model-driven image generation and media playback** NEW

45

Adds model-driven image generation and a media player to the desktop app supporting playback of generated audio artifacts and generated video.

– Describes what was added but not how it's triggered. desktop-v0.0.14-beta.1 desktop-beta

05 **Voice input and realtime voice sessions** NEW

35

Adds microphone voice input in the composer and realtime voice sessions with UI refinements and improved session lifecycle management.

– Names the feature but not mechanism or scope. desktop-v0.0.14-beta.1 desktop-beta

06 **Redesigned question prompts with animated disclosures**

IMPROVED

35

Redesigns the `AskQuestion` component in the UI and adds animated reasoning and tool disclosures.

– Names the component but gives no behavioural detail. desktop-beta desktop-v0.0.14-beta.1

07 **Avatar/pet overlay companion** NEW

30

Adds an avatar/pet overlay (preview) to the desktop app — a floating desktop companion that reacts to live session activity.

– Cosmetic feature described only briefly, no interaction detail. desktop-v0.0.14-beta.1 desktop-beta

WAS THIS USEFUL?  

OpenHands

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

OpenHands: AI-Driven Development

> OpenHands' latest releases add LLM cost tracking and CSV/JSON export to the Activity Log, a new Automate pane and Git Sync page for managing automations, LLM pre-flight validation, structured Canvas error display, and Kimi K3 as the new free default Canvas model.

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

01 **Activity Log cost tracking and export** NEW 73

The Activity Log now displays the LLM cost in USD to four decimal places for each completed run (`$0.0000` for zero cost, blank when unavailable). New export controls let you download run data as CSV or JSON, each including a raw numeric `cost` field per run that is exported as `null` when the backend reports no cost.

– Names formats and field but only a UI export button, no API product docs

THINNER COVERAGE BELOW

02 **Automate pane in sidebar** NEW 50

The Automate pane in the sidebar lists all automations on the active backend, showing each automation's name, trigger type, and enabled state, and surfaces an add option when no automations exist yet.

– Names UI location and fields shown but no config keys product docs

03 **LLM pre-flight validation** NEW 40

OpenHands now runs LLM pre-flight validation that checks LLM profiles and rejects misconfigured ones before they can be saved.

– Explains behaviour but no specific checks or errors named v1.14.0

04 **Kimi K3 as default free Canvas model** IMPROVED 40

Kimi K3 is now set as the default Canvas model and is tagged as free, giving users a no-cost starting point.

– Names the model but no override or pricing detail given v1.14.0

05 **Git Sync automation page** NEW 30

A new Git Sync page under Automations lets users manage repository synchronization workflows.

– Thin one-line description with no workflow mechanics v1.14.0

06 **Structured Canvas error display** IMPROVED 25

Canvas now presents structured error outcomes so failures surface with actionable detail rather than raw error text.

— No example of the structured error format shown v1.14.0

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗ CODE ↗

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

> Claude Code v2.1.234's biggest change is letting `/permissions`, `/add-dir` and several dialogs (`/autocompact`, `/theme`, `/help`, `/config`, `/advisor`) open mid-turn while an agent is working, alongside new GitLab MR status badges, automatic session continuation at usage limits, and two-way Remote Control state sync across phone, desktop and terminal — while removing the `/config` Default teammate model setting as a breaking change.

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

01 Mid-turn access to permissions and dialogs IMPROVED

83

`/permissions` can now be opened while Claude is actively working, with rule changes applying immediately to the rest of the current turn; `/add-dir <path>` can also now be invoked mid-turn; and `/add-dir`, `/autocompact`, `/theme`, `/help`, `/config`, and `/advisor` dialogs all now open mid-turn in the fullscreen TUI.

— Names every command/dialog affected, but no worked example v2.1.234

02 Configurable /goal background check-in timing NEW

83

Adds `CLAUDE_CODE_GOAL_CHECKIN_MINUTES` environment variable (set to `0` to opt out) to control how long `/goal` waits before checking in on background tasks that have been running 30+ minutes.

Opt out of the 30-minute background-task check-in so `/goal` never interrupts long-running agents.

```
$ export CLAUDE_CODE_GOAL_CHECKIN_MINUTES=0
claude
```

— Named env var, default threshold, and opt-out example given v2.1.234

03 CLAUDE_CODE_PROJECT_DIR_NAME environment variable NEW

80

Adds optional `CLAUDE_CODE_PROJECT_DIR_NAME` environment variable so hosts that give each session its own config directory can choose a short name for the per-project

transcript directory.

Give each CI job its own short transcript directory name instead of a long auto-generated one.

```
$ export CLAUDE_CODE_PROJECT_DIR_NAME=my-feature-branch
claude
```

— Named env var with a runnable example command v2.1.234

04 **GitLab merge request badge** NEW 65

Adds a GitLab merge request badge to the footer and statusline: repos with a GitLab remote and an authenticated `glab` CLI now show `MR !N` with draft/pending/green states.

— Clear UI location and states, no direct command v2.1.234

05 **Reduced context cost of claude-api skill** IMPROVED 65

Reduces context cost of loading the built-in `claude-api` skill from ~200k+ tokens to ~25k by loading reference docs on demand.

— Concrete before/after numbers but automatic, nothing to run v2.1.234

06 **Remote Control state sync across clients** IMPROVED 62

Remote Control: effort picks made on a phone or on claude.ai/code now apply to terminal- and Desktop/VS Code-hosted sessions, with the session publishing its effort level to connected clients; and Remote Control sessions from Claude Code Desktop or VS Code now keep phones and claude.ai/code updated on the session's permission mode and model as they change.

— Explains sync mechanism but no user-facing action v2.1.234

07 **Removed default teammate model setting** BREAKING 60

Removes the Default teammate model setting from `/config`; agent-team teammates now use the leader's model unless the spawn command names one.

— Named setting and new fallback behavior, migration note implicit v2.1.234

THINNER COVERAGE BELOW

08 **Truncated session list reporting in agent APIs** IMPROVED

56

`SendMessage` and `ListAgents` now report when your account's session list was too long to check completely, instead of treating unseen sessions as absent.

— Named APIs but behavior change only, no action for user v2.1.234

09 **Automatic session continuation at usage limit** NEW

55

Adds automatic session continuation when a claude.ai usage limit resets; toggle off via `/config` (Continue automatically at usage limit).

— Named config toggle gives a clear action v2.1.234

10 **/goal auto-clears on unrecoverable error** IMPROVED

55

`/goal` now clears itself with a notice when a turn ends on an unrecoverable error (revoked auth, exhausted credit balance, or context overflow) instead of staying armed.

— Specific triggers named but nothing for reader to do v2.1.234

11 **selection:clear keybinding action** NEW

45

Adds a `selection:clear` keybinding action, bindable to any key, that clears an in-app text selection in both the main view and the agents view.

— Named keybinding, but no default key or example given v2.1.234

12 **setup-token argument validation** IMPROVED

45

`claude setup-token` now rejects unexpected extra arguments instead of silently ignoring them.

— Named command, minimal detail on error behavior v2.1.234

13 **Improved auto-generated session titles** IMPROVED

30

Improves auto-generated session titles to read as short, specific names (e.g. Login button bug) rather than sentences restating the request.

— Cosmetic change described only in prose v2.1.234

└──▷ **BREAKING ON UPGRADE**

! The Default teammate model setting is removed from `/config` ; agent-team teammates now inherit the leader's model unless the spawn command explicitly names one.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI

Feedback ⓘ 1 RELEASE · 2026-08-17 NOTES ↗

CODE ↗

Lightweight coding agent that runs in your terminal

> Codex CLI's alpha release centers on deeper `codex doctor` network diagnostics and a new TypeScript SDK config-passthrough option, alongside a handful of smaller TUI and policy additions.

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

01 **Raw config passthrough in TypeScript SDK** NEW 90

Adds `CodexOptions.configOverrides` to the TypeScript SDK to pass ordered `--config key=value` arguments unchanged to the Codex CLI, applied after structured `config` values — useful for settings like permission-map paths that aren't representable via structured dotted-key config.

Pass a permission-map config key with a literal path (not representable via structured dotted-key config) when starting a Codex session from the TypeScript SDK.

typescript

```typescript
const session = new Codex({
  configOverrides: [
    '--config', 'sandbox.writable_roots./home/user/project=true'
  ]
});
```

— Named API field plus a runnable code example rust-v0.148.0-alpha.21

02 **Proxy and network diagnostics in codex doctor** IMPROVED 82

`codex doctor` now reports proxy-environment details, macOS system-proxy configuration status (automatic/PAC, manual, direct, unavailable), and detects proxy authentication challenges (HTTP 407).

— Names exact proxy modes and HTTP 407 detection, clear command to run

rust-v0.148.0-alpha.21

THINNER COVERAGE BELOW

03 **Guardian v2 transcript and defaults updates** IMPROVED 43

Guardian v2 transcripts now include images, and Guardian v2 defaults are sourced from the model catalog.

— Describes what changed but no configuration path `rust-v0.148.0-alpha.21`

04 **TUI usability additions** NEW 35

Adds working-directory commands to the TUI, compacts successful command activity in the history view for a cleaner display, and enables editing queued messages with Vim history-up in the composer.

— No specific commands, flags, or keybindings named `rust-v0.148.0-alpha.21`

05 **Per-environment shell and command policies** IMPROVED 31

Honors per-environment shell variable policies and enforces environment-specific command policies.

— No mechanism or config surface specified `rust-v0.148.0-alpha.21`

06 **Configurable skill catalog token budget** NEW 26

Adds a configurable skill catalog token budget.

— No config key or default value given `rust-v0.148.0-alpha.21`

07 **Experimental app-server project APIs** NEW 15

Adds experimental app-server project APIs.

— Bare mention with no detail on scope or usage `rust-v0.148.0-alpha.21`

WAS THIS USEFUL?  

AGENT

◆ AI Agent Frameworks

Hermes

Feedback ⓘ 2 RELEASES · 2026-08-17 → 2026-08-18 NOTES ↗

CODE ↗

The agent that grows with you

> Hermes launched its CLI and desktop installer with a full onboarding and configuration command set, added NVIDIA SkillEvaluator-backed advisory scanning and trust/quarantine controls for project-local skills, introduced `hermes peer` for cross-machine bot-to-bot messaging, and expanded Bot Mode with multi-machine group chats alongside new provider, plugin, and connection improvements.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 Project-local skills with trust and quarantine controls NEW

100

Skills load from `<project-root>/.hermes/skills/` (Hermes-native) or `<project-root>/.agents/skills/` (cross-tool convention) when Hermes runs inside a git checkout. `hermes skills trust` / `hermes skills untrust` grant or revoke trust for a repo, stored in `skills.trusted_project_dirs`; `skills.project_discovery: false` disables project-skill scanning and notices entirely. Project skills sit at the highest-precedence tier (`project → local (~/.hermes/skills/) → external_dirs`), tagged `[project]` in the skill index, and are scanned with the same security scanner used for Skills Hub installs, with results content-hash cached under `~/.hermes/cache/project_skill_scans/` and re-run on content change. Skills with a malicious verdict (prompt-injection directives, credential-exfiltration commands, hidden-text tricks) are quarantined, excluded from the skill index and slash commands, and refuse to load with an explanatory error; non-interactive surfaces (cron, API, ACP) inherit the interactive trust decision from the working directory without prompting or auto-trusting.

Trust a cloned repo so its `.hermes/skills/` and `.agents/skills/` project skills are loaded in sessions started inside it.

```
$ cd ~/myproject && hermes skills trust
```

Disable project-skill discovery globally so Hermes never scans repos or shows trust notices.

```yaml
skills.project_discovery: false
```

— Full mechanism, precedence order, config keys, and commands. v2026.8.18

02 **`hermes peer` cross-machine bot-to-bot messaging** NEW

95

`hermes peer add <name> --url <URL> [--key <KEY>] [--note TEXT]` registers or updates a remote Hermes gateway as a peer, storing the key as `HERMES_PEER_<NAME>_KEY`. `hermes peer list` shows registered peers and whether each has a key configured. `hermes peer dm <peer>[/<agent>] [message]` messages a remote agent's canonical Bot Chat over a peer gateway and prints the reply to stdout — `<peer>` targets the gateway's main agent, `<peer>/<agent>` targets a named profile via its `/p/<profile>/` mirror. `hermes peer remove <peer>` removes a peer (key entry preserved). Once a peer is registered, the `agent.bot_mode_protocol` taught to every canonical Bot Chat automatically includes the peer roster and the `hermes peer dm` pattern, so agents discover cross-machine teammates without manual SOUL edits.

— Every subcommand and flag named with exact syntax. v2026.8.18

03 **Installer deployment flags and dependency provisioning** NEW

90

Supports `--skip-browser` to bypass the Playwright/Chromium install step for headless deployments, and `--skip-computer-use` to opt out of pre-installing `cua-driver`. Adds a root-mode FHS install layout (`/usr/local/lib/hermes-agent/`, `/usr/local/bin/hermes`) for shared-machine or system-service deployments, with per-user config still stored under `~/.hermes/` or `$HERMES_HOME`. Non-sudo/unprivileged installs degrade gracefully — Chromium installs into the user's local Playwright cache and the required admin command is printed. The installer auto-provisions Python 3.11 (via `uv`), Node.js v22, `ripgrep`, and `ffmpeg` with no manual prerequisite installation.

Install on a headless service account that has no sudo access and does not need browser automation.

```
$ curl -fsSL https://hermes-agent.nousresearch.com/install.sh | bash -s -- --skip-browser
```

— Names every flag, path, and dependency; one runnable example. product docs

04 **NVIDIA SkillEvaluator advisory scanning on skill installs**

NEW 90

Skill installs run Tier 1 advisory checks via NVIDIA SkillEvaluator and SkillSpector, covering license and security checks; findings print with file and line before install confirmation, with credential-class findings (private keys, cloud access keys, tokens, credentialed connection strings) highlighted in red while PII findings are informational and never block install. The dashboard's Browse-hub scan button now returns this advisory data in a `tier1` field alongside the built-in scanner verdict.

Install the NVIDIA scanner binaries so Hermes runs Tier 1 advisory checks on every hub install.

```
$ uv tool install --python 3.13 "skillevaluator @
git+https://github.com/NVIDIA/SkillEvaluator.git@v0.1.0"
uv tool install
"git+https://github.com/NVIDIA/SkillSpector.git@v2.9.5"
```

— Names scanner tools, field, and finding categories with install command. v2026.8.18

05 **Tool Gateway via `hermes setup --portal`** NEW 85

`hermes setup --portal` logs in to Nous Portal, sets Nous as the LLM provider, and enables the Tool Gateway — web search, image generation, TTS, and a cloud browser — under a single Nous Portal subscription, all in one command.

Fastest onboarding path: authenticate with Nous Portal and enable the full Tool Gateway (web search, image generation, TTS, cloud browser) in one step.

```
$ hermes setup --portal
```

— Names command, provider, and bundled tools; runnable example. product docs

06 **Suite of CLI setup and management commands** NEW 80

Adds `hermes model` to choose and configure the LLM provider and model, `hermes tools` to configure enabled tools, `hermes gateway setup` to connect messaging platforms (Telegram, Discord, Slack, and others), `hermes config set` / `hermes config get` to read and write individual configuration values, a `hermes setup` wizard to configure all settings at once, `hermes desktop` to install and launch the desktop app after a CLI-only install, `hermes import` and `hermes profile import` for restoring a full backup or a single agent profile when migrating machines, `hermes doctor` to diagnose missing pieces and the detected install method, `hermes config check` / `hermes config migrate` to detect and migrate stale configuration after updates, and `hermes update` to auto-detect the install path (git installer, Docker, or NixOS) and print the matching update command.

Diagnose a broken or post-upgrade setup to see exactly what is missing and verify the detected install method.

```
$ hermes doctor
```

— Names every command but each individually thin. product docs

07 **Inline message components and widget interactivity** NEW

65

Plugins can render inline components inside assistant messages via `::name{...}` directives, including `::preview` which renders pages live inside the message. Widget clicks now reach the agent as hidden user turns, making widget interaction a first-class agentic response path.

— Names directive syntax and widget mechanism, no example given. v2026.8.18

08 **Bot Mode plugin and multi-machine group chats** NEW

65

Bot Mode (`hermes-bots`) ships as a built-in, default-on plugin providing the core teammate protocol for Bot Chat sessions. Adds cross-connection Bot Mode — a Create-on picker and multi-machine group chats — plus a `disband` command to delete group chats.

— Names plugin, picker, and disband command; no config detail. v2026.8.18 v2026.8.16.2

09 **Remote gateway connection improvements** IMPROVED 65

Adds support for remote gateway headers, carried through the connections registry, test probes, and Settings UI. Adds connection-aware plugin routing in the desktop app, scoping plugin socket connections to the active profile. Adds a status bar reconnect action for offline gateways, with self-healing for dropped SSH/HTTP registered remote connections.

— Names UI surfaces and connection types, no exact command. v2026.8.16.2

10 **Codex OAuth context window increase** IMPROVED 60

Raises the Codex OAuth context window to 900K tokens for the gpt-5.6 family and gpt-5.4 (subscription 1M rollout), and to a live-verified 350K for those same models.

— Concrete numbers and models, but no user action. v2026.8.16.2

THINNER COVERAGE BELOW

11 **Out-of-sync build warning in About** NEW 50

Adds an out-of-sync warning in the About screen when the app build diverges from the runtime, with a direct link to the installer.

— Names UI location and behaviour, a clear starting point. v2026.8.18

12 **Cua Driver 0.20 support with auto-repair** IMPROVED 50

Supports Cua Driver 0.20 runtime contracts for computer-use, with auto-repair of installed drivers that fail the runtime contract.

— Names version and behaviour, no manual steps given. v2026.8.16.2

13 **Meta Muse Spark provider support** NEW 45

Adds the Meta Model API (Muse Spark) as a bundled provider plugin, with Meta Muse Spark 1.2 available in the OpenRouter curated picker.

— Names provider and version but no usage steps. v2026.8.18

14 **context-manager protocol on SessionDB** NEW 40

Adds `context-manager` protocol support on `SessionDB` for safer session state management.

— Names protocol and component, no mechanism detail. v2026.8.16.2

15 **Per-plugin durable data directory** NEW 30

Adds a per-plugin durable data directory that survives plugin updates and removal.

— States behaviour only, no path or config named. v2026.8.18

16 **Browser Use mode A/B benchmark** NEW 30

Adds a Browser Use mode A/B benchmark (`feat(evals): add the Browser Use mode A/B benchmark`) to the evals suite.

— Names commit-style tag but no mechanism or usage. v2026.8.18

17 **UTF-16 file reads and Gemini 3 tool-call ID preservation** IMPROVED 30

Supports UTF-16 file reads and preserves Gemini 3 tool-call IDs.

— Bare statement, no mechanism or scope given. v2026.8.16.2

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Agentic Infrastructure

emisar

Feedback ⓘ 1 RELEASE · 2026-08-18 NOTES ↗

An MCP that lets AI tools securely connect to your infrastructure, write IaaS code, debug issues, and assist during incidents - without risking production stability. Built for security teams to approve and infrastructure teams to experience like magic.

> emisar v0.41.0 focuses on tightening access control - role-scoped consoles, immediate session reconnection on permission changes, and stricter checks on sensitive reads - while expanding its pack catalog to 100 packs and 1,689 actions with new Consul, GCP Monitoring, and HCP Terraform capabilities.

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

01 **Pack catalog expands to 100 packs with new Consul, GCP, and HCP Terraform actions**

NEW 75

The pack catalog expands to 100 packs and 1,689 actions. The Consul pack gains the ability to snapshot registration churn; the debugging pack adds bounded process context, environment-key, argv, and connection reads; the GCP Monitoring pack adds Cloud Logging name and entry reads; and the HCP Terraform pack exposes replacement paths in plan summaries.

— Concrete counts and named actions across four distinct packs v0.41.0

02 **Live session reconnection on permission changes** IMPROVED

60

Role and scope changes (member role, runner access, pack access) now take effect immediately in open web sessions: the session is reconnected before stale authority can be acted on.

— Clear mechanism and named scopes, but no user-facing action v0.41.0

THINNER COVERAGE BELOW

03 **Access control hardening for admins, reads, and staff console**

IMPROVED 55

Scoped admins are now blocked from delegating more reach than they hold or from triggering account-wide pack cleanup. Reads that expose every pack version or an exact run command now require a checked subject rather than relying on already-filtered callers, and the read-only staff console now requires MFA proof tied to the current enrollment.

— Three distinct security fixes named but described only in prose v0.41.0

04 **Role-scoped consoles for billing, operators, and general access**

NEW

50

Billing managers now receive a finance-only console, and operators can manage agents, approvals, and runbooks without gaining team or policy administration rights. Console navigation, actions, filters, and empty states now adapt to the member's role and access level, with role-aware explanations shown for restricted pack and action views.

— Names specific roles and console elements but no exact UI paths v0.41.0

05 **Bidirectional pagination for SSO group mappings** IMPROVED

25

SSO group mappings now paginate in both directions.

— Single-line change with no further mechanism or scope given v0.41.0

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Giskard

Feedback ⓘ 1 RELEASE · 2026-08-18 NOTES ↗

Open-Source Evaluation & Testing library for LLM Agents

> Giskard's v1.0.1rc1 release adds PEP 561 type stubs and tighter public type exports for downstream tooling, expands prompt-injection test coverage, and introduces two breaking changes to scenario persistence and check field naming.

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

01 **Check field renamed to target_key** BREAKING 65

The value-under-test field on every check is renamed to `target_key` — any code or serialized data referencing the old field name will break.

— Names exact field and scope, actionable migration note giskard-core/v1.0.1rc1

02 **PEP 561 type stubs across core packages** NEW 60

Ships `py.typed` marker files (PEP 561) for `giskard-core`, `giskard-llm`, `giskard-checks`, and `giskard-scan`, enabling static type checking in downstream projects that import these packages.

— Names exact packages and standard, no usage example given giskard-core/v1.0.1rc1

03 **Persisted-scenario tree rejects unknown fields** BREAKING 60

The persisted-scenario tree now rejects unknown fields; any saved scenario files containing unrecognized keys will fail to load after upgrading.

— Clear before/after and migration impact, no field examples giskard-core/v1.0.1rc1

THINNER COVERAGE BELOW

04 **Tightened public types in giskard.types** IMPROVED 43

Exports and tightens public `Literal` `/status` types from `giskard.types` for stricter API contracts.

— Names module and type kind but no concrete example giskard-core/v1.0.1rc1

05 **Expanded prompt-injection test suite** IMPROVED 28

Expands the prompt-injection test suite with additional scenarios and templates covering more injection attack patterns.

— No specific counts, templates, or scenario names given `giskard-core/v1.0.1rc1`

└──▷ **BREAKING ON UPGRADE**

! Persisted-scenario tree now rejects unknown fields — any saved scenario files containing unrecognized keys will fail to load after upgrading.

! The value-under-test field is renamed to `target_key` on every check — any code or serialized data referencing the old field name will break.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ 1 RELEASE · 2026-08-17 BLOG ↗

Cloudflare is a global network platform providing CDN, DDoS protection, DNS, security, and developer services for websites and applications.

> Cloudflare's window is led by four new Workers AI models (Grok 4.6, DeepSeek V4 Flash/Pro with million-token context, Qwen 3.8 27B) and package-ecosystem-aware Gateway HTTP policies, alongside new ways to attach Cloudflare Access directly to Workers, Artifacts data localization, Realtime SFU delivery controls, new Image Transformations/R2/MCP API endpoints, Certificate Transparency Monitoring GA, and scheduled WAF rules for HTTP/2 request smuggling and XSS event-handler coercion.

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

01 **New models added to Workers AI** NEW 90

Workers AI gains four new models: `grok-4.6` (xAI Grok 4.6), a flagship reasoning model supporting text and image inputs, function calling, and structured outputs for coding and agentic tasks; `@cf/deepseek-ai/deepseek-v4-flash-0731` (`deepseek-v4-flash-0731`), the official release superseding the preview with enhanced agentic capabilities; `@cf/deepseek-ai/deepseek-v4-pro-0813` (`deepseek-v4-pro-0813`), a high-capability reasoning model with a 1,048,576-token context window for long-horizon agentic workflows; and `@cf/qwen/qwen3.8-27b` (`qwen3.8-27b`), a 27-billion-parameter vision-language model from Alibaba with a 262,144-token context window, thinking mode, and function calling. DeepSeek and Qwen models are available via Workers AI binding, `POST /ai/run` , and (DeepSeek) the OpenAI-compatible endpoint; DeepSeek requires a Workers Paid plan or prepaid credits.

— Names all four models, IDs, context windows and access paths. product docs

02 **Package ecosystem selectors in Gateway HTTP policies** NEW 90

Cloudflare Gateway HTTP policies add `pkg.ecosystem` , `pkg.name` , `pkg.version` , `pkg.namespace` , and `pkg.purl` selectors so teams can allow or block software package downloads by ecosystem (npm, PyPI, RubyGems, Cargo, Go, Maven, NuGet), package name, version, or PURL; `pkg.purl` is API-only and the feature requires TLS decryption to be enabled.

Block downloads of a specific npm package from any registry (including corporate proxies) in a Gateway HTTP policy.

📍 # In Zero Trust dashboard: Settings › Gateway › HTTP Policies › Add a rule # Selector: Package Ecosystem = npm # Then add nested selector: pkg.name = lodash # Action: Block # Or via API — the selectors to reference are pkg.ecosystem and pkg.name

— Named selectors, ecosystems and TLS requirement with a UI/API example. product docs

03 **Data localization for Artifacts namespaces** NEW 90

A `jurisdiction` field on the Artifacts namespace creation API (`POST /client/v4/accounts/{account_id}/artifacts/namespaces`) lets teams pin repo data to `'eu'` or `'us'` for data localization; the jurisdiction is immutable after namespace creation.

Create an Artifacts namespace pinned to the EU for GDPR-scoped repo data.

```
$ curl --request POST \
  --url
"https://api.cloudflare.com/client/v4/accounts/$ACCOUNT_ID/artifacts/namespaces" \
  --header "Authorization: Bearer $CLOUDFLARE_API_TOKEN" \
  --header "Content-Type: application/json" \
  --data '{"namespace": "my-eu-namespace", "jurisdiction": "eu"}'
```

— Exact endpoint, field values, immutability note, and runnable curl example.

product docs

04 **Image Transformations settings API endpoints** NEW 90

New API endpoints manage Image Resizing configuration: `GET /accounts/{account_id}/settings/transformations` lists an account's Image Resizing configurations; `GET` / `PUT` `/zones/{zone_id}/settings/transformations_allowed_origins` read and change the Allowed Origins setting; and `GET` / `PUT` `/zones/{zone_id}/settings/transformations_c2pa` read and change the C2PA (Content Credentials) setting per zone.

Restrict Image Transformations to a specific allowed origin on a zone to prevent unauthorized image resizing from third-party domains.

```
$ curl -X PUT
'https://api.cloudflare.com/client/v4/zones/{zone_id}/settings/tr
ansformations_allowed_origins' \
  -H 'Authorization: Bearer <API_TOKEN>' \
  -H 'Content-Type: application/json' \
  --data '{"value": ["https://example.com"]}'
```

Enable C2PA content credentials on image transformations for a zone to embed provenance metadata in served images.

```
$ curl -X PUT
'https://api.cloudflare.com/client/v4/zones/{zone_id}/settings/tr
ansformations_c2pa' \
  -H 'Authorization: Bearer <API_TOKEN>' \
  -H 'Content-Type: application/json' \
  --data '{"value": "on"}'
```

— All five endpoints named with runnable curl examples for two of them. product docs

05 **Local Access simulation via wrangler.jsonc** NEW 86

An `access.dev` block in `wrangler.jsonc` lets developers locally simulate Cloudflare Access by setting an `aud` and a mock `identity` (e.g. `email`), so authenticated and unauthenticated flows can be tested without deploying.

Test an Access-protected Worker locally using a mock identity without deploying.

```json
{
  "access": {
    "dev": {
      "aud": "my-app",
      "identity": {
        "email": "admin@example.com"
      }
    }
  }
}
```

— Exact config key, file, and a runnable JSON snippet provided. product docs

06 **Configurable delivery modes for Realtime SFU DataChannels**

NEW 85

The Realtime SFU DataChannels API (`POST /v1/apps/{app_id}/sessions/{session_id}/datachannels/new`) adds `ordered` , `maxRetransmits` , and `maxPacketLifeTime` fields, enabling unordered and partially reliable delivery modes for latency-sensitive applications.

— Exact endpoint and field names, no worked example given. product docs

---

07 **New WAF rules for request smuggling and XSS coercion** NEW

75

Scheduled for 2026-08-24, the WAF adds rule `...1489d892` (HTTP/2 Request Smuggling - Request Body Anomaly), rule `...20646260` (XSS - JavaScript Event Handler Coercion - Headers), rule `...d706d517` (XSS - JavaScript Event Handler Coercion - Body), and rule `...660886c8` (XSS - JavaScript Event Handler Coercion - URI), all deploying in Log mode. The beta rule `...c293b926` (XSS, HTML Injection - Script Tag - Beta) is merged into the existing rule `...7b58420b` (XSS, HTML Injection - Script Tag) on the same date.

— Named rule IDs, mode and date, but change is automatic, not user-actioned.

WAF - WAF Release - Scheduled changes f…

---

08 **Access identity API for Workers** NEW

70

Workers can call ctx.access.getIdentity() on the `ctx.access` object to read an authenticated request's Cloudflare Access identity — email, name, and groups — without manually validating the Access JWT.

— Names exact method but no runnable example given. product docs

---

09 **Access policies attachable directly to Workers** NEW

70

Cloudflare Access policies can now be attached to a Worker directly so all its domains, routes, and preview URLs are protected automatically, including an account-wide 'protect all Workers by default' option with a per-Worker bypass to make individual Workers public.

.
.
.
.

— Clear mechanism and screenshots but no exact API/config path. product docs

---

10 **R2 catalog metadata deletion endpoint** NEW

70

Adds `DELETE /accounts/{account_id}/r2-catalog/{bucket_name}/delete` to remove R2 catalog metadata for a bucket.

Clean up stale R2 catalog metadata for a decommissioned bucket.

```
$ curl -X DELETE
'https://api.cloudflare.com/client/v4/accounts/{account_id}/r2-
catalog/{bucket_name}/delete' \
  -H 'Authorization: Bearer <API_TOKEN>'
```

— Exact endpoint and runnable curl example, but no other detail. product docs

THINNER COVERAGE BELOW

11 **Certificate Transparency Monitoring general availability** NEW

45

Certificate Transparency Monitoring is now generally available across all Cloudflare plans, with automatic filtering of certificates that Cloudflare issues on a zone's behalf.

— States scope and behaviour but no config surface named. product docs

12 **MCP Portal and Server management endpoints** NEW

32

Adds API endpoints to update and delete an MCP Portal, and to read details, update, and delete an MCP Server.

— No exact paths, methods, or examples given. product docs

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 18, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue