THE CYBER TOOLCHAIN

NO. 023

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 23, 2026 · EVERY DAY

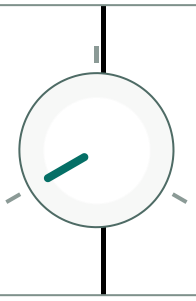
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 54 tools

OFFENSIVE SECURITY

vim ? help

◆ Reconnaissance & OSINT

Hunter

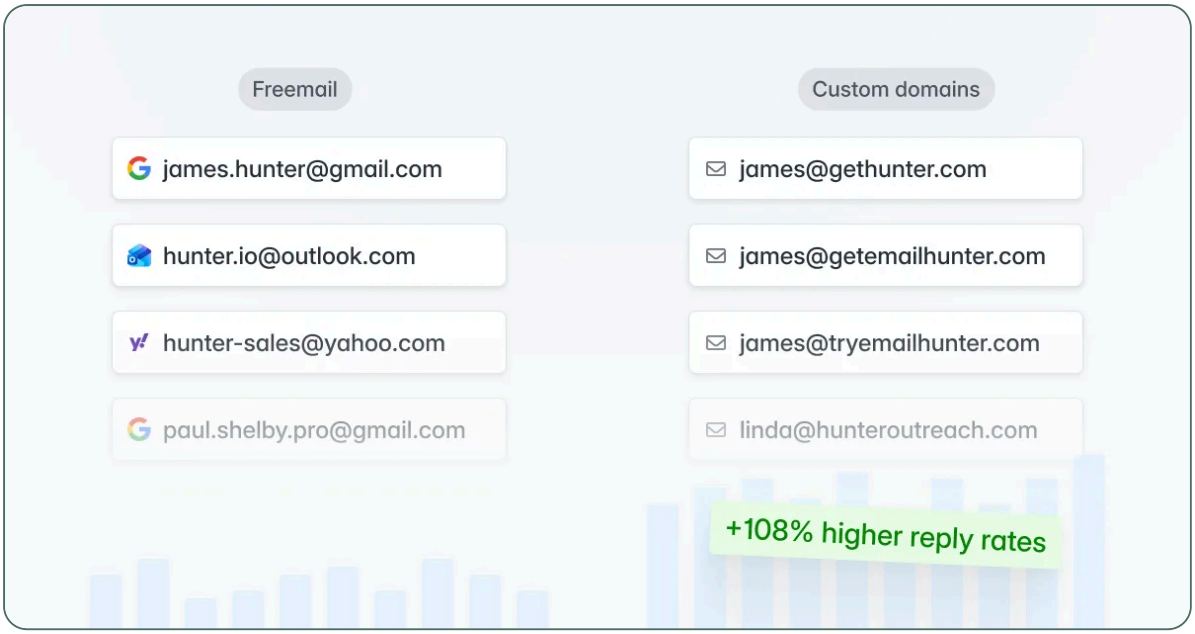
Feedback ⓘ

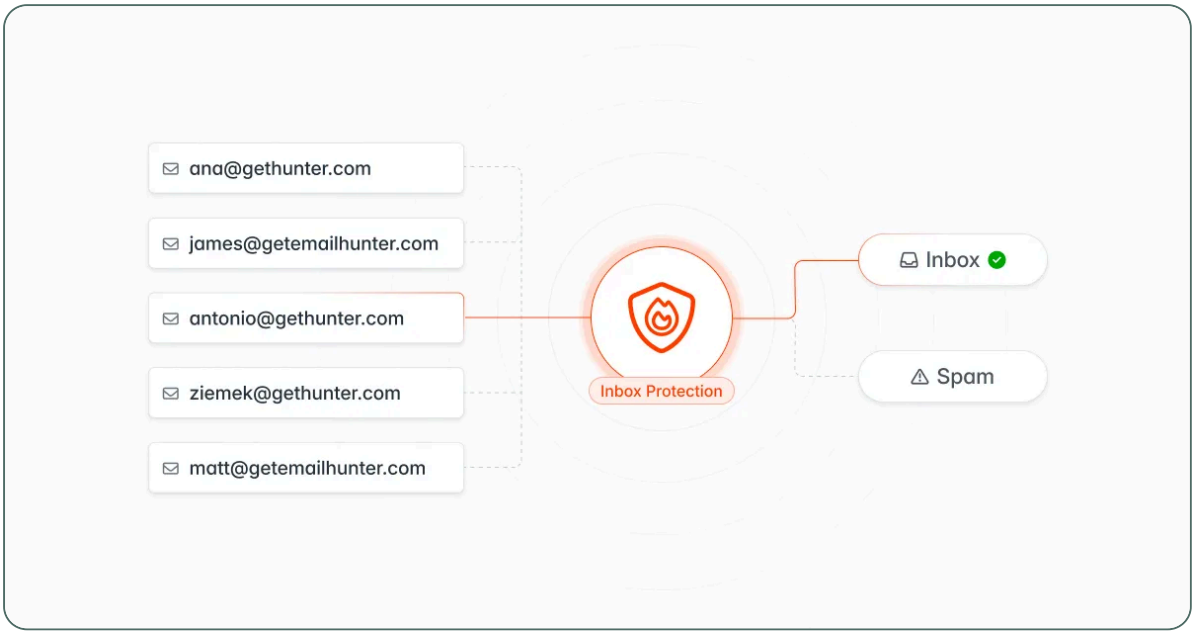
Buy domains, create email accounts, and protect your sender reputation

8 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Hunter is an email finding and verification tool that helps users discover professional email addresses associated with any domain or company.

Hunter adds managed email account provisioning and automated inbox warm-up to support full outreach setup.





└──▷ HOW TO FIND IT

Set up a new outreach domain end-to-end without leaving Hunter — useful when launching a cold outreach campaign from a fresh domain.

📍 In the Hunter console, go to Email accounts center › Buy a domain, choose your domain name, sender name(s), and number of accounts per sender. Accounts are ready within 1 hour with SPF, DKIM, and DMARC pre-configured.

Protect deliverability on a new or inactive account before sending outreach by enabling the Inbox Protection warm-up add-on.

📍 In the Hunter console, go to Email accounts center, select a connected email account, and enable Inbox Protection as an add-on to begin the 30-day automated warm-up.

Adds managed email account provisioning: buy domains and create pre-configured email accounts (with SPF, DKIM, and DMARC) directly within Hunter, ready within ~1 hour.

New Inbox Protection add-on automates a 30-day email warm-up, progressively increasing sending volume via real replies and positive engagement signals.

Inbox Protection dashboard tracks warm-up progress including interactions, open rates, reply rates, and spam placement.

After warm-up completes, Inbox Protection automatically switches to maintenance mode for continuous sender reputation protection.

Provisioned accounts marked Active are immediately available for use in Hunter Sequences.

▾ **7 more releases in this issue** · seen 2026-07-23

Company Logo API (free) NOTES ↗ STABLE

Hunter ships a free, no-auth Company Logo API returning logos for 16M+ companies by domain.

└──▷ **TRY IT**

Fetch a company logo by domain to enrich a CRM record or populate a dashboard — no authentication needed.

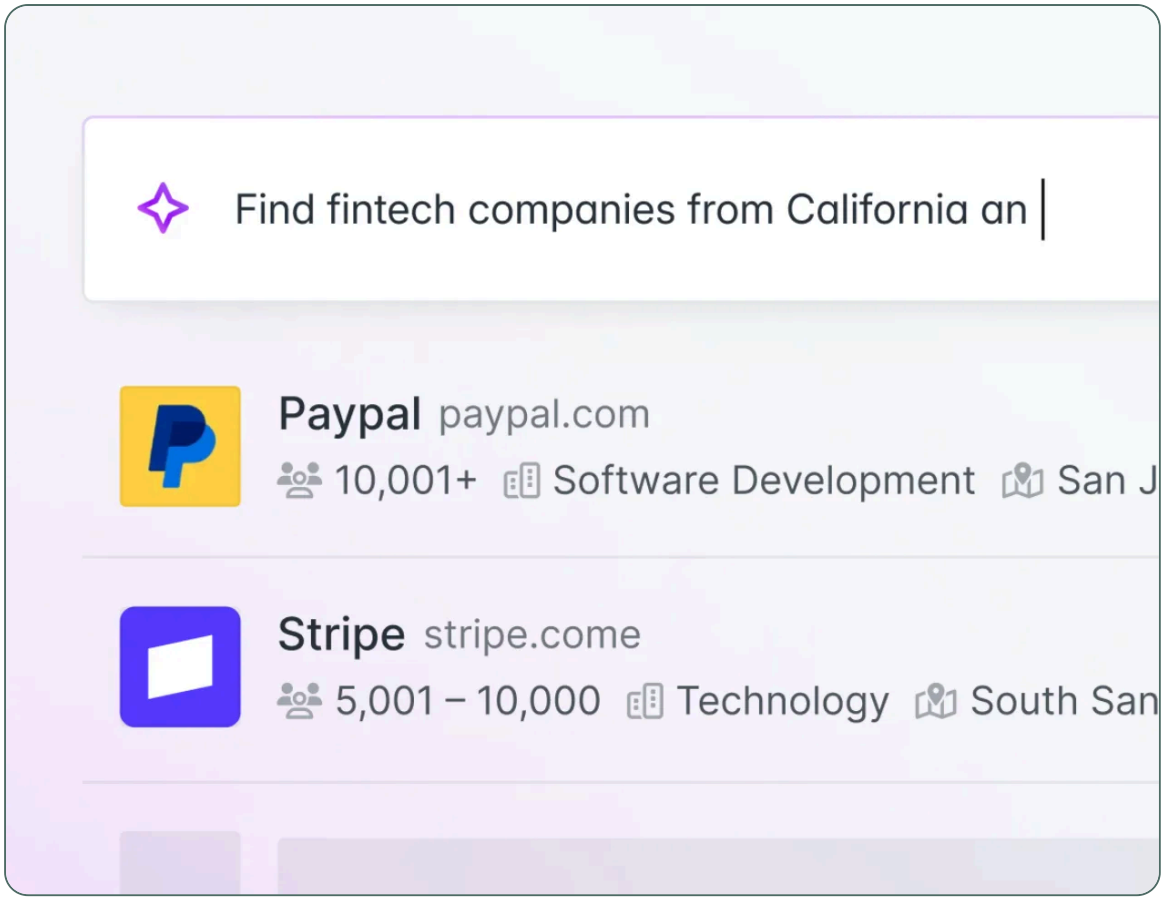
```
$ curl -L https://logos.hunter.io/stripe.com
```

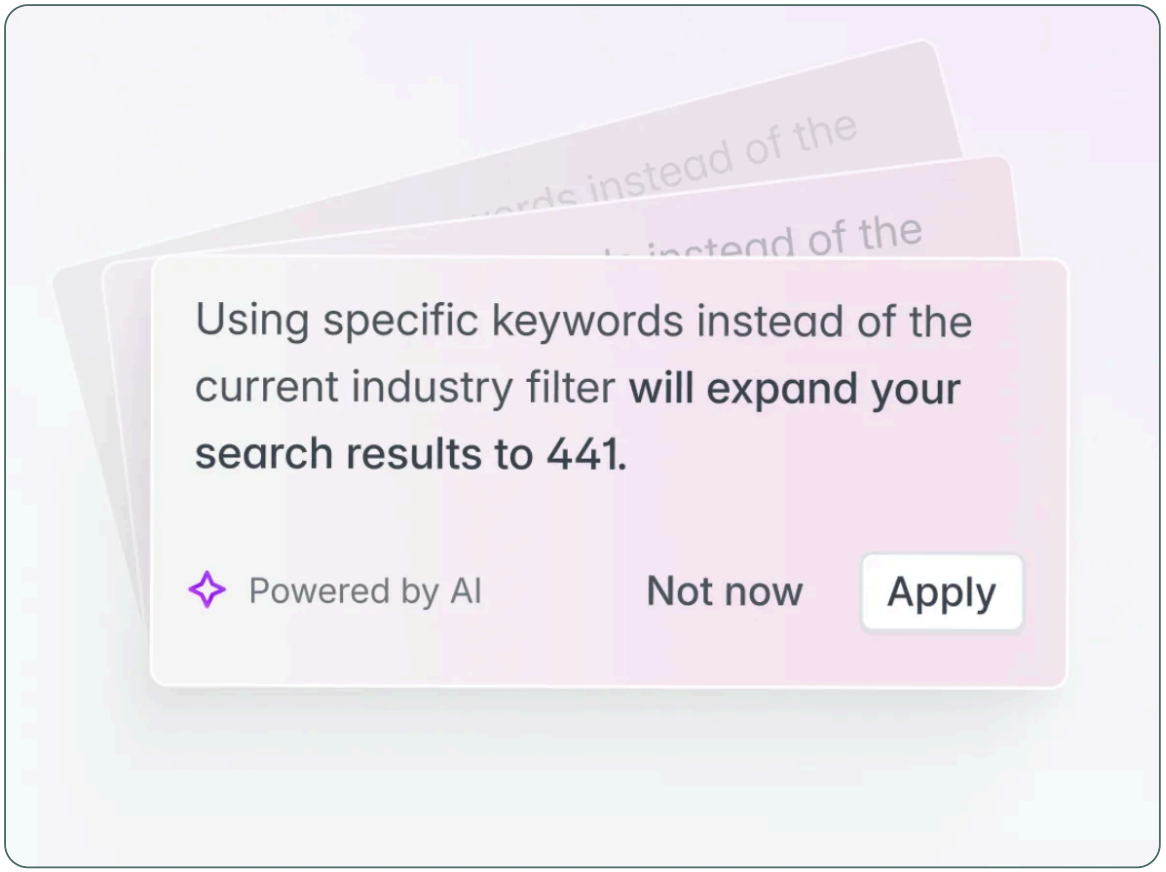
New free endpoint `https://logos.hunter.io/{domain}` retrieves a company logo via plain GET — no signup, no API key, no setup required.

Covers 16M+ company logos, enabling CRM enrichment, dashboard branding, and report generation without adding a separate provider.

Discover: smarter AI Assistant & saved companies filter NOTES ↗ STABLE

Hunter Discover gains a smarter AI Assistant with one-click refinements and a filter to exclude already-saved companies.







└──▷ **HOW TO FIND IT**

Find decision-maker contacts at matching companies without touching a single filter — describe your audience and let the Assistant do the rest.

> 📍 In Discover, open the AI Assistant and type: "Sales Managers at SaaS companies in Germany"

Prospect from your value proposition rather than a target persona — useful when you know your offer but not yet the exact buyer profile.

> 📍 In Discover, open the AI Assistant and type: "I am a Founder of a small engineering company helping with waste management in the United Kingdom"

Keep your Discover results fresh by hiding companies you've already added to your lead lists, so you never scroll past duplicates.

> 📍 In Discover, enable the `Exclude saved companies` filter to hide any company already saved to Leads from your search results.

Adds improved AI Assistant in Discover that translates plain-language audience descriptions into ready-to-use lead searches — no manual filter configuration required.

AI Assistant now jumps from company results to specific contacts when job titles are mentioned in the search query.

After each search, the Assistant analyzes results and surfaces one-click suggestions to narrow, expand, or sharpen targeting with filters like recent funding or job functions.

New filter in Discover to exclude companies already saved to Leads, ensuring only new opportunities appear in search results.

`Email Accounts Center & Sender Health` **NOTES ↗** **STABLE**

Hunter launches Email Accounts Center and Global Health Score for centralized sender account management and deliverability monitoring.

└──▷ **HOW TO FIND IT**

Check your overall sender health and act on Deliverability Guard warnings before launching a new outreach sequence.

> 📍 In Hunter Sequences, go to the Sequences navigation panel › Email Accounts Center. Review the Global Health Score and any Deliverability Guard warnings in the overview bar, then click into failing checks to follow the linked fix guidance.

Bulk-enable Inbox Protection and update sending limits across multiple sender accounts at once to tighten deliverability controls.

> 📍 In Hunter Sequences, go to Email Accounts Center › select multiple accounts using the checkboxes › choose `Enable Inbox Protection` or `Edit sending limits` from the bulk actions menu.

New Email Accounts Center consolidates all connected sender accounts in Sequences into a single view with status, daily sending progress, scheduled emails, health score, and Inbox Protection status.

Adds bulk actions across multiple accounts: run health checks, edit sending limits, update sender names, enable Inbox Protection, or disconnect accounts simultaneously.

New Domains tab provides per-domain health scores and detailed domain reports.

New Global Health Score (out of 100) aggregates sending health across all connected accounts, covering Authentication (SPF, DKIM, DMARC), Reputation (blacklists, domain age, website signals), and Email Account factors.

Health checks now run automatically when a new account is connected, with on-demand re-runs available from the accounts table.

`+1 more`

`Engagement rates are now recipient-based` **NOTES ↗** **STABLE**

Hunter Sequences now report open, click, reply, and unsubscribe rates per unique recipient, not per email sent.

Open, click, reply, and unsubscribe rates in Sequences now count distinct recipients who engaged divided by recipients who received, reflecting true audience reach.

Recipient-based rates surface across all reporting surfaces: Sequences list, sequence stats, CSV export, reporting overview and detail tabs, the API, and MCP tools.

└──▷ **BREAKING ON UPGRADE**

! For API users: per-step counts no longer sum to the recipient-based totals at the sequence level — any downstream logic that aggregated step-level counts to derive sequence-level rates will now produce incorrect results.

Progressive Sending: gradually scale your outbound volume NOTES ↗ STABLE

Hunter adds Progressive Sending to automatically ramp up daily email volume for new or reactivated inboxes without manual adjustments.

.

└──▷ HOW TO FIND IT

Enable Progressive Sending for a newly connected inbox to safely ramp up to 50 emails/day without touching settings again.

📍 In Hunter, go to your email account settings, select a newly connected inbox, enable Progressive Sending, and set your target daily sending limit (e.g. 50 emails/day). Hunter will handle all volume increases automatically.

New Progressive Sending feature automatically and gradually increases daily email volume in Sequences until a user-defined target limit is reached.

Supports newly connected email accounts and accounts inactive for 30+ days, covering the most deliverability-vulnerable ramp-up phase.

Volume progression is conditional: the daily cap only increases when the previous day's limit was hit and bounce rate remains healthy.

Available on all All-in-One Outreach plans, including free tier.

Reveal emails in Discover based on remaining credits NOTES ↗ STABLE

Hunter Discover now lets you partially reveal leads up to your available credit balance instead of blocking the entire dataset.

.

Enables partial lead reveals in Discover's people-finding workflow, revealing up to the exact number of leads your current credit balance allows rather than blocking the full dataset.

Adds a prompt in Discover to either proceed with the partial reveal or obtain more credits (upgrade or buy an extra credits pack) when a dataset exceeds available credits.

Sequences: A/B testing is here NOTES ↗ STABLE

Hunter Sequences gains A/B testing with automatic winner selection and per-variant step metrics.

.
.
.

└──▷ **HOW TO FIND IT**

When launching a new cold outreach sequence and unsure which subject line resonates, enable A/B testing on a step to let Hunter split recipients and surface the top performer automatically.

> 📍 In the console, open Sequences › select or create a sequence › edit a step › enable A/B Testing › define Variant A and Variant B (subject line and/or body) › save. Monitor results under Sequences › Reporting › A/B Results tab.

Adds A/B testing for each step in Sequences, letting you test subject line and email content variants against split recipient groups.

New step-level metrics surface opens, clicks, replies, and bounces per variant for direct comparison.

Enables automatic promotion of the winning variant to all remaining scheduled sends once a winner is selected.

New `A/B Results` tab in Sequences Reporting summarizes variant performance across all steps.

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

Dreadnode

Feedback ⓘ

AIRT Report Builder now generally available with templates, branding, and DOCX export

11 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Dreadnode is an AI-powered security research platform that helps teams build, run, and evaluate offensive AI agents for penetration testing and red teaming.

Dreadnode AIRT Report Builder goes GA with templates, DOCX export, new AI red-teaming skills, and a major CLI overhaul.

└──▷ TRY IT

Uninstall a capability you no longer need from your agent environment.

```
$ dn capability uninstall <capability-name>
```

Update the Dreadnode CLI to the latest version without leaving your terminal.

```
$ dn update
```

Build and download an AIRT report in DOCX format with custom branding for client delivery.

> 📍 In the platform, go to AIRT › Reports, select or create a report template, configure branding (logo, company name, confidentiality footer), then click Export › DOCX.

AIRT Report Builder is now generally available for all users, with saved templates (load/save/rename/delete), custom branding (logo, company name, title, confidentiality footer), DOCX export, and a live structural preview panel.

New trace-analysis advisor skill surfaces attack-effectiveness analysis, transform recommendations, and vulnerability pattern identification from historical run data directly inside the agent.

New GuardSessionPolicy gates every agent tool call through an LLM judge before execution, with a built-in safety rubric and optional custom rules.

GuardSessionPolicy supports five transcript strategies (`rubric_only` , `intent_only` , `intent_plus_calls` , `full` , and more) to control how much context is sent to the judge.

New export endpoint converts session transcripts to ATIF v1.7 or OpenAI Chat Completions format; new `dn session` CLI command group and a session UI download

menu enable browsing, inspecting, and exporting sessions.

+11 more

└──▷ **BREAKING ON UPGRADE**

! The `Export PDF Report` button on the AIRT overview page has been removed; the Reports tab now handles PDF and all other export formats.

! The default GuardSessionPolicy transcript strategy is now `intent_plus_calls` (previously unspecified/different), which changes what context is sent to the LLM judge for existing GuardSessionPolicy configurations.

▾ **10 more releases in this issue** · seen 2026-07-23

Agent Output available in beta for all orgs NOTES ↗ STABLE

Dreadnode adds ATLAS multi-agent attack framework, 170 multimodal red-team transforms, LiteLLM on-prem inference, E2B sandbox, and a full slash-command/@ palette in web chat.

└──▷ **TRY IT**

Cap token spend on an evaluation run for reproducible cost control across CI pipelines.

```
$ dn evaluation create --max-steps 50 --config evaluation.yaml
```

Adds ATLAS multi-agent attack framework (`dreadnode.airt.atlas`) running campaigns across eight attack modes using a MDP+Hedge router, with tool calls captured per-trial and surfaced in findings across the UI, API, and SDK.

Ships 130+ image, audio, and video perturbations for red teaming vision/audio/video models, covering ImageNet-C corruptions, AugLy augmentations, SpecAugment, FigStep, and DolphinAttack techniques.

Exposes 131 multimodal transforms in the AI red-teaming TUI (up from 26), including 57 image, 52 audio, and 22 video transforms.

Adds slash-command palette in the web chat composer — type `/` for a fuzzy-filtered overlay with built-in actions (`/new` , `/clear` , `/rewind` , `/model` , `/agent` , and more) plus dynamically-discovered skills.

Adds `@` `-mention` agent picker in the web chat composer, navigable by keyboard or mouse, inserting agents as inline mentions.

+16 more

└──▷ **BREAKING ON UPGRADE**

! The `Structured Output` nav item and route are renamed to `Agent Output` at `/agents/agent-outputs/`; any bookmarks or integrations referencing the old route or label will break.

`Analytics query and notebook UX overhaul` NOTES ↗ STABLE

Dreadnode adds web chat for project runtimes, persistent web-security agent memory, and a refreshed analytics UX.

└──▷ **HOW TO FIND IT**

When reviewing a web-security agent run, use example query chips and schema column insertion on the analytics page to quickly build queries against persisted memory artifacts without writing SQL from scratch.

📍 In the console, go to Analytics › Query, select a schema column from the column insertion panel, or choose an example query chip to pre-populate the editor with a query against your project's reconnaissance findings or discoveries.

Adds empty states, a stats summary strip, loading skeletons, inline error banners, schema column insertion, and example query chips to the analytics query and notebook pages.

Enables the web-security agent to persist reconnaissance findings, gadgets, leads, defense bypasses, and discoveries across sessions using structured project memory.

Allows revealed images in tool results and logged outputs to be re-hidden, freeing the decoded blob from memory.

`Android APK research capability and issue-tracker connectors for web-security` NOTES ↗ STABLE

Dreadnode adds Android APK static analysis, Jira/GitHub issue-tracker connectors, artifact logging, and admin model controls for web-security agents.

Adds `android-apk-research` capability bundle: a 10-tool MCP server with four skills for static semantic vulnerability research on Android APKs.

Adds Jira Cloud MCP connector for web-security agents to export validated findings directly as remediation tickets.

Adds GitHub MCP connector for web-security agents to file findings as GitHub issues and post follow-up comments from within the agent workflow.

Adds four artifact logging tools (`log_image_output`, `log_audio_output`, `log_video_output`, `log_file_artifact`) for attaching media and file artifacts to a web-security agent run.

Adds `waymore` to the web-security capability runtime, enabling historical URL and response retrieval from Wayback Machine, CommonCrawl, OTX, URLScan, and

VirusTotal.

+4 more

Binary analysis capability replaces windows-reversing with cross-platform PE/ELF/Mach-

NOTES ↗ STABLE

Dreadnode adds cross-platform binary analysis, Android APK research, AEM/Sling web testing, and bulk evaluation log downloads.

└──▷ **USE IT**

Enable cross-platform binary analysis (PE/ELF/Mach-O) in an agent capability config, replacing the old windows-reversing bundle.

yaml

```yaml
capability: binary-analysis
```

Run static vulnerability research against an Android APK by enabling the new android-apk-research capability bundle.

yaml

```yaml
capability: android-apk-research
```

Mark an MCP server header optional so OAuth runs as the default when the env var is unset.

yaml

```yaml
headers:
  - name: Authorization
    value: "Bearer ${MY_API_KEY}"
    optional: true
```

New `binary-analysis` capability supports PE, ELF, and Mach-O formats with an expanded Ghidra MCP server and a consolidated 8-phase analysis skill.

New `android-apk-research` capability ships a 10-tool MCP server and four skills for static semantic vulnerability research on Android APKs.

Automated web-security testing now covers Adobe Experience Manager and Apache Sling, surfacing selector-abuse, dispatcher-bypass, and XSS weaknesses.

Bulk evaluation log download lets you export all sample logs from an evaluation as a single ZIP bundle.

Dashboard workspace toggle lets users filter activity between personal workspace and full organization view.

+4 more

└──▷ **BREAKING ON UPGRADE**

! The `windows-reversing` capability bundle is replaced by `binary-analysis`; any configuration referencing `windows-reversing` must be updated to use `binary-analysis`.

`Claude Code engine — run Claude's agentic loop as a governed, session-backed agent`

NOTES ↗ STABLE

Dreadnode adds Claude Code engine, Task Sets, per-role environment model selection, eval sample search, and two new offensive security skills.

└──▷ **TRY IT**

Override the defender (environment) model per role independently of the solver and judge, e.g. to pit a weaker defender against a stronger attacker.

```
$ dn eval run --env-model defender=openai/gpt-4o --env-model
judge=anthropic/claude-3-5-sonnet <eval-target>
```

Group related tasks into a versioned Task Set and run them all as one evaluation target from the CLI.

```
$ dn task-set run <org>/<task-set-name>
```

Adds pluggable agent execution engines with Claude Code as the first built-in option, running Claude's agentic loop as a governed, session-backed agent with full trace viewer support.

Introduces Task Sets — named, versioned collections of tasks runnable as a single evaluation target via the API, `dn task-set` CLI commands, or the Environments UI.

Enables user-controlled task-environment models: choose which model each environment role uses via `--env-model role=id` (CLI), `model_overrides` in `evaluation.yaml`, or the API, independently of solver and judge models.

Adds free-text search across all samples in an evaluation, with match highlighting, per-sample match counts, and a Matches only filter.

Adds a blind SQL injection extraction skill supporting boolean/timing-based oracle identification, WAF bypass patterns, and automated char-by-char extraction.

+5 more

└──▷ **BREAKING ON UPGRADE**

! The `--session <name>` flag is now required for concurrent web-security agents; omitting it will cause agents to clobber each other's browser state.

`Compound eval verification — script_and_judge and flag_and_judge` NOTES ↗

STABLE

Dreadnode adds compound eval verification, agentic outcome judging, 6 new web pentesting skills, and expanded AIRT goal categories.

└──▷ **TRY IT**

Run an agentic outcome judge against an agent trajectory to get a pass/fail verdict without writing a custom scoring script.

```
$ dn outcome-judge run
```

New `script_and_judge` and `flag_and_judge` verification methods chain a mechanical check (script or flag) with an LLM judge to close reward-hacking gaps in eval tasks — no custom rubric required.

New `outcome_judge` verification method lets eval tasks use an agentic LLM to inspect agent trajectories and emit pass/fail verdicts, exposed via the new `dn outcome-judge run` CLI subcommand.

New opt-in `intent_plus_outputs_summary` guard strategy replaces raw tool outputs with LLM-generated summaries in judge prompts, reducing prompt-injection surface while preserving tool-call context.

Six new web app pentesting skills added: ESI injection, gRPC-Web pentest, H2C/WebSocket smuggling, HTTP connection contamination, timing-attack recon, and XSLT injection.

Adds `waymore` to the web-security capability runtime for historical URL and response retrieval from Wayback Machine, CommonCrawl, OTX, URLScan, and VirusTotal.

+4 more

└──▷ **BREAKING ON UPGRADE**

! The AIRT workspace path structure has changed to `~/.dreadnode/airt/[org]/[workspace]/`, replacing the previous env-var-based path system — existing setups relying on the old env-var paths will need to be updated.

`End-to-end multimodal AI red teaming — probe vision, audio, and video models with full`

NOTES ↗ STABLE

Dreadnode adds end-to-end multimodal AI red teaming, custom HTTP targets, Task Sets, and a SecurityContext MCP server.

└──▷ **USE IT**

Point a PAIR attack at a private Azure OpenAI endpoint using Azure AD auth via the new custom_http target mode.

```python
build_target(TargetSpec(mode="custom_http", url="https://<your-resource>.openai.azure.com/openai/deployments/<deployment>/chat/completions", auth="azure_ad"))
```

View the README for a capability in the CLI before running it, to understand inputs, outputs, and auth requirements.

```
$ dn capability info <capability-name> --readme
```

New multimodal AI red teaming: probe LLMs with text, image, audio, and video inputs; apply per-modality transforms; score with the new `multimodal_judge` scorer; view full message parts inline in findings, traces, and parquet exports (schema v3 with base64 media columns).

Cross-product `prompt_matrix` mode runs N prompts × M media items = N×M trials for systematic multimodal attack coverage.

Nova Sonic speech-to-speech target support added to the TUI capability picker and `build_target` SDK factory.

New `custom_http` target mode and build_target(TargetSpec) factory let TAP, PAIR, and other attacks hit any endpoint (Azure, Vertex, SageMaker, Bedrock, OpenAI-compatible) with declarative auth strategies (API key, Bearer, AWS SigV4, Azure AD, GCP).

Per-modality response scores in findings and Assessment Findings tables, with the highest-scoring modality emphasized.

+15 more

└──▷ **BREAKING ON UPGRADE**

! The `items` key in capability manifests is replaced by a unified `produces` key that configures built-in item types, custom types, and disablement — existing manifests using `items` will need to be updated.

Free-text search in evaluation trajectories NOTES ↗ STABLE

Dreadnode adds free-text search in evaluation trajectories, session rewind, cross-org share links, and sandbox dependency auto-install.

└──▷ **HOW TO FIND IT**

Quickly locate a specific tool call or output string inside a long evaluation run without manually scrolling through every trajectory step.

> 📍 In the console, open an evaluation trajectory, then type your search term into the free-text search box — matching content is highlighted inline and tool calls containing the term expand automatically.

Roll back a TUI conversation to an earlier point to re-run from a known-good message without losing the full history.

```
$ /rewind
```

Adds free-text search in hosted evaluation trajectories — matches are highlighted inline and tool calls auto-expand when their content matches the search term.

Adds session rewind via `/rewind` or double-Esc in the TUI, rolling back a conversation to any prior user message while preserving full session history.

Enables cross-org capability share links so Capability Hub links automatically redirect recipients to the correct page regardless of their org slug.

Sandbox runtimes now auto-install Python packages and system dependencies declared in capability manifests at load time.

AIRT CLI now auto-discovers attacks, transforms, and scorers at runtime (61 attacks, 547 transforms, 141 scorers) with accurate capability counts in the hub.

+3 more

Session and Trace Views Now Display Agent-Captured Images, Audio, and Video

NOTES ↗ STABLE

Dreadnode adds media playback in session views, ASM capability, HTTP/2 WAF bypass, GraphQL pentesting, and a sandbox runtime hierarchy.

└──▷ **TRY IT**

Run a smoke test that mirrors production runtime and validates judge-backed tasks before submitting — catches multi-file solution issues early.

```
$ dn task validate --smoke
```

Review agent-captured media artifacts from a session without auto-decoding potentially hostile content.

> 📍 In the platform UI, open a session transcript or trace details view, then click `Show media` on any image, audio, or video item to decode and display it.

Adds inline display of agent-captured images, audio, and video in session transcripts and trace details, with an explicit `Show media` click to prevent auto-decoding of hostile artifacts.

New attack surface management capability combining BBOT scanning, Shodan enrichment, Neo4j graph analysis, screenshot triage, and a multi-agent ASM pipeline.

New h2-waf-bypass skill covering 6 HTTP/2 WAF bypass classes: delayed DATA frame timing, body size truncation, Extended CONNECT, ForwardAuth body stripping, path normalization, and JSON content-type gaps.

New graphql-pentest skill covering endpoint discovery, introspection mining, resource abuse (CWE-400, CWE-674), content-type CSRF, and capability matrix testing.

New standalone vuln-assessment-methodology capability providing a reusable severity matrix, disprove-first rules, and reporting standards any security capability can load.

+8 more

Tool media evidence in session transcript NOTES ↗ STABLE

Dreadnode adds ProjectMemory, inline media in transcripts, new DoS/path-traversal/EXIF skills, and a unified LiteLLM admin proxy.

└──▷ **TRY IT**

Inject secrets or config into a challenge service during smoke testing without modifying docker-compose.yaml.

```
$ dn task validate --smoke -e API_KEY=abc123 -e
TARGET_HOST=http://localhost:8080
```

Adds ProjectMemory tools so agents can persist and retrieve durable project-scoped memory across runs.

Screenshots and media logged via dn.log_output() inside tools now render inline in the session transcript next to the producing tool call.

New unified LiteLLM admin proxy lets admins manage model deployments and credentials directly from the platform; platform models are now always org-allowlist-

scoped.

Adds Ctrl+K shortcut in the TUI to open the full model browser.

New app-layer-dos skill covers ReDoS, decompression bombs, server delay exploitation, and cross-protocol amplification for application-layer DoS testing.

+11 more

WAS THIS USEFUL? 👍 👎

Strix Feedback ⓘ v1.3.0 NOTES ↗

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

Strix v1.3.0 adds a Feedback & support tab and a prominent scan switcher with pentest terminology to the local viewer UI.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.3.0 https://github.com/usestrix/strix.git
# already have the repo? check out this version:
$ git checkout v1.3.0
```

Adds a Feedback & support tab to the local viewer UI.

Adds a prominent scan switcher to the local viewer with renamed `pentest` terminology.

Shrinks the sandbox Docker image from 7.2 GB to 3.8 GB via cache cleanup, multi-stage Go build, and dropping ZAP — enabling use on disk-constrained environments.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian SaaS

Feedback ⓘ changelog · Jul 23, 2026

CHANGELOG ↗

GitGuardian SaaS is a secrets detection platform that scans code repositories and developer workflows to identify and remediate exposed credentials and sensitive data.

└──▷ HOW TO FIND IT

Loop in a teammate or team directly on an incident without leaving GitGuardian — they receive an email linking straight to the incident.

📍 In any incident, open the Notes or Feedback panel, type @ and select a member (@name) or a team (@team-name) from the searchable list, then submit.

Filter the incident list to set aside noise from archived or deleted sources, then save it as a persistent view.

📍 On the Perimeter page or Incident list, open the source status filter and select `Archived` or `Deleted`, then save as a view.

Reduce false positives on a Custom Source install by restricting scans to detectors that can actually verify secret validity.

📍 In the Custom Source install settings, enable the 'use only detectors with validators' option for the target source.

Adds @mention support in incident notes and feedback — type @name or @team-name to notify individual members or entire teams via email, linked directly to the incident.

WAS THIS USEFUL? 👍 👎

GitGuardian Detection Engine

Feedback

changelog · Jul 23, 2026

CHANGELOG ↗

GitGuardian Detection Engine scans code repositories and developer environments to automatically detect secrets, credentials, and sensitive data exposures in real time.

Updated Hume AI API key checker endpoint.

WAS THIS USEFUL?

◆ Web Application Security

Cloudflare

Feedback ⓘ changelog · Jul 23, 2026 CHANGELOG ↗

Cloudflare is a cloud-based platform that provides website performance optimization, DDoS protection, CDN services, and application security to safeguard and accelerate internet traffic.

Cloudflare adds Agents SDK MCP controls, Code Mode direct APIs, Devin Outposts sandboxes, and new WAF rules for Next.js/ColdFusion/WordPress CVEs.

└──▷ USE IT

Suppress Think's automatic MCP tool injection when you're managing MCP tools yourself through Code Mode, to avoid duplicate tool registration.

javascript

```
import { Think } from "@cloudflare/think";

const result = await Think({
  includeMcpTools: false,
  waitForMcpConnections: true,
  // ... other Think options
});
```

Allow an MCP server or custom host to discover and invoke a connector method on the Code Mode durable runtime without using the AI SDK.

javascript

```
const matches = await runtime.search("create issue");
const docs = await runtime.describe(matches.results[0].path);
const outcome = await runtime.execute({
  code: `async () => github.create_issue({ title: "Bug" })`
});
```

Adds `includeMcpTools` option to `@cloudflare/think` to opt out of Think's automatic MCP tool exposure when tools are managed externally (e.g., via Code Mode).

Agents SDK MCP clients now reuse converted input/output schemas across model turns within a live connection, eliminating redundant JSON Schema-to-Zod conversions.

Adds execute(), search(), and describe() methods to the `@cloudflare/codemode` durable runtime handle, enabling MCP servers and non-AI-SDK hosts to invoke Code Mode directly.

Introduces Devin Outposts: run Devin agent sessions on Cloudflare-managed, container-isolated infrastructure via the Sandbox SDK.

Deploys new WAF rules covering Next.js DoS (CVE-2026-64641), Next.js SSRF in Server Actions (CVE-2026-64649), Next.js information disclosure (CVE-2026-64643), Adobe ColdFusion path traversal (CVE-2026-48276, CVE-2026-48282), and WordPress unauthenticated SQL injection (CVE-2026-60137).

+2 more

└─▷ **BREAKING ON UPGRADE**

! The Account Roles API is deprecated and replaced by the Permission Groups API; integrations must obtain Permission Group IDs from the new API and use them in Account Members API policies — legacy Role IDs must be remapped.

! The `PermissionGroup` response replaces the legacy `Role` top-level `description` and `permissions` object (keyed by resource type with edit/read flags) with a `meta` object containing `label` and `scopes`; individual permissions are no longer returned as part of the permission group.

WAS THIS USEFUL? 👍 👎

Levo

Feedback ⓘ Release Notes â Q4 2025

6 RELEASES · SEEN 2026-01-01 → 2026-07-23 NOTES ↗ LATEST STABLE

Levo is an API security testing tool that automatically discovers, observes, and tests APIs to detect vulnerabilities and security misconfigurations in development and production.

Levo Q4 2025 adds web app scanning, AI security capabilities, a rebuilt notification system, and expanded API discovery.

New refreshed dashboard experience for improved operational visibility.

Completely rebuilt notification system for security alerts and events.

First wave of web application scanning capabilities.

New AI security capabilities for detecting AI-related risks.

Expanded API discovery to surface more of the attack surface.

+2 more

▾ **5 more releases in this issue** · seen 2026-01-01 → 2026-07-23

> Release Notes â May 2026 NOTES ↗ STABLE
>
> Levo May 2026: redesigned Policy Hub, AI traffic tagger with agent attribution, and Streamable HTTP for MCP server.
>
> Redesigned Policy Hub with identity-aware policies for finer-grained security controls.
>
> Next-generation AI traffic tagger attributes traffic to individual agents and sessions.
>
> New Agent Detail view surfaces per-agent activity and session context.
>
> Adds Streamable HTTP transport support for the MCP server.
>
> Deeper API discovery from source code expands coverage of detected endpoints.

> Release Notes â April 2026 NOTES ↗ STABLE
>
> Levo adds AI agent/model/MCP discovery, AI Policies with guardrail alerting, and AI-driven DAST authentication.
>
> Discovers AI agents, models, and MCP tools from cloud providers in the API catalog.
>
> Introduces AI Policies with guardrail alerting for AI security enforcement.
>
> Adds AI-driven authentication support for web application (DAST) scanning.
>
> Expands reporting capabilities.
>
> Ships new sensor deployment options.

Release Notes â March 2026 NOTES ↗ STABLE

Levo adds scheduled DAST scans, AI-driven auth, live AI Gateway guardrails, and a new application comparison report.

Adds scheduled and on-demand DAST scan modes to the web application scanning engine.

Introduces AI/LLM-driven authentication for automated crawling of login-protected applications.

Enables smarter crawling in the DAST engine for broader application coverage.

Unifies AI Gateway policies with live guardrail enforcement.

Delivers a new application comparison report via a faster async reporting pipeline.

+3 more

Release Notes â February 2026 NOTES ↗ STABLE

Levo February 2026: CLI-driven DAST, AI Guardrails policies, agentless API discovery, and tagging across findings

Adds CLI-driven web application (DAST) scans for triggering scans from pipelines or the command line.

Includes evidence capture in DAST findings for proof-of-exploitation documentation.

Supports DOCX report export for DAST scan results.

Introduces a policy framework for AI Guardrails to define and enforce guardrail rules.

Enables multi-tenant and parallelized scanning in AI Guardrails for larger-scale coverage.

+3 more

Release Notes â January 2026 NOTES ↗ STABLE

Levo ships Vigil AI firewall, DAST scanning, and external attack surface management in January 2026

Introduces Vigil, an AI-native firewall for AI apps and MCP servers.

Launches web application scanning (DAST) capabilities.

Adds external attack surface management (EASM) capabilities.

Ships AI guardrails with end-to-end coverage.

Delivers AI discovery working end to end.

WAS THIS USEFUL?  

c/side

Feedback ⓘ Cmd+K command palette 8 RELEASES · SEEN 2026-07-23

NOTES ↗ LATEST STABLE

c/side is a browser-side security tool that detects and monitors malicious third-party scripts on websites to protect users from client-side attacks.

c/side adds a Cmd+K command palette for keyboard-first navigation across domains, features, and settings.

└──▷ HOW TO FIND IT

Jump directly to any domain or setting without mouse navigation — useful when triaging alerts across many monitored domains.

> 📍 Press Cmd+K from anywhere in the dashboard, then type the domain name or setting to jump to it instantly.

New Cmd+K command palette enables instant keyboard-driven search and navigation to any domain, feature, or setting in the dashboard.

▾ 7 more releases in this issue · seen 2026-07-23

Fingerprint Dashboard is now generally available

NOTES ↗ STABLE

c/side's Fingerprint Dashboard is now GA — monitor device fingerprints and filter AI agent traffic from a single view.

└──▷ HOW TO FIND IT

Separate bot and AI agent traffic from legitimate users to investigate suspicious script activity.

> 📍 In the console, go to the Fingerprint Dashboard and apply the AI agent signals filter to view only automated visitors.

Adds generally available Fingerprint Dashboard for monitoring device fingerprints across all traffic.

Enables filtering by AI agent signals to distinguish automated visitors from real users.

Provides unified single-view management of fingerprint data, billed within existing c/side subscriptions.

JA4 TLD fingerprinting

NOTES ↗ STABLE

c/side adds JA4 TLS fingerprinting to detect VPN clients masquerading as browsers

Adds per-connection JA3/JA4 TLS fingerprinting using the full ClientHello (cipher suites, extensions, ALPN, elliptic curves) to detect VPN client software impersonating web browsers.

New Notifications Engine supporting Linear and Jira NOTES ↗ STABLE

c/side adds Linear and Jira integrations to automatically create issues from security alerts.

.

└──▷ HOW TO FIND IT

Route c/side security alerts directly into your Linear or Jira project to keep remediation in your existing issue tracker.

📍 In the console, go to Settings › Notifications, connect your Linear or Jira account, and configure which alert types create issues.

New Notifications Engine lets you configure which alerts route to which destinations, giving granular control over alert delivery.

Integrates with Linear to automatically create issues from c/side alerts (Business and Enterprise plans).

Integrates with Jira to automatically create issues from c/side alerts (Business and Enterprise plans).

Onboarding that matches your use case NOTES ↗ STABLE

c/side adds use-case-guided onboarding for PCI compliance, fraud prevention, AI agent detection, and device fingerprinting.

Script Ad Injection Detection NOTES ↗ STABLE

c/side adds ad injection detection, flagging client-side scripts that trigger 5+ redirects or suspicious iframe behavior.

New detections identify ad injection via client-side scripts, including iframe data population and URL resolution.

Alerts customers when scripts exhibit suspicious redirect behavior (5 or more redirects).

Self-service SSO for enterprise customers NOTES ↗ STABLE

c/side enterprise adds self-service SSO setup for Okta, Microsoft Entra ID, and Duo Security.

.

└──▷ HOW TO FIND IT

Onboard your organization's identity provider without waiting on support — configure SSO directly in the dashboard.

> 📍 In the cside dashboard, go to Settings › SSO and select your identity provider (Okta, Microsoft Entra ID, or Duo Security) to complete setup.

Supports self-service SSO configuration for Okta, Microsoft Entra ID, and Duo Security directly in the cside dashboard — no support ticket required.

We now detect when bots try to hide NOTES ↗ STABLE

c/side now automatically detects bots that spoof browser automation properties to appear as legitimate users.

Detects bots that spoof browser automation properties to evade fingerprinting, surfacing the signal in fingerprint data.

WAS THIS USEFUL? 👍 👎

PortSwigger Burp Suite Feedback ⓘ

Professional / Community 2026.7.1 (2026-07-23) BLOG ↗

Burp Suite is a web security testing platform that intercepts, analyzes, and manipulates HTTP/S traffic to identify vulnerabilities in web applications.

Burp Suite 2026.7.1 introduces Burp AT, an agentic AI-driven testing capability for Burp Suite Professional.

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

DeepSource

Feedback ⓘ AI Review on more programming languages

NOTES ↗

DeepSource is an automated code review platform that analyzes code for bugs, security vulnerabilities, anti-patterns, and style issues across multiple languages.

DeepSource adds MCP Server, BYOK inference, Kubernetes Gateway API, monorepo SCA, and expanded AI Review language support.

.
.
.
.

└──▷ HOW TO FIND IT

Enable Kubernetes Gateway API on Enterprise Server so DeepSource creates HTTPRoutes instead of Ingress resources, keeping traffic within your existing Gateway.

> 📍 In the Admin Console, go to Config, set 'Use Gateway API?' to Yes, then provide your Gateway name and namespace.

Turn on AI Review for a newly supported language that has no DeepSource static analyzer.

> 📍 In the DeepSource dashboard, go to Settings → Code Review → Analyzers, pick the target language, and click Save changes.

Adds DeepSource MCP Server exposing 30 tools across 8 categories so AI coding agents can access code review results, vulnerability data, and repository metrics via the Model Context Protocol.

Supports Bring Your Own Key (BYOK) model provider credentials for Enterprise Server AI Review, routing inference directly to Azure OpenAI or GCP Vertex AI without passing through DeepSource Cloud.

Expands AI Review to ten additional programming languages that lack a DeepSource static analyzer.

Introduces two AI Review tiers (Standard and Advanced), with Team plan users defaulted to Standard.

Enables AI Review to run on-demand only when @deepsourcebot is mentioned in PR comments.

+5 more

WAS THIS USEFUL? 👍 👎

Codacy Feedback ⓘ

Adding GolangCI-Lint as new supported tool January, 2026 NOTES ↗

Codacy is an automated code quality and security analysis platform that reviews pull requests, enforces coding standards, and tracks technical debt across repositories.

Codacy adds GolangCI-Lint as a supported static analysis tool for Go repositories.

WAS THIS USEFUL? 👍 👎

Cloudsmith Feedback ⓘ

Cooldown policies now support Maven (2026-07-23) BLOG ↗

Cloudsmith is a cloud-native, universal package management platform that enables teams to securely host, manage, and distribute software artifacts across multiple formats.

Cloudsmith cooldown policies now extend to Maven packages for supply-chain protection.

Supports cooldown policies for Maven packages, blocking newly published versions from being consumed until a configured delay period has elapsed.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ DAST updates (2026-07-22)

2 RELEASES · 2026-07-22 BLOG ↗ LATEST STABLE

Veracode is a cloud-based application security platform that helps developers identify and remediate vulnerabilities in their code through static, dynamic, and software composition analysis.

Veracode DAST adds custom host-to-IP mappings, directory restrictions, enhanced scheduling, and open-port scanning with new detections.

Adds custom host-to-IP address mapping configuration for DAST scans when DNS lookup is unavailable or not desired.

Adds Directory Restrictions in DAST target configuration to control scan scope.

Enhances DAST Schedule feature with advanced scheduling capabilities and greater configuration flexibility; all existing schedules are automatically migrated.

Adds open-port scanning support for web applications in Dynamic Analysis and DAST Enterprise mode.

Adds new detections for HTTP headers, TLS, and fingerprinting to the DAST scan engine.

└──▷ **BREAKING ON UPGRADE**

! The DAST REST API scheduling endpoints have changed; integrations that use the REST API to manage schedules must be updated to use the new endpoints.

▾ **1 more release in this issue** · 2026-07-22

Static Analysis updates (2026-07-22) BLOG ↗ STABLE

Veracode SAST adds PySpark/Flask/Django/FastAPI Python support, Pipeline Scan exploit levels, and Android 17 detection in July 2026 updates.

└──▷ **USE IT**

Exclude a bundled third-party .jar from static analysis to eliminate noise from non-application code.

json

```
{
  "exclude": [
    "lib/third-party-sdk.jar"
  ]
}
```

Exclude a specific Java class (e.g., generated code) from Static Analysis results via veracode.json.

```json
{
  "exclude": [
    "com/example/generated/AutoGeneratedClient.class"
  ]
}
```

Adds enhanced Python scanner with improved taint tracking, cross-file analysis, and pattern recognition — applies automatically to all Python scans from July 22, 2026.

Adds Python framework support for PySpark, Flask, Django, and FastAPI in the enhanced scanner.

Adds Python support for httplib to detect CWE-295 and CWE-297 SSL/TLS flaws in v2026.07.

Adds support for Dart 3.12 and Flutter 3.44 in v2026.06.

Adds support for Kotlin 2.4 in v2026.06.

+23 more

WAS THIS USEFUL? 👍 👎

npq

Feedback ⓘ v3.23.0 4 RELEASES · 2026-07-07 → 2026-07-17 NOTES ↗

CODE ↗ LATEST STABLE

safely install npm packages by auditing them pre-install stage

npq now suggests a proven-older semver when a package version fails the 30-day recency check.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.23.0 https://github.com/lirantal/npq.git
# already have the repo? check out this version:
$ git checkout v3.23.0
```

Adds an alternative-version suggestion to version-maturity failures, surfacing the highest older semver published outside npq's 30-day recency window so practitioners have an actionable candidate to audit next.

Renders the suggestion in rich, plain, and JSON output without affecting issue counts, and suppresses it automatically for packages flagged as malicious.

▾ 3 more releases in this issue · 2026-07-07 → 2026-07-17

v3.22.0 NOTES ↗ CODE ↗ STABLE

npq now corroborates expired-domain warnings with authoritative RDAP lookups to reduce false positives.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.22.0 https://github.com/lirantal/npq.git
# already have the repo? check out this version:
$ git checkout v3.22.0
```

Adds RDAP corroboration for expired-domain warnings: a warning is only emitted when DNS returns NXDOMAIN *and* the authoritative RDAP service confirms no active registration, reducing false positives from DNS-only failures.

Routes RDAP lookups through the official IANA DNS bootstrap registry with HTTPS enforcement, 3-second timeouts, no redirects, and process-lifetime caching.

Reports inconclusive RDAP evidence (unsupported TLDs, timeouts, rate limits, multipart suffixes, malformed names) as `not evaluated` rather than blocking installs.

v3.21.0 NOTES ↗ CODE ↗ STABLE

npq v3.21.0 adds authenticated custom registry support, routing all audits through your configured npm registry.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.21.0 https://github.com/lirantal/npq.git
# already have the repo? check out this version:
$ git checkout v3.21.0
```

└──▷ TRY IT

Audit and install a private scoped package from an Artifactory registry, passing registry auth inline — useful in CI where you can't rely on a pre-configured .npmrc.

```
$ npq install @company/tool --
registry=https://artifactory.example.com/artifactory/api/npm/npm-
virtual/
```

Supports authenticated custom registries: reads standard npm registry configuration (auth, proxies, TLS, client certs) from `.npmrc` and routes all package metadata, manifest, signing-key, and attestation lookups through a shared authenticated registry client.

Adds `--registry`, `--userconfig`, and `--globalconfig` CLI flags to explicitly set the registry and npm config paths per invocation.

Reports signing, provenance, and download-count capabilities as `NOT EVALUATED` (rather than failures) when the selected registry does not expose those services.

v3.20.0 NOTES ↗ CODE ↗ STABLE

npq v3.20.0 refreshes its typosquatting detection dataset for more accurate supply-chain threat coverage.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.20.0 https://github.com/lirantal/npq.git
# already have the repo? check out this version:
$ git checkout v3.20.0
```

Updates the typosquatting package dataset to replace stale entries, improving detection of malicious lookalike packages.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

Feedback ⓘ changelog · Jul 23, 2026 CHANGELOG ↗

ElcomSoft develops forensic and password recovery tools that help law enforcement and security professionals access encrypted data, devices, and accounts.

Elcomsoft Phone Breaker 11.03 adds universal binary for macOS and restores SMS-based iCloud two-factor authentication.

Ships as a universal binary for macOS, running natively on both Apple Silicon and Intel Macs.

└──▷ **BREAKING ON UPGRADE**

! The pre-login 2FA status check is removed, as Apple no longer supports it.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Fastly

Feedback ⓘ Fastly Terraform Provider 9.4.0 (2026-07-22)

2 RELEASES · 2026-07-21 → 2026-07-22 BLOG ↗ LATEST STABLE

Fastly is a cloud edge platform that accelerates, secures, and delivers websites, apps, and APIs by caching content closer to users globally.

Fastly Terraform Provider 9.4.0 exposes computed `certificate_id` on TLS subscriptions and enriches API error messages with subscription context.

└──▷ USE IT

Chain a resource that depends on a freshly issued certificate in a single apply, using the new computed `certificate_id` output from the validation resource.

hcl

```
resource "fastly_tls_subscription_validation" "example" {
  subscription_id = fastly_tls_subscription.example.id
}

resource "fastly_tls_activation" "example" {
  certificate_id  =
fastly_tls_subscription_validation.example.certificate_id
  domain          = "example.com"
  configuration_id = fastly_tls_subscription.example.configuration_id
}
```

Exposes computed `certificate_id` on `fastly_tls_subscription_validation`, enabling issuance-dependent resources to be chained in a single `terraform apply`.

Includes subscription ID and domains in `fastly_tls_subscription` API error messages, making it easier to identify failing resources when managing many subscriptions at once.

▾ 1 more release in this issue · 2026-07-21 → 2026-07-22

> Add HTTP and TLS metrics to historical and real-time Domain Inspector APIs (2026-07-21
>
> BLOG ↗ STABLE
>
> Domain Inspector APIs now expose TLS and HTTP version distribution metrics for incoming requests.

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Sysdig Secure

Feedback ⓘ changelog · Jul 23, 2026 CHANGELOG ↗

Sysdig Secure is a cloud-native security platform that detects threats, manages vulnerabilities, ensures compliance, and responds to incidents across containers, Kubernetes, and cloud environments.

Sysdig Secure adds multi-bucket Custom Storage support and a redesigned Posture Overview & Findings experience.

└─▷ HOW TO FIND IT

Configure a MinIO bucket for use with Captures and Rapid Response via the new Custom Storage settings.

> 📍 In Sysdig Secure, go to Data Sources › Storage, click Add Bucket, select S3-compatible endpoint, enter your MinIO endpoint URL and credentials, then click Test Connection and Save.

Identify the weakest posture areas and assign remediation to platform teams using the new Findings workspace.

> 📍 In Sysdig Secure, navigate to Posture › Findings, apply filters for high-severity controls, group results by resource or category, and use the Assign Remediation action to route issues to the appropriate team.

Adds Custom Storage: connect multiple S3 or S3-compatible buckets (MinIO, Google Cloud Storage, etc.) for use across Captures, Response Actions, and Rapid Response — replacing the previous single shared S3 Capture Storage configuration.

Supports AWS access keys, AWS role delegation (AssumeRole), and S3-compatible endpoints for bucket authentication, with connection testing on save and per-bucket failure history.

Introduces a new Posture Overview page surfacing passing scores, 30-day trends, and hotspots across weakest categories, resources, and high-severity controls.

Introduces a new Posture Findings workspace with filtering, grouping, and remediation assignment to drive faster issue resolution.

└─▷ BREAKING ON UPGRADE

! The previous single shared S3 Capture Storage configuration is replaced by the new Custom Storage system; existing Capture Storage settings will need to be reconfigured under Data Sources > Storage.

WAS THIS USEFUL? 👍 👎

Plexicus

Feedback ⓘ

Plexicus 2026.7.17: Multifactor Priority Scoring v2, Registry Scanning Activated & Retic

NOTES ↗

Plexicus is an AI-powered application security platform that automatically detects, prioritizes, and remediates vulnerabilities across code, APIs, and cloud environments.

Plexicus adds Multifactor Priority Scoring v2, activates Registry Scanning, and overhauling Reticulum accuracy.

Introduces Multifactor Priority Scoring v2 for vulnerability prioritization.

Activates Registry Scanning as a new scanning capability.

Delivers a Reticulum accuracy overhaul improving detection precision.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Rapid7 InsightCloudSec

Feedback ⓘ

changelog · Jul 23, 2026 DOCS ↗ BLOG ↗

Rapid7 InsightCloudSec is a cloud-native security platform that continuously monitors, detects, and automatically remediates security risks and compliance violations across multi-cloud environments.

InsightCloudSec 26.7.21 adds Runtime Sensor health monitoring, unified compliance rule mapping, new API endpoints, and a CIS Azure Storage pack.

└──▷ TRY IT

Retrieve all Insight-to-rule mappings for a specific compliance pack to audit current coverage programmatically.

```
$ curl -X GET 'https://<insightcloudsec-
host>/api/pack/<pack_id>/insight-rule-mappings' \
  -H 'Authorization: Bearer <api_token>' \
  -H 'Content-Type: application/json'
```

Map one or more Insights to a specific compliance rule within a custom pack in a single API call.

```
$ curl -X POST 'https://<insightcloudsec-host>/api/map-to-compliance-
rule' \
  -H 'Authorization: Bearer <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"pack_id": "<pack_id>", "rule_id": "<rule_id>", "insight_ids":
["<insight_id_1>", "<insight_id_2>"]}'
```

Create a custom compliance rule on a pack to enforce an internal policy control.

```
$ curl -X POST 'https://<insightcloudsec-
host>/api/pack/<pack_id>/rules/create' \
  -H 'Authorization: Bearer <api_token>' \
  -H 'Content-Type: application/json' \
  -d '{"name": "Internal Policy 2.1"}'
```

Adds real-time Runtime Sensor health status (Healthy, Disconnected, or Degraded) on the Kubernetes Clusters page and Resources Listing for customers with a Runtime

license.

Introduces a unified `Map to Compliance Pack` row and bulk action in the Insights Library, replacing the separate `Edit Metadata` and `Add to Custom Pack` actions.

Enables cascading dropdown selection of one or more compliance rules when mapping an Insight to a compliance pack.

Supports user-defined compliance rules (e.g., Custom Control 1.3) on custom packs, with a new Compliance Rules tab for adding, renaming, and deleting rules.

Supports rule inheritance from base packs; inherited rules appear read-only on the child pack's Compliance Rules tab.

+5 more

└──▷ **BREAKING ON UPGRADE**

- ! The `Edit Metadata` and `Add to Custom Pack` actions in the Insights Library have been replaced by the single `Map to Compliance Pack` action; any workflows or documentation referencing those actions must be updated.
- ! The query filter `Snapshot Accessible To Public` is deprecated (renamed to `Snapshot Accessible To Public (Deprecated)`); its functionality is now split across `Snapshot Public Access` and `Database Snapshot Public Access`.
- ! The Insight `Snapshot Available to the Public` is deprecated (targeted removal in version 27.1); it should be replaced with `Snapshot Is Accessible to the Public` and `Database Snapshot Is Accessible to the Public`.

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

NVIDIA OpenShell

Feedback ⓘ v0.0.90 NOTES ↗

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

OpenShell v0.0.90 adds config key support to provider create/update forms in the TUI.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v0.0.90 https://github.com/NVIDIA/OpenShell.git
# already have the repo? check out this version:
$ git checkout v0.0.90
```

WAS THIS USEFUL? 👍 👎

NoLabs nono

Feedback ⓘ v0.69.0 NOTES ↗ CODE ↗

Sandbox any AI agent in seconds - zero setup, zero latency.

nono v0.69.0 adds SPIFFE/SPIRE workload identity auth for upstream proxy routes and profile-declared no_proxy bypass support.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.69.0 https://github.com/nolabs-ai/nono.git
# already have the repo? check out this version:
$ git checkout v0.69.0
```

└──▷ **USE IT**

Bypass the nono proxy for internal hosts (e.g. a local Redis or internal registry) that should be reached directly.

```json
{
  "no_proxy": ["redis", "registry.internal"]
}
```

Adds SPIFFE/SPIRE workload identity authentication for upstream proxy routes, supporting both JWT-SVID bearer injection and OAuth2 client assertion (RFC 7523) flows; all flows are fail-closed if the SPIRE socket is unreachable.

Adds profile-declared `no_proxy` bypass support, letting profiles specify hosts that skip the nono proxy.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

dockerscan

Feedback ⓘ v2.1.0 NOTES ↗ CODE ↗

The Most Comprehensive Docker Security Scanner

DockerScan v2.1.0 overhauls CVE detection with MITRE cvelistV5 (~125K CVEs), dpkg-accurate version matching, and optional NVD CVSS enrichment.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v2.1.0 https://github.com/cr0hn/dockerscan.git
# already have the repo? check out this version:
$ git checkout v2.1.0
```

└──▷ TRY IT

Fetch the new ~125K-CVE database immediately after upgrading so scans reflect the MITRE cvelistV5 data.

```
$ dockerscan update-db
```

Switches CVE database source to MITRE cvelistV5 daily snapshots, eliminating NVD API keys and rate limits while growing coverage from ~1,500 to ~125,000 CVEs.

Rewrites package-to-CVE matching with a dpkg-style version comparator (Debian Policy 5.6.12) and per-(vendor, product) version ranges for accurate results.

Adds `vendor_verified` metadata to findings, distinguishing vendor-confirmed matches from heuristic name-only matches (which now include explicit remediation caveats).

Adds optional, non-blocking CVSS enrichment from the NVD API via `--enrich-from-nvd`, `--enrich-only`, and `--enrich-timeout` flags — NVD downtime never fails the build.

Expands vendor-verified package alias coverage with 11 new entries (gnu/coreutils, gnu/tar, gnu/sed, gnu/grep, gnu/gzip, gnu/findutils, busybox, musl-libc/musl, and more).

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Detection Engineering

Tenzir

Feedback ⓘ changelog · Jul 23, 2026 CHANGELOG ↗

Tenzir is an open-source security data pipeline platform that helps teams collect, route, shape, and analyze large-scale security telemetry efficiently.

Adds a Splunk search input operator, enabling Tenzir nodes to pull data directly from Splunk searches.

WAS THIS USEFUL? 👍 👎

◆ SIEM

Google SecOps

Feedback ⓘ changelog · Jul 23, 2026 CHANGELOG ↗

Google SecOps is a cloud-based security operations platform that helps organizations detect, investigate, and respond to cybersecurity threats using AI-powered analytics and threat intelligence.

Google SecOps adds Advanced Filtering with dynamic YARA-L query variables in dashboards, plus Security Tokens for agentic consumption metering.

└──▷ HOW TO FIND IT

Create a dashboard filter with a multi-value token to dynamically scope a YARA-L detection query by hostname at runtime, rather than hardcoding values.

📍 In the Google SecOps console, open a dashboard and select Advanced Filtering. Define a Token Variable (e.g., `target_hostname`) with alphanumeric/underscore characters, set Filter Value Generation to YARA-L query results, enable Multi-Select, and set the delimiter to `|`. Reference the token in your YARA-L query to inject the selected values at runtime.

Adds Advanced Filtering in dashboards (Public Preview): security analysts can define Token Variables to inject dynamic values, complex regular expressions, or boolean logic directly into YARA-L queries at runtime.

Token values can be generated dynamically from YARA-L query results or entered as a static manual list.

Customizable prefix/suffix wrappers on token values enable advanced patterns such as inline regular expressions within dashboard queries.

Multi-select support for token variables allows combining multiple values in a query using a configurable delimiter (e.g., `|`).

Introduces Security Tokens for metering agentic consumption within Google SecOps, consumed by generally available security agents.

└──▷ BREAKING ON UPGRADE

! From October 26, 2026, new Google SecOps instances will no longer accept legacy Backstory API or Ingestion API calls; only Chronicle API endpoints will be supported.

! From July 20, 2027, all requests to legacy Backstory API and Ingestion API endpoints will fail for all existing instances as the legacy APIs are completely turned down.

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Hudson Rock

Feedback ⓘ { New } - Company Search Endpoint

5 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Hudson Rock is a cybercrime intelligence platform that provides developers and security teams with access to data stolen by infostealer malware to detect compromised credentials.

Hudson Rock adds a new Company Search endpoint to its API.

▾ 4 more releases in this issue · seen 2026-07-23

{ New } - Search by PC Name Endpoint NOTES ↗ STABLE

Hudson Rock adds a Search by PC Name endpoint to query compromised credentials tied to a specific machine.

└──▷ TRY IT

Query Hudson Rock's API to check whether a specific machine name appears in compromised-credentials data.

```
$ curl "https://cavalier.hudsonrock.com/api/json/v2/osint-tools/search-by-computer-name?computer_name=DESKTOP-ABC123" -H "api-key: <your-api-key>"
```

New API endpoint enables searching the compromised-credentials database by PC name to identify infected machines.

{ New } - Search by Term Endpoint NOTES ↗ STABLE

Hudson Rock adds a new Search by Term endpoint to its API.

{ New } - Subdomain Filtering Feature NOTES ↗ STABLE

Hudson Rock adds subdomain filtering to narrow infostealer data queries by subdomain.

Adds subdomain filtering capability to scope intelligence queries to specific subdomains.

{ New} - Directory Tree Field NOTES ↗ STABLE

Hudson Rock adds a Directory Tree field to expose file system structure from infostealer compromises.

WAS THIS USEFUL?  

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ ADE customization for Mac: Managed profile delivery

8 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Manage devices and everything they touch. See how security, IT, and compliance get easier when unified around devices with Iru.

Iru now delivers managed configuration profiles (PPPC, background items) during Setup Assistant for ADE Mac enrollments.

▾ 7 more releases in this issue · seen 2026-07-23

> Continued improvements to Apple beta release management NOTES ↗ STABLE
>
> Iru extends Apple beta program enrollment and version enforcement to iOS and iPadOS, with new device-record visibility.
>
> Adds Apple beta program enrollment and version enforcement for iOS and iPadOS in Managed OS and ADE Library Items.
>
> Surfaces beta program enrollment status on the device record details tab for any enrolled Apple device.
>
> New optional Prism column displays Apple beta program enrollment status across device listings.

> Public IP address now reported for Apple devices NOTES ↗ STABLE
>
> Iru now reports public IP addresses for Apple devices via device records, Prism, and the enterprise API.

> Request AppleCare enhanced logging from iOS and iPadOS devices testing Apple betas
>
> NOTES ↗ STABLE
>
> Iru now supports AppleCare enhanced log collection requests from iOS and iPadOS 27 beta devices.



Restrictions improvements for Apple beta releases NOTES ↗ STABLE

Iru adds new macOS 27 beta restrictions and expands visionOS 27+ support for four existing restrictions.

Adds `ForceCaptivePortalConnectionFromLockScreen` restriction for devices running the macOS 27 Apple beta.

Adds `ForceWifiConfigurationOnLockScreen` restriction for devices running the macOS 27 Apple beta.

Extends `allowAutomaticAppDownloads`, `allowCamera`, `allowListedAppBundleIDs`, and `blockedAppBundleIDs` restrictions to visionOS 27 and later.

Vulnerability Management Enhancements NOTES ↗ STABLE

Iru Vulnerability Management adds Bulk Device Exclusion to remove multiple devices from scope in a single action.

└──▷ **HOW TO FIND IT**

Temporarily take a batch of devices out of scope (e.g. machines pending maintenance) so they stop inflating vulnerability counts and triggering alerts.

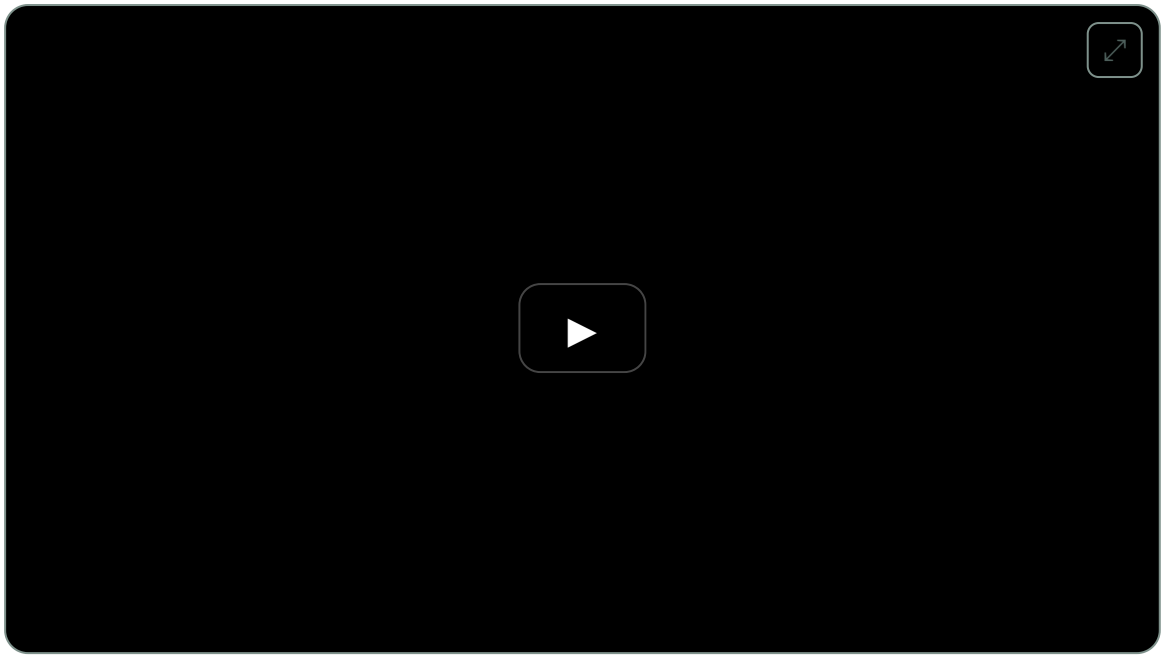
> 📍 In the console, go to Vulnerability Management › Devices, select the target devices, and choose Bulk Device Exclusion to exclude them in a single action.

Adds Bulk Device Exclusion to exclude multiple devices from Vulnerability Management at once, hiding them from all vulnerability views, detection counts, and alerts.

Supports bulk removal of exclusions with the same single-action workflow.

What's New in Iru | July 10th NOTES ↗ STABLE

Iru adds Windows Managed OS, three new Compliance Automation frameworks, and proactive ServiceNow updates.



Adds Managed OS support for Windows, expanding automated OS lifecycle management beyond Apple devices.

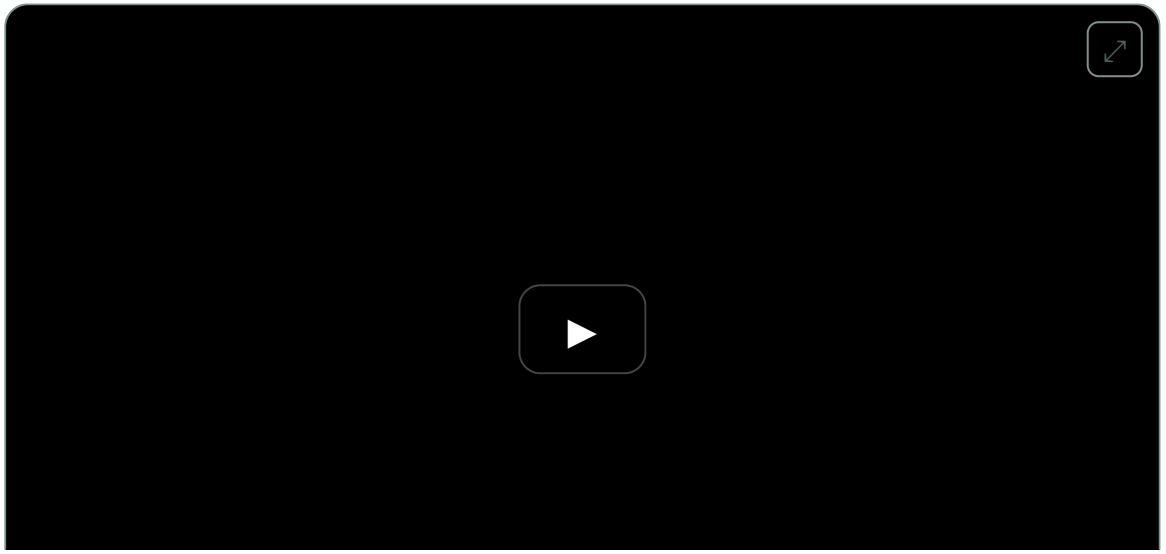
Introduces updates to Managed OS for Apple Library Items.

Adds three new Compliance Automation frameworks for automated compliance coverage.

Enables proactive ServiceNow updates, pushing device or compliance data to ServiceNow automatically.

What's New in Iru | July 17th NOTES ↗ STABLE

Iru adds MSP multi-client management, a Policy Management module, and EDR wildcard support.





New purpose-built features for MSPs managing multiple clients.

New Policy Management module within the Compliance Automation suite.

Adds wildcard support for Endpoint Detection and Response.

WAS THIS USEFUL? 👍 👎

XFA

Feedback ⓘ Claude Code 'Auto' mode flagged as unsafe AI configuration

15 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

XFA is a developer security tool that scans and monitors open-source dependencies for vulnerabilities, helping teams manage software supply chain risks efficiently.

XFA now detects Claude Code's `auto` and `bypassPermissions` modes as unsafe AI configurations alongside the legacy flag.

Extends Unsafe AI Mode detection to flag Claude Code's `auto` and `bypassPermissions` values in `permissions.defaultMode` within `~/.claude/settings.json` (or `~/.claude/settings.local.json`).

Devices configured with either flagged mode now appear in the Unsafe AI Mode check with remediation guidance surfaced to end users in the in-app guide.

Treats `bypassPermissions` as equivalent to `--dangerously-skip-permissions` and `auto` as the same risk class — both auto-approving all tool calls including shell commands and network requests.

▾ **14 more releases in this issue** · seen 2026-07-23

Clearer policy settings and a redesigned Discovery

NOTES ↗ STABLE

XFA adds live Enforcement policy timing previews and a redesigned Discovery connection flow with provider picker.

└──▷ **HOW TO FIND IT**

When configuring an Enforcement policy, use the live timing summary and suggestion chips to confirm warn/block windows before saving — eliminating guesswork about when devices will be affected.

> 📍 In the Dashboard, go to Policies › Enforcement, create or edit a policy, and observe the threshold timeline and timing summary as you adjust warn and block windows. Select a suggestion chip (e.g. a common warn window) to apply it, then verify the exact warn and block times shown before saving.

When onboarding Discovery for the first time, use the provider picker to connect an identity provider and start seeing all devices in your organization.

> 📍 In the Dashboard, go to Discovery, click Add connection, select Microsoft, Google, or Okta from the provider picker, and follow the step-by-step setup page for the chosen provider.

Adds a threshold timeline and live timing summary to Enforcement policy editing, so

you can see exactly when a device would be warned or blocked before saving.

Adds suggestion chips for common warn and block windows when configuring Enforcement policies.

Redesigns the Discovery overview with a two-column layout and a provider picker that routes directly to a step-by-step setup page for Microsoft, Google, or Okta.

Devices table on the Dashboard now shows whether a device's Last Seen came from the XFA agent or from Discovery.

Enables disconnecting a Microsoft Teams integration from the Integrations page, matching existing Slack disconnect capability.

+1 more

Compliance Status Bar NOTES ↗ STABLE

XFA dashboard gains a compliance status bar for at-a-glance policy-based device security health.

└──▷ HOW TO FIND IT

Review overall policy health and spot devices needing attention without drilling into individual records.

📍 In the XFA console, open the Dashboard and locate the Compliance Status Bar to see compliant, warned, blocked, and uncovered device counts at a glance.

New compliance status bar on the dashboard surfaces compliant, warned, blocked, and uncovered devices at a glance.

Enables admins to identify devices with missing information that require verification before compliance status is finalized.

Custom onboarding and sign-in messages NOTES ↗ STABLE

XFA now supports custom messages on onboarding and sign-in cards, configurable per application from the Dashboard.

└──▷ HOW TO FIND IT

Reinforce your acceptable-use policy or link to an internal runbook exactly where end users will see it during device onboarding or sign-in.

📍 In the Dashboard, go to Applications › [your application] and set the onboarding message and sign-in message fields at the bottom of the card configuration.

Adds custom text messages to the bottom of XFA's device onboarding and sign-in cards, with separate messages for each flow.

Enables configuration of onboarding and sign-in messages from the Applications page in the Dashboard.

Device Snooze NOTES ↗ STABLE

XFA adds Device Snooze to pause checks, notifications, and optionally MFA for individual devices with a time-limited exception.

└──▷ HOW TO FIND IT

Grant a temporary exception to a device under repair so it stops triggering checks and alerts until a set date.

📍 In the console, go to Devices › (select device) › Action menu › Snooze device, choose a preset or custom end date/time, optionally add a reason, then click Confirm.

Review all currently snoozed devices that are hidden from the default Devices overview.

📍 In the console, go to Devices overview › enable the Show snoozed devices filter.

Pause MFA enforcement for a device during the snooze window, in addition to pausing device checks.

📍 In the console, go to Devices › (select device) › Action menu › Snooze device, enable Include MFA in snooze, set the end date/time, then click Confirm.

Adds Device Snooze to grant individual devices a temporary exception, pausing device checks and awareness notifications for a chosen preset or custom date/time window.

Snoozed devices are hidden from overviews and analytics by default, with a Show snoozed devices filter to surface them on demand.

Supports optional reason logging per snooze, and allows updating or removing a snooze from the same device action menu.

Enables explicit MFA inclusion in the snooze to also pause MFA for Enforcement sign-ins during the exception window.

Easier Linux installation NOTES ↗ STABLE

XFA Linux installation now requires a single curl command instead of manual package selection.

└──▷ TRY IT

Install XFA on a Linux host without navigating package dropdowns — paste and run.

```
$ curl -fsSL https://distribution.xfa.tech/xfa-native-desktop-application/install.sh | sh
```

Adds a one-line bootstrap installer for Linux, eliminating the need to manually select and download a package.

`Guided setup, device posture at a glance, and rebuilt Statistics trends`

NOTES ↗ STABLE

XFA adds guided setup, device posture overview cards, rebuilt Statistics trends, and unattended remote access detection.

:
:
:

└──▷ **HOW TO FIND IT**

Review your organization's posture across all three pillars after onboarding to identify which areas still need configuration.

> 📍 In the XFA dashboard, go to Devices. The Discovery, Awareness, and Enforcement cards at the top of the page show current status and counts for each pillar.

Check whether any enrolled devices allow unattended remote access and enforce a remediation policy against it.

> 📍 In the XFA dashboard, open a device's detail view and review the `Unattended remote access` check result, then go to your device policy to add or adjust the corresponding rule.

New Getting Started page walks through four tracked setup steps — from device discovery to compliance enforcement — so admins always know what to configure next.

Devices page now opens with Discovery, Awareness, and Enforcement status cards showing real-time counts (devices discovered, alerts sent, successful sign-ins) at a glance.

Statistics Trends rebuilt out of BETA: four metric tiles as tabs, unified chart with period/OS filter/chart-type controls inline, and custom date range support.

Device detail view now reports whether a device allows unattended remote access, with an explanation of what the check looks for and a link to act on it in policy.

Disconnected Discovery integrations now surface a red alert and a Reconnect button on the dashboard, and count toward the Discovery badge.

Lost connection status for devices NOTES ↗ STABLE

XFA adds a Lost connection device status with filtering and verification requests to surface endpoints that stopped reporting.

└──▷ HOW TO FIND IT

Identify and remediate endpoints that have gone silent — filter the Devices view and trigger re-verification without leaving the dashboard.

📍 In the console, go to Devices › filter by status Lost connection, select affected devices, then click Request verification from the Dashboard to prompt users to reconnect the XFA app.

New Lost connection status highlights devices where the XFA agent has stopped reporting a verified status in the Devices overview.

Filter by Lost connection status in both the Devices and People views to quickly identify affected endpoints.

Request verification directly from the Dashboard to prompt users to reconnect the XFA app and resume device security status sharing.

Microsoft Teams (awareness) connection NOTES ↗ STABLE

XFA adds Microsoft Teams integration to deliver device risk alerts, enrollment invites, and sign-in verification directly to end users.

└──▷ HOW TO FIND IT

Connect your Microsoft Teams tenant to start delivering device risk and enrollment notifications directly to users.

📍 In the XFA Dashboard, go to Integrations, select Microsoft Teams, and connect your tenant. Then toggle each notification category (device risk alerts, new-device enrollment invitations, sign-in verification requests) on or off and use Send test message to verify delivery.

Connects XFA to a Microsoft Teams tenant so end users receive device risk alerts, new-device enrollment invitations, and sign-in verification requests as direct messages — no email required.

Enables per-category notification toggles so admins can independently enable or disable each message type.

Supports sending a test message per notification section to verify delivery before rolling out to users.

Delivers messages localized in English, Dutch, French, German, and Spanish.

OPAL/SED hardware disk encryption detection on Linux NOTES ↗ STABLE

XFA Linux disk-encryption checks now recognize TCG OPAL 1.x/2.x self-encrypting drives alongside LUKS.

Detects TCG OPAL 1.x and 2.x self-encrypting drives (via sedutil-cli) as encrypted on Linux when drive locking is enabled, enabling accurate compliance reporting for hardware-encrypted devices.

Policy Compliance Summary NOTES ↗ STABLE

XFA policy detail pages now show a live compliance summary sidebar alongside policy settings.

Screen lock timeout enforcement NOTES ↗ STABLE

XFA adds screen lock timeout enforcement so admins can cap idle-lock delays on macOS, Windows, and Linux.

└──▷ HOW TO FIND IT

Tighten idle-device security by capping how long employees can set their screen lock timeout — useful for regulated environments requiring short inactivity windows.

📍 In the XFA console, go to your organization's policy settings and set the maximum screen lock timeout (in minutes) under the screen lock enforcement section.

Adds maximum screen lock timeout enforcement: admins can set a maximum allowed timeout duration (in minutes) in their organization's policy, and XFA verifies the device's configured timeout does not exceed it.

Supports platform-specific timeout detection on macOS, Windows, and Linux.

Combines timeout check with existing screen lock enable/disable check — both must pass for a device to be compliant.

Slack (awareness) connection NOTES ↗ STABLE

XFA adds Slack integration to deliver device risk alerts, enrollment invites, and sign-in verification requests directly to end users.

└──▷ HOW TO FIND IT

Enable the Slack integration so end users get device risk alerts and enrollment invites as DMs instead of email.

📍 In the Dashboard, go to Integrations › Slack, connect your workspace, then toggle on the desired notification categories (device risk alerts, new-device enrollment, sign-in verification, safe-browsing policy) and set delivery frequency per section.

Connects XFA to a Slack workspace so end users receive device risk alerts as direct messages without email.

Sends new-device enrollment invitations via Slack DM to streamline onboarding.

Delivers sign-in verification requests and safe-browsing policy notifications through Slack.

Admins can toggle each notification category on/off and configure per-section delivery frequency from the Integrations page.

Slack messages are localized in English, Dutch, French, German, and Spanish.

`Stronger mobile checks at sign-in, and a smoother dashboard` NOTES ↗ STABLE

XFA now warns or blocks risky mobile devices (jailbroken, rooted, dev-mode, rebooted) at sign-in, matching existing desktop policy enforcement.

└──▷ **HOW TO FIND IT**

If your policy already blocks or warns on jailbroken/rooted desktop devices, navigate to your device policy to confirm the same posture now applies to mobile at sign-in.

> 📍 In the dashboard, go to your Device Policy settings and review the conditions for jailbroken, rooted, developer-mode, and recently-rebooted devices to ensure your desired warn/block action is set for mobile.

Extends jailbreak, root, developer-mode, and recent-reboot signals from mobile devices into the access decision at sign-in, so risky mobile devices are warned or blocked per policy — matching existing desktop behaviour.

Adds a `Watch demo` action in the Dashboard for Discovery and Enforcement sections, opening a guided tour without leaving the dashboard.

`VeraCrypt system encryption detection (Windows)` NOTES ↗ STABLE

XFA now detects VeraCrypt system-drive encryption on Windows, reporting those devices as compliant automatically.

Supports VeraCrypt system (boot-drive) encryption on Windows as a compliant disk-encryption state, alongside BitLocker.

WAS THIS USEFUL? 👍 👎

Devicie

Feedback ⓘ App Protection Policies Go Live

5 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Devicie is a Microsoft Intune-based device management platform that automates Windows device configuration, security policy deployment, and compliance monitoring for enterprises.

Devicie introduces App Protection Policies, now live for all users.

Adds App Protection Policies, enabling enforcement of data protection rules on managed and unmanaged apps.

▾ **4 more releases in this issue** · seen 2026-07-23

Bespoke App Request Form (In-Portal) NOTES ↗ STABLE

Devicie adds an in-portal bespoke app request form for streamlined custom app submissions.

└──▷ **HOW TO FIND IT**

Submit a custom app request without leaving the portal.

📍 In the Devicie portal, locate and complete the Bespoke App Request Form to submit a custom app request.

CIS 4.0 Reporting NOTES ↗ STABLE

Devicie adds CIS 4.0 Reporting support for up-to-date benchmark compliance visibility.

Supports CIS Benchmark 4.0 as a reporting framework for compliance assessments.

Device Spotlight to Intune Deep Link NOTES ↗ STABLE

Devicie adds a direct deep link from Device Spotlight to the corresponding device record in Microsoft Intune.

Multi-Account Login & Fast Switching NOTES ↗ STABLE

Devicie adds multi-account login so practitioners can switch between accounts without signing out.

Supports simultaneous login to multiple Devicie accounts with fast switching between them.

WAS THIS USEFUL?  

◆ Security Data Pipeline

Cribl

Feedback ⓘ

Architecting for the speed of "vibe coding": Introducing Apps (2026-07-23)

2 RELEASES · 2026-07-22 → 2026-07-23 BLOG ↗ LATEST STABLE

Cribl is an observability pipeline platform that helps organizations collect, process, route, and transform machine data from any source to any destination.

Cribl introduces Apps: a sandboxed iframe-based framework letting customers build custom UIs and workflows directly on the Cribl platform.

Introduces Apps, a new sandboxed application framework that lets customers build and run custom UIs and workflows on the Cribl platform without requiring internal platform privileges.

Provides built-in managed backend primitives — including state persistence and secure API proxies — so custom apps need no separate backend infrastructure.

Enforces a strict trust boundary: Apps interact with the Cribl platform only through existing public APIs, never with internal auth tokens or API keys.

Runs customer-written app code in a zero-privilege sandboxed iframe, isolating failures from the core platform's latency, security, and availability.

Enables local development and packaging of Apps using standard web technologies, eliminating the need to learn proprietary platform DSLs.

▾ 1 more release in this issue · 2026-07-22 → 2026-07-23

What's New! We heard you like AI. (2026-07-22) BLOG ↗ STABLE

Cribl 4.19 adds an app framework, embedded MCP server, multi-region Workspaces, and new Snowflake/BigQuery/Bedrock destinations.

└──▷ HOW TO FIND IT

Reduce top-of-hour CPU spikes when many scheduled searches fire simultaneously by enabling Schedule Jitter on a recurring search.

📍 In the Cribl Search UI, open a scheduled search, locate the Schedule settings, and enable Schedule Jitter with a configurable window (e.g., 15 minutes) to introduce a random delay within that window before execution.

Stream security or operational data directly into a Snowflake table without Kafka or S3 staging, using the new native Destination.

📍 In Cribl Stream, go to Routing › QuickConnect (or a Pipeline's Output), add a new Destination, select Snowflake, and configure the Snowflake Destination per target table with your account credentials.

Collect Amazon Bedrock model invocation and CloudTrail audit logs into Cribl Stream for centralized analysis.

📍 In Cribl Stream, add a new Source, select the Amazon Bedrock tile, configure S3/SQS ingestion using the cribl-aws-bedrock-io Pack, and enable dedicated inputs for invocation logs and CloudTrail data.

New app framework lets teams build and install custom UIs, dashboards, and workflow shortcuts directly inside Cribl using Cribl APIs and AI-assisted tools like Cursor, importable via files, URLs, or Git repos.

Cribl-managed MCP server is now embedded in the Cribl Leader as a first-class platform feature, enabling MCP-compatible AI tools and agents to connect using existing auth and access controls — no separate Docker deployment required.

Cribl Insights now supports monitoring, metrics, and alerting for non-main Workspaces (dev, test, and production environments), not just the main Workspace.

Multi-region Workspace provisioning allows Workspaces to be created in any supported region independently of the Organization region, supporting data residency, sovereignty, and compliance requirements.

Centralized credentials support lets teams define credentials once in HashiCorp Vault or Cribl's built-in secret store and reuse them across Stream and Edge, reducing credential sprawl.

+6 more

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Qualys

Feedback ⓘ changelog · Jul 22, 2026 BLOG ↗

Qualys is a cloud-based security and compliance platform that continuously monitors, detects, and helps remediate vulnerabilities across IT assets and infrastructure.

Qualys releases QID coverage for CVE-2026-64600 (RefluXFS), a critical XFS kernel privilege-escalation vulnerability affecting RHEL, Oracle Linux, Amazon Linux, and Fedora.

└──▷ HOW TO FIND IT

Identify Linux systems running an unpatched kernel version exposed to RefluXFS (CVE-2026-64600) across your asset inventory.

> 📍 In the Qualys console, go to Vulnerabilities › Search and query for QID 387841 to surface hosts vulnerable to the RefluXFS XFS reflink race condition, then prioritize patching for internet-facing and multi-tenant systems.

Determine the active kernel version on scanned Linux hosts to assess whether a patched kernel has been applied and rebooted into.

> 📍 In the Qualys console, go to Vulnerabilities › Search and query for QID 45097 to view the Linux kernel version active at scan time, confirming whether vendor-fixed kernels (per RHSA-2026:39179, ELSA-2026-39494, FEDORA-2026-085c9e92a4, etc.) have taken effect after reboot.

New QID 387841 detects the RefluXFS local privilege escalation (CVE-2026-64600) via XFS reflink direct-I/O race condition.

New QID coverage added across VULNSIGS-2.6.651-2 through 2.6.656-2 for vendor kernel advisories on RHEL, Oracle Linux, AlmaLinux, Rocky Linux, and Fedora addressing CVE-2026-64600.

Existing QID 45097 (Linux Kernel Version Running) can be used to identify the active kernel version and assess exposure to RefluXFS across Linux endpoints.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

CrowdStrike Falcon Exposure Management

Feedback

ⓘ undated-41a947f6 BLOG ↗

CrowdStrike Falcon Exposure Management continuously discovers and prioritizes attack surface vulnerabilities across an organization's assets, helping security teams reduce risk and remediate exposures faster.

Falcon Exposure Management adds KEV-native prioritization, automated exploit detection, attack path analysis, and an AI Exposure Analyst Agent for BOD 26-04 compliance.

Adds continuous 24/7 discovery and risk scoring of internet-facing assets, shadow IT, and cloud workloads via external attack surface management (EASM).

Integrates the CISA KEV catalog natively with endpoint telemetry, automatically flagging affected hosts with remediation guidance and zero-config KEV visibility.

Adds Network Vulnerability Assessment (NVA) leveraging existing Falcon sensors to assess unmanaged devices without dedicated scanning appliances.

Adds Secure Configuration Assessment (SCA) that continuously evaluates systems against CIS and other benchmarks while ingesting third-party vulnerability data for unified exposure visibility.

Adds Exploit Automatability Detection using behavioral AI and indicators of attack (IOAs) to detect automated exploitation attempts in real time across endpoints, cloud, and identity.

+3 more

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Auth0

Feedback ⓘ changelog · Jul 23, 2026 CHANGELOG ↗

Auth0 is a cloud-based identity platform that provides authentication and authorization services, enabling developers to securely add login and access control to applications.

Auth0 adds Anonymous Sessions beta and self-service SSO third-party access delegation for customer admins.

└──▷ HOW TO FIND IT

Enable anonymous sessions for your tenant and a specific application so guest activity is tracked before login.

> 📍 In the Auth0 Dashboard, go to Settings › (Tenant Settings) and enable Anonymous Sessions length configuration, then navigate to Applications › <your app> › Settings and enable Anonymous Sessions for that client.

New Anonymous Sessions (beta, paid plans): create and manage stateless sessions for unauthenticated users via a new `POST /anonymous/token` endpoint supporting guest checkout, cart persistence, and pre-login personalization.

Anonymous sessions can be linked to authenticated users through pre-registration and post-login actions, preserving cross-device activity on the authenticated profile.

Anonymous session length is now configurable as a tenant-level setting, and anonymous session support can be enabled per Application and API.

Self-Service Enterprise Configuration tickets can now delegate third-party application access decisions to the customer admin, surfacing a new wizard step during SSO setup rather than requiring the setting at ticket creation time.

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ changelog · Jul 23, 2026

CHANGELOG ↗

Nudge Security discovers and secures all SaaS apps used across an organization, helping IT teams manage shadow IT and enforce security policies through automated nudges.

Nudge Security adds admin audit log export, inherent/residual risk scores, and automatic business criticality tiers for SaaS and AI apps.

└──▷ HOW TO FIND IT

Audit who revoked OAuth grants or changed roles over a given period — useful after a security incident or for compliance reviews.

> 📍 In the console, go to Settings › Audit event, filter events by date or action type, then click Export to download the log as a CSV.

Identify your highest-risk AI vendors by combining residual risk score with recommended remediation actions.

> 📍 In the app inventory, filter by Risk Score = Critical or High, then open any app overview and click the Risk tab to view the factor breakdown and ranked recommended actions.

Focus security posture reviews on apps that handle sensitive data by filtering the inventory by data type and criticality tier.

> 📍 In the app inventory, use the Data Types filter to select a sensitive category (e.g., financial data) and the Business Criticality filter set to Critical or High to surface the apps warranting immediate attention.

Exports a complete admin activity log as CSV, capturing OAuth revocations, role changes, field/setting updates, and app connect/disconnect events (Settings > Audit event).

Introduces inherent and residual risk scores (Low/Medium/High/Critical) for every SaaS and AI vendor, with transparent factor breakdowns covering vendor security posture, internal deployment risk, and compensating controls — now in open beta.

Continuously recalculates risk scores as conditions change (e.g., new MCP server connections, vendor breach disclosures) and surfaces recommended actions ranked by projected residual risk reduction.

Automatically assigns business criticality tiers (Low/Medium/High/Critical) to every app based on inferred data types, enabling inventory filtering across 27 distinct data types with no manual research — now in open beta.

Allows customization of data sensitivity tiers per organization's data governance model, with business criticality recalculating automatically on change (Settings > Risk).

└──▷ **BREAKING ON UPGRADE**

! Previously tagged data types are no longer in the main Data Types field; they now appear under Data Types (Legacy), separate from the new automatically populated Data Types field.

WAS THIS USEFUL? 👍 👎

Ory

Feedback ⓘ New rate limit model for Ory Network

7 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Ory is an open-source identity and access management platform providing authentication, authorization, and user management infrastructure for developers building secure applications.

Ory Network moves to fine-grained, bucket-based rate limiting instead of fixed per-endpoint-path limits.

Introduces bucket-grouped project rate limits, replacing fixed per-endpoint-path rate limiting for more granular control.

▾ **6 more releases in this issue** · seen 2026-07-23

Ory Hydra, Ory Kratos v26.2.16 released NOTES ↗ STABLE

Ory Hydra now accepts OpenID Connect `prompt=select_account` in OAuth 2.0 authorization requests.

└──▷ **TRY IT**

Trigger account selection during an OAuth 2.0 authorization flow so users can switch between identities.

```
$ GET /oauth2/auth?response_type=code&client_id=
<client_id>&redirect_uri=
<redirect_uri>&scope=openid&prompt=select_account
```

Ory Network v26.2.21 released NOTES ↗ STABLE

Ory Network now accepts workspace API keys on GET `/identity-schemas`, enabling programmatic schema listing.

└──▷ **TRY IT**

List all identity schemas from a CI pipeline or backend service using a workspace API key instead of a user session.

```
$ curl -H 'Authorization: Bearer <workspace-api-key>' https://<your-
project>.projects.oryapis.com/identity-schemas
```

Supports workspace API keys on `GET /identity-schemas`, allowing automated/service-account access to list identity schemas without a user session.

Ory Network, Ory Hydra, Ory Kratos, Ory Keto, Ory Oathkeeper v26.3.2 released

NOTES ↗ STABLE

Ory Network v26.3.2 enforces DNS-label boundaries on wildcard return URLs and CORS origins.

└──▷ **BREAKING ON UPGRADE**

! Wildcard `*` in return URLs and CORS origins is now honored only when the `*` is within a single DNS label — previously accepted wildcards that spanned multiple DNS labels will no longer match, potentially blocking requests that worked before.

Ory Network, Ory Kratos v26.2.17 released NOTES ↗ STABLE

Ory Network SCIM endpoints now support enterprise user attribute provisioning.

Ory Network, Ory Kratos, Ory Polis v26.2.19 released NOTES ↗ STABLE

Ory Network adds per-identity data region selection directly from the Console identity details page.

└──▷ **HOW TO FIND IT**

Reassign a specific identity to a different data region without API calls, directly from the Console.

📍 In the Ory Console, navigate to Identity details for the target identity, then use the data region selector (above the organization field) to choose the desired region and save.

Ory Network, Ory Kratos, Ory Polis, Ory Elements v26.2.18 released NOTES ↗

STABLE

Ory Network event streams can now be paused and resumed without losing configuration.

└──▷ **HOW TO FIND IT**

Check or manage the status of an event stream to temporarily halt event delivery without tearing down the stream configuration.

📍 Check the `status` field on an event stream response — it will be either `active` or `paused`.

Adds pause and resume support for event streams, with a new `status` field (`active` or `paused`) indicating stream state — a paused stream retains its full configuration.

WAS THIS USEFUL?  

Castle

Feedback ⓘ Additional Authentication Method types

8 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Castle is a security platform that helps developers detect and block account takeover, fraud, and bot attacks in real time.

Castle adds Proof-of-Personhood and Payment authentication method types for richer signal classification.

Adds Proof-of-Personhood (`pop`) authentication method type for events verified through PoP providers (e.g., Worldcoin, Humanode, etc.).

Adds Payment authentication method type for events tied to a payment method or provider (e.g., card, bank, wallet variants).

▾ **7 more releases in this issue** · seen 2026-07-23

CSV export for Lists NOTES ↗ STABLE

Castle Lists can now be exported to CSV, with background processing and email delivery for large datasets.

└──▷ **HOW TO FIND IT**

Download all entries from a Castle List for offline analysis or reporting.

> 📍 In the Castle console, open the target List, then click Export in the list toolbar to download the CSV (large lists will be emailed as a download link when ready).

Adds CSV export for List entries — click Export in the list toolbar to download; large lists run in the background and deliver a link via email.

Dark mode NOTES ↗ STABLE

Castle adds dark mode with automatic system-preference detection and manual Light/Dark/System override.

└──▷ **HOW TO FIND IT**

Manually switch between Light, Dark, or System (auto) display modes without relying on OS preference.

> 📍 In the Castle console, open the user menu and select your preferred appearance: Light, Dark, or System (auto).

Fraudulent Email Domain signal NOTES ↗ STABLE

Castle adds a Fraudulent Email Domain signal to detect domains registered specifically for fraud and bot signups.

New Fraudulent Email Domain signal detects email domains confirmed as actively used in fraud and abuse campaigns, distinct from disposable email services like Mailinator or TempMail.

IOS SDK 4.0.0 Release NOTES ↗ STABLE

Castle iOS SDK 4.0.0 adds expanded device signals and SDK protection in a major update.

Collects additional device signals for broader fraud and bot detection coverage.

Adds SDK protection capabilities.

Multiple API keys NOTES ↗ STABLE

Castle now supports multiple API keys per account, enabling zero-downtime key rotation and per-integration secrets on Enterprise plans.

└──▷ HOW TO FIND IT

When rotating a compromised or expiring key, create a new one in the dashboard, label it by purpose, migrate traffic gradually, then delete the old key — all without downtime.

📍 In the Castle console, go to Settings › API Keys, click `Create new key`, enter a label (e.g. `Production Web`), copy the new Publishable API Key and API Secret, update your integration, then delete the old key once traffic has migrated.

Create multiple Publishable API Keys and API Secrets within a single Castle account.

Rotate or revoke individual keys without disrupting other active integrations.

Label each key by purpose (e.g., Production Web, Mobile iOS) for clear per-integration tracking.

Supports gradual traffic migration: create a new key, shift traffic incrementally, then delete the old one.

Phone call detection NOTES ↗ STABLE

Castle adds phone call detection to catch social engineering attacks during login and transactions.

.

Detects when a user is on a phone call during sensitive actions (e.g., login, transactions) to identify potential social engineering attacks.

Exposes three new fields for phone call state in Castle events.

Webhook API Key selection NOTES ↗ STABLE

Castle lets you choose which API Secret signs each webhook, enabling independent secret rotation per integration.

└──▷ HOW TO FIND IT

Assign a specific API Secret to a webhook endpoint to isolate it from other integrations and rotate its secret without affecting others.

> 📍 In the Castle dashboard, navigate to your Webhook configuration, select the endpoint, and choose the desired API Secret from the signing key selector.

Enables per-webhook API Secret selection for signing, so integrations can be isolated and secrets rotated independently (Enterprise plans).

WAS THIS USEFUL? 👍 👎

---

iDenfy Feedback ⓘ snapshot-20260723 NOTES ↗

iDenfy is an identity verification platform that uses AI-powered document scanning and biometric facial recognition to verify user identities and prevent fraud.

iDenfy adds Arabic/RTL IDV support, KYC multi-theme branding, KYB dynamic workflows, Companies House integration, and new Digital ID support across sprints.

Adds Arabic language and RTL layout support for identity verification (IDV) flows (Sprint 127).

Adds KYC multi-theme branding, enabling different visual themes per verification workflow (Sprint 127).

Adds KYB sole proprietorship identity verification (Sprint 127).

Adds dynamic workflow decline action for KYB (Sprint 127).

Adds new Digital ID support: Austrian Handy Signatur, Freja eID, and Mobile ID (Sprint 126).

+34 more

WAS THIS USEFUL? 👍 👎

Ping Identity

Feedback ⓘ 21 Jul 2026 (2026-07-22) BLOG ↗

Ping Identity is an identity security platform that provides authentication, single sign-on, multi-factor authentication, and access management solutions for enterprises.

Advanced Identity Cloud adds emailService binding for next-gen scripts, managed secret references in custom nodes, and a new secret validation endpoint.

└──▷ **TRY IT**

Validate that a secret mapping resolves correctly in the configured secret store before relying on it in production.

```
$ curl -X POST \
  'https://<tenant>.forgeblocks.com/am/json/realms/root/realm-
config/secrets/stores/<store-id>/mappings/<mapping-id>?
_action=validate' \
  -H 'Authorization: Bearer <access-token>' \
  -H 'Content-Type: application/json'
```

Custom scripted nodes can now reference managed secrets from a secret store using a `SECRET` property type, keeping API keys and passwords out of script code.

Custom scripted nodes now support combining `multivalued: true` with `options` to create multi-select list properties in the journey designer.

New `validate` action on the secrets endpoint lets operators confirm whether a secret mapping resolves to a valid secret in the configured secret store.

The Enable Connection via Credential toggle for the PingOne Worker service now automatically enables or disables relevant connection fields based on selection.

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ Configurable default OAuth scopes (2026-07-22)

2 RELEASES · 2026-07-22 BLOG ↗ LATEST STABLE

Clerk is a user authentication and identity management platform that provides developers with pre-built UI components, APIs, and SDKs for sign-up, sign-in, and user management.

Clerk lets you configure default OAuth scopes for dynamic client registration.

▾ **1 more release in this issue** · 2026-07-22

> Connect your AI tools to Clerk MCP with one command (2026-07-22) BLOG ↗
>
> STABLE
>
> Clerk MCP server now installs into 10 AI clients with a single command — no manual config required.
>
> └──▷ **TRY IT**
>
> Register the Clerk MCP server in all supported AI clients at once — no manual JSON or config editing needed.
>
> ```
> $ clerk mcp install
> ```
>
> New `clerk mcp install` command registers the Clerk MCP server across 10 AI clients (including Claude Code, Cursor, and VS Code) via a built-in bridge with zero manual configuration.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Pomerium

Feedback ⓘ Configurable DNS Resolvers

8 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Pomerium is an identity-aware proxy that provides secure, clientless access to internal applications by enforcing context-based authentication and authorization policies.

Pomerium v0.32.0 adds configurable DNS resolvers with custom servers and tunable refresh rates for upstream hostname resolution.

Supports specifying custom DNS servers for resolving upstream routes, independent of the host system's resolver.

Enables tunable DNS refresh rates so Pomerium picks up record changes on a schedule matching your environment.

Delivers consistent, centrally configured name resolution across all Pomerium components for dynamic and clustered deployments.

▾ 7 more releases in this issue · seen 2026-07-23

Faster Config Updates & a New Databroker Query Engine NOTES ↗ STABLE

Pomerium v0.32.0 adds incremental route diffing, a new Databroker query/index engine, and Pebble-backed in-memory storage.

Adds incremental config diffing for routes — only changed routes are recomputed and applied, reducing churn on large route sets.

Introduces Databroker indexing and query support, including field-based record lookup, a new GetOptions API, additional filter operators, and order-by.

Replaces the in-memory store with Pebble, aligning it with the persistent backend and setting a default path for the file backend.

Sync gains a `latest stream` option and safely aborts calls with invalid record versions, improving resilience for large deployments.

Hosted Identity Provider NOTES ↗ STABLE

Pomerium v0.32.0 adds a hosted identity provider so teams can authenticate without standing up their own OIDC server.

└──▷ HOW TO FIND IT

Enable the Pomerium-hosted IdP as your identity source instead of wiring up an external OIDC provider — ideal for evaluations or small deployments.

📍 In the Pomerium console, go to Settings › Identity Provider and select the Pomerium-hosted IdP as your identity source.

New hosted identity provider type lets Pomerium manage authentication without requiring teams to configure an external OIDC provider.

Built-in refresh handling authenticates session refresh requests and correctly hashes the derived IdP configuration for clean session renewal.

`MCP: Refresh Tokens, Dynamic Client Registration & Friendlier Errors` **NOTES ↗**

STABLE

Pomerium MCP gateway gains OAuth refresh tokens, dynamic client registration, and customizable error pages.

Adds OAuth refresh token support for MCP sessions, keeping long-running AI agents and clients connected without repeated interactive logins.

Supports dynamic client registration via client ID metadata documents, letting MCP clients register with Pomerium automatically.

Enables customizable MCP 401 error pages with clearer messaging for unauthorized client domains.

Makes allowed client domains optional in MCP configuration, simplifying onboarding.

Sets a proper WWW-Authenticate header on unauthorized MCP responses.

`MCP: Upstream OAuth, Dynamic Client Registration & Route Discovery` **NOTES ↗**

STABLE

Pomerium's MCP gateway gains upstream OAuth with token caching, Dynamic Client Registration, and route auto-discovery.

Adds upstream OAuth support for MCP routes: Pomerium discovers upstream OAuth metadata, completes authorization flows, caches client and token state, and injects upstream access tokens into MCP traffic.

Adds Dynamic Client Registration (DCR) as a fallback for upstream OAuth flows, enabling tools like MCP Inspector that lack Client ID Metadata Document support.

Includes the `iss` parameter in MCP authorization flows, returning Pomerium's host as the issuer to align with the upcoming 2026-07-28 MCP spec.

Enables auto-discovery for MCP connect, disconnect, and authorize endpoints, reducing static upstream OAuth configuration for compatible servers.

MCP metadata discovery now uses SSRF-safe HTTP fetching with stricter protected-resource metadata handling, including path-prefix validation.

+3 more

Native SSH: Authorization Code Flow, Reverse Tunnels & Jump-Host Mode NOTES ↗

STABLE

Pomerium v0.32.0 adds OAuth Authorization Code flow, reverse tunnels, jump-host mode, and an interactive routes portal for native SSH.

└──▷ USE IT

Enable reverse tunneling so SSH servers behind NAT or on edge networks can dial out to Pomerium and receive authorized inbound sessions.

```yaml
ssh_upstream_tunnel: true
```

Enable jump-host (ProxyJump) mode so clients can tunnel straight through Pomerium to an upstream SSH server without requiring the upstream to trust Pomerium's User CA.

```yaml
ssh_allow_direct_tcpip: true
```

Enable the interactive routes portal so users who SSH to Pomerium without specifying a route get a picker of all destinations they're authorized to reach.

```yaml
ssh_routes_portal: true
```

Switches SSH authentication to the OAuth Authorization Code flow (replacing Device Code), delivering a faster browser login with full telemetry observability.

Adds reverse tunneling via `ssh_upstream_tunnel` flag, letting upstream SSH servers dial out to Pomerium — ideal for edge devices, NAT'd networks, and isolated environments.

Adds jump-host (ProxyJump) mode via `ssh_allow_direct_tcpip` flag, letting clients connect through Pomerium to upstreams without requiring the upstream to trust Pomerium's User CA.

Adds an interactive routes portal via `ssh_routes_portal` flag, presenting authorized destinations as a picker when a client SSHes to Pomerium without specifying a route.

Adds per-key session management and revocation in the Pomerium UI, binding each SSH key to an OAuth user with policy evaluated on every connection.

+1 more

└──▷ **BREAKING ON UPGRADE**

! SSH authentication now uses the Authorization Code grant; upstream IdPs must have the Authorization Code grant enabled (v0.31.x and earlier used Device Code, which is no longer used).

Revoke and Block: Stronger Session Revocation (Enterprise) NOTES ↗ STABLE

Pomerium Enterprise adds `Revoke and block future authentication` to permanently lock out users during incident response.

└──▷ **HOW TO FIND IT**

During an incident response, permanently lock out a compromised user so they cannot silently re-authenticate through SSO after session revocation.

📍 In the Enterprise Console, go to Sessions, find the target user's session, open the revocation dialog, and select `Revoke and block future authentication`.

Adds a `Revoke and block future authentication` option to the session revocation dialog, preventing SSO re-authentication loops after revocation.

Enforces blocks via a deny/revocation list so blocked users are rejected on all subsequent requests, not just the current session.

SSH Session Recording (Enterprise) NOTES ↗ STABLE

Pomerium Enterprise adds SSH session recording with full keystroke capture, S3-compatible storage, and a Console playback UI.

└──▷ **HOW TO FIND IT**

Review a full audit trail of a privileged SSH session after the fact — filter by user or host, then replay terminal output.

📍 In the Enterprise Console, go to Recordings, filter by SSH user or target host, select a session, and click Play to replay the recorded terminal output.

Adds full SSH session capture and playback: every keystroke and terminal output for privileged SSH connections is recorded and replayable directly from the Enterprise Console.

Supports flexible recording storage: point recordings at a local on-disk directory or any S3-compatible bucket, configured per cluster.

Introduces a blob-URI builder in the Enterprise Console so admins can configure recording storage without hand-writing connection strings, with support for Go CDK driver options and advanced parameters.

> New Recordings UI lets teams browse, sort, and filter SSH session recordings by route, SSH user, target host, duration, and size.
>
> Recordings are stored in open, portable asciinema-style (.asciicast) format for easy archiving, parsing, and third-party replay.

WAS THIS USEFUL? 👍 👎

Cisco Duo Desktop

Feedback ⓘ snapshot-20260723 NOTES ↗

Cisco Duo Desktop is an endpoint security application that verifies device health and security posture before granting users access to protected corporate resources.

Duo Desktop adds Microsoft Defender for Endpoint detection on Linux, Ubuntu 26.04 LTS support, and user-level preferences persistence on macOS.

Adds detection for Microsoft Defender for Endpoint on Linux (v4.6.0).

Supports Ubuntu 26.04 LTS on Linux (v4.6.0).

Adds a configuration option to restrict the app's interaction to certain account keys on Linux (v4.5.0).

User-level preferences now persist on uninstall on macOS Public Beta (v7.19.3.0).

└──▷ **BREAKING ON UPGRADE**

! Dropped support for Ubuntu 23.04 and 20.04 LTS on Linux (v4.6.0).

WAS THIS USEFUL? 👍 👎

Zscaler

Feedback ⓘ changelog · Jul 22, 2026 BLOG ↗

Zscaler is a cloud-based security platform that protects users, devices, and data by routing internet traffic through its global network for inspection and threat prevention.

Zscaler adds Agentic SecOps with closed-loop remediation and ZDX Network Intelligence for per-hop ISP visibility.

Introduces Agentic SecOps solution (limited availability) — an AI-driven security operations layer that connects telemetry across the Zscaler platform, endpoints/EDR, identity, and cloud to detect threats and orchestrate automated closed-loop remediation.

Adds ZDX Network Intelligence dashboard with global route visualization color-coded by severity, enabling NetOps teams to pinpoint last-mile and transit ISP degradation across hybrid workforces.

Enables BGP Autonomous System Number (ASN) `-level` aggregation of probe telemetry within ZDX, surfacing per-ASN latency and packet loss to identify which carrier is responsible for user experience degradation.

Adds hop-level drill-down within BGP ASNs in ZDX, showing per-hop latency and packet loss to identify specific routers or peering points causing problems.

Introduces ZDX Peer Impact Analysis, providing cross-customer visibility into whether other Zscaler customers traversing the same ISP path are experiencing the same anomaly — confirming whether an issue is carrier-wide or organization-specific.

WAS THIS USEFUL? 👍 👎

DNSFilter

Feedback ⓘ

07/22 July 22 – macOS Roaming Client v2.4.4 Beta (2026-07-22) BLOG ↗

PRE-RELEASE

DNSFilter is a cloud-based DNS security platform that blocks malicious websites, filters unwanted content, and protects networks from cyber threats in real time.

DNSFilter macOS Roaming Client v2.4.4 Beta adds live local domain and resolver sync without daemon restarts.

└──▷ HOW TO FIND IT

Install the beta to get live config sync — navigate to the beta channel in the dashboard to download.

> 📍 In the dashboard, go to Deployments → Roaming Clients → Beta Channel and download the latest beta installer.

Enables live sync of local domain and resolver configuration changes from the dashboard without requiring a daemon restart.

> └──▷ BREAKING ON UPGRADE
>
> ! macOS 12 Monterey is no longer supported; v2.4.4 requires macOS 13 Ventura or later.

WAS THIS USEFUL? 👍 👎

ZeroTier

Feedback ⓘ June 15, 2026

3 RELEASES · 2026-04-15 → 2026-06-15 NOTES ↗ LATEST STABLE

ZeroTier is a virtual networking tool that creates secure, encrypted peer-to-peer networks connecting devices globally as if on a local LAN.

ZeroTier New Central gains the full rules engine: micro-segmentation via capabilities, tags, and custom flow rules.

└──▷ HOW TO FIND IT

When you want to enforce micro-segmentation on a ZeroTier network in New Central — e.g., allow only tagged `servers` to reach tagged `databases`.

> 📍 In New Central, go to your Network › Flow Rules, select `Custom` mode, and enter rules using capabilities and tags (e.g., `accept` / `drop` statements against `tag` and `cap` matchers) in the syntax-highlighted editor, then click Save.

Adds Custom Flow Rules to New Central, enabling network micro-segmentation policies using capabilities, tags, and the full ZeroTier rule definition language (Scale and Enterprise plans).

New Flow Rules interface supports both templated and custom rule modes with a rule editor featuring syntax highlighting and validation.

▾ 2 more releases in this issue · 2026-04-15 → 2026-06-15

May 6, 2026 NOTES ↗ STABLE

ZeroTier Central adds webhooks for real-time notifications on network, member, and organization events.

└──▷ HOW TO FIND IT

Enable webhooks in ZeroTier Central to receive real-time alerts when members join or networks change.

> 📍 In the ZeroTier Central console, go to your Organization settings, select Webhooks, and create a new webhook by entering your endpoint URL and choosing the event types to monitor (e.g. network changes, member activity, organization updates).

April 15, 2026 NOTES ↗ STABLE

ZeroTier Central API v2 launches with full programmatic access to orgs, networks, members, and new service accounts for CI/CD.

> New Central API v2 provides RESTful programmatic access to organizations, network groups, networks, members, and access controls.
>
> Adds Service Accounts with scoped API tokens for automated workflows and CI/CD pipelines, with customizable token expiration dates.

WAS THIS USEFUL?  

◆ Data Security & DLP

Immuta

Feedback ⓘ snapshot-20260723 NOTES ↗

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

Immuta adds separate access/masking request forms, managed groups for policies, Redshift viewless discovery, and OAuth 2.0 for Alation.

└──▷ HOW TO FIND IT

Configure separate request forms for data access vs. masking exceptions on a data product so stewards can ask context-appropriate questions for each request type.

> 📍 In the Immuta UI, navigate to your data product's settings, open the Request Forms section, and configure one form under Data Access Requests and a separate form under Masking Exception Requests — each with its own questions and data use agreement.

Get answers about Immuta features without leaving the product during an investigation or policy authoring session.

> 📍 In the Immuta UI, click the Documentation icon in the top menu, then type your question in natural language into the AI-powered documentation chatbot.

Adds separate access and masking exception request forms for data products, allowing stewards to configure independent questions and data use agreements per request type.

Introduces three built-in system-managed groups (All Agents, All Users, All Identities) for writing subscription and data policies without manual membership management.

Enables automated sensitive data discovery for Amazon Redshift viewless integration, supporting regex, dictionary, and column name regex identifiers with automatic tag application.

Adds OAuth 2.0 authentication support for the Alation integration, eliminating the need to manually rotate expiring access tokens.

Expands Databricks Unity Catalog tag ingestion to support foreign catalogs (created via Lakehouse Federation) by falling back to system.information_schema tag views.

+4 more

└──▷ BREAKING ON UPGRADE

! Immuta will update Snowflake policy role names used to grant users access to Snowflake data sources as part of the architecture rollout beginning July 23, 2026; policy roles will be recreated under new names (user role names remain unchanged).

! Integration error banners for Snowflake integration validation failures will no longer be displayed in the UI; error messages will instead be available through the connections experience.

WAS THIS USEFUL? 👍 👎

Ubiq Security

Feedback ⓘ snapshot-20260723 NOTES ↗

Ubiq Security is a developer-focused encryption platform that provides APIs and SDKs enabling applications to easily encrypt and decrypt sensitive data at the field level.

Ubiq overhauls dataset access control with Access Groups, adds SIEM streaming, SSO, and expands SDK type support.

└──▷ **HOW TO FIND IT**

Enable SSO for the Ubiq Dashboard using an existing enterprise SAML/OIDC provider so users authenticate without separate Ubiq credentials.

📍 In the Ubiq Dashboard, go to Account Administration › SSO and configure your SAML or OpenID provider (Google Workspace, Microsoft Entra, or Okta). See SSO for the Ubiq Dashboard for the full setup guide.

Warm the SDK key cache on startup to eliminate first-use latency in latency-sensitive Java services.

java

```
ubiqClient.loadCache();
```

Introduces Access Groups to replace Dataset Groups for identity-to-dataset access control, separating organization from access management.

Adds Dataset Tagging as a flexible, non-hierarchical replacement for Dataset Groups, supporting grouping by environment, application area, or custom structure.

Adds advanced Dataset Configuration wizard enabling policy-level definition of encryption behavior for dates, numerics, short fields, and large character sets without custom code.

Updates Identities section to consolidate API Keys and IDP-managed users with an access graph for at-a-glance visibility.

Enables Event Streaming (SIEM Log Forwarding) for Enterprise customers to push activity and usage events to Splunk.

+17 more

└──▷ **BREAKING ON UPGRADE**

! Dataset access is now managed through Access Groups instead of Dataset Groups; the identity-to-dataset relationship has changed from `Dataset → Dataset Group → Identity` to `Identity → Access Group → Dataset`.

! Existing Dataset Groups have been converted to Tags; workflows relying on Dataset Groups as an access control mechanism must be updated to use Access Groups.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Third-Party & Vendor Risk (TPRM)

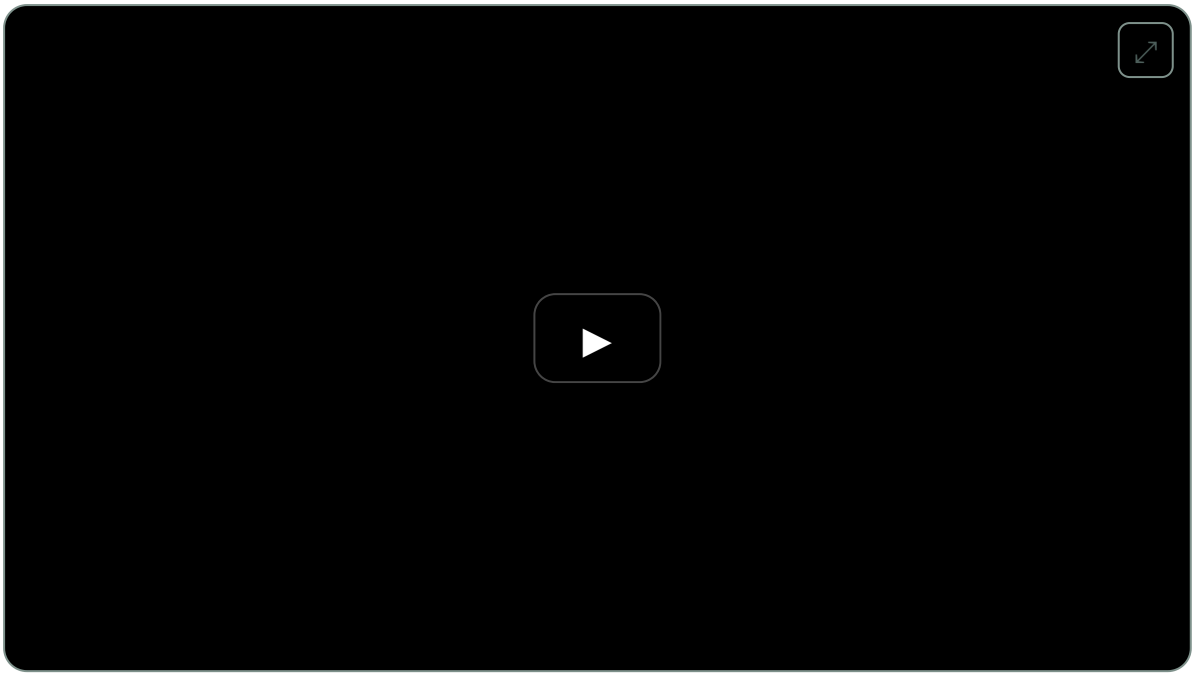
Panorays

Feedback ⓘ Q1'23 12 RELEASES · SEEN 2026-07-23 NOTES ↗

LATEST STABLE

Panorays is a third-party security risk management platform that helps organizations evaluate, manage, and continuously monitor the cybersecurity posture of their vendors and suppliers.

Panorays Q1'23 adds third-party inventory, endpoint detection, questionnaire collaboration, risk policy controls, remediation management, and automated reevaluation.



Adds Third-Party Inventory: a single consolidated view to classify, manage, and upload security documentation for all third parties regardless of tier or evaluation type.

Introduces Endpoint Detection as a new external attack surface assessment category, giving visibility into third-party endpoint security posture and its impact on cyber risk ratings.

Adds Questionnaire Conversations, enabling real-time in-platform collaboration with third parties and internal team members on individual questionnaire questions before submission.

Introduces Correlated Insights, displaying questionnaire results and external attack surface assessment findings side-by-side for cross-validation in a single pane of glass.

Adds customizable Risk Policy controls allowing admins to adjust weighting between questionnaires and external assessments, factor in critical tests and expired questionnaires, and modify the risk rating matrix.

+4 more

▾ **11 more releases in this issue** · seen 2026-07-23

Q1’24 NOTES ↗ STABLE

Panorays Q1'24 adds AI-powered answer validation, asset-based segmentation, richer supply chain discovery, and automatic remediation reminders.

New Smart Validation AI automatically parses third-party documents (questionnaire attachments or Files & Documentation uploads), correlates them with specific questions, and generates confidence scores plus supporting/contradicting evidence.

New Segments feature lets you group third parties into subsets by domain or geolocation assets, generating a focused attack surface assessment for those specific assets.

Supply Chain Discovery now surfaces detailed evidence and navigation controls for direct and unregistered Nth parties, with the ability to select which to add for assessment and continuous monitoring.

Remediation Management adds automatic email reminders to third parties 3 days before open remediation task due dates, removing the need for manual follow-up.

Q1’25 NOTES ↗ STABLE

Panorays Q1'25 adds supply-chain risk insights, customizable executive reports, a new dashboard widget, and custom integrations.

New Risk Insights and Response Portal gives a single-page view of cyber event impact across direct and indirect third parties for faster, data-driven response.

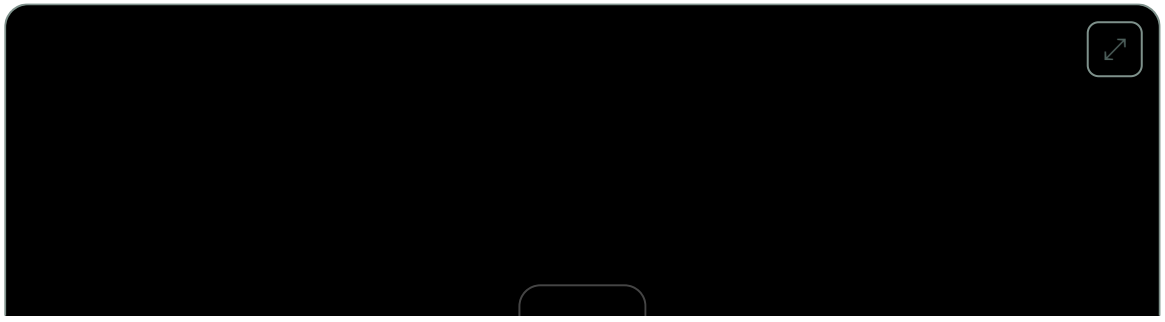
Adds customizable executive reports with mix-and-match widgets and export capability, tailored for board and management audiences.

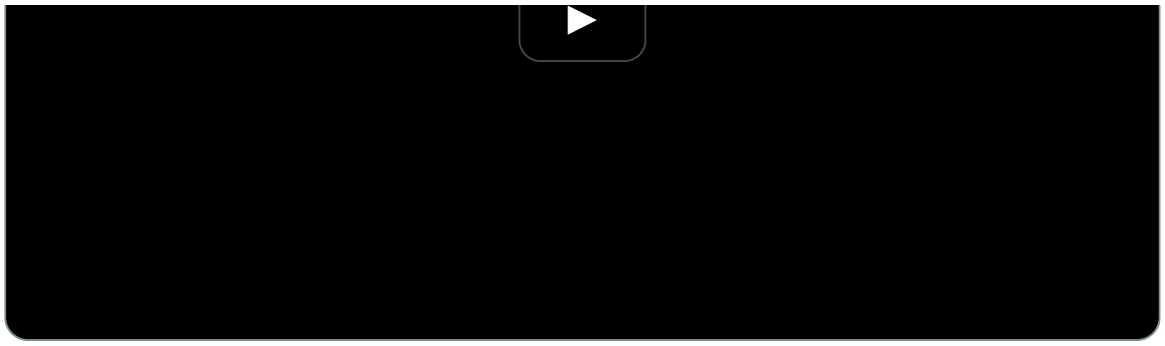
New Live Dashboards Criterion Summary widget surfaces unresolved criteria with criticality, sources, and affected third parties in one consolidated view.

Introduces custom integrations to connect Panorays features with existing security systems for streamlined data exchange.

Q2’23 NOTES ↗ STABLE

Panorays Q2'23 adds Supply Chain Discovery, Risk Insights portal, KEV prioritization, and enriched Asset Details.





New Supply Chain Discovery maps direct and indirect suppliers automatically, surfaces Shadow IT, assigns cyber posture and Commonality ratings, and lets users add discovered suppliers to their inventory in one click.

New Risk Insights and Response Portal delivers real-time alerts on vulnerabilities, breaches, news, and zero-day events across the digital supply chain, identifies affected direct and indirect suppliers, and enables immediate breach-questionnaire dispatch.

CISA Known Exploited Vulnerabilities (KEVs) are now ingested and automatically prioritized as Critical findings within technology-vulnerability detection.

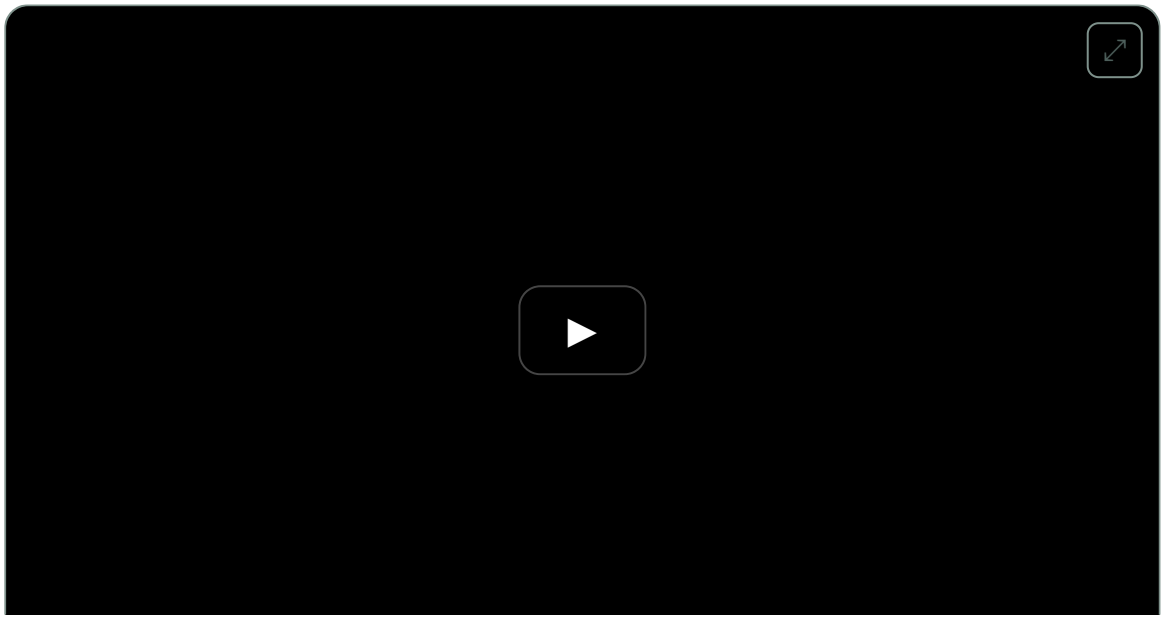
New Asset Details sidebar exposes enriched DNS, Network, Registration, and Web data for first- and third-party discovered assets, including discovery path and criticality context.

New Bulk Task Creation lets users multi-select findings and questions simultaneously and consolidates all resulting remediation tasks into a single outbound email to the third party.

+3 more

Q2’23 AI Release NOTES ↗ STABLE

Panorays Q2'23 adds Smart Match, an AI feature that suggests pre-answered responses to streamline questionnaire completion.



└──▷ **HOW TO FIND IT**

Speed up questionnaire submissions by letting Smart Match pull from your answered-question library instead of drafting responses from scratch.

> 📍 In the questionnaire submission flow, open a question field and select a suggested answer from the AI-generated library ranked by similarity score.

Adds Smart Match, an AI-powered capability that surfaces previously answered questions ranked by similarity score while completing a questionnaire, accelerating response time and ensuring answer consistency.

Q2'24 NOTES ↗ STABLE

Panorays Q2'24 adds internal questionnaires, AI answer lookup, SEC 8-K cyber news, CSP findings, live dashboards, and custom risk factors.

└──▷ **HOW TO FIND IT**

Tune a specific supplier's risk score to reflect offline assurance evidence (e.g., a completed SOC 2 report) without waiting for an automated scan update.

> 📍 In the console, open a supplier profile › Risk Policy › Custom Factors, add a point adjustment for SOC2 completion, and optionally mark evidence as mandatory.

Reduce third-party questionnaire turnaround time by letting vendors search prior answers — third parties access the Answer Lookup search bar in the questionnaire portal to find previously answered similar questions before typing a new response.

> 📍 In the vendor-facing questionnaire portal, use the Answer Lookup search bar to enter a question keyword; review the returned list of similar past questions, their answers, source questionnaires, and similarity match scores.

New Internal Questionnaires feature lets admins circulate questionnaires within their own organization to gather third-party information from relationship owners at any assessment stage (due diligence, SLA reviews, etc.).

New Risk Policy Custom Factors lets admins add or reduce points from a supplier's bottom-line risk rating based on external factors (offline questionnaires, SOC2 completion, trusted vendor status), with per-supplier configuration and mandatory evidence requests.

New Live Dashboards enable customizable reporting dashboards built from selectable data-visualization widgets, tailored for management and board reporting.

New centralized Findings Management page provides enriched finding details and in-place actions such as confirming resolutions or disputing findings.

New AI-powered Answer Lookup search bar helps third parties surface past questionnaire responses, showing similar questions, their answers, source questionnaires, and a similarity match score.

+2 more

Q2’25 NOTES ↗ STABLE

Panorays Q2'25 adds Gemini-powered answer suggestions, simplified supplier onboarding, richer AI business snapshots, and Cyber News webhooks.

New AI Answer Advisor powered by Google Gemini generates intelligent questionnaire answer suggestions drawn from publicly available data, complementing SmartMatch's existing scan of prior responses.

Redesigned Add Supplier flow enables onboarding suppliers without prior background knowledge, with automatic cyber posture assessment initiated upon addition.

Enhanced Business Snapshot aggregates a broader spectrum of public data sources using advanced AI to enrich supplier profiles with company identifiers, financial attributes, and improved compliance accuracy.

Public API now supports Cyber News webhooks, delivering real-time threat intelligence for custom integrations with existing security systems.

Q3’23 NOTES ↗ STABLE

Panorays Q3'23 adds Workflow Automation, Smart Match questionnaire uploads, and Internal References for reviewers.

New Workflow Automation feature lets you define preset rules to automatically approve or reject third parties, generate remediation tasks, and assign responsibilities to multiple stakeholders — triggered by Questionnaire Ratings, Cyber Posture Ratings, Risk Ratings, business impact, and third-party risk status.

Smart Match now accepts uploaded questionnaires from outside Panorays, improving auto-answer accuracy for new questionnaires based on previously completed responses.

New Internal References capability lets reviewers leave internal notes on questionnaires to surface company policy and compliance context, accelerating the review process.

Q3’24 NOTES ↗ STABLE

Panorays Q3'24 adds DORA reporting, third-party portfolios, activity logging, dark web CTI alerts, and task management.

└──▷ HOW TO FIND IT

When you need to submit your DORA Register of Information to a Supervisory Authority, use the DORA File Generator to produce a pre-coded, ready-to-send file instead of manually filling the Excel sheet.

> 📍 In the Panorays console, navigate to the DORA module, select `DORA File Generator`, complete the built-in questionnaire templates, and export the Register of Information file for submission to your Supervisory Authority.

When managing a large third-party portfolio, create segmented groups to focus review resources on your highest-criticality or most sensitive suppliers.

> 📍 In the Panorays console, go to Portfolios, select `Create Portfolio`, then filter and assign third parties by criteria such as criticality, sensitivity, geolocation, or company department.

When investigating a supplier incident or preparing an audit, export a filtered activity log from the Activity Center for the relevant supplier and time range.

> 📍 In the Panorays console, open Activity Center, apply filters for the target supplier and date range, then select `Export` to download the log for reporting.

New DORA File Generator add-on module generates a complete DORA Register of Information report as a ready-to-send file for Supervisory Authorities, with built-in questionnaire templates pre-coded for DORA's Excel format.

New Portfolios feature lets you organize third parties into distinct groups segmented by criticality, sensitivity, geolocation, company department, and more for risk-based prioritization.

New Activity Center provides a consolidated log of all supplier activities on Panorays, with filtering and export for reporting and audit trails.

Enhanced Dark Web monitoring surfaces mentions relevant to your company and third parties with added metadata fields (Source Type, Site, and more) and CTI-backed alerting.

New Task Management: General Tasks enables opening customized tasks for internal teams or third parties from a single centralized location.

`+2 more`

`Q3'25` `NOTES ↗` `STABLE`

Panorays Q3'25 adds supplier inventory management, per-supplier Live Dashboard reports, enhanced tag management, and questionnaire re-submission.

└──▷ HOW TO FIND IT

Classify and monitor suppliers by region or department for targeted assessment plans.

📍 In the console, go to Tags Management, create a tag (e.g., `Region: EU` or `Dept: Finance`), apply it to the relevant vendors, then use the Tags Management page to tier and segment those vendors for a custom assessment plan or report.

Build a focused cybersecurity report tailored to executive stakeholders for a single supplier.

📍 In the console, navigate to Live Dashboards, select a supplier, add and configure interactive widgets relevant to your use case, then save the dashboard as a named report version for that supplier.

New Inventory view provides a centralized location to manage your supplier list, view automatic risk indicators, and move suppliers into tiered assessments based on risk level or priority.

Live Dashboards now support single-supplier reports, letting you build and save multiple tailored report versions per supplier using interactive widgets.

New Tags Management page lets you group and label vendors by characteristics such as location, department, or risk level, and manage all tagged vendors from a single centralized view.

Segments now support configurable Questionnaire Rating calculation — choose to mirror the parent company's questionnaires or create dedicated questionnaires per segment.

Internal Questionnaire responders can now update answers on already-submitted questionnaires without requiring a new submission; senders are automatically notified of any changes.

Q4'23 NOTES ↗ STABLE

Panorays Q4'23 adds supplier intake forms, archival status, AI-powered cyber news, internal notes, and an expanded API v2.0.

└──▷ **HOW TO FIND IT**

Archive an inactive supplier to stop ongoing cyber assessments while retaining their historical documents.

📍 In the Panorays console, open the third party's profile, go to their status settings, and change the status to `Archived`.

New Intake Form lets internal stakeholders add suppliers via a one-time verification process without requiring a Panorays sign-in.

New `Archived` status pauses cyber assessments for inactive or terminated third parties while preserving their files and documents.

Cyber News enriched with additional data sources and an AI-based categorization and classification system for business-specific risk exposure.

Risk Insight and Response Portal gains an Internal Note feature for team communication, memos, and decision-rationale documentation within the portal.

API v2.0 expands access to questionnaire templates, remediation data retrieval, and Portfolio assignment to third parties, with faster response times.

Q4’24 NOTES ↗ STABLE

Panorays Q4'24 adds SIG 2025 support, task consolidation, drill-down dashboards, and security event audit trails.

Adds SIG 2025 evaluation with mapping to DORA and NIS2, per-supplier scoping by material risk, and seamless third-party SIG uploads across all SIG versions.

Adds Inherent Risk Questions to questionnaires, enabling business impact assessments with internal and external stakeholder collaboration.

Adds Task Management Consolidation — a single centralized view to find, track, and manage all tasks assigned to suppliers and internal teams.

Adds drill-down capability to Live Dashboard widgets, surfacing detailed supplier breakdowns for report types and tracking critical findings and remediation task status.

Adds a security events category to the Activity Center audit trail, tracking user login activities per supplier and across all suppliers.

WAS THIS USEFUL? 👍 👎

Conveyor

Feedback ⓘ

📄 Auto-Generate Bridge Letters in Your Trust Center with Smart Documents

7 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Conveyor is a security tool that automates vendor security reviews by enabling companies to share and manage security documentation and questionnaire responses efficiently.

Conveyor Trust Center can now auto-generate signed, dated bridge letters on demand via Smart Documents.

└──▷ HOW TO FIND IT

Set up a bridge letter that self-dates on every download, eliminating manual paperwork for SOC 1/SOC 2 coverage gaps.

> 📍 In the Trust Center console, go to Documents › Add Documents, then select Smart Document, paste your bridge letter template with a date placeholder, and save.

Adds Smart Documents: admins create a template with a date placeholder and Conveyor auto-generates a fresh, signed, and dated bridge letter each time a visitor downloads it.

Supports folder organization, access controls, and other existing document settings for Smart Documents.

Built-in preview lets admins verify the generated output before making a Smart Document live.

▾ 6 more releases in this issue · seen 2026-07-23

📜 New Integration: Automate Trust Center Access with Ironclad NOTES ↗

STABLE

Conveyor Trust Center now integrates with Ironclad to auto-approve access and sync signed NDAs based on contract records.

└──▷ HOW TO FIND IT

Enable the Ironclad integration to start automating Trust Center access decisions based on existing contract records.

> 📍 In the Conveyor console, go to Trust Center › Integration Settings and connect Ironclad.

New Ironclad integration: auto-approve Trust Center access requests when a requester's email or domain matches an existing Ironclad contract record.

Enables automatic NDA gate bypass for counterparties with a signed NDA already on file in Ironclad.

Writes signed clickwrap NDAs back to Ironclad as new records with the PDF attached, keeping contract history in sync.

📝 Draft Changes to Your Trust Center Before They Go Live NOTES ↗ STABLE

Conveyor Trust Center now supports draft mode — stage layout changes privately and publish when ready.

└──▷ HOW TO FIND IT

Preview and publish Trust Center changes without taking the live site offline.

📍 In the Trust Center designer, open your Trust Center, select Create a draft, make layout or content changes, share the preview link for internal review, then publish when ready.

🔌 Assign Questions and Write Answers Through the API NOTES ↗ STABLE

Conveyor's v2 questionnaires API now supports assigning questions and writing/updating answers in bulk.

Assigns questions to individual users or teams via the v2 questionnaires API, individually or in bulk.

Writes and updates questionnaire answers through a new API endpoint, with full edit history preserved.

Exposes questionnaire assignment and answer-writing through the MCP interface for AI agent workflows (e.g., Claude).

🛡️ Topic Preferences: Set Guardrails for How ConveyorAI Answers Sensitive Topics

NOTES ↗ STABLE

Conveyor adds Topic Preferences so admins can set per-topic guardrails controlling how ConveyorAI answers sensitive questions.

└──▷ HOW TO FIND IT

Restrict ConveyorAI from giving generic answers on legal or financial topics by creating a guardrailed topic that flags or overrides responses whenever those subjects come up.

📍 In the Conveyor console, go to Knowledge Library › Topics and click Create Topic. Choose the topic (e.g., Legal Liability), select the desired action ("flag whenever this comes up" or answer this verbatim), and save.

New Topic Preferences feature lets Knowledge Editors, Admins, and Superadmins define per-topic guardrails for ConveyorAI, with actions such as answering verbatim, using specified sources, or flagging whenever a topic surfaces.

Richer downvote feedback is now available on any ConveyorAI answer, independent of Topic Preferences, enabling more granular quality signals.

🤖 Auto-Return Completed Questionnaires, No Human Touch Required NOTES ↗

STABLE

Conveyor can now automatically return completed questionnaires to the source system when they meet configurable criteria.

└──▷ **HOW TO FIND IT**

Eliminate analyst sign-off on already-complete questionnaires by configuring auto-return thresholds in your Intake workflow.

📍 In the Conveyor console, go to your Intake workflow settings and set your return criteria (e.g., minimum answered rate, no flagged or unanswered questions) to enable automatic return.

Adds auto-return for Intake questionnaires: set criteria such as minimum answered rate or no flagged/unanswered questions, and Conveyor returns the questionnaire to the source system without manual review.

🤖 Bring Your Customer Trust Program To Any AI Assistant with the AI Connector

NOTES ↗ STABLE

Conveyor's AI Connector brings your trust knowledge base into Claude, ChatGPT, Gemini, and any MCP-compatible assistant.

└──▷ **HOW TO FIND IT**

Connect Conveyor to your MCP-compatible AI assistant so sales can self-serve approved security answers without leaving their AI tool.

📍 Add `https://mcp.conveyor.com/mcp` as a custom connector in your AI assistant (Claude, ChatGPT, Gemini, etc.) and authenticate via OAuth with your Conveyor credentials.

New AI Connector integrates Conveyor with Claude, ChatGPT, Gemini, and any MCP-compatible AI assistant via a custom connector endpoint.

Enables sales and trust teams to query approved security answers, pull analytics, and update knowledge in plain language directly from their AI tool.

Preserves per-user Conveyor permissions through OAuth, so the AI assistant only

> surfaces data each user is already authorized to see.
>
> Supports org-wide admin registration for enterprise clients, allowing a single setup to cover all users.

WAS THIS USEFUL?  

◆ Compliance Automation & Trust Management

Mondoo

Feedback ⓘ Executive Reports is generally available

11 RELEASES · SEEN 2025-08-01 → 2026-07-23 NOTES ↗ LATEST STABLE

Mondoo is a security platform that continuously scans and assesses infrastructure, cloud environments, and software supply chains for vulnerabilities and misconfigurations.

Mondoo's Executive Reports is now GA — export live VM data as board-ready PowerPoint or PNG slides.

└──▷ HOW TO FIND IT

Generate a board-ready VM report and export it as a PowerPoint deck for stakeholder presentations.

> 📍 In the Mondoo console, navigate to Executive Reports, select your report, and choose Export › PowerPoint or Export › PNG.

Enables Executive Reports for all customers, turning live Mondoo vulnerability management data into board-ready presentations.

Supports export of Executive Report slides to PowerPoint and PNG formats.

▾ 10 more releases in this issue · seen 2025-08-01 → 2026-07-23

Executive Reports is in Preview NOTES ↗ STABLE

Mondoo previews Executive Reports: board-ready ROI slides exportable to PowerPoint or PNG.

└──▷ HOW TO FIND IT

When you need to present vulnerability program ROI to leadership, use Executive Reports to generate and export polished slides directly from live Mondoo data.

> 📍 In the Mondoo console, navigate to Executive Reports (Preview) and select Export to download slides as PowerPoint or PNG.

New Executive Reports feature (Preview) converts live Mondoo vulnerability management data into board-ready ROI presentations.

Supports export of individual report slides to PowerPoint or PNG formats.

Mondoo Release Highlights June 2026 NOTES ↗ STABLE

Mondoo June 2026 brings smarter internet-exposure detection, AI infrastructure inspection, six new platforms, and 145 new checks.

Adds internet-exposure detection across cloud environments to surface attacker-visible attack surface.

Adds AI infrastructure inspection capabilities.

Supports six new platforms for security policy coverage.

Adds 145 new security checks across supported platforms.

Mondoo Release Highlights May 2026 NOTES ↗ STABLE

Mondoo May 2026 debuts a redesigned UI with Risk Dimensions scoring, remediation context, new dashboards, and executive reports.

Introduces Risk Dimensions, a transparent five-dimension scoring model for evaluating security risk.

Adds remediation context alongside findings to help practitioners prioritize and act on issues.

Delivers redesigned dashboards surfacing risk data through the new scoring model.

New executive reports provide leadership-ready summaries of security posture.

Mondoo Release Highlights April 2026 NOTES ↗ STABLE

Mondoo April 2026 adds shadow AI detection, seven new platforms, expanded IaC scanning, and the largest MQL update yet.

Adds shadow AI detection capabilities.

Supports seven brand-new platforms for asset scanning.

Expands infrastructure as code (IaC) scanning coverage.

Deepens cloud and network resource coverage.

Delivers the largest MQL (Mondoo Query Language) expansion to date.

Mondoo Release Highlights March 2026 NOTES ↗ STABLE

Mondoo March 2026 adds network security scanning, hundreds of new cloud resources/checks, IaC scanning for Terraform and CloudFormation variants, and new CIS benchmarks.

Adds network security scanning capabilities.

Introduces hundreds of new cloud resources and security checks.

Supports IaC scanning for Terraform and CloudFormation variants.

Expands remediation content across policies.

Adds new CIS benchmarks.

Mondoo Release Highlights February 2026 NOTES ↗ STABLE

Mondoo February 2026 brings broader vulnerability detection, new policy checks, resource-level cloud findings, and MQL expansion.

Expands vulnerability detection coverage across infrastructure.

Adds hundreds of new policy checks.

Introduces resource-level findings for cloud assets.

Delivers the largest-ever expansion of MQL (Mondoo Query Language).

Mondoo Release Highlights November 2025 NOTES ↗ STABLE

Mondoo November 2025 adds new security advisories, broader vuln detection, NPM supply chain coverage, and network security capabilities.

Adds new Mondoo Security Advisories to surface emerging threats.

Expands vulnerability detection coverage to additional applications and operating systems.

Includes guided remediation steps alongside vulnerability findings to accelerate patching.

Enhances NPM package detection to strengthen software supply chain security.

Introduces new network security detection capabilities.

Mondoo Release Highlights October 2025 NOTES ↗ STABLE

Mondoo October 2025 adds end-of-life OS detections, a known ransomware use risk factor, and new CIS benchmark policies.

Adds detections for end-of-life operating systems to surface legacy assets in security posture.

Introduces a new known ransomware use risk factor for prioritizing vulnerability remediation.

Delivers new and enhanced CIS benchmark policies.

Expands vulnerability detections and remediations.

Mondoo Release Highlights September 2025 NOTES ↗ STABLE

Mondoo September 2025 adds new OS/app vuln detections, DORA & BSI SYS 1.5 compliance, and experimental MCP security.

Adds vulnerability detection support for new operating systems.

Adds vulnerability detection for an expanded list of third-party applications.

Supports EU/German compliance frameworks DORA and BSI SYS 1.5.

Introduces experimental MCP (Model Context Protocol) security scanning.

Streamlines listing and filtering of findings.

+1 more

Mondoo Release Highlights August 2025 NOTES ↗ STABLE

Mondoo August 2025 expands vulnerability detections, adds remediation snippets for Ansible/Bash/Terraform, and ships new CIS benchmarks.

Expands vulnerability detection coverage with new findings.

Adds Ansible, Bash, and Terraform remediation snippets to speed up risk resolution.

Includes latest CIS benchmarks for updated security and compliance coverage.

Enhances Mondoo Query Language (MQL) to simplify creation of custom policies.

WAS THIS USEFUL? 👍 👎

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ CLI 6.37.45 (2026-07-22) BLOG ↗

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

RegScale 6.37.45 adds eMASS full-system sync and AWS Inspector tag-based exclusion filtering.

└──▷ **TRY IT**

Perform a full one-shot sync of all eMASS structural records (assets, controls, POAMs, milestones, artifacts) for an SSP without manually ordering dependent calls.

```
$ regscale emass sync_system
```

Run AWS Inspector sync while excluding resources tagged for non-production environments to keep findings scoped to production assets only.

```
$ regscale aws inspector sync --exclude_tags
Environment=dev,Environment=staging
```

New eMASS `sync_system` command synchronizes all structural records (assets, controls, POAMs, milestones, artifacts) between a RegScale SSP and its linked eMASS system in a single dependency-ordered invocation.

New `--exclude_tags` option for AWS Inspector sync to skip resources carrying specified tags.

WAS THIS USEFUL? 👍 👎

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v3.20.2 NOTES ↗

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

CISO Assistant v3.20.2 adds full GDPR processing register export/import and CyFun 2025 Excel import support.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.20.2 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v3.20.2
```

Adds full export/import support for GDPR processing records, enabling backup and migration of processing register data.

Supports importing CyFun 2025 assessments via Excel, extending the list of supported frameworks for ingestion.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Cursor

Feedback ⓘ Cursor Router NOTES ↗

Cursor is an AI-powered code editor that helps developers write, edit, and debug code faster using integrated large language model assistance.

Cursor Router brings intelligent per-request model routing to Auto mode with Cost/Balance/Intelligence optimization modes and admin controls.

└──▷ HOW TO FIND IT

Surface the routed model on each request so your team can audit which model is being used for compliance or cost-tracking purposes.

> 📍 In the Cursor dashboard, go to Admin › Cursor Router settings and set `Routed model visibility` to displayed (visible to members).

Introduces Cursor Router: an intelligent model router that classifies each request by task type and complexity and dispatches it to the appropriate model automatically.

Adds three optimization modes under Auto — Cost, Balance, and Intelligence — letting users trade off token spend against model capability on a per-session basis.

Adds admin controls to enable Cursor Router per team or group, restrict which optimization modes members can use, set the default mode, and maintain model allow/block lists.

Supports toggling visibility of the routed model (hidden by default), giving teams the option to surface which model handled each request.

Adds soft and hard enforcement options so admins can standardize users on Auto mode.

+2 more

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ v2.1.218 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 v2.1.218 runs `/code-review` as a background subagent, adds MCP error details, and expands auto-mode intelligence.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.1.218 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.218
```

└──▷ **USE IT**

Opt a forked skill back into foreground execution when it must block the conversation before continuing.

```yaml
# In your skill's frontmatter:
---
context: fork
background: false
---
```

Runs `/code-review` as a background subagent so review output no longer floods the conversation and stacked slash commands stay on target.

Adds HTTP status codes and error text to `claude mcp list` and `/mcp` when a server fails to connect, plus a warning for MCP config values with hidden leading or trailing whitespace.

Adds screen-reader announcements of deleted text for word and line deletions (`Option+Delete`, `Ctrl+W`, `Cmd+Backspace`, `Ctrl+U`, `Ctrl+K`) in `--ax-screen-reader` mode.

Announces fast-mode changes when switching models via `/config model=<x>` or Remote Control.

Changes `/deep-research` to start only on explicit invocation — Claude no longer launches it autonomously.

+6 more

└──▷ **BREAKING ON UPGRADE**

! Agent markdown files with names containing `:` are now rejected; any agent whose name includes `:` must be renamed before it will load.

! Skills with `context: fork` now run in the background by default; skills that must run in the foreground need `background: false` added to their frontmatter.

WAS THIS USEFUL?  

GOVERN

◆ AI/LLM Security

Daytona

Feedback ⓘ Docs Search, Git Clone & API 400s

3 RELEASES · SEEN 2026-07-23 NOTES ↗ LATEST STABLE

Daytona is a self-hosted development environment manager that automates provisioning, configuration, and management of standardized, secure dev environments across any infrastructure.

Daytona adds docs search, Git clone support, and improved API 400 error handling.

Adds documentation search capability.

Adds Git clone support.

Adds API 400 error responses for improved developer feedback.

▾ **2 more releases in this issue** · seen 2026-07-23

> SDK Simplification & Per-Sandbox Resource Limits NOTES ↗ STABLE
>
> Daytona introduces SDK simplification and per-sandbox resource limits for finer workload control.
>
> Simplifies the SDK surface to reduce integration complexity.
>
> Adds per-sandbox resource limits enabling granular CPU/memory control per sandbox.

> Sandbox Forking SDK & Org Metrics NOTES ↗ STABLE
>
> Daytona adds sandbox forking via SDK and organization-level metrics.
>
> Adds sandbox forking capability accessible through the SDK.
>
> Introduces organization-level metrics for monitoring sandbox usage.

WAS THIS USEFUL? 👍 👎

Backslash

Feedback ⓘ Block MCP action

5 RELEASES · SEEN 2026-07-23

NOTES ↗ LATEST STABLE

Backslash is an application security platform that analyzes code and dependencies to identify, prioritize, and remediate reachable vulnerabilities in cloud-native applications.

Backslash adds a Block MCP action to control MCP activity.

New Block MCP action enables practitioners to block MCP (Model Context Protocol) actions within Backslash.

▾ **4 more releases in this issue** · seen 2026-07-23

> Claude Code support NOTES ↗ STABLE
>
> Backslash adds support for Claude Code.

> PR Event Notifications NOTES ↗ STABLE
>
> Backslash adds PR event notifications to alert teams of security findings at pull request time.

> Remove PR comment after successful scans NOTES ↗ STABLE
>
> Backslash now removes PR comments automatically after scans come back clean.
>
> Automatically removes previously posted PR comments when a subsequent scan returns no findings, keeping pull requests free of stale security noise.

> Selected repos shown in top NOTES ↗ STABLE
>
> Backslash now surfaces selected repositories at the top of the repo list for faster navigation.
>
> Selected repositories are now pinned to the top of the repository list for quick access.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 23, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue