THE CYBER TOOLCHAIN

NO. 044

# Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 13, 2026 · EVERY DAY



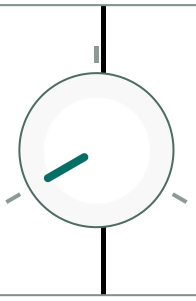
*The daily firehose — everything the toolchain shipped today, already filtered.*

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

**scroll**



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 20 tools

## OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

# Dradis

Feedback ⓘ 1 RELEASE · 2026-08-13 BLOG ↗

Dradis is a collaboration platform that aggregates security findings and vulnerabilities from multiple tools into centralized reports for teams.

> Dradis Pro v5.3 centers on reusable ServiceNow and CSV field mappings in the Mappings Manager, Gateway HTML/PDF export directly from projects, and multiplayer Echo chat on Issues, alongside smaller integration, API, and dashboard updates.

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

01 **ServiceNow field mapping in Mappings Manager** NEW 85

Adds `ServiceNow` to the Mappings Manager: define how Dradis fields map to ServiceNow fields, with Liquid templating support, across the Application, Container, and Vulnerability Response modules, so mapped fields auto-populate whenever a Vulnerable Item is created. Configure via `Mappings Manager` → `ServiceNow` → `New ServiceNow Mapping`.

Set up a reusable ServiceNow mapping with Liquid so your team's standard Dradis fields populate automatically every time a Vulnerable Item is created.

📍 In Dradis Pro, go to `Mappings Manager` → `ServiceNow` → `New ServiceNow Mapping`, then define each Dradis field (e.g. using Liquid like {{ `issue.fields.severity` }}) mapped to the target ServiceNow field across the Application, Container, or Vulnerability Response module.

Set up reusable ServiceNow field mappings with Liquid so every Vulnerable Item export is pre-populated consistently.

📍 In Dradis Pro, go to Mappings Manager → ServiceNow → New ServiceNow Mapping, then define each field mapping using Liquid (e.g. {{ issue.title }}) to control exactly what is sent to each ServiceNow field across the Application, Container, and Vulnerability Response modules.

Set up a reusable ServiceNow field mapping with Liquid so every Vulnerable Item created from Dradis populates the correct fields automatically.

📍 In Dradis Pro, go to Mappings Manager → ServiceNow → New ServiceNow Mapping. Define how each Dradis field (e.g. Title, CVSS Score) maps to the target ServiceNow module field, using Liquid syntax such as {{ `issue.fields.title` }}

where you need dynamic content. Save the mapping; it will be applied automatically each time a Vulnerable Item is created.

— Names config path, modules and Liquid mechanism `New in Dradis Pro v5.3 (2026-08-13)`

---

## 02 Gateway HTML/PDF export from project Export tab NEW

83

Adds an `HTML / PDF` Gateway export option directly from a project's `Export` tab, generating themed HTML output (including the Athena theme, recommended for print-to-PDF) without requiring access to the Results Portal, alongside the existing Word, Excel, and CSV export options.

Generate a polished, print-ready PDF report directly from a project without sharing it through the Results Portal.

> 📍 In Dradis Pro, open your project, go to the Export tab, select `HTML / PDF`, choose the `Athena` Gateway theme, then print the resulting browser view to PDF.

Generate a print-ready PDF deliverable from a project without sharing it through Gateway.

> 📍 Open your project → Export tab → select HTML / PDF → choose the Athena Gateway theme → open the generated HTML in your browser → File → Print → Save as PDF.

Produce a polished PDF deliverable from a project without sharing it through Gateway.

> 📍 In Dradis Pro, open your project → Export tab → select `HTML / PDF` → choose your Gateway theme (e.g. Athena) → open the generated HTML in your browser → use browser print-to-PDF to save the report.

Export a polished PDF deliverable from a project without sharing it via the Results Portal.

> 📍 Open your project in Dradis Pro, go to the Export tab, select HTML / PDF, choose the Athena Gateway theme, and generate the report. Open the resulting HTML in your browser and use the browser print dialog to save as PDF.

Generate a polished PDF deliverable from a project without sharing it through the Results Portal.

> 📍 In your project, go to the Export tab, select `HTML / PDF`, and choose the `Athena` Gateway theme. Open the generated HTML in your browser and use browser print-to-PDF — print-specific styles remove interactive elements such as the comment box automatically.

— Exact UI path and theme named, runnable steps given New in Dradis Pro v5.3 (2026-08-13)

03 **Multiplayer Echo chat on Issues** NEW 70

Adds multiplayer Echo chat on Issues via the Echo tab: start a conversation from a saved prompt, send follow-up messages, and watch Roslin's replies stream live to all project members viewing the same Issue.

Collaborate with your team to refine a finding using Echo multiplayer chat without losing context between messages.

> 📍 In Dradis Pro, open an Issue → Echo tab → select a saved prompt → send your initial message. Teammates viewing the same Issue see Roslin's reply stream live and can add follow-up messages to the same thread.

Collaborate with teammates on refining a finding using multiplayer Echo chat on an Issue.

> 📍 Open an Issue in Dradis Pro, navigate to the Echo tab, select a saved prompt, and send your first message to Roslin. Follow up with additional messages in the same thread — teammates viewing the same Issue will see replies stream live and can add their own messages.

— Describes live-streaming mechanism and shared-session scope

New in Dradis Pro v5.3 (2026-08-13)

04 **CSV import mapping in Mappings Manager** NEW 66

Adds CSV import support to the Mappings Manager: column-to-Issue/Evidence field mappings are defined once per Report Template and applied automatically to future uploads that use the same column headers.

Eliminate repetitive CSV column mapping on recurring imports by saving a mapping for a project's Report Template.

> 📍 In Dradis Pro, import a CSV into a project that has a Report Template assigned → map columns to Issue/Evidence fields when prompted → Dradis saves the mapping. On the next upload with the same column headers, the saved mapping is applied automatically.

— Names Report Template mechanism but light on limits New in Dradis Pro v5.3 (2026-08-13)

THINNER COVERAGE BELOW

05 **API author attribution and Content Block activity logging**

IMPROVED 50

Sets the `author` field to the authenticated user automatically when creating notes and evidence via the REST/JSON API, and adds Recent Activity entries for API edits made to Content Blocks.

– Names exact API field and activity type, no example call

New in Dradis Pro v5.3 (2026-08-13)

06 **Burp and VSTS integration additions** NEW 45

Adds Type as an available Issue field in the Burp integration, and adds Work Item attachment support to the VSTS integration.

– Names both fields but no mechanism or scope detail New in Dradis Pro v5.3 (2026-08-13)

07 **Author dashboard redesign** NEW 45

Redesigns the author dashboard with contextual widget panels, a What's New feed, CTA cards, and an onboarding view for authors without active projects.

– Lists UI elements but no navigation or behaviour detail

New in Dradis Pro v5.3 (2026-08-13)

08 **Line-break preservation in streamed Ollama responses** IMPROVED 35

Preserves line breaks in streamed Ollama responses for models that emit them as standalone tokens.

– Names cause but no user-facing steps or scope product docs

09 **Move task card to different list** NEW 25

Adds an option to move a task card to a different list.

– Bare description, no location or mechanism given New in Dradis Pro v5.3 (2026-08-13)

10 **Admin 1:1 training requests** NEW 25

Allows Admins to request 1:1 training for new users.

– Bare description, no process or location given New in Dradis Pro v5.3 (2026-08-13)

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

## APPLICATION SECURITY

◆ SAST / SCA / Code Security

# npq

Feedback ⓘ 1 RELEASE · 2026-08-13 NOTES ↗ CODE ↗

safely install npm packages by auditing them pre-install stage

> npq v3.26.0 expands its detection of AI coding-agent environments to better tailor install-time auditing when running under such tools.

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

01 **Expanded coding-agent detection signals** IMPROVED 26

npq now expands its coding-agent detection signals to more accurately identify when it is running inside an AI coding agent environment.

— No mechanism, signals, or config named beyond prose description. v3.26.0

WAS THIS USEFUL? 👍 👎

## Vercel deepsec

Feedback ⓘ changes since 2026-08-09 CODE ↗

Deepsec is a security harness for finding vulnerabilities in your codebase powered by coding agents

deepsec adds a `local` model route so machines with an existing claude/codex login skip all credential setup and env-var checks.

└──▷ **GET THIS VERSION**

```
$ git clone --branch commits-2026-08-09 https://github.com/vercel-labs/deepsec.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-09
```

└──▷ **TRY IT**

Bootstrap a workspace on a machine already logged in to the claude CLI, skipping all cloud credential setup.

```
$ pnpm deepsec --model-auth local
```

Persist the local-subscription route in config so every subsequent preflight skips env-var checks automatically.

deepsec.config.yaml

```
ai:
  mode: local
  provider: local
```

Adds `--model-auth local` CLI flag (and interactive Use local subscriptions as option 4 in the model-access prompt) to configure a `local` model route, delegating model auth entirely to the machine-wide `claude` / `codex` login — no credential is resolved, injected, or verified.

Sandbox commands under the local route still require `AI_GATEWAY_API_KEY` (an ambient key is accepted), preserving the firewall token-brokering requirement.

Documents `--scaffold-only` in `docs/getting-started.md`, including what it writes, what it skips, and how to resume the automated flow later, plus a Using local subscriptions section across models.md, FAQ, CLI help, and the scaffolded workspace README.

WAS THIS USEFUL? 👍 👎

## NETWORK SECURITY

◆ Firewalls & Network Defense

### Palo Alto Networks

Feedback ⓘ 1 RELEASE · 2026-08-13 DOCS ↗

BLOG ↗

Palo Alto Networks is a cybersecurity company providing cloud-native security platforms, firewalls, and threat protection solutions.

> Palo Alto Networks brought a previously unavailable OpenAI model to its Unit 42 Frontier AI Exposure Analysis for automated exploit validation, and added new response and integration capabilities to Device Security.

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

01 **GPT-5.6 Daybreak in Frontier AI Exposure Analysis** NEW

65

Unit 42 Frontier AI Exposure Analysis now deploys GPT-5.6 Daybreak — previously unavailable for commercial use — directly in customer environments to find exposures, test exploitability, validate attack paths, and prioritize remediation.

— Names model and use case but no setup steps 2026-08-13-3fb889a9

THINNER COVERAGE BELOW

02 **Action Center and IGEL UMS integration in Device Security**

NEW

50

Device Security adds an Action Center enabling automated response actions against IoT/OT devices, and gains an integration with IGEL UMS for unified endpoint and device security management.

— Names two additions but no config or navigation detail 2026-08-13-3fb889a9

WAS THIS USEFUL? 👍 👎

## CLOUD & CONTAINER SECURITY

◆ CNAPP

# Palo Alto Prisma Cloud

Feedback ⓘ 1 RELEASE · 2026-08-13

DOCS ↗ BLOG ↗

Palo Alto Prisma Cloud is a cloud security platform that protects cloud infrastructure, applications, and data across multiple cloud providers.

> Prisma Cloud's AIRS Runtime API now natively integrates with OpenAI Codex Enterprise for org-wide prompt inspection and DLP, alongside a new AI Discovery capability via Cortex AISPM.

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

## 01 Prisma AIRS integration with OpenAI Codex Enterprise NEW

81

The Prisma AIRS Runtime API now integrates natively with OpenAI Codex Enterprise: pasting the Prisma AIRS API key and endpoint into the Codex Global Admin Console routes all user prompts across an organization through real-time inspection with no client-side agents required. This also extends Enterprise DLP controls to Codex, scanning developer inputs — including logs, config snippets, and code functions — for proprietary data and credentials before they reach the AI model.

Enable org-wide Prisma AIRS scanning for all Codex users without client-side agents — useful for security teams who need DLP and runtime threat inspection across an entire developer organization.

> 📍 1. In the Prisma AIRS management console, generate your API key and endpoint. 2. In the Codex Enterprise Management UI, open the Global Admin Console. 3. Paste the Prisma AIRS API key and endpoint into the Codex configuration fields and save. 4. All user prompts across your organization are now automatically routed through the Prisma AIRS Runtime API for real-time inspection.

— Full setup steps and DLP scope given; no API endpoint path named. 2026-08-13-a70faccf

THINNER COVERAGE BELOW

## 02 AI Discovery via Cortex AISPM NEW

23

Adds AI Discovery capability via Cortex AISPM, surfaced as a new documentation area alongside the Codex integration.

— Only a name is given, no mechanism or usage detail. 2026-08-13-a70faccf

WAS THIS USEFUL?  

◆ Cloud Security Posture

# Prowler

Feedback ⓘ 1 RELEASE · 2026-08-13 NOTES ↗ CODE ↗

Prowler is the Open Cloud Security Platform trusted by thousands to automate security and compliance in any cloud environment. With thousands of ready-to-use checks and compliance frameworks, delivers real-time, customizable monitoring and seamless integrations, making cloud security simple, scalable, cost-effective for organizations of any size.

> Prowler 5.39.0 introduces CMMC 2.0 as a universal compliance framework and substantially expands Attack Paths with 20 new AWS privilege-escalation queries and a redesigned graph, alongside a Compliance Watchlist, multi-domain SAML SSO, guided Azure Management Group onboarding, a new Jira Cloud integration framework, and a Lighthouse AI Skills menu for findings.

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

## 01 Jira Cloud integration class and developer framework NEW

95

A new `Jira` integration class (in `prowler/lib/outputs`) exposes send_findings(findings, project_key, issue_type), test_connection(redirect_uri, client_id, client_secret, raise_on_exception), get_projects(), get_available_issue_types(project_key), get_access_token(), refresh_access_token(), auth_code_url(), and input_authorization_code() to push findings into Jira Cloud via the Atlassian REST API. The OAuth 2.0 authflow uses `AUTH_URL` (`https://auth.atlassian.com/authorize`), `TOKEN_URL` (`https://auth.atlassian.com/oauth/token`), and `API_TOKEN_URL` (`https://api.atlassian.com/oauth/token/accessible-resources`) with scopes `read:jira-user`, `read:jira-work`, and `write:jira-work`; a developer guide documents building further custom integrations (e.g. Slack, AWS Security Hub) in the same `prowler/lib/outputs` folder.

Authenticate to Jira Cloud and push a list of Prowler findings as individual Jira issues under a specific project and issue type.

```python
from prowler.lib.outputs.jira import Jira

jira = Jira(
    redirect_uri="http://localhost:8080",
    client_id="client_id",
    client_secret="client_secret",
)
jira.send_findings(findings=findings, project_key="SEC",
issue_type="Bug")
```

Validate Jira Cloud credentials before configuring the integration in a pipeline, returning a Connection object with error details on failure.

```python
from prowler.lib.outputs.jira import Jira

conn = Jira.test_connection(
    redirect_uri="http://localhost:8080",
    client_id="client_id",
    client_secret="client_secret",
    raise_on_exception=False,
)
print(conn.is_connected, conn.error)
```

— Names every method, OAuth URL and scope, with runnable code. product docs

02 **AWS privilege-escalation Attack Paths and graph redesign** NEW

85

The `iam_policy_allows_privilege_escalation` check gains 22 additional escalation combinations, and 20 new AWS privilege-escalation Attack Paths queries cover `iam:PassRole` paths across AWS Batch, Braket, Cognito Identity, ECS, EMR, EMR Serverless, GameLift, Glue, EC2 Image Builder, Kinesis Analytics, HealthOmics, EventBridge Scheduler, Systems Manager, and Step Functions, plus existing-resource abuse, permissions-boundary removal, role assumption, and IAM Identity Center permission-set policy injection. Each query now exposes an outcome category (code execution, privilege escalation, public exposure, resource inventory), and the Attack Paths graph was redesigned to group same-class resources into expandable nodes with counts and terminate every path in an explicit outcome node naming the destination impact.

— Named check and services, but only UI-level actionability. 5.39.0

03 **Grouped Jira dispatch via MCP tool** NEW

75

The `prowler_send_findings_to_jira` MCP tool gains Cloud-only grouped dispatch: it selects failed findings by check IDs against the latest completed scan and sends one Jira work item per check listing up to 50 affected resources, with per-group error reporting.

— Names the MCP tool and dispatch logic, no invocation example. 5.39.0

04 **Active tenant field on users/me endpoint** IMPROVED 70

`GET /api/v1/users/me` membership relationships now identify the active tenant with `meta.active` for both JWT and API key authentication.

— Exact endpoint and field named, directly callable. 5.39.0

05 **Supply-chain security hardening for container builds** IMPROVED 65

Prowler API, UI, SDK, and MCP container images (and Prowler Cloud/Private Cloud production images) now publish per-architecture SBOMs and build-provenance attestations. SDK and API container builds also verify checksums of downloaded PowerShell, Trivy, and zizmor binaries before installation, and Grype now complements Trivy across container-image security gates, detecting components and vulnerabilities manifest-based scanners can miss and blocking fixable high and critical findings.

— Internal build security named in detail, no reader action. product docs

06 **AWS hardcoded-secret detection checks** NEW 65

Adds `batch_job_definition_no_secrets`, scanning Batch job definition environment variables and command parameters for hardcoded secrets, and `awslambda_layer_no_secrets_in_content`, scanning Lambda layer package content for hardcoded secrets.

— Named check IDs, runnable directly via Prowler. 5.39.0

07 **CIS Microsoft 365 Foundations Benchmark checks** NEW 60

Adds 12 new checks for CIS Microsoft 365 Foundations Benchmark v7.0.0 across Admin Center, Defender, Entra ID, Exchange Online, and Microsoft Teams, followed by 20 new Microsoft 365 Entra ID checks for the same benchmark covering password protection, default user permissions, Conditional Access, sessions and authentication methods, and PIM and access reviews.

— Named benchmark and areas, no individual check IDs. 5.39.0

08 **CMMC 2.0 compliance framework** NEW 60

Adds CMMC 2.0 as a universal compliance framework covering all 149 requirements across Levels 1, 2, and 3, mapped to checks across AWS, Azure, GCP, Alibaba Cloud, Oracle Cloud, and Microsoft 365.

— Framework scope well specified, no usage command. 5.39.0

09 **Azure Management Group onboarding wizard** NEW 60

A guided wizard discovers Azure subscriptions within Management Groups and selectively onboards them through a single workflow. In Prowler Cloud/Private Cloud, entering the Microsoft Entra tenant ID and authenticating with a single tenant-wide service principal onboards all subscriptions under that tenant at once.

— UI wizard steps named, no API endpoint given. 5.39.0

10 **Lighthouse AI Skills menu and permission control** NEW 60

Adds a Lighthouse AI Skills menu on every finding in Prowler Cloud, offering Contextual Fix, Triage Decision, Systemic Scope, and a free-form prompt, with follow-up actions including creating a Jira issue or muting the finding (Cloud/Private Cloud only). A new Manage Lighthouse AI role permission in role forms and role details lets this configuration permission be granted or restricted independently of other permissions.

— Named skills and permission, UI-only actionability. 5.39.0

THINNER COVERAGE BELOW

11 **Compliance Watchlist for pinned frameworks** NEW 55

Organizations can pin compliance frameworks from any compliance view, manage them through a searchable catalog, and filter the Compliance section to show only pinned frameworks; the Overview page reports the latest score for every pinned framework.

— Describes UI behavior without named endpoints or config. product docs

12 **Resolved triage status for MANUAL findings** NEW 55

Adds a `Resolved` triage status for `MANUAL` findings in Prowler Cloud, recording written evidence and verifying the finding as an effective PASS for 90 days or until a real failure is detected (Cloud/Private Cloud only).

— Clear behavior and 90-day duration, no config key named. 5.39.0

13 **Pre-configured credential links in add-provider wizard** NEW 55

The add-provider wizard now surfaces pre-configured credential creation links: Cloudflare exposes a User API Token template and an Account-Owned template pinned

to the entered Cloudflare Account ID, and GitHub exposes a personal-repositories template and an organization-scanning template pinned to the entered identifier.

— Named provider templates in wizard, no API endpoint. 5.39.0

14 **Organization-level MCP onboarding tools** NEW 55

Adds eight organization MCP tools to the Prowler Cloud MCP server, enabling agents to onboard and manage entire AWS Organizations, GCP organizations, and Azure tenant root management groups end to end (Cloud/Private Cloud only).

— Named cloud providers covered, but no individual tool names. 5.39.0

15 **Multi-domain SAML SSO and sign-in method visibility** NEW 45

SAML SSO now supports one configuration authorizing a primary email domain plus up to 19 additional domains through a single Identity Provider and one stable Assertion Consumer Service (ACS) URL. The Users table now shows each account's sign-in methods as tags, including email/password, Google, GitHub, SAML with linked domains, and Partner SSO.

— Config behavior described, no exact setting names. product docs

16 **Severity scaling for organization_repository_creation_limited** IMPROVED 40

Scales the `organization_repository_creation_limited` GitHub check severity to low when repository creation is limited to private or internal visibility.

— Named check, minor severity tweak with no further detail. 5.39.0

WAS THIS USEFUL? 👍 👎

## DETECTION & RESPONSE

◆ SIEM

# Google SecOps

Feedback ⓘ 1 RELEASE · 2026-08-12 BLOG ↗

Google SecOps is a unified security operations platform combining threat detection, investigation, and response capabilities across cloud and enterprise environments.

> Google SecOps added a public-preview Cloud Logging integration for diagnosing feed and ingestion pipeline issues, and rolled out an upgraded rich-text editor across its SOAR case and reporting surfaces.

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

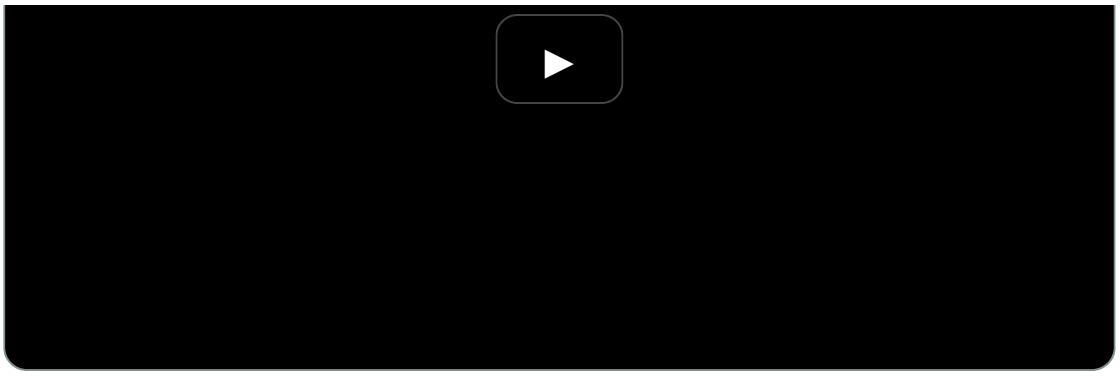
## 01 Cloud Logging integration for feed and pipeline diagnostics

NEW 90

Google SecOps adds a public preview Cloud Logging integration (requires a Bring Your Own Project (BYOP) Google Cloud project) to monitor, debug, and troubleshoot SIEM ingestion pipelines and feeds via Logs Explorer, supporting both push- and pull-based ingestion mechanisms. A new Debug with logs option on the Feed management page opens Logs Explorer pre-filtered for a specific feed for fast per-feed diagnosis. Exclusion filters can be configured in the Log Router to exclude specific logs (e.g., Storage Transfer Service (STS) logs) from being routed to Cloud Logging, and Gemini Cloud Assist can be used to investigate logging and metrics telemetry directly from the Google SecOps console.







Quickly isolate delivery failures for a specific feed without manually constructing Logs Explorer queries.

> 📍 In the Google SecOps console, go to Settings › Feed management, select a feed, and click `Debug with logs` to open Logs Explorer pre-filtered for that feed.

— Names BYOP, Logs Explorer, Log Router, STS, Gemini Cloud Assist, and gives UI steps.

August 12, 2026 (2026-08-12)

---

## 02 Upgraded rich-text editor across SOAR surfaces

BREAKING

70

The rich-text editor used in Cases Wall, the Use Case Upload dialog, the Report Template dialog, and the Dashboard Editor widget has been upgraded with semantic font sizes (Small, Normal, Large, Huge) and simplified table insertion/removal. As a result, the new editor no longer supports formatting inside table cells, and the Cut, Copy, and Paste buttons have been removed from the toolbar, though standard OS keyboard shortcuts still work.

— Names affected surfaces and font options but no exact navigation steps. product docs

### └──▷ ALSO FROM THESE RELEASES

### └──▷ BREAKING ON UPGRADE

! The upgraded rich-text editor no longer supports formatting inside table cells.

! The Cut, Copy, and Paste buttons have been removed from the rich-text editor toolbar (standard OS keyboard shortcuts remain supported).

---

WAS THIS USEFUL? 👍 👎

# Elastic Security Labs

Feedback ⓘ 1 RELEASE · 2026-08-11

BLOG ↗

Elastic Security Labs publishes security research, threat analysis, and vulnerability findings to help organizations understand emerging cyber threats.

> Elastic Security Labs shipped a full pipeline for auditing every Cursor AI coding agent action, spanning a cross-platform hook collector, an ECS-mapped logging schema, agent/IDE detection, MDM-free fleet deployment, and Elastic Agent ingestion with ES|QL querying.

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

## 01 Cross-platform hook collector for Cursor tool calls NEW

91

A 280-line dependency-free bash hook collector script ( `log-tool-calls.sh` ), published in the `elastic/elasticsearch-labs` repository with a PowerShell port for Windows, wires into Cursor's `hooks.json` (covering `sessionStart` , `beforeShellExecution` , `afterShellExecution` , `beforeMCPExecution` , `afterMCPExecution` , `postToolUse` , `postToolUseFailure` , `afterFileEdit` , `beforeReadFile` , `subagentStart` , `subagentStop` , `sessionEnd` , `stop` ) to record every Cursor AI agent tool call as JSONL to a date-rotated log file readable only by the owner ( `chmod 0600` ). It covers both the IDE and the headless `cursor-agent` CLI (used in CI and scripted runs) through the same `hooks.json` , firing identical hook events, and enforces local retention of 30 days with opportunistic deletion of older files while Elasticsearch serves as the system of record.

Deploy the hook collector system-wide on macOS so both the Cursor IDE and `cursor-agent` CLI log every tool call without MDM.

```
$ sudo mkdir -p /usr/local/share/ai-hooks
sudo cp log-tool-calls.sh /usr/local/share/ai-hooks/log-tool-
calls.sh
sudo chmod +x /usr/local/share/ai-hooks/log-tool-calls.sh
sudo tee "/Library/Application Support/Cursor/hooks.json" <<'EOF'
{
  "version": 1,
  "hooks": {
    "sessionStart":          [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "beforeShellExecution":  [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "afterShellExecution":   [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "beforeMCPExecution":    [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "afterMCPExecution":     [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "postToolUse":           [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "postToolUseFailure":    [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "afterFileEdit":         [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "beforeReadFile":        [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "subagentStart":         [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "subagentStop":          [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "sessionEnd":            [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }],
    "stop":                  [{ "command": "/usr/local/share/ai-
hooks/log-tool-calls.sh" }]
  }
}
EOF
```

Verify the CLI surface is captured separately from the IDE after deployment by running a headless agent call and checking the log.

```
$ cursor-agent --print "say hello"
tail -1 ~/.config/ai-hooks/logs/tool-calls-$(date -u +%F).jsonl
```

— Exact script, hooks.json, and file permissions given, fully runnable

13 million tool calls: auditing every A…

02 **Agent and IDE surface detection** NEW 75

Each event is tagged with a detected `agent` and `ide` value ( `cursor` , `cursor-cli` , or `remote` ) by inspecting inherited environment variables `CURSOR_VERSION` , `CURSOR_CODE_REMOTE` , `VSCODE_PID` , `VSCODE_CWD` , and `VSCODE_IPC_HOOK` , so IDE and CLI sessions are distinguishable within a single log stream.

Verify the CLI surface is captured separately from the IDE after deployment by running a headless agent call and checking the log.

```
$ cursor-agent --print "say hello"
tail -1 ~/.config/ai-hooks/logs/tool-calls-$(date -u +%F).jsonl
```

— Named env vars and a validation command confirm behaviour

13 million tool calls: auditing every A…

03 **Structured logging schema with ECS-conformant fields** NEW 60

Every log entry promotes top-level structured fields — `hook_event_name` , `tool_name` , `command` , `file_path` , `mcp_server` , `model` , `session_id` , `duration` , `duration_ms` — plus an ECS-conformant `event` object ( `event.category` , `event.type` , `event.outcome` ) so entries correlate with existing Elastic security indices.

— Field names listed but no query or command shown

13 million tool calls: auditing every A…

04 **MDM-free deployment via Cursor Cloud Distribution** NEW 60

The collector can be delivered to devices outside MDM via Cursor's Cloud Distribution console, which pushes a self-bootstrapping hook command carrying the collector script gzipped and base64-encoded, and verifies it against a pinned SHA-256 hash on every invocation.

— Mechanism described but deployment is via console UI, no command given

13 million tool calls: auditing every A…

05 **Elastic Agent filestream ingestion and ES|QL querying** NEW 60

An Elastic Agent filestream integration configuration parses the JSONL logs into structured fields, enabling ES|QL queries across all agent activity — for example, identifying which hosts ran an agent that read a `.pem` file in a given time window.

— Concrete query example given but no exact ES|QL syntax or config shown

13 million tool calls: auditing every A…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

## IDENTITY, ACCESS & DATA PROTECTION

◆ **Secure Access (VPN / SASE / ZTNA)**

### Zscaler

Feedback ⓘ 1 RELEASE · 2026-08-13 BLOG ↗

Zscaler is a cloud-based security platform providing zero-trust network access, threat prevention, and data protection for organizations.

> Zscaler launched Zero Trust Gateway (ZTGW) for Google Cloud, a native integration that transparently steers GCP workload traffic to Zscaler for inspection without touching applications or network routes.

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

01 **Zero Trust Gateway for Google Cloud** NEW 75

Zscaler Zero Trust Gateway (ZTGW) integrates with Google Cloud's Network Security Integration (NSI) to transparently steer workload traffic to Zscaler for deep inspection without application changes, manual tunnels, or route table modifications, and is deployable in under 10 minutes. It supports Single VPC, Multiple VPC, and Shared VPC Google Cloud topologies, uses Google Cloud Security Profiles and Profile Groups to control traffic steering and enable policy-driven enforcement, and provides unified visibility across GCP, other clouds, and on-premises environments through the Zscaler Admin Portal.

— Names integration mechanism, topologies, and console but no exact setup steps.

2026-08-13-d8c47b24

WAS THIS USEFUL? 👍 👎

# Zscaler Internet Access

Feedback

1 RELEASE · 2026-08-13

BLOG ↗

Zscaler Internet Access is a cloud-based security platform that inspects and controls user internet traffic to block threats and enforce policies.

> Zscaler launched Zero Trust Gateway for Google Cloud, bringing inline SSL/TLS inspection with threat protection, secure access controls for AI development tools, and consolidated multi-cloud firewall policy.

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

01 **Advanced Threat Protection with inline SSL inspection** NEW

53

Zero Trust Gateway for Google Cloud adds Advanced Threat Protection with inline SSL/TLS inspection and sandboxing to detect and block malware and ransomware in cloud traffic.

— Names inspection and sandboxing mechanism but no config surface 2026-08-13-d0653cd9

02 **Secure access controls for AI development tools** NEW

46

Zero Trust Gateway for Google Cloud enables secure access controls for AI development tools such as Devin and Cursor without blocking developer workflows.

— Names specific tools but no policy mechanism detail 2026-08-13-d0653cd9

03 **Multi-cloud firewall policy consolidation** NEW 44

Zero Trust Gateway for Google Cloud consolidates virtual firewall infrastructure across multi-cloud environments under a single policy framework to reduce sprawl.

— Describes consolidation goal without specific setup steps 2026-08-13-d0653cd9

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

# Zscaler Data Protection

Feedback ⓘ 1 RELEASE · 2026-08-13

BLOG ↗

Zscaler Data Protection is a cloud-based security service that prevents data loss and enforces policies across web, cloud, and email channels.

> Zscaler introduced a Zero Trust Gateway for Google Cloud that adds inline DLP and lets policy be scoped using Google Cloud resource metadata.

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

## 01 Zero Trust Gateway for Google Cloud with inline DLP NEW

55

Zscaler Zero Trust Gateway for Google Cloud adds inline Data Loss Prevention (DLP) to inspect and block sensitive data in transit across Google Cloud workloads. It supports Google Cloud resource labels and network tags as policy metadata, enabling context-aware DLP enforcement scoped to specific workloads (e.g., stricter rules for Production vs. Development environments) tied to the Google Cloud resource hierarchy.

— Describes mechanism and scope but no config/UI steps given. 2026-08-13-b5d3c8a8

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

# Auth0

Feedback ⓘ 1 RELEASE · 2026-08-13 BLOG ↗

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

> Auth0 brought Flexible Password Policy to general availability for database connections, consolidating password rules into a single configurable object and tightening the connections API to reject mixed legacy/new password configs.

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

## 01 Flexible Password Policy GA for database connections NEW

90

Auth0's Flexible Password Policy is now generally available, replacing legacy password fields ( `passwordPolicy` , `password_complexity_options` , `password_history` , `password_no_personal_info` , `password_dictionary` ) with a single `options.password_options` object on database connections. It adds composition controls (minimum length, required character types, blocking of sequential or repeated characters), password history enforcement against a configurable number of previous passwords, dictionary blocking via a 10,000- or 100,000-word built-in list plus custom entries, and profile-data blocking that rejects passwords containing values from a configurable list of user profile fields. New database connections created as of July 2026 default to this policy with a 15-character minimum length and no required character types, aligning with NIST guidance, and can be created via Dashboard, Management API, Deploy CLI, Terraform, or an Auth0 SDK.

— Names config object, sub-controls and defaults but no runnable example

Flexible Password Policy is now general…

## 02 Connections API rejects mixed password field payloads

85

BREAKING

`PATCH /api/v2/connections/{id}` now returns `400 invalid_body` if a request contains both `password_options` and any legacy password field ( `passwordPolicy` , `password_complexity_options` , `password_history` , `password_no_personal_info` , `password_dictionary` ) simultaneously. Automation that reads a connection and writes it back with legacy fields appended will break and must be updated to send only `password_options` .

— Exact endpoint, method, error code and required migration step

Flexible Password Policy is now general…

└──▷ ALSO FROM THESE RELEASES

:
:
:
:

└──▷ **BREAKING ON UPGRADE**

! `PATCH /api/v2/connections/{id}` returns `400 invalid_body` if a request contains both `password_options` and any legacy password field ( `passwordPolicy` , `password_complexity_options` , `password_history` , `password_no_personal_info` , `password_dictionary` ) simultaneously — automation that reads a connection and writes it back with legacy fields appended will break and must be updated to send only `password_options` .

! New database connections created without an explicit password configuration now default to `options.password_options` with a minimum password length of 15 (up from legacy defaults) — pipelines that create connections and expect legacy defaults will receive the new NIST-aligned defaults instead.

---

WAS THIS USEFUL?  

---

**BUILD**

◆ AI Coding Agents

# Kiro

Feedback ⓘ 1 RELEASE · SEEN 2026-08-13 NOTES ↗

> Kiro's CLI adds on-device voice dictation, a dedicated spec review screen for staging line comments, workspace-wide AGENTS.md steering, and makes cloud sessions opt-in by default.

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

01 **Voice dictation via /voice command** NEW 80

Adds the `/voice` command (also triggerable via `Ctrl+O` or by holding `Space`) for on-device speech-to-text prompt dictation using Whisper, with no audio leaving the machine and no cloud API key required. A `--continuous` flag enables hands-free back-and-forth mode that auto-starts a new recording each turn.

Dictate a prompt hands-free during a long agentic run without touching the keyboard.

```
$ /voice --continuous
```

— Names command, shortcuts, flag and on-device mechanism.

Voice Mode, Opt-In Cloud Sessions, and …

02 **Spec review screen at checkpoints** NEW 80

`Ctrl+X` at spec phase checkpoints opens a spec review screen where you can read phase documents and stage line comments; these staged comments are batched into a single revision request when the checkpoint question is answered.

Review a generated spec document and queue targeted line edits before the agent continues.

> 📍 At a spec phase checkpoint, press Ctrl+X to open the review screen, annotate the lines you want changed, then answer the checkpoint question to send all staged comments as one revision request.

— Names shortcut and exact workflow, no numeric limits given.

Voice Mode, Opt-In Cloud Sessions, and …

03 **Nested AGENTS.md steering context** IMPROVED 65

`AGENTS.md` files now load as steering context from anywhere in the workspace tree, not just the workspace root or `~/.kiro/steering/`, enabling co-location of steering instructions with the code they describe.

— Names file and paths but no example of use. Voice Mode, Opt-In Cloud Sessions, and …

THINNER COVERAGE BELOW

04 **Cloud sessions now opt-in by default** BREAKING 55

Cloud sessions are now disabled by default unless an administrator explicitly enables them via the Cloud Sessions toggle in the Kiro console; organizations that never configured this toggle will find cloud sessions disabled after upgrading until an admin turns it on.

— States exact upgrade impact and admin action needed.

Voice Mode, Opt-In Cloud Sessions, and …

└──▷ **BREAKING ON UPGRADE**

! Cloud sessions (v3) are now opt-in: organizations that never configured the Cloud Sessions toggle in the Kiro console will find cloud sessions disabled after upgrading until an admin explicitly enables them.

WAS THIS USEFUL?  

# OpenAI Codex CLI

Feedback ⓘ 1 RELEASE · 2026-08-13 NOTES ↗

CODE ↗

Lightweight coding agent that runs in your terminal

> Codex CLI's alpha.12 release focuses on session resilience and sandbox hardening: it adds turn recovery after interruptions, routes blocked network requests through the shared approval pipeline, and isolates inline visualization viewer documents from sandboxed sessions.

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

01 **Turn recovery for interrupted sessions** NEW 75

Adds `RecoverTurnRequest` and `CodexThread::recover_turn_if_idle` to resume an interrupted regular turn with its existing turn ID and updated thread settings, including Plan-mode recovery without injecting an empty user message. Turn recovery is rejected while another turn is already active, without applying settings or queuing input.

— Names the API surfaces but gives no usage example rust-v0.148.0-alpha.12

02 **Inline visualization viewer sandbox isolation** IMPROVED 75

Materializes inline visualization viewer documents in a dedicated cache under `CODEX_HOME`, keyed by source and artifact thread IDs, instead of alongside visualization artifacts, preventing sandboxed sessions from modifying viewer documents before they open in a browser. Visualization links are now created only when the active filesystem policy cannot write to the viewer cache, disabling links for full-disk-write sessions and applying the policy consistently when loading transcript history.

— Names CODEX_HOME cache location but no direct command rust-v0.148.0-alpha.12

THINNER COVERAGE BELOW

03 **Network requests routed through approval pipeline** IMPROVED 56

Blocked network requests now route through the shared approval pipeline so permission hooks, automatic review, and user review all use the common approval flow, including for background terminals started by an earlier turn. The final applied network decision is recorded in tool telemetry without exposing the destination or assigning an approval source.

— Explains mechanism but no flag or command surfaced rust-v0.148.0-alpha.12

WAS THIS USEFUL?  

## Cline

Feedback ⓘ 3 RELEASES · 2026-08-13 NOTES ↗ CODE ↗

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

> Cline shipped a directional Hub-daemon upgrade flow with automatic CLI self-update, tightened Claude Code provider integration and telemetry accuracy in the SDK, and overhauled the desktop session transcript with per-tool-call rows alongside new Vertex AI model support.

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

### 01 Directional Hub daemon upgrades with CLI self-update IMPROVED

85

Managed Hub daemons now upgrade directionally — when another Cline install ships a newer Hub build, hosts attach to the newer daemon and prompt to update and restart (via `onHubUpdateRestart`) instead of two installs repeatedly retiring each other's daemons; yolo and sandbox sessions are never interrupted by this prompt. After the user accepts the dialog, the CLI self-updates automatically, running `cline update` under the hood and printing instructions to reconnect.

— Names exact hook and update command, but flow is automatic `cli-v3.0.54` `sdk/sdk/v0.0.74`

### 02 Claude Code provider workspace and permission handling

IMPROVED

80

The Claude Code provider now declares its own native tools, anchors the session on the workspace directory, and loads `~/.claude` plus project settings so user-configured permission rules apply; file edits under the workspace are auto-approved while command execution remains gated by Claude settings.

— Names exact config path and permission scoping `sdk/sdk/v0.0.74`

03 **Session transcript and diff rendering overhaul** IMPROVED

75

Every tool call now gets its own row in the session transcript with its own icon, status indicator, and expandable detail — commands render like a terminal (e.g. `$ bun test`) with captured output on expand, and edits show diffs inline per hunk. Running tool rows are highlighted in brand violet and settle to gray on completion, with error rows staying red. File diffs in both chat rows and the diff panel now render through a shared syntax-highlighted renderer that follows the app theme, and the thinking indicator stays visible during quiet stretches of a turn (such as while tool arguments are streaming) so turns no longer appear frozen.

— Details UI mechanics but no config or API surface `desktop-v0.0.12`

04 **Vertex AI Fable 5 and custom model support** NEW

65

Vertex AI provider adds Fable 5 and support for custom model IDs, and the global-region picker no longer hides models from the live catalog.

— Names new model and fixes picker, UI-level action `desktop-v0.0.12`

05 **Compaction honors Max Output Tokens setting** IMPROVED

65

Compaction now respects the Max Output Tokens setting instead of a hardcoded 1024-token cap, enabling reasoning models to produce summaries instead of skipping compaction.

— Names exact setting and removed hardcoded cap `desktop-v0.0.12`

06 **Disjoint token telemetry buckets** IMPROVED

60

Token telemetry now reports disjoint per-request buckets — uncached input, cache reads, and cache writes — eliminating ~5x inflation of per-task sums on cache-heavy sessions.

— Explains mechanism and scale but no user action `sdk/sdk/v0.0.74`

THINNER COVERAGE BELOW

07 **Explicit surfacing of involuntary logouts** IMPROVED

50

Involuntary Cline logouts due to rejected refresh tokens are now surfaced explicitly rather than silently clearing credentials; transient network failures on startup no longer register as logouts.

— Describes behavior change, no config surface named `sdk/sdk/v0.0.74`

08 **Reduced telemetry volume from stream deltas** IMPROVED

40

Per-token stream deltas are no longer mirrored into telemetry, eliminating ~97% of agent event volume that had no analytical value.

— States numeric impact but no actionable surface `sdk/sdk/v0.0.74`

---

WAS THIS USEFUL? 👍 👎

# OpenHands

Feedback ⓘ 1 RELEASE · 2026-08-13 NOTES ↗

OpenHands: AI-Driven Development

> OpenHands launched its `openhands-sdk` / `openhands-tools` framework for building custom AI agents with LLM, tool, and conversation classes, alongside a new Agent Canvas UI, a Sandbox Server REST API, local LLM integrations, org-level spend budgets, and a Skills system, plus v1.13.0 UI additions like conversation archiving and a context-window usage meter.

└──▷ **WHAT SHIPPED · 14 FEATURES** most completely described first ? what's the number?

## 01 Built-in tools and Agent/Conversation composition NEW

86

Built-in `TerminalTool`, `FileEditorTool`, and `TaskTrackerTool` are importable from `openhands.tools.terminal`, `openhands.tools.file_editor`, and `openhands.tools.task_tracker`. The `Agent` class composes an agent from an `llm` and a list of `Tool` objects, and the `Conversation` class manages the interaction lifecycle via send_message() and run() within a given `workspace` path.

Compose and run an agent programmatically with terminal and file-editing capabilities in the current working directory.

python

```
import os
from openhands.sdk import LLM, Agent, Conversation, Tool
from openhands.tools.file_editor import FileEditorTool
from openhands.tools.terminal import TerminalTool

llm = LLM(
    model=os.getenv("LLM_MODEL", "anthropic/claude-sonnet-4-5-
20250929"),
    api_key=os.getenv("LLM_API_KEY"),
    base_url=os.getenv("LLM_BASE_URL", None),
)
agent = Agent(llm=llm, tools=[Tool(name=TerminalTool.name),
Tool(name=FileEditorTool.name)])
conversation = Conversation(agent=agent, workspace=os.getcwd())
conversation.send_message("List all Python files and count lines
of code.")
conversation.run()
```

— Names classes, methods, and import paths with a full working example. product docs

02 **LLM class with LiteLLM and OpenHands Cloud provider support**

NEW 81

The `LLM` class accepts `model`, `api_key`, and `base_url` constructor arguments and can be configured via the `LLM_MODEL`, `LLM_API_KEY`, and `LLM_BASE_URL` environment variables to reach any LiteLLM-compatible provider. It also accepts `openhands/` `-prefixed` model names (e.g. `openhands/claude-sonnet-4-5-20250929`) when using OpenHands Cloud as the LLM provider.

Create a minimal agent that writes project facts to a file, using OpenHands Cloud as the LLM backend.

```
$ export LLM_MODEL="openhands/claude-sonnet-4-5-20250929"
export LLM_API_KEY="your-openhands-api-key"
uv run python examples/01_standalone_sdk/01_hello_world.py
```

Compose and run an agent programmatically with terminal and file-editing capabilities in the current working directory.

```python
import os
from openhands.sdk import LLM, Agent, Conversation, Tool
from openhands.tools.file_editor import FileEditorTool
from openhands.tools.terminal import TerminalTool

llm = LLM(
    model=os.getenv("LLM_MODEL", "anthropic/claude-sonnet-4-5-20250929"),
    api_key=os.getenv("LLM_API_KEY"),
    base_url=os.getenv("LLM_BASE_URL", None),
)
agent = Agent(llm=llm, tools=[Tool(name=TerminalTool.name), Tool(name=FileEditorTool.name)])
conversation = Conversation(agent=agent, workspace=os.getcwd())
conversation.send_message("List all Python files and count lines of code.")
conversation.run()
```

— Names class args, env vars, and model prefix with runnable examples. product docs

03 **ChatGPT subscription login for Codex models** NEW 75

LLM.subscription_login(vendor, model) authenticates with a ChatGPT Plus/Pro subscription so agents can access Codex models (e.g. `gpt-5.2-codex`) without consuming API credits.

Authenticate with a ChatGPT Plus/Pro subscription to run an agent on Codex without using API credits.

```python
from openhands.sdk import LLM
llm = LLM.subscription_login(vendor="openai", model="gpt-5.2-codex")
```

— Names method signature, model, and benefit with example. product docs

04 **SDK packages and source installation** NEW 70

OpenHands ships as the `openhands-sdk` and `openhands-tools` PyPI packages, which must be installed together at matching versions (e.g. `pip install "openhands-sdk==1.22.1" "openhands-tools==1.22.1"`); optional `openhands-workspace` and `openhands-agent-server` packages add sandboxed workspaces in Docker or on a remote server. The SDK can also be built from source via `git clone https://github.com/OpenHands/software-agent-sdk.git` and `make build`.

— Names packages and install commands but no runtime behaviour. product docs

05 **Agent Canvas: architecture, remote/Docker backends, ACP CLI support, and LLM profiles**

NEW 60

Agent Canvas connects to execution, automation, and sandbox services with documented architecture, and can run Claude Code, Codex, or Gemini CLI inside it via the Agent Client Protocol (ACP). It can connect to a remote Agent Server backend on another machine or container, run with Docker for a sandboxed backend with a mounted project workspace, and manage LLM profiles to configure and reuse models across conversations.

— Names ACP, CLIs, and backend options but no exact setup steps. product docs

THINNER COVERAGE BELOW

06 **Local LLM server integrations** NEW 45

Adds support for connecting OpenHands to local LLM servers including LM Studio, Ollama, vLLM, and SGLang.

— Names four supported local servers, no config detail. product docs

07 **Context window usage meter and manual compaction** NEW

45

Adds a context window usage meter, usage drawer, and manual compaction control so users can monitor and manage LLM context consumption mid-session.

— Names three UI elements but no thresholds or commands. v1.13.0

08 **Runnable SDK examples library** NEW

35

Ships 24+ runnable examples under `examples/01_standalone_sdk/`, including custom tools (`02_custom_tools.py`) and microagent skills (`03_activate_microagent.py`).

— Names example files but no description of what they teach. product docs

09 **Sandbox Server REST API v1** NEW

35

A new Sandbox Server REST API (V1) exposes REST endpoints for conversations and sandboxes.

— Names the API surface but not exact endpoint paths or methods. product docs

10 **Org-level spend budgets** NEW

35

Adds org-level spend budgets to cap AI spending for organizations and their members.

— States the capability but no configuration mechanism. product docs

11 **Client-side conversation archive** NEW

35

Adds a client-side conversation archive, letting users archive and manage past conversations directly in the UI.

— Names the UI capability with no navigation detail. v1.13.0

12 **Skills system for reusable agent instructions** NEW

30

A new Skills system lets practitioners give OpenHands reusable instructions, domain knowledge, workflows, and supporting resources.

— Describes purpose only, no config or file format named. product docs

13 **Inline markdown artifact previews** NEW

30

Adds inline markdown artifact previews in the chat panel, rendering markdown output without leaving the conversation.

— Describes behaviour but no exact UI location. v1.13.0

14 **Issue readiness gate for dev handoff** NEW

30

Adds a ready-for-dev issue readiness gate with type-specific criteria to validate issues before they are handed off to the agent.

— States purpose but not the specific criteria. `v1.13.0`

---

WAS THIS USEFUL? 👍 👎

# Claude Code

Feedback ⓘ 1 RELEASE · 2026-08-12 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's v2.1.229 release adds plugin marketplace command sources, Remote Control session resume and status labels, a workflow fan-out prompt-prefix staggering control, SSE keepalives for long-running gateway streams, and VSCode session groups — while tightening self-hosted runner startup and commit-push-pr auto-approval defaults as breaking changes.

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

## 01 Remote Control session resume and status labels IMPROVED

72

`claude remote-control --continue` resumes the most recent Remote Control session without specifying a session ID, useful for reconnecting after a network drop. `ListAgents` now marks disconnected Remote Control sessions as `offline` and labels cloud sessions as `cloud`, giving clearer session inventory.

Resume the most recent Remote Control session without looking up its session ID — useful when reconnecting after a network drop.

```
$ claude remote-control --continue
```

— Names exact command and new status labels, no deeper mechanism given v2.1.229

## 02 Workflow fan-out prompt-prefix staggering NEW

72

Introduces workflow fan-out staggering that lets same-prefix sibling agents share a cached prompt prefix; it can be disabled by setting `CLAUDE_CODE_WORKFLOW_PREFIX_STAGGER_MS=0`, e.g. for latency-sensitive pipelines where agents should not wait before starting.

Disable workflow fan-out staggering in a latency-sensitive pipeline where you do not want agents to wait before starting.

```
$ CLAUDE_CODE_WORKFLOW_PREFIX_STAGGER_MS=0 claude
```

— Env var and runnable command given, mechanism briefly explained v2.1.229

## 03 Self-hosted runner Windows requires --base-dir

BREAKING

70

Self-hosted runner Windows startup now requires an explicit `--base-dir`; there is no default checkout directory on Windows, so launches that relied on a default will fail without it.

— Exact flag and migration consequence stated v2.1.229

04 **Removes silent approval of dangerous git/gh flags** BREAKING

70

`/commit-push-pr` no longer auto-approves git/gh commands with dangerous flags such as `--force`, `--amend`, `--no-verify`; workflows relying on silent approval of these flags will now be prompted for confirmation.

— Named flags and command with clear before/after behavior v2.1.229

05 **Plugin marketplace command sources** NEW

62

Adds plugin marketplace `command` sources: a local command (e.g. an IDE) prints the plugin directory, which is re-resolved each session without a restart. `mode: "link"` uses the directory in place rather than copying it.

— Mechanism and config mode named, no runnable example v2.1.229

THINNER COVERAGE BELOW

06 **Session groups in VSCode sidebar** NEW

58

Adds session groups in the VSCode sidebar — right-click to create, rename, or delete a group, and Cmd/Ctrl- or Shift-click to move multiple sessions at once.

— UI path given, no config or command surface v2.1.229

07 **SSE keepalives on gateway streaming** IMPROVED

43

Adds SSE keepalive pings to gateway streaming responses during long thinking pauses, preventing idle-timeout disconnects on Vertex and Bedrock upstreams.

— Explains behavior and named upstreams but no user-facing action v2.1.229

08 **Resizable /btw side-question panel** IMPROVED

43

Makes the `/btw` side-question panel resizable by dragging its boundary, in both side-docked and stacked layouts, in VSCode.

— Simple UI tweak, minimal mechanism beyond drag-to-resize v2.1.229

09 **Server-supplied hooks for self-hosted runners** NEW

30

Adds server-supplied Claude Code hook support for self-hosted runner sessions, matching managed-environment behavior.

— Bare description of parity feature, no mechanism or config detail `v2.1.229`

└──▷ **BREAKING ON UPGRADE**

! Self-hosted runner Windows startup now requires an explicit `--base-dir`; launches that previously relied on a default checkout directory will fail without it.

! `/commit-push-pr` no longer auto-approves git/gh commands that include dangerous flags (`--force`, `--amend`, `--no-verify`, etc.); workflows depending on silent approval of these flags will now prompt for confirmation.

---

WAS THIS USEFUL? 👍 👎

---

## GOVERN

◆ AI/LLM Security

# Daytona

Feedback ⓘ 1 RELEASE · SEEN 2026-08-13 NOTES ↗

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

> Daytona 0.201.0 shipped pre-signed file transfer URLs, typed SDK error codes, and lifecycle-aware sandbox filtering.

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

01 **autoDestroyAt range filter for sandbox listing** NEW 50

Adds `autoDestroyAt` range filter to the sandbox list API, enabling lifecycle-aware sandbox queries.

– Names the field and API but no exact endpoint path. snapshot-20260813

02 **Pre-signed file upload/download URLs** NEW 45

Adds pre-signed sandbox file download and upload URLs for direct, credential-free file transfers to/from sandboxes.

– Describes benefit but no endpoint or command shown. snapshot-20260813

03 **Typed SDK error codes** NEW 45

Adds structured error responses with typed error codes to the SDK, making programmatic error handling against the API practical.

– States purpose but no code names or sample. snapshot-20260813

WAS THIS USEFUL? 👍 👎

# NVIDIA SkillSpector

Feedback ⓘ 1 RELEASE · 2026-08-13

NOTES ↗

Security scanner for AI agent skills. Detect vulnerabilities, malicious patterns, security risks, prompt injection, data exfiltration, and supply-chain risks in Claude Code, Codex, and MCP skills before you install them.

> SkillSpector v2.9.4 hardens its MCP HTTP transport against local-path abuse, adds a new supply-chain finding for shipped Python bytecode, and improves prompt-injection detection for whitespace-padding evasion.

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

## 01 HTTP MCP transport rejects local filesystem targets

BREAKING

85

The HTTP MCP transport ( `skillspector mcp` with `http` transport) now rejects local filesystem scan targets and local YARA-rule directories, enforcing a security boundary between remote callers and scanner-host paths. The `stdio` transport retains full local-scan capability; HTTP clients must switch to a remote repository or URL, or use `stdio` for local scans.

Run SkillSpector's MCP server over stdio (trusted local transport) to retain the ability to scan local filesystem paths after the HTTP restriction.

```
$ skillspector mcp --transport stdio
```

— Names transport modes and flag; includes runnable stdio command. v2.9.4

## 02 SC8 finding for shipped Python bytecode NEW

60

Adds a HIGH-severity `SC8` finding when a skill ships Python bytecode or `__pycache__` content, expanding supply-chain coverage.

— Names finding ID and trigger but no remediation steps. v2.9.4

THINNER COVERAGE BELOW

## 03 Whitespace-padding prompt-injection detection NEW

40

Adds detection for whitespace-padding techniques used to hide prompt-injection instructions in skill content.

— Describes what's detected but not mechanism or finding ID. v2.9.4

04 **allowed-tools recognised as least-privilege guidance** IMPROVED

35

Recognises `allowed-tools` as valid least-privilege permission guidance in remediations and documentation.

— Names the field but gives no usage detail. v2.9.4

05 **Skill Inspector companion skill guide** NEW

15

Ships a new Skill Inspector companion skill guide.

— Bare mention with no detail on content or use. v2.9.4

└──▷ **BREAKING ON UPGRADE**

! HTTP MCP clients can no longer scan local filesystem paths or provide local YARA-rule directories; switch to a remote repository or URL for HTTP requests, or use `stdio` transport for local scans.

WAS THIS USEFUL?  

## PLATFORMS & UTILITIES

◆ General Utilities & Libraries

# Obsidian

Feedback ⓘ 1 RELEASE · 2026-08-13 CHANGELOG ↗

Obsidian is a personal knowledge management app that lets users create, link, and organize markdown notes locally with powerful search and graph visualization.

> Obsidian's 1.13.5–1.13.7 updates overhaul image handling in Live Preview and remove the URI action approval dialog in favor of existing guardrails, alongside smaller polish to tab behavior, plugin API, and settings UI.

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

01 **SettingDefinitionBase#disabled plugin API** NEW 65

Adds `SettingDefinitionBase#disabled` to the plugin API, letting plugin authors programmatically disable individual settings controls.

– Names exact API field but no usage example given changelog-20260813-20918661

02 **Revamped Live Preview image editing** IMPROVED 60

Images in Live Preview can now be selected via keyboard, no longer auto-expand to show filenames, and pressing Backspace / Delete removes the embedded image when selected.

– Names exact keys and behavior changes, no version-specific scope beyond that changelog-20260813-20918661

THINNER COVERAGE BELOW

03 **URI action approval dialog removed** BREAKING 40

Obsidian removes the URI action approval dialog, now relying on guardrails introduced in 1.13 (e.g. Mermaid diagrams off by default) instead of a confirmation prompt.

– Explains what changed but not full guardrail mechanism changelog-20260813-20918661

04 **Shift-Tab focus navigation fix** IMPROVED 35

Shift-Tab now moves focus to the next focusable element instead of staying locked to the current row.

– Names exact key combo and behavior change changelog-20260813-20918661

05 **Undo close tab behavior improved** IMPROVED 33

Undo close tab now reopens tabs in the correct tab group and no longer reuses the active New tab page.

– Describes before/after but no concrete surface changelog-20260813-20918661

06 **Mobile toolbar settings icons** IMPROVED 30

Mobile toolbar settings page now shows icons next to each command.

— Names UI location but minimal further detail changelog-20260813-20918661

07 **Installer updated to Electron 43.3.0** IMPROVED 30

Installer updated to use Electron v43.3.0.

— Names exact version but no functional impact described changelog-20260813-20918661

08 **Granular startup stats window** IMPROVED 20

The startup stats window now shows more granular information about startup speed.

— Bare description with no specifics on metrics shown changelog-20260813-20918661

09 **Accent-colored settings sliders** IMPROVED 18

Sliders in settings now use the accent color.

— Purely cosmetic, minimal detail changelog-20260813-20918661

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

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

▷ Subscribe | Feedback on this issue