THE CYBER TOOLCHAIN

NO. 036

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 5, 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 28 of 29 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

Horizon3 NodeZero

Feedback ⓘ 1 RELEASE · SEEN 2026-08-05

NOTES ↗

Horizon3 NodeZero is an autonomous security tool that continuously identifies and exploits vulnerabilities in networks to assess organizational risk.

> Horizon3 NodeZero added a public GraphQL API for programmatic access to the Portal, overhauled Rapid Response with asset-level exploitability detail and a redesigned Real-Time View, and replaced Asset Groups with reusable Scopes in External Pentest workflows.

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

01 **Rapid Response coverage for new CVEs** NEW 50

Rapid Response coverage was added for Starlette/LiteLLM BadHost, Apache Solr, Palo Alto Networks PAN-OS, Nginx UI, Ivanti Sentry, Check Point Remote Access VPN, Splunk Enterprise, Cisco Unified Communications Manager, and Squid.

– Lists specific products covered but no CVE IDs or test details snapshot-20260805

02 **External Pentest Scopes replace Asset Groups** BREAKING 45

The External Pentest workflow was rebuilt so reusable Scopes replace Asset Groups, with improved bulk asset management and uploads for large external attack surfaces.

– Names the replacement concept but no migration steps or config given snapshot-20260805

03 **Public GraphQL API for Horizon3 Portal** NEW 40

A new public GraphQL API provides programmatic access to a subset of Horizon3 Portal functionality, enabling automation, integration, and advanced workflows.

– Names GraphQL API but no endpoint, schema, or auth details given api-docs-cfcbf60ea3eb

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Portkey

Feedback ⓘ 1 RELEASE · SEEN 2026-08-05 NOTES ↗

Portkey is an AI gateway providing observability, security, and reliability features for LLM applications with multi-model support.

> Portkey's March 2026 release adds vault-backed Secret References, weekly rate-limit windows, cross-cloud GCS log delivery via GCP Workload Identity Federation, four new guardrail integrations, and expanded provider support across DeepInfra, DeepSeek, Bedrock, Vertex AI, and Azure AI Foundry.

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

01 **DeepSeek tool calling and reasoning mode** NEW 75

Adds `tools`, `tool_choice`, and `stream_options` for `deepseek-chat`; maps `reasoning_effort` to thinking mode and surfaces `reasoning_content` in streams for `deepseek-reasoner`.

– Parameters and models named, no runnable example snapshot-20260805

THINNER COVERAGE BELOW

02 **Vertex AI resource label mapping** NEW 30

Adds Vertex AI resource label support by mapping Portkey metadata to Vertex resource labels.

– No field names or example given snapshot-20260805

WAS THIS USEFUL? 👍 👎

Semgrep

Feedback ⓘ 1 RELEASE · SEEN 2026-08-05 NOTES ↗

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

> Semgrep introduces a public beta of AI-enabled Agentic Workflows for multi-step detection pipelines, alongside refinements to license compliance blocking, Supply Chain incident notifications, and policy management UX.

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

01 **Agentic Workflows public beta** NEW 60

Semgrep Agentic Workflows enters public beta, running multi-step AI-enabled detection pipelines from the Semgrep AppSec Platform that combine program analysis, deterministic tools, and constrained AI-based reasoning to produce reviewable security issues.

– Explains mechanism and platform but no exact steps or API named. snapshot-20260805

02 **Default Supply Chain incident notification policy** NEW 60

A default Supply Chain incident notification policy is now available to all Supply Chain customers using Unified policies, with optional Slack channel notifications for ongoing incidents based on the most recent SBOM.

– Names Unified policies and Slack integration but no setup steps. snapshot-20260805

THINNER COVERAGE BELOW

03 **License compliance block-all logic** IMPROVED 45

License compliance can now be configured to block a dependency only when all of its licenses are blocked, rather than when any single license is blocked.

– No named config key or flag for the toggle. snapshot-20260805

04 **Confirmation warning for project tag changes** NEW 40

Semgrep now shows a confirmation warning when adding or removing project tags attached to policies, during bulk edit and from a project's settings page.

– Names UI locations but no further mechanism detail. snapshot-20260805

05 **Detailed webhook validation error messages** IMPROVED 20

Webhook validation and test failure error messages now include specific failure reasons.

– Bare description with no example of message content. snapshot-20260805

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback

2 RELEASES · 2026-08-04 → 2026-08-05

BLOG ↗

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

> StepSecurity's Harden Runner now supports AWS CodeBuild-hosted runners, opt-in container process tracing for ARC pods, versioned Runner Job Hooks, and automatic allowance of single-label hostnames in block mode, while threat intel notifications gain a tenant-wide channel configuration page and three match-granularity levels.

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

01 **Threat intel match-granularity levels** NEW 90

Introduces three match-granularity levels for threat intel notifications — All threat intel incidents, Affected packages, and Exact version only — selectable from the Threat intel notifications dialog; alerting is an explicit opt-in per organization via the `Notify when StepSecurity threat intel flags a compromised component` event in organization notification settings.

Limit threat intel noise for an application team by alerting only when the exact compromised package version is present in their environment.

📍 In organization notification settings, enable the 'Notify when StepSecurity threat intel flags a compromised component' event, open the Threat intel notifications dialog, and select Exact version only.

— Names exact dialog, event, and levels with a runnable configuration path.

Threat Intel Notifications: Tenant-Wide…

02 **Container process tracing for ARC pods** NEW 85

ARC Harden Runner `v2.18.6` adds opt-in container process tracing for ARC-hosted pods, enabled via the Helm value `env: enableContainerProcessTracing: "true"`.

Enable container process tracing on ARC-hosted pods to capture process events from inside containers.

yaml

```yaml
env:
  enableContainerProcessTracing: "true"
```

— Exact Helm config value provided, ready to apply.

Harden Runner Releases- August 5, 2026 …

03 **Tenant-wide threat intel notification channels** NEW

85

Adds a Notifications page under Admin console settings to configure email, Slack webhook, and Teams webhook once at the tenant level, replacing per-organization channel setup; granularity can also be set independently at the tenant level so a tenant can subscribe to the full incident feed regardless of individual organization configuration.

Configure a single Slack webhook for all organizations in your tenant so every team receives threat intel alerts without repeating channel setup per org.

> 📍 In the Admin console, go to Settings › Notifications and enter your Slack webhook URL under the Slack webhook field, then save.

— Exact admin console navigation and field given to configure.

Threat Intel Notifications: Tenant-Wide…

04 **Versioned Runner Job Hooks v1.0.0** NEW

75

Publishes Runner Job Hooks `v1.0.0` at `step-security/gha-runner-job-hooks`, letting runners pin to a specific hook version for pre-job and post-job execution; the hooks were rewritten from shell script to TypeScript/JavaScript.

Configure versioned Runner Job Hooks on self-hosted runners via the StepSecurity app UI.

> 📍 In the StepSecurity app, go to Settings → Harden-Runner Installation → Runner Job Hooks and follow the installation instructions to reference a specific version from step-security/gha-runner-job-hooks.

— Repo and version named; setup is via UI navigation, not a command.

Harden Runner Releases- August 5, 2026 …

05 **Implicit single-label hostname allowance in block mode** IMPROVED

70

Single-labeled internal hostnames (no dot) are now implicitly allowed when `egress-policy: block` is set, across GitHub Action `v2.20.1`, ARC `v2.18.6`, and Self-Hosted VM `v1.8.14`, removing the need to list them in `allowed-endpoints`.

— Behaviour change is automatic; no action required from users.

Harden Runner Releases- August 5, 2026 …

THINNER COVERAGE BELOW

06 **AWS CodeBuild runner support in Harden-Runner** NEW

35

Harden-Runner GitHub Action `v2.20.1` adds support for GitHub Actions runners hosted on AWS CodeBuild.

— Names the version and platform but no setup steps given.

Harden Runner Releases- August 5, 2026 …

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

infisical

Feedback ⓘ 7 RELEASES · 2026-07-06 → 2026-08-04 NOTES ↗

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

> Infisical shipped a major Privileged Access Management overhaul alongside a wave of new secrets-sync destinations (Spacelift, Rundeck, Hasura Cloud), Cloudflare secret rotation, HSM-backed internal certificate authorities, native Windows ADCS support, and cross-project secret sharing.

└──▷ WHAT SHIPPED · 28 FEATURES most completely described first ? what's the number?

01 **Certificate sync external identifier exposed** IMPROVED 55

Exposes certificate sync external identifier in the API and UI, making it easier to correlate Infisical-managed certificates with external systems.

— Names API and UI exposure, purpose is clear v0.161.13

02 **HSM-backed internal certificate authorities** NEW 55

Adds HSM-backed internal certificate authorities, enabling hardware-protected signing for internally issued certificates.

— Explains mechanism (hardware-protected signing) clearly v0.161.13

03 **Public certificate authority get-cert endpoint** IMPROVED 50

Makes the certificate authority `get-cert` endpoint publicly accessible.

— Names the exact endpoint made public v0.162.3

04 **Native Windows ADCS certificate support** NEW 50

Adds native Windows Active Directory Certificate Services (ADCS) support for certificate management.

— Names the exact Windows service integrated v0.161.14

05 **Hasura Cloud secrets sync** NEW 50

Supports Hasura Cloud as a secret sync destination, pushing secrets directly to Hasura Cloud environments.

— Names destination and behaviour but no config path v0.161.13

06 **Cloudflare API token and R2 key rotation** NEW 45

Adds secret rotation support for Cloudflare API tokens and R2 access keys.

— Names exact secret types rotated, no flag/endpoint given v0.162.16

07 **JSON and plain text formats in Cloudflare Workers sync** IMPROVED 45

Supports JSON and plain text variable formats in Cloudflare Workers sync.

— Names exact formats and integration destination v0.162.3

08 **Cross-project secret sharing** NEW 40

Secrets can be shared across different Infisical projects; this capability is gated by license entitlement.

— Names entitlement gating but no config surface given v0.162.16 v0.161.13

09 **PKI post-sync commands for agents** NEW 40

Adds PKI post-sync command support for Windows and Linux agents.

— Names platforms but not the command syntax v0.162.16

10 **Azure CLI access in PAM** NEW 40

Adds Azure CLI access to Privileged Access Management (PAM).

— Names the specific CLI target but no setup steps v0.162.3

11 **Rundeck secrets sync integration** NEW 40

Adds Rundeck as a supported app connection for secret sync.

— Names the connection type, no setup path given v0.162.0

12 **Spacelift secrets sync integration** NEW 35

Adds Spacelift as a new secrets sync integration.

— Named integration but no setup detail v0.162.16

13 **403 forbidden-access error page** NEW 35

Adds a dedicated 403 forbidden-access error page in the UI.

— Names the HTTP status handled, no navigation path v0.162.16

14 **Guided self-hosted bootstrap onboarding** NEW 35

Adds guided self-hosted bootstrap onboarding for admins.

— Names admin audience, no flow steps described v0.162.16

15 **PAM assumer role removed** BREAKING 35

The PAM assumer role has been removed.

— Breaking change named but no migration guidance v0.162.3

16 **Revamped Privileged Access Management (PAM) experience** IMPROVED 35

Revamps the PAM (Privileged Access Management) module with a redesigned experience.

— Repeated across releases but no specifics on redesign v0.162.2 v0.162.0

17 **Machine identity credential expiry alerts** NEW 30

Adds expiry alerts for machine identity authentication credentials.

— States what it does, no mechanism or threshold v0.162.16

18 **Staleness check for Windows PAM discovery** NEW 30

Adds a staleness check for Windows PAM discovery.

— Bare description, no detail on how staleness is flagged v0.162.16

19 **Gateway and relay UI migrated to v3** IMPROVED 30

Migrates the gateway and relay UI to v3.

— Names version but not what changed v0.162.16

20 **Expanded onboarding flow with product selection** IMPROVED 30

Expands the onboarding flow with product selection and metrics.

— Names added steps, no further mechanism v0.162.16

21 **Slack notifications for access approvals** NEW 30

Adds Slack notifications for access approvals.

— Names Slack channel but no configuration detail v0.162.3

22 **Certificate manager UI revamp** IMPROVED 25

Revamps the certificate manager UI.

— Only a bare name, no specifics on changes v0.162.16

23 **Organization name shown on invite signup** NEW 25

Shows organization name on the invite signup screen.

— Small UI change, minimal detail v0.162.16

24 **Personal settings UI revamp** IMPROVED 20

Revamps personal settings UI.

— Bare mention with no scope detail v0.162.16

25 **Secrets activation modal** NEW 20

Adds an activation modal for secrets management workflows.

— Vague description of purpose and trigger v0.162.4

26 **Certificate upload support** NEW 20

Supports certificate upload.

— Extremely thin, one-line description v0.162.3

27 **Expanded framework and infrastructure integrations** IMPROVED 15

Revamps and expands coverage of framework integrations and infrastructure integrations.

— No named integrations or scope given v0.162.4

28 **Improved metering UI usage visibility** IMPROVED 15

Improves metering UI to surface more usage information for users.

— Generic improvement, no specific metrics named v0.161.14

└──▷ **BREAKING ON UPGRADE**

! The PAM assumer role has been removed

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

PortSwigger Burp Suite Feedback ⓘ

DAST 2026.8 (2026-08-05) BLOG ↗

Burp Suite is a web application security testing platform for finding vulnerabilities through intercepting and analyzing HTTP traffic.

Burp Suite DAST 2026.8 gives every vulnerability a persistent identity that carries across scans.

Assigns each vulnerability a persistent identity that carries across scans, enabling lifecycle tracking of findings over time rather than treating them as isolated scan results.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

NVIDIA OpenShell Feedback ⓘ 1 RELEASE · 2026-08-04 NOTES ↗

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

> OpenShell v0.0.98 adds OTLP trace export for VM-backed sandbox drivers, giving observability into sandbox execution.

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

01 **OTLP trace export for VM sandboxes** NEW 40

OpenShell exports driver traces from VM-backed sandboxes over OTLP, providing observability into sandbox execution.

— Names OTLP export but no config key or endpoint given. v0.0.98

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ 1 RELEASE · SEEN 2026-08-05 NOTES ↗

> Iru added the ability to exclude and re-include specific users from a policy assignment, plus styling updates to the generated-policy editor.

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

01 **Exclude users from policy assignments** NEW 45

Admins can exclude specific users from a policy assignment when publishing or updating a policy, to handle exceptions to broad policy coverage. Admins and auditors can later re-include a previously excluded user in that assignment.

— Describes the capability and who can use it, no UI path given

Policy assignment exclusions and genera…

02 **Generated-policy editor styling update** IMPROVED 5

The generated-policy editor received styling changes, per the release title; no further detail on what changed is provided.

— Only named in the release title, no description given

Policy assignment exclusions and genera…

WAS THIS USEFUL? 👍 👎

Microsoft Security

Feedback ⓘ 1 RELEASE · 2026-08-04 BLOG ↗

Microsoft Security is a comprehensive platform providing threat protection, vulnerability management, and security monitoring across devices, identities, and cloud environments.

> Microsoft added a DevSecOps pillar and an AI-focused Zero Trust Assessment to its Zero Trust Workshop, alongside expanded AI pillar guidance and new published guidance on securing AI agents and the software supply chain.

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

01 **DevSecOps pillar in Zero Trust Workshop** NEW 70

A new DevSecOps pillar with 15 control groups and 91 tasks covers developer platforms, CI/CD pipelines, source repositories, dependencies, artifacts, and infrastructure-as-code. It also adds four AI-assisted development tasks: code governance, tool allowlisting, data protection, and AI/ML pipeline supply-chain security.

– Concrete counts and named task areas but no exact navigation path 2026-08-04-0378d1c4

THINNER COVERAGE BELOW

02 **AI pillar guidance on AI Memory framework** IMPROVED 50

The AI pillar in the Zero Trust Workshop is improved with guidance based on the Microsoft AI Memory framework, treating memory as a governed security boundary with intent, provenance, lifecycle visibility, and user control.

– Names the framework and its principles but no actionable steps 2026-08-04-0378d1c4

03 **AI agent security pattern and practice guidance** NEW 40

New pattern and practice guidance covers least privilege for AI agents, Zero Trust for source code access, memory safety in agentic systems, software supply chain protection, and security adoption for development.

– Lists topic areas but no specifics on the guidance content 2026-08-04-0378d1c4

04 **AI-focused Zero Trust Assessment** NEW 35

A new Zero Trust Assessment experience establishes a baseline across Zero Trust pillars, including AI and DevSecOps scenarios.

– States purpose but not how the assessment is run or accessed 2026-08-04-0378d1c4

05 **Zero Trust for AI e-book** NEW 25

Publishes a new e-book, "Zero Trust for AI: Rebuilding security controls for autonomous and agentic systems", covering AI agents, tools, memory, data, and runtime operations.

— Names the resource and its scope but no operational detail `2026-08-04-0378d1c4`

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ 1 RELEASE · 2026-08-03 BLOG ↗

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

> Trend Vision One's AI Scanner adds new export formats, an attack-objective filter, and a way to upload scan reports directly into the console for air-gapped environments.

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

01 Multiple export formats for AI Scanner reports NEW

70

AI Scanner scan reports can now be exported as `JSON`, `CycloneDX 1.6` (VEX form), and self-contained sortable `HTML`, in addition to the existing `CSV` option; all formats respect the active dashboard filters.

Export a CycloneDX 1.6 VEX report scoped to a specific attack objective for ingestion into an SBOM or vulnerability management pipeline.

> 📍 In AI Scanner, set the Objective filter to the desired attack objective, then choose Export → CycloneDX 1.6 (VEX).

— Names formats and versions but only a UI export path, not an API.

Export, filter, and upload scan reports…

02 Direct upload of AI Scanner HTML reports NEW

70

AI Scanner HTML reports (up to 100 MB) can be uploaded directly into the TrendAI Vision One console via AI Security → AI Application Security, supporting on-premise and air-gapped deployments.

Get scan results from an air-gapped system into TrendAI Vision One without direct network access.

> 📍 In the console, go to AI Security → AI Application Security and use the upload control to import an AI Scanner HTML report (up to 100 MB).

— Gives exact console path and size limit plus air-gapped use case.

Export, filter, and upload scan reports…

THINNER COVERAGE BELOW

03 Objective filter for AI Scanner results NEW

55

An Objective filter narrows AI Scanner widgets, the findings table, and exported reports to one or more attack objectives, defaulting to all selected.

Export a CycloneDX 1.6 VEX report scoped to a specific attack objective for ingestion into an SBOM or vulnerability management pipeline.

> 📍 In AI Scanner, set the Objective filter to the desired attack objective, then choose Export → CycloneDX 1.6 (VEX).

— Names the filter and its scope but only a UI navigation step.

`Export, filter, and upload scan reports…`

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Cotool

Feedback ⓘ 1 RELEASE · SEEN 2026-08-05 NOTES ↗

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

> Cotool v0.62.0 makes agent runs durable across worker restarts, adds full-text Linear search and GitHub sandbox file uploads, expands Huntress, Darktrace, and Kolide integrations, and introduces guided model upgrades and skill version history for agents.

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

01 **Full-text Linear issue search with filters** NEW 50

Adds full-text Linear issue search with team, status, archive, and comment filters.

– Names filters and integration but no query syntax or path v0.62.0

02 **Darktrace model breach actions** IMPROVED 50

Expands Darktrace support with actions to acknowledge, reopen, and comment on model breaches.

– Names concrete actions but no invocation method shown v0.62.0

03 **Huntress MSP/MSSP account support** IMPROVED 45

Expands Huntress support for MSP and MSSP accounts with optional organization scoping.

– Names account types and scoping option, no setup steps v0.62.0

04 **Durable agent execution across worker restarts** IMPROVED 40

Agent runs, delegated sub-agents, tool calls, waiting prompts, and handoffs are now durable across worker restarts, meaning these states survive a restart rather than being lost.

– Explains scope of durability but no config or trigger detail v0.62.0

05 **Skill version history and restore** NEW 40

Adds skill version history with the ability to inspect and restore earlier versions.

– Clear capability but no location or command given v0.62.0

06 **Guided model upgrades for agents** NEW 35

Adds guided, user-approved model upgrades for agents with supporting research for each recommendation.

— Describes approval flow but no UI or command detail v0.62.0

07 **GitHub sandbox file uploads** NEW 35

Adds support for GitHub tools to upload sandbox-generated files.

— Names GitHub tools and sandbox files but lacks mechanism v0.62.0

08 **Kolide multi-value report query handling** IMPROVED 35

Expands Kolide report queries with improved multi-value parameter handling.

— Names integration and improvement area only briefly v0.62.0

09 **Instruction-aware integration recommendations in builder** NEW 30

Adds instruction-aware integration recommendations directly in the agent builder.

— Names UI location but no detail on recommendation logic v0.62.0

10 **API vs interactive execution audit distinction** IMPROVED 25

Distinguishes API agent executions from interactive runs in audit logs.

— States change but no log field or query example v0.62.0

11 **Provider-native conversation compaction** IMPROVED 25

Improves long conversations with provider-native compaction that preserves more useful context.

— Vague mechanism, no limits or provider names given v0.62.0

WAS THIS USEFUL?  

Tenzir

Feedback ⓘ 1 RELEASE · 2026-08-05 BLOG ↗

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

> Tenzir Ship v2.0.0 changes how pre-1.0 projects handle breaking-change version bumps and adds explicit controls for declaring a first stable release.

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

01 **Pre-1.0 breaking-change version bump logic** BREAKING

80

Projects on major version zero no longer automatically bump to v1.0.0 when unreleased entries include breaking changes; instead they receive a minor bump, e.g. a breaking entry after `v0.4.2` now produces `v0.5.0` instead of `v1.0.0`. Reaching v1.0.0 now requires an explicit invocation rather than happening automatically.

— Clear before/after mechanism with version numbers, but no runnable check command.

Tenzir Ship v2.0.0: Safe pre-1.0 breaki…

02 **--major flag on release create** NEW

80

New `--major` flag on `tenzir-ship release create` explicitly declares the first stable (1.0.0) release, and also accepts `1.0.0` as an explicit version argument. Example: `tenzir-ship release create --major --yes`.

Declare the first stable 1.0.0 release for a project that is still on major version zero.

```
$ tenzir-ship release create --major --yes
```

— Exact flag and runnable command given directly. product docs

THINNER COVERAGE BELOW

03 **stats ambiguous version warning** IMPROVED

45

`stats` now warns and reports no next version instead of failing when multiple outstanding release-candidate series make the target version ambiguous.

— Names the command but no reproduction steps or example. product docs

└──▷ **BREAKING ON UPGRADE**

! Projects on major version zero will no longer automatically bump to v1.0.0 on a breaking change — they receive a minor bump instead (e.g., v0.4.2 → v0.5.0). Reaching v1.0.0 now requires an explicit `tenzir-ship release create --major --yes` invocation.

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

Cribl

Feedback ⓘ 1 RELEASE · 2026-07-31 BLOG ↗

Cribl is a data processing and routing platform that collects, transforms, and delivers security and observability data from multiple sources.

> Cribl launched a new App for AI Observability, giving teams a single pane to search, investigate, and report on AI telemetry across models and providers, backed by dashboards for usage/cost, security/compliance, sensitive-data trace tracking, and a conversational AI Insights agent.

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

01 **Sensitive Data Traces for PII and secrets** NEW 65

Sensitive Data Traces pinpoint exactly where sensitive data (PII, credentials, API keys) appears within a trace or session, with analysis by category, user, model, and tool over time.

– Describes mechanism and multiple analysis dimensions

What is your AI actually doing? Introdu…

02 **Sessions and Traces view** NEW 65

The Sessions and Traces view provides end-to-end trace visibility with filterable session lists, span timelines, and full prompt, response, and tool-call payloads.

– Names concrete UI elements: filters, timelines, payloads

What is your AI actually doing? Introdu…

03 **Usage and Cost Dashboard** NEW 60

The Usage and Cost Dashboard surfaces spending trends, token consumption, top users and models, LLM efficiency anti-patterns, and an activity heatmap.

– Named dashboard with several concrete metrics but no exact navigation

What is your AI actually doing? Introdu…

04 **Security and Compliance Dashboard** NEW 60

The Security and Compliance Dashboard tracks sensitive data exposure rates, findings by type and user, tool-call audit, and content risk indicators.

– Named dashboard with several concrete tracked elements

What is your AI actually doing? Introdu…

THINNER COVERAGE BELOW

05 **Goals and Monitors against budgets** NEW 55

Goals and Monitors continuously evaluate live AI telemetry against user-defined targets, such as month-to-date token spend vs. budget.

— Explains mechanism and gives concrete example target

What is your AI actually doing? Introdu…

06 **Cribl App for AI Observability** NEW 50

New Cribl App for AI Observability delivers a single pane to search, investigate, and report on AI telemetry across models, tools, and environments, built on Cribl's AI Platform for Telemetry.

— Names the platform and scope but no concrete UI path or command

What is your AI actually doing? Introdu…

07 **Org View spend roll-up** NEW 50

Org View provides a hierarchical roll-up of spend, tokens, and requests across an organization.

— Named view with clear metrics, no deeper mechanism given

What is your AI actually doing? Introdu…

08 **AI Insights conversational agent** NEW 45

AI Insights offers a conversational AI agent for querying usage, cost, and security data across the organization in plain language.

— Describes capability but no query syntax or entry point given

What is your AI actually doing? Introdu…

09 **Usage Efficiency tooling** NEW 35

Usage Efficiency tooling exposes LLM efficiency issues to improve token efficiency and model effectiveness.

— Thin prose description with no named surface or metric

What is your AI actually doing? Introdu…

10 **AI-generated attention briefs** NEW 25

AI-generated briefs automatically surface what needs attention so teams are alerted to problems without manual hunting.

.

— Generic description with no mechanism or named surface

What is your AI actually doing? Introdu…

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ 1 RELEASE · 2026-08-05 BLOG ↗

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

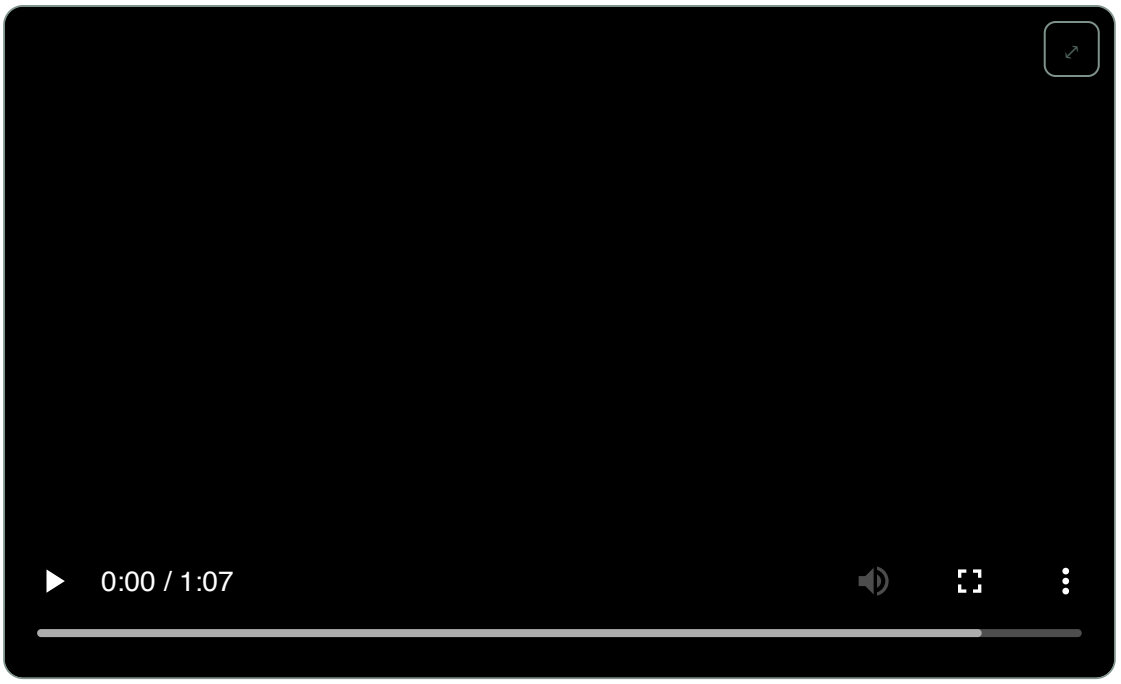
> Tines reworked the case list layout, giving users more control over what information is shown and how it is displayed.

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

01 **New case list layout options** IMPROVED 30

The case list now offers new layout options that give users more control over what information is shown and how it is displayed.





— Describes change but no specific display options or config named.

New layout options for the case list (2…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta

Feedback ⓘ 1 RELEASE · SEEN 2026-08-05 NOTES ↗

Immuta is a data governance platform that controls access to sensitive data and ensures compliance through automated policies and masking.

> Immuta's latest release overhauls the Connections experience with health dashboards, per-object policy sync controls and Redshift viewless auto-discovery, while introducing several breaking changes around Snowflake role naming, Databricks query auditing and data-source grant revocation.

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

01 Databricks query_text masking breaks Unity Catalog query audit

75

BREAKING

Starting August 26, 2026, Databricks will mask `query_text` in the `query_history` system table by default, causing Immuta's Databricks Unity Catalog query audit to lose query text with affected events showing `Indeterminant`, unless the Immuta service principal is added to the `databricks_pii_access` group.

– Names table, field, group and exact remediation action. snapshot-20260805

THINNER COVERAGE BELOW

02 Connection health dashboard and integration settings NEW

55

Connections now display a per-connection health state of `ACTIVE`, `INACTIVE`, or `UNHEALTHY`, and integration settings for Databricks Unity Catalog and Snowflake have been moved directly into Connections.

– Names the three health states but not the exact navigation. snapshot-20260805

03 Snowflake policy role renaming under new integration architecture

55

BREAKING

As part of a new Snowflake integration architecture rollout beginning July 23, 2026, Immuta will rename Snowflake policy role names and recreate existing policy roles under new names; user role names remain unchanged and there is no downtime.

– Gives timing and impact scope but no migration steps. snapshot-20260805

04 Separate access and masking exception request forms per data product

NEW

50

Data products can now have separate access request forms and masking exception request forms, each with independent questions and data use agreements.

— Explains the split but not exact configuration surface. snapshot-20260805

05 **Combined manual resync for Snowflake data and subscription policies**

IMPROVED

45

Manual data policy resync for Snowflake data sources now resyncs both data policies and subscription policies simultaneously.

— Clear behavior change, no command or UI path named. snapshot-20260805

06 **Faster masking exception queries for Snowflake** IMPROVED

30

Query performance for Snowflake securables with advanced masking exceptions is improved via a more efficient comparison method.

— Mechanism described only vaguely as 'more efficient'. snapshot-20260805

└──▷ BREAKING ON UPGRADE

! Immuta will rename Snowflake policy role names as part of the new Snowflake integration architecture rollout (beginning July 23, 2026); existing policy roles will be recreated under new names. User role names remain unchanged and there is no downtime.

! Starting August 26, 2026, Databricks will mask `query_text` in the `query_history` system table by default; Immuta's Databricks Unity Catalog query audit will lose query text and all affected events will show as `Indeterminant` unless the Immuta service principal is added to the `databricks_pii_access` group.

! Immuta will no longer display integration error banners when integration validation fails; error messages will only be available through the new Connections experience.

! When a data source is disabled, Immuta now revokes data owners' Immuta-managed grants in the data platform (behavior change release; on by default August 1, enabled for all accounts September 1).

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Microsoft Entra

Feedback ⓘ 1 RELEASE · 2026-08-04 BLOG ↗

Microsoft Entra is a cloud-based identity and access management platform that secures user authentication and authorization across applications.

> Microsoft Entra expanded its Zero Trust Assessment with three new posture pillars covering AI, Security Operations, and Infrastructure.

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

01 **New Zero Trust Assessment pillars** IMPROVED 45

The Zero Trust Assessment adds AI, Security Operations, and Infrastructure pillars, expanding automated posture evaluation beyond the existing Identity, Devices, Network, and Data pillars.

— Names the pillars but no mechanism or scoring detail given. 2026-08-04-1bf83279

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

JumpCloud

Feedback ⓘ 1 RELEASE · 2026-08-04 BLOG ↗

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

> JumpCloud's agent update (2.152.1) refines the Self Service App Catalog and extends the updated login window to Entra ID-joined Windows devices.

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

01 **Self Service App Catalog filters unsupported apps** IMPROVED

35

The Self Service App Catalog now filters unsupported package-manager applications so users only see installable apps in the JumpCloud app.

— Describes behaviour but no config or scope detail.

JumpCloud Agent Release Notes for Augus…

02 **Updated login window on Entra ID-joined devices** IMPROVED

30

The updated login window is now enabled on Microsoft Entra ID-joined devices, expanding supported sign-in coverage on Windows.

— States rollout scope but no mechanism or steps.

JumpCloud Agent Release Notes for Augus…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

Clerk

Feedback ⓘ 1 RELEASE · 2026-08-04 BLOG ↗

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

> Clerk extended its <SignIn /> component to support the combined sign-in-or-up flow even when strict user enumeration protection is enabled.

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

01 **Sign-in-or-up flow under strict enumeration protection** NEW

58

The `<SignIn />` component now supports the combined sign-in-or-up flow when strict user enumeration protection is active — visitors first verify their identity, and Clerk then either signs them in or creates a new account based on that verification.

— Names the component and behaviour but no config/API detail.

Sign-in-or-up now works with strict enu…

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ 1 RELEASE · 2026-08-04 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's v3.20.8 release adds AI-assisted risk assessment import via chat and expands its framework library with an ANSSI industrial cybersecurity standard, alongside translation updates.

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

01 **ANSSI industrial cybersecurity framework library** NEW 55

Adds the ANSSI 'la cybersécurité des systèmes industriels - mesures détaillées' framework library, extending the platform's set of 150+ supported global frameworks with automatic control mapping.

— Names the framework but not how to enable or map it v3.20.8

02 **AI chat risk assessment import** NEW 35

CISO Assistant now supports importing risk assessments through the AI chat interface.

— No detail on import mechanism, format or limits v3.20.8

03 **Dutch and Ukrainian translation updates** IMPROVED 15

Updates Dutch (NL) and Ukrainian translations.

— Bare mention with no scope or content details v3.20.8

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Claude Code

Feedback ⓘ 1 RELEASE · 2026-08-04 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 v2.1.222 tightens permission enforcement across multi-agent workflows and restricts where Remote Control auto-start can be configured, while also removing the ultraplan feature.

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

01 **Remote Control auto-start scoping restricted** BREAKING

85

Remote Control auto-start can no longer be enabled via repo-local `.claude/settings.json` or `.claude/settings.local.json`; it must now be enabled at user scope via `/config`. Repo-local settings can still be used to turn auto-start off.

— Names exact config files and replacement command v2.1.222

02 **Diffs now render raw git blob content** IMPROVED 65

The `/diff` view, Remote Control workspace diff, and file-edit diffs in web sessions now use raw git blob content, ignoring workspace-configured diff drivers and `textconv`.

— Names affected surfaces and the textconv behavior bypassed v2.1.222

THINNER COVERAGE BELOW

03 **Permission classifier now vets inter-agent messages** IMPROVED

55

In auto mode, messages dispatched to other agent sessions via `SendMessage` are now evaluated by the permission classifier before delivery, improving safety of multi-agent workflows.

— Names SendMessage mechanism but no reader action v2.1.222

04 **disable-model-invocation refusal prompts direct skill run** IMPROVED 50

When `disable-model-invocation` refusal behavior triggers, Claude now prompts the user to run the skill directly rather than attempting to replicate its workflow.

— Names the setting but behavior change is brief v2.1.222

05 **ultraplan feature removed** BREAKING 30

The `ultraplan` feature has been removed from Claude Code; any configuration or workflows depending on it will break.

— States removal only, no detail on what it did v2.1.222

└──▷ BREAKING ON UPGRADE

! Remote Control auto-start can no longer be enabled from `.claude/settings.json` or `.claude/settings.local.json`; configurations relying on repo-local settings to turn it on will stop working — use `/config` at user scope instead.

! The `ultraplan` feature has been removed; any configuration or workflows depending on it will break.

WAS THIS USEFUL? 👍 👎

Cline

Feedback ⓘ 4 RELEASES · 2026-08-05 NOTES ↗ CODE ↗

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

> Cline enforced hard, tool-level blocking of file-editing shell commands in Plan mode across the CLI, SDK, and desktop app, while also shipping session auto-approval, CLI theming, a universal macOS build, and a batch of provider reliability and compatibility fixes.

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

01 **TUI color themes via `/theme` command** NEW 95

Adds a `/theme` command, a command palette, and a Theme row in `/settings` to select the CLI's TUI color theme with live preview; the `CLINE_THEME` environment variable overrides the persisted choice at startup. Built-in themes include Auto, Cline Dark, Cline Light, Tokyo Night, Gruvbox Dark, Nord, Dracula, Catppuccin Mocha, One Dark, Solarized Dark, and Solarized Light.

— Names exact command, settings row, env var, and theme list cli-v3.0.50

02 **Plan mode hard-blocks file-editing shell commands** NEW 86

Plan mode now rejects file-manipulation shell commands at the enforcement/tool level rather than relying on prompting: in-place editors (`sed -i` , `perl -i`), file redirections, mutating git subcommands, package installs, and nested command strings (`sh -c` , `eval` , `sudo` , `xargs`) are blocked in `run_commands` on Linux, Windows, and PowerShell, with a tool error returned; read-only investigation via `run_commands` still works.

— Names exact blocked command patterns and the run_commands tool, but no invocation example

sdk/sdk/v0.0.70 cli-v3.0.50 desktop-v0.0.9 v4.1.4

03 **Auto-approve tool execution without prompts** NEW 85

Adds an `--auto-approve true` CLI flag to start a session with tool auto-approval enabled from the start, an in-client Auto-approve tools per-session toggle for ACP clients like Zed, and an `auto_approve` session config option (`AUTO_APPROVE_CONFIG_ID`) now honored over ACP.

— Runnable flag and named config key with an example command cli-v3.0.50

04 **Automatic context-window overflow recovery** IMPROVED 65

Context-window overflow is now detected and recovered from automatically: the run force-compacts using a deterministic strategy requiring no extra LLM call and retries

once, surfacing an actionable message for terminal cases instead of a raw provider error.

– Explains mechanism in detail but nothing for reader to configure `sdk/sdk/v0.0.70` `cli-v3.0.50` `desktop-v0.0.9`

05 **AI SDK 7 model layer upgrade** IMPROVED 60

Upgrades the model layer to AI SDK 7 and switches Ollama to the native AI SDK provider, with canonical AI SDK 7 file parts for images so image-bearing requests no longer log deprecation warnings.

– Names the SDK version and provider change, but internal, not user-actionable

`sdk/sdk/v0.0.70` `desktop-v0.0.9`

06 **Claude 4.6+/5.x thinking compatibility fix** IMPROVED 60

Claude 4.6+ and 5.x models are no longer rejected with "`thinking.type.enabled` is not supported": adaptive-era models are no longer sent the manual thinking wire shape, the model catalog now carries reasoning metadata, and unlisted or user-typed adaptive IDs are inferred automatically.

– Names exact error and models fixed, but no user action needed `sdk/sdk/v0.0.70` `desktop-v0.0.9`

07 **`mode` field on StartSessionInput** NEW 60

Adds a `mode` field to `StartSessionInput` with values `user`, `automation`, `subagent`, or `team`, alongside the existing `source` field, to record how a session began.

– Names the exact field, its enum values, and containing type `sdk/sdk/v0.0.70`

08 **Bedrock Converse cachePoint markers** IMPROVED 60

The Bedrock provider now emits Converse `cachePoint` markers instead of Anthropic `cache_control`, restoring prompt cache reads and writes and eliminating stray `cache_control` fields in the request body.

– Names the exact field change but is internal to the provider integration

`sdk/sdk/v0.0.70`

THINNER COVERAGE BELOW

09 **Universal macOS binary** IMPROVED 55

Ships a single universal macOS download that runs natively on both Apple Silicon and Intel; existing per-architecture installs migrate automatically.

– Explains behavior and migration but no install command `desktop-v0.0.9`

10 **Session history filter by origin** NEW 55

A new filter control in the sidebar lets users filter session history by origin — Desktop, CLI, extension, or scheduled.

— Names the UI control and its filter values `desktop-v0.0.9`

11 **Context usage ring in composer** NEW 50

The composer displays a token usage ring showing context-window consumption and cumulative cost, changing color as the limit approaches.

— Describes the widget's behavior but no config options `desktop-v0.0.9`

12 **Empty model response retries on all providers** IMPROVED 50

Empty model responses are now retried on every provider — OpenRouter, Cline, and OpenAI-compatible endpoints — rather than only Ollama, which previously failed outright with `Model returned empty response`.

— Named providers and error message, but nothing actionable to do `sdk/sdk/v0.0.70` `cli-v3.0.50` `desktop-v0.0.9`

13 **Skills surfaced in slash command menu** IMPROVED 45

Skills now appear alongside workflows in the slash command menu, and commands that share a name are disambiguated instead of one silently shadowing the other.

— Names the UI surface but no config or command details `desktop-v0.0.9` `v4.1.4`

14 **Lazy root-session persistence** IMPROVED 45

Root-session persistence is now lazy: starting a runtime allocates the session ID in memory without writing a database row, so abandoned sessions no longer leave empty history entries.

— Explains mechanism but is internal SDK behavior, not user-facing action `sdk/sdk/v0.0.70`

15 **`max_completion_tokens` for reasoning models** IMPROVED 45

Reasoning models on OpenAI-compatible endpoints now receive `max_completion_tokens` instead of the previously rejected `max_tokens`.

— Names both parameters but is an internal request-shape fix `sdk/sdk/v0.0.70`

16 **Model-initiated plan-to-act switching removed** BREAKING 40

The model can no longer transition itself from plan mode to act mode mid-turn; mode transitions are now user-driven only.

— Clear behavior change but no migration steps given `v4.1.4`

17 **Infomaniak and SCX.ai added to provider catalog** NEW

40

Adds Infomaniak and SCX.ai to the bundled provider and model catalog.

— Names two providers but no configuration steps `desktop-v0.0.9` `cli-v3.0.50`

18 **Telegram slash command delivery fix** IMPROVED

40

Telegram slash commands such as `/clear` are now delivered to the connector command host instead of being silently dropped.

— Names a specific command and fix but limited scope detail `cli-v3.0.50`

19 **Chutes provider support** NEW

35

Adds Chutes as a recognized model provider.

— Bare provider name with no setup detail `v4.1.4`

20 **Bedrock geo inference profiles** IMPROVED

35

Bedrock foundation models are now routed through geo inference profiles.

— Names the mechanism but gives no further detail on effect `sdk/sdk/v0.0.70`

21 **Diagnostics and telemetry improvements** IMPROVED

30

Application errors are now reported in diagnostics, and telemetry configuration is baked into the sidecar at build time.

— Vague description with no detail on what diagnostics show `desktop-v0.0.9`

22 **Image content substitution for non-image models** IMPROVED

30

Requests to models without image support now substitute image content instead of failing outright.

— Describes fallback behavior with no mechanism detail `sdk/sdk/v0.0.70`

└──▷ **BREAKING ON UPGRADE**

! Model-initiated switching from plan mode to act mode is removed; the model can no longer transition out of plan mode mid-turn — only the user can.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

CrewAI

Feedback ⓘ 1 RELEASE · 2026-08-05 NOTES ↗

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

> CrewAI 1.15.11 adds a built-in Db2 search tool and lays groundwork for tying open-source usage to enterprise accounts, alongside expanded telemetry for coding agents.

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

01 **project_id enterprise account linking** NEW 50

Adds `project_id` support to link open-source CrewAI usage to an enterprise account.

— Names concrete field but no file or API location given 1.15.11

02 **AMP surfaced in AGENTS.md** IMPROVED 45

Surfaces CrewAI AMP in `AGENTS.md` and adds detection of coding agents in telemetry.

— Names file and mechanism but lacks further detail 1.15.11

03 **IBM Db2 search tool** NEW 40

CrewAI adds an IBM Db2 search tool, enabling agents to query Db2 databases as a built-in tool.

— Names the tool but no config or usage detail 1.15.11

04 **Interception-hook dispatch telemetry** IMPROVED 30

Tracks interception-hook dispatches in telemetry for improved observability of agent execution.

— Describes what's tracked but no mechanism or access point 1.15.11

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Agentic Infrastructure

emisar

Feedback ⓘ 5 RELEASES · 2026-07-12 → 2026-07-16 NOTES ↗

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

> This window hardened emisar's MCP layer (typed argument coercion, clearer faults, trimmed and enriched run summaries) and its release/infrastructure trust chain (digest-pinned, SBOM-attested portal images, immutable signed binaries, a dedicated pack registry, and a public status page), while shipping new dispatch-log operator tooling (`emisar doctor` , `state check-dispatch-log` , `QUARANTINE_DISPATCH_LOG`) and closing a string of security gaps (step-up 2FA disable, Secure cookies, Redis ACL redaction, SCIM rate limiting, TLS 1.2 floor).

└──▷ WHAT SHIPPED · 37 FEATURES most completely described first ? what's the number?

01 `QUARANTINE_DISPATCH_LOG` install.sh environment variable NEW

90

Adds `QUARANTINE_DISPATCH_LOG=1` environment variable to `install.sh` to automatically quarantine an unreadable dispatch log instead of leaving a host crash-looping after an upgrade.

Automatically quarantine a corrupt dispatch log during an unattended host upgrade instead of crash-looping.

```
$ QUARANTINE_DISPATCH_LOG=1 ./install.sh
```

— Exact env var and runnable install command given v0.28.0

02 `emisar doctor` and `state check-dispatch-log` commands NEW

85

Adds `emisar doctor` checks for the dispatch log, degraded packs, and last cloud rejection; adds `emisar state check-dispatch-log` for a scriptable verdict on dispatch-log health.

Check runner health — dispatch log integrity, degraded packs, and last cloud rejection — before or after an upgrade.

```
$ emisar doctor
```

Script a CI/CD gate that fails a pipeline if the dispatch log is unreadable or in a bad state.

```
$ emisar state check-dispatch-log && echo "dispatch log OK" ||
exit 1
```

— Runnable commands with CI-gating example given v0.28.0

03 **Dedicated pack registry hostname** NEW 75

Pack registry now serves from its own hostname at `registry.emisar.dev/v1/catalog.json`, with anonymous access narrowed to exact object reads.

— Exact endpoint URL given, navigable directly v0.24.0

04 **Bridge version visibility and staleness flag** NEW 70

Console LLM agents list now groups connections by the operator behind each key, displays the `emisar-mcp` bridge version inline per connection, and flags out-of-date bridges with a one-step path to the upgrade command.

— Names the bridge and UI grouping, no exact upgrade command v0.26.0

05 **Specific MCP argument fault messages** IMPROVED 65

MCP argument faults now name the specific field and type mismatch — e.g. `limit must be a JSON integer from 1 to 100; it was sent as a string.` — replacing generic contract errors; enum faults list allowed values and unknown arguments are spelled out against the allowed set.

— Exact error text quoted, mechanism explained v0.29.0

06 **Risk level raised for secret-dense config dumps** IMPROVED 65

Cloud-init user/vendor data, cloud-config, pm2 env, nginx config dumps, and caddy config dumps moved from risk level `low` to `medium`, matching `linux.sudoers_dump`, removing them from the ungated low-read path.

— Names exact packs and risk levels moved v0.29.0

07 **MCP scalar argument type coercion** IMPROVED 60

Scalar MCP arguments sent as canonical strings (e.g. `"limit": "50"`, `"dry_run": "true"`) now coerce to their expected types instead of failing the call.

— Concrete field examples given but no runnable command v0.29.0

08 **MCP run summary enrichment and trimming** IMPROVED 60

MCP run summaries now include a bounded terminal failure message when the runner supplies one and a bounded policy reason for denied runs, while omitting zero-information fields: `stdout` / `stderr`, byte-count, and truncation fields are absent when a stream produced no bytes, `output_complete` appears only when false, and per-stream SHA-256 digests are removed from the MCP projection (retained on the run detail page and audit record).

— Field-level detail across three releases but no command v0.29.0 v0.28.0 v0.27.0

09 **Runner v0.12.0 dispatch journal migration break** BREAKING

60

Upgrading to runner-v0.12.0 starts a new durable dispatch journal and does not import older unacknowledged entries; downgrading to an earlier runner after v0.12.0 has written new dispatch records is unsupported.

— Clear migration note naming exact version boundary v0.27.0

10 **Attested, digest-pinned portal release images** NEW 60

Portal releases now publish the exact image exercised by CI, pinned by digest, with provenance and SBOM attestations; fixable HIGH/CRITICAL vulnerabilities block publication.

— Explains full mechanism and gating criteria v0.24.0

THINNER COVERAGE BELOW

11 **Redis ACL password hash redaction** IMPROVED 55

Redis ACL password hashes are redacted from `acl_list` and `acl_getuser` output, preventing ungated reads from leaking crackable material.

— Names exact commands affected v0.29.0

12 **Exec-argument leading-dash flag protection** IMPROVED

55

Exec-style arguments across docker, systemctl, podman, minio, and related packs now anchor their first character to an alphanumeric, blocking hostile values from being parsed as leading-dash flags.

— Names affected packs and the exact mechanism v0.29.0

13 **install.sh pre-flight dispatch state verification** IMPROVED

55

`install.sh` now verifies dispatch state with the staged binary before touching a running service, aborting with remediation options on failure.

— Named script and behavior, no exact command shown v0.28.0

14 **Degraded pack visibility in MCP and portal** NEW 55

Degraded packs now surface as named load failures in MCP `list_runners` responses and as an amber notice with reasons and reinstall remedy on the runner detail page.

– Names the endpoint and UI surface, no exact steps v0.28.0

15 **`docs/COMPATIBILITY.md` runner-state compatibility surface** NEW 55

Adds `docs/COMPATIBILITY.md` naming on-host runner state (dispatch log, installed pack trees) as a compatibility surface with documented migration rules.

– Names exact doc path readers can consult v0.28.0

16 **Public status page with on-call escalation** NEW 55

Public status page available at `status.emisar.dev`, with external availability monitoring and on-call escalation.

– Exact URL given, directly navigable v0.24.0

17 **Hardened hosted control-plane infrastructure** IMPROVED 55

Hosted control plane moves to hardened infrastructure with zero-downtime rolling deploys, a private-network database with TLS pinned to a private CA, and no public compute surface.

– Detailed architecture change, nothing for reader to act on v0.24.0

18 **Runner v0.13.0 with SLSA-3 provenance** IMPROVED 45

Runner binary ships as `runner-v0.13.0` with SLSA-3 build provenance and `SHA256SUMS`.

– Names exact version and artifacts shipped v0.28.0

19 **Hosted MCP server listed in official registry** NEW 45

Hosted MCP server (`dev.emisar/emisar`) is now published to the official MCP registry.

– Names the hosted server identifier v0.24.0

20 **Natural multi-term queries in `find_actions`** IMPROVED 40

`find_actions` supports natural multi-term queries.

– Names the function but no example query v0.28.0

21 **Operator-run GDPR account erasure task** NEW 40

Account and user erasure (GDPR delete) available as an operator-run release task.

— Names the task type but not exact invocation v0.28.0

22 **Corrupt dispatch state halts runner pre-connect** NEW

40

Corrupt durable dispatch state now halts the runner before it connects to the control plane, preventing a misconfigured runner from silently operating.

— Describes behavior, no exact command or flag v0.27.0

23 **Single Terraform apply gate for superseded plans** IMPROVED

40

A newer non-destructive production plan supersedes an older pending plan, with Terraform confirmation remaining the single human apply gate.

— Names Terraform gate but no exact workflow v0.27.0

24 **Step-up challenge required to disable 2FA** IMPROVED 35

Disabling two-factor authentication now requires a fresh step-up challenge, preventing a hijacked session from quietly weakening an account.

— Clear behavior but no config or flag named v0.29.0

25 **Full route coverage for SSO/MFA enforcement** IMPROVED

35

Sign-in enforcement for Require SSO and Require MFA now covers every controller route.

— States scope change but no config surface named v0.28.0

26 **Rate limiting on SCIM and admin endpoints** NEW 35

SCIM, audit-export, and runner-register endpoints are now rate-limited.

— Names endpoints but no limit values or config v0.28.0

27 **Direct alerts for critical infra conditions** NEW 35

Adds direct alerts for database-down, zero-healthy-backend, and portal VM saturation conditions.

— Names conditions monitored, no config shown v0.28.0

28 **OAuth account-chooser consent for connectors** IMPROVED

35

OAuth consent now lets a user choose which account to connect during connector setup.

— UI behavior described, no exact flow steps v0.27.0

29 **Immutable, reproducible signed binary releases** IMPROVED

35

Binary releases are now immutable and reproducible, with signed and verified release tags.

— States property, no verification steps given v0.24.0

30 **Secure flag on production session cookies** IMPROVED 30

Production session cookies always carry the `Secure` flag.

— Names the flag but nothing to act on v0.29.0

31 **Console-created runbooks executable via MCP** NEW 30

Console-created runbooks are now visible and executable over MCP.

— Brief statement, no mechanism or example v0.28.0

32 **TLS 1.2 minimum floor for outbound connections** IMPROVED 30

Outbound connections now assert a TLS 1.2 minimum floor.

— States the floor version, nothing to configure v0.28.0

33 **In-flight run recovery on runner reconnect** IMPROVED 30

Control plane recovers runs that were in flight when a runner drops and reconnects, instead of stranding them.

— Behavior described without mechanism detail v0.26.0

34 **Connector commands excluded from shell history** IMPROVED 25

Key-bearing connector commands are kept out of shell history.

— Brief statement with no mechanism given v0.27.0

35 **Redaction of sensitive args in failure causes** IMPROVED 25

Failure causes redact sensitive action arguments before they reach local audit records or the portal.

— States scope but no field-level detail v0.27.0

36 **Signing nonce state skipped for unsigned installs** IMPROVED

25

Unsigned installations no longer open or lock signing nonce state during startup.

— Terse internal change, no user-facing action v0.27.0

37 **Per-page runner upgrade notice scoping** IMPROVED

25

Runner version-upgrade notice now scopes its count to the current page rather than reading as a single runner across a multi-runner list.

— Minor UI fix, no config or command v0.26.0

└──▷ **BREAKING ON UPGRADE**

! Upgrading to runner-v0.12.0 starts a new durable dispatch journal and does not import older unacknowledged entries; downgrading to an earlier runner after v0.12.0 has written new dispatch records is unsupported.

WAS THIS USEFUL? 👍 👎

◆ Local LLM Runtimes

vMLX

Feedback ⓘ 1 RELEASE · 2026-07-13 NOTES ↗

vMLX - JANGTQ Uber Compressed MLX Models - L2 Disk Cache (survives restart) + L1 Paged (super fast ttft) + Hybrid SSM Scheduler + Cont Batching + etc!

> vMLX 1.6.7 flips two major defaults: paged KV cache now turns on automatically for supported text model families, and reasoning mode is now enabled by default across all reasoning-capable models.

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

01 Paged KV cache default for text model families IMPROVED

88

Paged KV cache now defaults ON for autodetected text families (Qwen3.5/3.6, Zaya, Nemotron-H, LFM2.5, Laguna, Hy3, Step-3.7) when continuous batching and prefix cache are active. The default can be overridden with the `--use-paged-cache` / `--no-paged-cache` flags, e.g. `vmlx serve mlx-community/Qwen3-8B-4bit --no-paged-cache`.

Override the new paged-cache default to explicitly disable it for a model family that behaves unexpectedly with block-based caching.

```
$ vmlx serve mlx-community/Qwen3-8B-4bit --no-paged-cache
```

Explicitly opt in to paged KV cache on a model where you want to confirm block-based deduplication is active.

```
$ vmlx serve mlx-community/Qwen3-8B-4bit --use-paged-cache
```

— Names exact families, gating conditions, and runnable override flags. v1.6.7

THINNER COVERAGE BELOW

02 Reasoning mode default ON for reasoning-capable models IMPROVED

40

Reasoning mode now defaults ON for every reasoning-capable model family, with `max_thinking_tokens` honored and gated on both streaming and non-streaming paths.

— Names the config field but gives no flag or command to control it. v1.6.7

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Promptfoo

Feedback ⓘ 1 RELEASE · SEEN 2026-08-05 NOTES ↗

Promptfoo is a testing framework for evaluating and benchmarking language model prompts and AI applications across multiple providers.

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

01 **numTests cap for red team strategies** NEW 70

The `numTests` config key caps the number of test cases generated per red team strategy, as in a `promptfooconfig.yaml` snippet setting `numTests: 20` under a strategy.

Cap test case generation per strategy and load config with env variable path substitution to keep CI scans bounded.

yaml

```yaml
# promptfooconfig.yaml snippet
redteam:
  strategies:
    - id: jailbreak
      numTests: 20
prompts:
  - file://$PROMPTS_DIR/system.txt
```

— Named config key with a working YAML example snapshot-20260805

02 **Model Audit cloud sharing flags** NEW 60

New `--share` and `--no-share` flags for Model Audit control whether results are shared to the cloud.

— Exact flag names given but no example command snapshot-20260805

THINNER COVERAGE BELOW

03 **Environment variable substitution in file paths** NEW 55

`$VAR` syntax is now supported for environment variables in file paths within configuration, as shown loading a prompt via `file://$PROMPTS_DIR/system.txt`.

Cap test case generation per strategy and load config with env variable path substitution to keep CI scans bounded.

```yaml
# promptfooconfig.yaml snippet
redteam:
  strategies:
    - id: jailbreak
      numTests: 20
prompts:
  - file://$PROMPTS_DIR/system.txt
```

— Named syntax with a config example, no further mechanism detail snapshot-20260805

04 **Code scanning for fork PRs and comment triggers** IMPROVED

40

Code scanning adds support for fork pull requests and comment-triggered scans.

— States the two new triggers with no setup detail snapshot-20260805

└──▷ ALSO FROM THESE RELEASES

WAS THIS USEFUL? 👍 👎

NVIDIA SkillSpector

Feedback ⓘ 1 RELEASE · 2026-08-04

NOTES ↗

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

> SkillSpector v2.5.2 adds MCP registry posture scanning and hardens its ingest pipeline with hard size and entry-count caps that fail closed on breach.

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

01 **Ingest size and entry-count caps** NEW 83

SkillSpector now enforces `INGEST_MAX_BYTES` (100 MiB) on streamed URL downloads, total uncompressed zip size, and post-clone Git repo disk usage, plus `INGEST_MAX_ZIP_MEMBERS` (10,000) on the number of zip entries. Breaching either limit fails closed and raises `IngestLimitExceededError`.

— Names exact env vars, defaults, and error type v2.5.2

THINNER COVERAGE BELOW

02 **MCP registry posture scanning** NEW 23

SkillSpector adds scanning of MCP registry posture as part of its security checks.

— Only a one-line mention with no mechanism or scope detail v2.5.2

WAS THIS USEFUL? 👍 👎

garak

Feedback ⓘ 1 RELEASE · 2026-08-04 NOTES ↗

the LLM vulnerability scanner

> garak v0.16.0 introduces a Context Aware Scanning framework for technique/intent-driven evaluation, unifies probe and buff selection into a single run.spec grammar, and adds a native Anthropic generator alongside several detector hardening improvements.

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

01 **Unified run.spec selection grammar** BREAKING 74

garak consolidates the old `probe_spec` and `buff_spec` configuration keys into a single `run.spec` definition within the `run` config, and adds support for selecting _intents_ alongside probes and buffs. Existing config files using top-level `probe_spec` or `buff_spec` will break, and the equivalent CLI plugin-selection options are now deprecated in favour of the unified syntax (old options still accepted but flagged deprecated).

— Names the exact config keys and migration path clearly. v0.16.0

02 **report.jsonl structure changes with probe_summary entry** 66

BREAKING

`report.jsonl` output gains a new `probe_summary` entry type, and the `digest` format is expanded with probe summary data; `digest` values tied to the old `*_spec` configuration keys are removed. Downstream parsers of `report.jsonl` must be updated for the new structure.

— Names exact output file and entry type, migration impact noted. v0.16.0

03 **Context Aware Scanning (CAS) framework** NEW 65

New framework adding technique and intent annotations on probes, a `technique_intent_matrix` in reports enriched with names and descriptions, and a new `IntentProbe` base class for intent-driven evaluation. Scan reports also gain new technique and intent tabs/lists for browsing CAS results.

— Names classes and report fields but no runnable command given. v0.16.0

04 **suppressed_params for BedrockGenerator** NEW 62

Adds `suppressed_params` support to `BedrockGenerator` for controlling which parameters are passed to the Bedrock API.

— Names exact config field and generator class. v0.16.0

THINNER COVERAGE BELOW

05 **Unicode normalisation in StringDetector** IMPROVED 46

Adds Unicode normalisation to `StringDetector` for more robust string matching across character representations.

— Names the detector and change but not the mechanism. v0.16.0

06 **Hardened web_injection markdown exfiltration detectors** IMPROVED 42

Hardens `web_injection` markdown exfiltration detectors against domain, extension, and parameter evasion techniques.

— Names the detector and evasion types but no usage detail. v0.16.0

07 **Native Anthropic generator plugin** NEW 36

New native Anthropic generator plugin for direct Anthropic API access.

— Only a name and one-line description, no config details. v0.16.0

08 **Simple adaptive attacks probe** NEW 28

New simple adaptive attacks probe added to the `probes` plugin.

— No probe name, mechanism, or usage given. v0.16.0

└──▷ BREAKING ON UPGRADE

! `probe_spec` and `buff_spec` configuration keys move into the `run` definition as part of the new unified spec syntax — existing config files using top-level `probe_spec` or `buff_spec` will break.

! CLI plugin selection options for probes and buffs are now deprecated in favour of the unified spec syntax (old options still accepted but marked deprecated).

! The `report.jsonl` output structure changes: new `probe_summary` entry type is added, the `digest` format is expanded, and `digest` values tied to the old `*_spec` configuration keys are removed — any downstream parsers of `report.jsonl` must be updated.

WAS THIS USEFUL? 👍 👎

PLATFORMS & UTILITIES

◆ General Utilities & Libraries

Obsidian

Feedback ⓘ 5 RELEASES · 2026-07-14 → 2026-07-30 BLOG ↗

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

> Obsidian's 1.13 cycle overhauls image viewing with a full-screen lightbox, keyboard-driven Live Preview image editing, and a new Vim `:image` command, while adding a searchable Settings panel on both desktop and mobile, an iOS Share Sheet integration, and tightened URI/HTML security — alongside breaking changes to the `--callout-color` CSS variable and the minimum supported iOS version.

└──▷ WHAT SHIPPED · 25 FEATURES most completely described first ? what's the number?

01 **Live Preview keyboard image selection and editing** IMPROVED 95

In Live Preview, images can now be selected via keyboard; `Backspace` / `Delete` deletes the selected image, `Ctrl/Cmd-C` / `Ctrl/Cmd-X` copies or cuts it, `+` / `-` resizes it, `0` resets its size, `Enter` edits the image, and `Tab` manually edits the image size. Live Preview also gained a new visual selection state for image embeds and no longer preserves the editor line height to match image height when an image is the only content on a line.

— Exact keybindings for every action, fully actionable

Obsidian 1.13 Desktop (Public) (2026-07… Obsidian 1.13.3 Desktop (Early access) … Obsidian 1.13.2 Desktop (Early access) …

02 **Vim `:image` Ex command** NEW 90

New Vim Ex command `:image` lets users grow, shrink, or reset a selected image via `:image grow`, `:image shrink`, and `:image reset`, and accepts a count prefix such as `:5image grow` to repeat the action five times.

— Exact command syntax given, fully runnable Obsidian 1.13 Desktop (Public) (2026-07… Obsidian 1.13.2 Desktop (Early access) …

03 **Full-screen image lightbox** NEW 85

Clicking an image (or the new Zoom button added to image embeds in the editor) opens a full-screen image lightbox; users can pan by click-and-drag, navigate between all images embedded in the current file with arrow keys, and open it from Live Preview by selecting an image and pressing Space. On mobile, swipe down dismisses the zoomed image.

— Full interaction detail but no config surface named

Obsidian 1.13 Desktop (Public) (2026-07… Obsidian 1.13.3 Desktop (Early access) …

Obsidian 1.13.2 Desktop (Early access) … Obsidian 1.13.3 Mobile (Early access) (… Obsidian 1.13 Mobile (Public) (2026-07-…

04 **`--callout-color` now accepts any CSS color format** BREAKING

85

The `--callout-color` CSS variable now accepts any valid CSS color value (e.g. `#ff0000`, rgb(255,0,0)) instead of a bare RGB triplet (e.g. `255,0,0`); existing theme or snippet overrides using the old format will break.

– Names exact variable and value formats, migration implied

Obsidian 1.13 Desktop (Public) (2026-07…

05 **Searchable Settings panel** NEW

75

Settings now include a search box for finding options by name or description, navigable with `Ctrl/Cmd-F` to re-focus; search currently covers core settings and core plugins, and community plugins can migrate to the new Settings API to appear in results. Keyboard navigation lets users move between items with arrow keys and open a tab with `Enter`, including Vim key binding support. On mobile, a new searchable Settings panel was added for faster navigation across all settings.

– Names shortcut and API but no exact menu path Obsidian 1.13 Desktop (Public) (2026-07… Obsidian 1.13 Mobile (Public) (2026-07-…

06 **Settings open in a new window** NEW

75

Settings now open in a new window by default; this can be disabled via `Settings → Interface → Open settings in new window`.

– Exact settings path given, minimal mechanism detail

Obsidian 1.13 Desktop (Public) (2026-07…

07 **Sync sidebar view improvements** IMPROVED

75

The Sync sidebar view gains search, multi-selection (`Shift` `-arrow` keys or `Alt` `-click`), drag-into-other-views support, file-level context menu items, and `Delete` / `Backspace` to delete files.

– Names exact keybindings, no config surface Obsidian 1.13 Desktop (Public) (2026-07… Obsidian 1.13.2 Desktop (Early access) …

08 **iOS Share Sheet integration** NEW

75

New iOS Share Sheet lets users save content from Safari, YouTube, and other apps directly to Obsidian without opening the app; supports customizable Locations to define the vault, folder, and template used.

– Named mechanism and config options, no exact path

Obsidian 1.13 Mobile (Public) (2026-07-…

09 **Sized external images without alt text in reading mode** NEW

65

Reading mode now supports external images with a size but no alt text, e.g. ![200](https://www.path.to/image.png).

– Exact syntax example given, no further mechanism

Obsidian 1.13 Desktop (Public) (2026-07… Obsidian 1.13.2 Desktop (Early access) …

10 **Minimum iOS version raised to 15** BREAKING

65

The minimum supported iOS version was raised from iOS 14.5 to iOS 15.

– Exact version numbers given, clear impact Obsidian 1.13 Mobile (Public) (2026-07-…

11 **Drag Base links into File Explorer to move files** NEW

60

Adds support for dragging links from a Base into the File Explorer to move the file to a new location.

– Clear UI action, no further mechanism detail Obsidian 1.13 Desktop (Public) (2026-07… Obsidian 1.13.2 Desktop (Early access) …

THINNER COVERAGE BELOW

12 **Obsidian URI confirmation dialog** NEW

50

Obsidian URIs now show a confirmation dialog before firing, with a per-action allow list manageable from Settings.

– No exact settings path or allow-list mechanism detail

Obsidian 1.13 Desktop (Public) (2026-07…

13 **ConfirmationModal plugin API component** NEW

50

Adds a new `ConfirmationModal` component for plugin developers.

– Names the API component but no usage details Obsidian 1.13 Desktop (Public) (2026-07…

14 **Exit Restricted Mode without re-enabling plugins** NEW

50

Adds the ability to exit Restricted Mode without re-enabling plugins, for vault debugging.

– Clear purpose, no exact settings path Obsidian 1.13 Desktop (Public) (2026-07…

15 **File Recovery keyboard navigation** IMPROVED

50

The File Recovery view now supports keyboard navigation via arrow keys through file history records.

– Clear interaction, no further mechanism Obsidian 1.13 Desktop (Public) (2026-07…

Obsidian 1.13.3 Desktop (Early access) …

16 **Bases table column resize on mobile** NEW 50

Bases gains a new menu item to resize columns in table views.

— Names feature area but scant mechanism Obsidian 1.13 Mobile (Public) (2026-07-…

17 **Note Composer relative link rewriting on extraction** IMPROVED 45

When extracting a section with Note Composer, links in that section are now rewritten to be relative to the new file location.

— Describes behavior change, no exact command path

Obsidian 1.13 Desktop (Public) (2026-07…

18 **Folder drag-and-drop import preserving structure** IMPROVED 45

Dragging a folder into the desktop app now imports the entire folder and preserves its directory structure.

— Clear behavior, minimal further detail Obsidian 1.13 Desktop (Public) (2026-07…

19 **Tablet press-and-hold split resizing** NEW 45

On tablet, press-and-hold now resizes splits and pinned sidebars.

— Clear gesture named, minimal further detail Obsidian 1.13 Mobile (Public) (2026-07-…

20 **Consolidated Interface settings page** IMPROVED 40

New Interface settings page consolidates several settings that were previously located elsewhere.

— No list of which settings moved Obsidian 1.13 Desktop (Public) (2026-07…

22 **Drag-and-drop tab reordering** NEW 40

The Tab switcher now supports drag and drop to reorder tabs.

— Clear UI action, no further mechanism Obsidian 1.13 Mobile (Public) (2026-07-…

23 **One-time Mermaid rendering confirmation banner** NEW 35

Adds a new one-time banner confirming rendering of Mermaid code blocks in the vault.

— Bare notice description, no configuration detail

Obsidian 1.13 Desktop (Public) (2026-07…

24 **Network-drive HTML resource warning** NEW 30

Adds a warning before loading HTML resources when the source points to a network drive.

— Bare description of warning, no configuration detail

Obsidian 1.13 Desktop (Public) (2026-07…

25 **Sync plugin-sync warning dialog** NEW 30

Adds a warning dialog when choosing to sync plugins via Obsidian Sync.

— Bare description, no further mechanism given Obsidian 1.13 Desktop (Public) (2026-07…

26 **Bookmarks view search** NEW 30

Adds search to the Bookmarks view.

— Bare feature name only Obsidian 1.13 Desktop (Public) (2026-07…

└──▷ BREAKING ON UPGRADE

! The `--callout-color` CSS variable now expects a valid CSS color (e.g. `#ff0000` or rgb(255,0,0)) instead of a bare RGB triplet (e.g. `255,0,0`); existing theme/snippet overrides using the old format will break.

! Minimum iOS version raised from iOS 14.5 to iOS 15.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 5, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue