THE CYBER TOOLCHAIN

NO. 045

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 14, 2026 · EVERY DAY

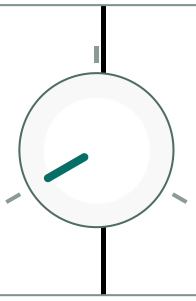
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 34 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

Prescient Security Cait

Feedback ⓘ

1 RELEASE · SEEN 2026-08-14 NOTES ↗

> Cait expanded its platform with Attack Surface Management and new asset testing capabilities in this release.

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

01 **Attack Surface Management added to Cait** NEW 15

Cait's platform now includes Attack Surface Management, extending discovery and monitoring capabilities beyond what was previously supported.

— No mechanism, scope, or interface detail given snapshot-20260814

02 **New asset testing capabilities** NEW 5

Cait introduces new asset testing capabilities within its service.

— Bare mention with no detail on scope or mechanism snapshot-20260814

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

dirsearch

Feedback ⓘ 1 RELEASE · 2026-08-14 NOTES ↗

Web path scanner

> dirsearch's v0.5.0 release introduces a Rust-native scanning and wordlist backend, category-based wordlists with template placeholders, ffuf-style match/filter controls, multi-format reporting with direct MySQL/PostgreSQL output, and JSON-based resumable sessions, while accompanying config.ini documentation details new scan-calibration keys and native-backend proxy/redirect support.

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

01 Native Rust backend for scanning and wordlists NEW

95

dirsearch adds a Rust-based native backend selectable via `--request-backend native` and `--wordlist-backend native` (also configurable as the `request-backend`, default `python`, and `wordlist-backend`, default `auto`, keys in `config.ini`), with `--sync` to fall back to the legacy threaded stack. The native backend now supports HTTP proxying, follows HTTP redirects during enumeration, bounds response body memory to prevent unbounded consumption on large scans, and ships with parity guardrails and CI integration to keep its behaviour aligned with the standard backend.

Run a scan using the Rust native backend for maximum throughput on a large wordlist.

```
$ dirsearch -u https://target -e php,html,js --request-backend native --wordlist-backend native -w /path/to/wordlist
```

— Names flags, config keys, defaults and native-backend mechanics. v0.5.0

02 Multi-format reporting with SQL database output NEW

90

`output-formats` in `config.ini` supports `simple`, `plain`, `json`, `xml`, `md`, `csv`, `html`, and `sqlite`, writable together in a single run; `output-file` accepts template variables `{extension}`, `{format}`, `{host}`, `{scheme}`, `{port}`, `{date}`, and `{datetime}` for dynamic naming. New `mysql-url` and `postgres-url` keys write scan results directly to MySQL or PostgreSQL, with the destination table controlled by `output-sql-table` (default `{scheme}_{host}:{port}`).

Write scan results to PostgreSQL and an HTML report simultaneously, using dynamic file naming based on host and date.

```ini
# config.ini
output-formats = html
output-file = reports/{host}/{date}.{extension}
postgres-url = postgres://user:password@localhost/dirsearch
```

— Names every output key, format and default value. v0.5.0

---

03 **Category wordlists and template placeholders** NEW 85

`--wordlist-categories` selects category-based bundled wordlists (PHP, Python, Java, .NET, Node, ColdFusion, and infrastructure), matched by the `wordlist-categories` key in `config.ini` (e.g. `common`, `conf`, `web`). `--wordlist-status` and `--wordlist-max-size` (config key `wordlist-max-size`, default `500000`) control size, and `wordlist-backend` (config default `auto`, values `auto` / `python` / `native`) selects the generation engine. Template wordlists support placeholders `%SUBJECT%`, `%CRUD_OP%`, `%AUTH_OP%`, `%DATE%`, `%API_VERSION%`, `%CATEGORY:name%`, and `%EXT%`.

Scan a target using category-based wordlists to focus on PHP and Node.js paths without maintaining a manual wordlist.

```
$ dirsearch -u https://target -e php,js --wordlist-categories
php,node
```

— Lists flags, config keys, defaults and placeholder syntax. v0.5.0

---

04 **ffuf-style match and filter controls** NEW 85

config.ini and the CLI gain ffuf/wfuzz-style opt-in matching via `matcher-mode` and `filter-mode` (value `or`), plus `match-status`, `filter-status`, `match-size`, `filter-size`, `match-words`, `filter-words`, `match-lines`, `filter-lines`, `match-regex`, `filter-regex`, `match-time`, and `filter-time` keys for fine-grained response filtering. Additional keys `filter-threshold`, `exclude-redirect`, `exclude-response`, and `exclude-texts` (e.g. `'Not found'`, `'404'`) filter by numeric threshold, redirect pattern, reference response file, or response body text.

Enable ffuf-style response matching to surface only responses with 100-2000 bytes and 10-100 words, filtering noise from large or empty replies.

```ini
# config.ini
matcher-mode = or
match-size = 100-2000
match-words = 10-100
filter-size = 0
```

— Enumerates every match/filter key with example usage. v0.5.0

05 **Scan control, calibration and logging options** NEW

85

config.ini adds `skip-on-status` (default `429`) to auto-skip rate-limited targets, `target-max-time` for a per-target time cap, and `auto-calibration` to force extra wildcard calibration up front. Recursive scans gain `recursion-status` (default `200-399,401,403`) and `max-recursion-depth` (default `0`); other new keys are `network-interface` (e.g. `eth0`) to bind scans to a NIC, `replay-proxy` to forward matches through a secondary proxy, and `log-file` / `log-file-size` for persistent capped logging. Release v0.5.0 adds the `--disable-cli` flag, Nmap XML target import, and `--verbose` output showing response elapsed time and content type.

Automatically skip rate-limited targets and cap total scan time per host to avoid stalling CI pipelines.

```ini
# config.ini
skip-on-status = 429
target-max-time = 120
auto-calibration = True
```

— Names every scan-control key, flag and default value. v0.5.0

06 **Database extras and installer changes** BREAKING

70

Optional MySQL and PostgreSQL report dependencies are now explicit extras — `dirsearch[mysql]`, `dirsearch[postgresql]`, and `dirsearch[db]` — instead of mandatory runtime dependencies, so existing installs relying on auto-installed DB drivers must reinstall with the appropriate extra. The old dependency installer path was removed.

— Gives exact extras names and required reinstall action. v0.5.0

07 **JSON-based resumable sessions** BREAKING

65

Resumable scan sessions now store JSON controller state instead of pickle, with `--list-sessions` to enumerate saved sessions and resume by ID. Legacy pickle sessions are no longer loaded directly and produce a warning on upgrade.

— Names the flag and format change but not migration steps. v0.5.0

---

08 **Python API for programmatic scanning** NEW 65

Adds a Python API with `FuzzerConfig`, `DirsearchFuzzer`, structured results, wordlist helpers, and template wordlist support, plus a `python -m dirsearch` entrypoint for running scans without the standalone CLI.

— Names the API classes but gives no usage example. v0.5.0

---

09 **Cross-platform builds and Docker images** NEW 65

PyInstaller release builds now cover Windows x64, Linux x64, Linux ARM64, macOS Intel, and macOS Apple Silicon across `async`, `threaded`, and `native-rust` stack variants. Linux x64 Docker images are published to GitHub Container Registry for each release stack, and portable CPython archives with compiled dependencies are offered as an alternative for users hit by PyInstaller antivirus false positives.

— Names platforms and stacks but no pull command. v0.5.0

THINNER COVERAGE BELOW

10 **Expanded bundled discovery data** IMPROVED 40

Bundled discovery wordlists are expanded with new framework paths, WordPress and WooCommerce entries, API/Ollama/OpenAI endpoints, and traversal-oriented entries.

— Names data categories added but no counts or files. v0.5.0

---

11 **Safer XML parsing against malicious input** IMPROVED 25

Hardens XML parsing with safe XML handling to protect against malicious input.

— Only states the hardening exists, no mechanism given. v0.5.0

└──▷ **BREAKING ON UPGRADE**

! Resumable sessions now use JSON controller state; legacy pickle sessions are no longer loaded directly and will produce a warning on upgrade.

! Optional DB dependencies (MySQL, PostgreSQL) are now explicit extras (`dirsearch[mysql]`, `dirsearch[postgresql]`, `dirsearch[db]`) rather than mandatory runtime dependencies — existing installs that relied on auto-installed DB drivers must reinstall with the appropriate extra.

! The old dependency installer path was removed.

WAS THIS USEFUL?  

Dalfox

Feedback

1 RELEASE · 2026-08-13 NOTES ↗ CODE ↗

Dalfox is a powerful open-source XSS scanner and utility focused on automation.

> Dalfox hardened its server mode with cross-site and DNS-rebinding protections, host/origin allow-lists and memory caps, introduced a full `dalfox payload` subcommand with JSON output and a wide set of named payload-category selectors, and shipped scanning accuracy improvements including modern JS framework detection, wider rate limiting, better cookie injection and clearer panic reporting.

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

01 **Server hardening against cross-site and DNS-rebinding attacks**

NEW 95

`dalfox server` now blocks browser-driven cross-site and DNS-rebinding requests by validating the Origin, Sec-Fetch-Site, and Host headers, rejecting `Sec-Fetch-Site: cross-site` or `same-site` requests with `403` while leaving non-browser clients (curl, CLI, CI) unaffected. The `--allowed-hosts <names>` flag lists extra hostnames accepted in the Host header beyond the bind host, `localhost`, and IP literals (for reverse-proxy setups), while `--allowed-origins` permits specific web UI origins through the cross-site gate without disabling CORS enforcement. The server also prints a startup warning when JSONP is enabled without an API key, since JSONP bypasses the cross-site gate and would let any site launch scans and read results, and `--max-retained-scans` caps how many finished scan results are retained to prevent unbounded memory growth.

Run Dalfox server behind a reverse proxy that forwards a public hostname, so the Host-header gate doesn't reject legitimate requests.

```
$ dalfox server --allowed-hosts
"dalfox.internal,scan.corp.example"
```

Allow a specific internal web UI to call the Dalfox server API while keeping the cross-site gate active for all other origins.

```
$ dalfox server --allowed-origins "https://app.example.com"
```

Prevent unbounded memory growth in long-running server mode by capping how many finished scan results are retained.

```
$ dalfox server --max-retained-scans 200
```

— Named flags, headers and status codes with runnable commands v3.2.1

02 **Payload subcommand with named category selectors** NEW

85

`dalfox payload` gains selectors for `payloadbox` (fetches remote payloads from PayloadBox), `all` (prints every local selector under `# name` headers, remote excluded), `event-handlers` (DOM event handler attributes like `onclick`, `onmouseover`), `useful-tags` (tags like `script`, `img`, `svg`), `uri-scheme` (scheme-based payloads like `javascript:`, expanded from 5 to 11 entries), `dom-clobbering`, `mxss` (mutation-XSS/sanitizer-bypass), `blind` (skeletons where `{}` is substituted with an OOB callback URL), a special-characters/encoded-variants selector (also expanded), a visibly-confirmable-sinks selector with filter-surviving variants, and a polished alert(document.domain) PoC selector for screenshots. The command also gained per-selector item counts in its summary output and closest-selector suggestions (Levenshtein distance ≤ 2) when an unknown selector name is given.

Quickly dump every local payload category in one shot to pipe into a fuzzer or review offline.

```
$ dalfox payload all
```

Generate blind-XSS skeletons pre-populated with your OOB callback to drop straight into a payload list.

```
$ dalfox payload blind
```

— Every selector named with commands demonstrating the pattern v3.2.1

03 **INTERNAL_ERROR status for scanner panics** IMPROVED

65

Introduces an `INTERNAL_ERROR` status code on `target_summary` entries so a scanner panic is reported as `skipped` rather than silently falling through to `clean`.

— Named field and status code with clear before/after behavior v3.2.1

THINNER COVERAGE BELOW

04 **JSON output for payload command** NEW

55

The `--json` flag prints selected payloads as a JSON array instead of one item per line, giving machine-readable payload lists out of `dalfox payload`.

— Named flag but no usage example given v3.2.1

05 **Rate limiting extended to discovery and mining phases**

IMPROVED 55

REST and MCP `rate_limit` now applies to the discovery and mining phases, not only scanning, so it throttles requests across the whole workflow.

— Names the config surface but gives no example v3.2.1

06 **Cookie parameters injected as cookies** IMPROVED 55

Cookie parameters are now injected as actual cookies rather than as same-named headers, raising cookie recall from 80.0% to 87.5%.

— Concrete before/after metric but no actionable flag v3.2.1

07 **Modern JavaScript framework detection** NEW 25

Dalfox now detects modern JavaScript frameworks during scanning.

— No mechanism, frameworks, or scope given v3.2.1

WAS THIS USEFUL?  

Acunetix 360

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

Acunetix 360 is a cloud-based web application security scanning platform that identifies vulnerabilities through automated dynamic and static analysis.

> Acunetix 360's v26.8.0 release extends scan data retention and adds support for .NET 10.

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

01 **Extended scan data retention to 730 days** IMPROVED 35

Scan data retention is now extended to 730 days.

— States new limit but no mechanism or config detail.

v26.8.0 – 13 August 2026 (2026-08-13)

02 **.NET 10 support added** NEW 30

Acunetix 360 adds support for .NET 10.

— Names the version but no scope or usage detail. v26.8.0 – 13 August 2026 (2026-08-13)

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Wiz Code

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

> Wiz Code extends secret detection beyond organizational boundaries by inventorying developers' personal public GitHub repositories, validating credential exploitability with an AI-powered Red Agent, and mapping full attack paths on the Security Graph.

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

01 **Red Agent exploitability validation for exposed credentials**

NEW 60

Wiz Red Agent, described as an AI-powered attacker, actively tests whether an exposed credential found in a personal repo is live and what it can reach, going beyond traditional validity checks.

– Names the Red Agent and its behavior but no invocation shown 2026-08-13-8ecb9077

02 **Identity-attached secret findings with commit and owner data**

NEW 60

Findings from personal repositories surface the exact commit, GitHub user, and email that pushed a secret, enabling immediate owner-assigned remediation without manual investigation.

– Names surfaced fields, giving a clear remediation starting point 2026-08-13-8ecb9077

THINNER COVERAGE BELOW

03 **Remediation workflow integration for personal-repo findings**

NEW 55

Personal-repo findings integrate into existing remediation workflows including Jira ticket creation, Slack developer notification, and tracked token rotation and revocation.

– Names Jira and Slack integrations but no config keys or endpoints 2026-08-13-8ecb9077

04 **Security Graph attack path for personal-repo secrets** NEW

50

Secrets found in personal repositories are enriched onto the Wiz Security Graph to show the full attack path and blast radius of a compromised credential, including downstream resource access.

– Names the Security Graph mechanism but no query or navigation given 2026-08-13-8ecb9077

05 **Personal GitHub repo inventory and identity correlation** NEW

Wiz Code adds automated inventory and correlation of developer identities to their personal public GitHub repositories outside org boundaries, closing the visibility gap where unmanaged secret exposure lives.

— Explains the gap it closes but no command or config to invoke it `2026-08-13-8ecb9077`

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ 1 RELEASE · 2026-08-13 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 Secure Registry adds Go module proxying (Beta), rounding out its policy-controlled proxy support alongside npm, PyPI, Maven, and NuGet, with a full set of Go-specific security controls and checksum handling.

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

01 **Go-specific policy controls: tags, cooldown, compromised packages, block list**

NEW 80

Released Tags Only rejects Go module resolutions at commit hashes, branch names, and pseudo-versions such as `v0.0.0-20241114170450-2d3c2a9cc518`, permitting only published release tags. Cooldown Period holds newly published Go module versions for a configurable window, with per-module exemptions scoped to the full module path. Compromised Packages blocks Go module versions flagged by StepSecurity's SOC or the wider security community. Custom Block List bans a Go module or version range org-wide using `module@version` form with glob support.

— Enumerates every control and its exact syntax Secure Registry Now Supports Go (2026-0…

02 **Proxy Checksum Database for Go modules** NEW 75

A new Proxy Checksum Database setting under Go Registry Settings relays `sum.golang.org` requests through Secure Registry, removing the need to disable `GOSUMDB` when runners cannot reach the default checksum database.

— Names exact setting and env var it replaces Secure Registry Now Supports Go (2026-0…

03 **Go module proxy support in Secure Registry** NEW 70

Secure Registry now proxies Go modules (Beta) at `https://registry.stepsecurity.io/go`, joining existing npm, PyPI, Maven, and NuGet proxying; every module request is evaluated against your Go policy before a response is returned. Integration paths for Go proxy support cover JFrog Artifactory, Sonatype Nexus Repository, and Direct.

⋮

— Names endpoint and integration paths but lacks setup steps

Secure Registry Now Supports Go (2026-0…

THINNER COVERAGE BELOW

04 **Policy Evaluations logging for Go requests** NEW 45

Every Go module request is logged in the Policy Evaluations log as Allowed, Modified, or Blocked, attributed to the developer machine or workflow run that made the request.

— Describes log categories but no navigation or query detail

Secure Registry Now Supports Go (2026-0…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Aviatrix CoPilot

Feedback ⓘ 1 RELEASE · SEEN 2026-08-14

NOTES ↗

Aviatrix CoPilot is an AI-powered assistant that helps operators manage, troubleshoot, and optimize cloud network infrastructure across multiple cloud providers.

> Aviatrix CoPilot's biggest addition this window is a new Cloud Routes diagnostics module for multicloud routing and gateway troubleshooting, alongside an Egress Security Score dashboard, a VPN Client 3.0.2 refresh (Universal macOS binary, FIPS 140-3/SAML support, and a breaking CA key-length requirement), and a batch of smaller CoPilot 4.36.0 additions spanning gateway management, network segmentation, Kubernetes visibility, notifications, and serverless permission gaps.

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

01 FIPS build requires RSA 2048-bit UserVPN CA

BREAKING

90

The FIPS build of VPN Client 3.0.2 requires the VPN gateway's UserVPN CA to use an RSA key of at least 2048 bits. Gateways on Controller versions earlier than 8.2, or Controller 8.2 gateways whose UserVPN CA has not been rotated since upgrade, may still use a 1024-bit CA (common on Controllers originally initialized on Aviatrix 7.1 or earlier) and will fail client certificate verification; operators must rotate the UserVPN CA to RSA 2048-bit before upgrading.

— Exact migration action and affected controller versions given product docs

02 Cloud Routes diagnostics and multicloud routing view

NEW

85

New Cloud Routes feature at `CoPilot > Diagnostics > Cloud Routes` gives a central view of routing across the Aviatrix transit network, multicloud and on-prem Site2Cloud connections, without logging into individual cloud consoles. It includes a Gateway Routes tab (tunnel info, route state up/down, routing tables, filterable by gateway name, subnet, or IP), a VPC/VNet Routes tab (routing tables per VPC/VNet/VCN, filterable by route table name, subnet, or IP), an External Connections tab (data center tunnel status, connected gateway, remote IP, tunnel type, linked to Topology and FlowIQ), and a BGP Info tab (advertised/learned routes, connection map with ASN and IP). A gateway diagnostics panel adds PING, TRACEROUTE, TRACEPATH, TRACELOG (uploads directly to Aviatrix Support), Test Connectivity (TCP/UDP to a host and port), ACTIVE SESSIONS (filterable by port), INTERFACE STATS, PACKET CAPTURE (filter by

host, port, time, address, protocol, flags; download as pcap for Wireshark), and SERVICE ACTIONS (check/restart gateway services).

— Exact UI path and every diagnostic tool named, but no runnable command product docs

03 **Egress Security Score dashboard** NEW 80

New `Security > Egress > Overview` tab (requires Controller 8.0 and DCF feature enabled) shows an Egress Security Score card with a score-calculation explainer dialog and a Manage VPC/VNets link to the Egress VPC/VNets tab, plus a Protected VPC/VNets card breaking cloud-account VPC/VNets into Unprotected, Monitored, Partially Protected, Protected, No Egress, and Unknown states.

— Exact UI path, prerequisites, and state categories named product docs

04 **FIPS 140-3 encryption and SAML authentication in VPN Client** NEW 60

FIPS 140-3 compliant encryption is now supported in the VPN Client 3.0.2, with SAML-based authentication (Okta, Duo, and other SAML identity providers) available across Windows, macOS, and Ubuntu.

— Names providers and platforms but no configuration steps product docs

THINNER COVERAGE BELOW

05 **VPN Client 3.0.2 Universal macOS binary and Ubuntu 26.04 support** NEW 55

The VPN Client 3.0.2 macOS installer is now a Universal binary running natively on both Intel and Apple Silicon (ARM64), eliminating the Rosetta 2 dependency. Ubuntu 26.04 is now supported alongside existing Ubuntu 22.04 and 24.04.

— Names exact platforms and versions, no setup steps needed product docs

06 **Missing permissions visibility for serverless functions** NEW 40

New Missing Permissions Visibility for Serverless Functions surfaces IAM/permission gaps for serverless workloads directly in CoPilot.

— Names the gap it surfaces but not the UI location New and Enhanced Features in 4.36.0

07 **VRF support in Network Segmentation** NEW 35

CoPilot 4.36.0 adds VRF support in Network Segmentation, enabling VRF-aware policy across segmented networks.

— States capability but no configuration path or example

New and Enhanced Features in 4.36.0

08 **Kubernetes intra-cluster orchestration status view** NEW

35

New Kubernetes Intra-Cluster Orchestration Status view surfaces the health and state of intra-cluster traffic orchestration.

— Describes purpose only, no UI path or detail New and Enhanced Features in 4.36.0

09 **Gateway management controls for AZ affinity and sizing** NEW

30

CoPilot 4.36.0 adds AZ Affinity, Symmetric Routing, and T-Shirt Instance Sizing controls for Gateway Management.

— Names three controls but no mechanism or location New and Enhanced Features in 4.36.0

10 **FIPS 140-3 compliance indicators in CoPilot UI** NEW

30

CoPilot 4.36.0 adds FIPS 140-3 compliance indicators in the CoPilot UI.

— Names the standard but not where indicators appear New and Enhanced Features in 4.36.0

11 **Ubuntu 20.04 support deprecated in VPN Client** DEPRECATED

25

Ubuntu 20.04 support is deprecated as of VPN Client 3.0.2.

— Bare deprecation notice with no timeline or replacement product docs

12 **Updated Controller bell notifications** IMPROVED

25

Updated Controller Bell Notifications bring alerting improvements to the CoPilot notification center.

— Vague 'improvements' with no specifics given New and Enhanced Features in 4.36.0

└──▷ ALSO FROM THESE RELEASES

└──▷ BREAKING ON UPGRADE

! The FIPS build of VPN Client 3.0.2 requires the VPN gateway's UserVPN CA to use an RSA key of at least 2048 bits — gateways on Controller versions earlier than 8.2, or Controller 8.2 gateways whose UserVPN CA has not been rotated since upgrade, may use a 1024-bit CA (common on Controllers originally initialized on Aviatrix 7.1 or earlier) and will fail client certificate verification. Rotate the UserVPN CA to RSA 2048-bit before upgrading.

! Ubuntu 20.04 support is deprecated in VPN Client 3.0.2.

WAS THIS USEFUL?  

Fastly

Feedback ⓘ CLI v16.0.0 (2026-08-13) BLOG ↗

Fastly is a content delivery network and edge computing platform that accelerates website performance and provides security services globally.

Fastly CLI v16.0.0 now offers all available starter kits for each language during compute init.

WAS THIS USEFUL?  

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ 1 RELEASE · SEEN 2026-08-14 NOTES ↗

> Iru launched a Windows system tray app that lets end users install pending app updates themselves without IT involvement.

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

01 **Windows system tray update app** NEW 55

A new Iru system tray app on Windows displays all pending Auto App and Custom App updates, letting end users install one at a time or trigger Update all without IT involvement. The update list automatically clears completed entries as installs succeed, and the app handles closing open apps when required to finish an install.



— Describes behaviour well but no config keys or flags named.

Update Windows apps from the system tray

WAS THIS USEFUL? 👍 👎

Trend Vision One

Feedback ⓘ 1 RELEASE · 2026-08-13 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 now lets administrators reset MFA for individual user accounts directly from the console, removing the need to file a support case.

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

01 **Admin-level MFA reset for user accounts** NEW 72

Administrators can reset MFA for individual user accounts using the Admin-level MFA reset option on the User Accounts page, eliminating the need to open a support case when a user loses their MFA device or recovery code. After the reset, the user is prompted to set up MFA again at their next sign-in.

When a user is locked out after losing their MFA device, reset their MFA without a support case so they can re-enroll at next sign-in.

📍 In the Trend Vision One console, go to Administration → User Accounts, locate the target account, click the options icon, then click Admin-level MFA reset. The user will be prompted to set up MFA at their next sign-in.

— Names exact UI option and console path but no scope/limits.

Administrators can now reset MFA for us…

WAS THIS USEFUL? 👍 👎

Microsoft Sysinternals

Feedback

1 RELEASE · 2026-08-12

BLOG ↗

Microsoft Sysinternals is a suite of advanced system utilities for Windows diagnostics, monitoring, and troubleshooting at the kernel level.

> Microsoft Sysinternals shipped Process Tree PID filtering in Process Monitor, Azure Virtual Desktop and Dev Box support in RDCMan, and a new DemoMirror feature in ZoomIt, alongside smaller CPU-detection updates to CoreInfo and Process Explorer and a full Rust rewrite of NotMyFault.

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

01 **Process Tree PID filtering in Process Monitor** NEW

50

Process Monitor adds Process Tree PID filtering to match events based on any ancestor process ID in the process tree, useful for tracing all activity spawned by a parent process.

— Explains mechanism but no exact filter syntax given

NotMyFault, RDCMan, ZoomIt, CoreInfo, P…

02 **AVD and Dev Box support in RDCMan** NEW

45

RDCMan v3.20 adds Azure Virtual Desktop workspace feed support and Microsoft Dev Box support.

— Names the two supported services, no setup steps

NotMyFault, RDCMan, ZoomIt, CoreInfo, P…

03 **DemoMirror in ZoomIt** NEW

45

ZoomIt v12.2 introduces DemoMirror for mirroring the screen, a screen region, or a window onto a second monitor.

— Describes function, no configuration detail included

NotMyFault, RDCMan, ZoomIt, CoreInfo, P…

04 **Intel LP-E core support in CoreInfo** IMPROVED

35

CoreInfo v4.02 adds Intel LP-E cores support and improves CPU vendor name detection.

— Names two changes, no further mechanism detail NotMyFault, RDCMan, ZoomIt, CoreInfo, P…

05 **NotMyFault rewritten in Rust** IMPROVED

35

NotMyFault v4.6 is rewritten in Rust and ships with a modern UI.

— States rewrite and UI change, no specifics on new behavior

NotMyFault, RDCMan, ZoomIt, CoreInfo, P…

06 **Intel core type info in Process Explorer** NEW 30

Process Explorer v17.13 adds core type information for Intel CPUs.

— Bare mention of new CPU info field NotMyFault, RDCMan, ZoomIt, CoreInfo, P…

WAS THIS USEFUL? 👍 👎

◆ **Security Data Pipeline**

Cribl

Feedback ⓘ 1 RELEASE · 2026-08-14 BLOG ↗

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

> Cribl launched an Apps platform that lets users access pre-built tools and build custom data experiences on top of existing Cribl infrastructure using an AI-driven conversational interface.

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

01 **Apps platform for custom data experiences** NEW 60

A new Apps button appears in the upper-right corner of every Cribl screen, giving access to pre-built apps that add new views, processes, workflows, and capabilities built around existing Cribl data. Users can also build custom apps through an AI-driven conversational flow, while Cribl handles the underlying hosting, identity, RBAC, and data access/governance layers.

— Names UI entry point and AI flow but no config or API detail.

Cribl Apps: Do the fun part of developm…

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

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

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

> Tines published its first public REST API (369 endpoints across 71 areas) and rolled out a wave of platform additions: Tines 3B workflows as case actions, AI Agent skill extensions, a self-hosted LDAP configuration UI, Workbench credential restrictions, customizable AI credit alerts, per-workflow savings estimation, and improved change-control reviews.

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

01 **Public REST API across 71 areas** NEW 85

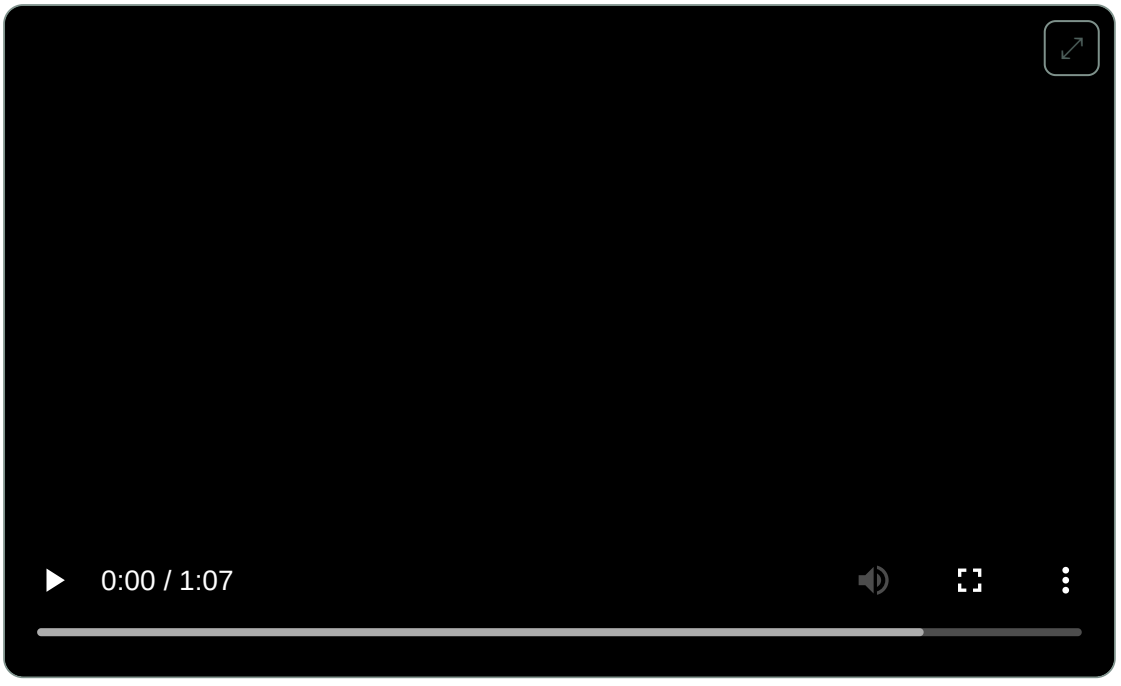
Tines published a full REST API covering 369 endpoints across 71 areas, each supporting create, read, update and delete operations. The largest areas are Actions (22 endpoints), `Stories > Notes` (12), Stories (11), Teams (v1) (11), `Admin > Users` (10), Global Resources (v1) (10), `Cases (v1) > Case Templates` (9) — which specifically enables managing case templates programmatically — and `Admin > Sync Destinations` (8), plus 63 additional areas spanning Cases, Records, Admin, Scim, Audit Logs, Workbench, Reporting, Folders (v1), Run Script Environments and more.

— Full endpoint areas and counts given, but no example calls 0.0.1

02 **Tines 3B workflows as case actions** NEW 65

Cases can now trigger published Tines 3B workflows directly as case actions. After configuring a 3B service account connection, users can select a published 3B workflow when adding a case action.





— Clear setup steps and UI flow named `Add Tines 3B workflows as case actions …`

THINNER COVERAGE BELOW

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

Adds a Self-hosted LDAP configuration UI for self-hosted deployments, giving admins a dedicated interface to set up LDAP.

— UI location named, no field-level detail `product docs`

04 **Workbench credential restrictions** NEW 45

Adds an admin control to Control which credentials can be used in Workbench, restricting which credentials are available within the Workbench AI environment.

— Admin control named but no exact config path product docs

05 **Customizable AI credit usage alerts** NEW 45

Lets teams Customize AI credit usage alerts by setting thresholds for AI credit consumption notifications.

— Threshold concept named, no exact settings path product docs

06 **AI Agent skill extensions** NEW 40

Introduces Extend AI Agents with skills, a new skill extension surface for AI Agents accessible via Actions and Storyboard.

— Names the surface but not skill mechanics product docs

07 **Per-workflow savings estimation** NEW 35

Adds configuration for detailed savings estimates on a per-workflow basis, letting teams quantify the value delivered by each automated workflow.



— No config key or UI path named product docs

08 **Improved group change reviews** IMPROVED 30

Introduces Improved group change reviews under Change Control for reviewing group-level changes to stories.

— Bare mention, no mechanism or UI detail product docs

WAS THIS USEFUL?  

◆ SIEM

Sumo Logic

Feedback ⓘ 2 RELEASES · 2026-08-13 BLOG ↗

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

> Sumo Logic shipped a new public API for managing AWS S3 archive destinations, a Mobot skill that builds dashboard panels from plain-language prompts, and a batch of schema and parameter updates across its otCollectors, plan, and sourceTemplate endpoints.

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

01 **otCollectors, plan and sourceTemplate API updates** IMPROVED 71

Response schemas changed for `GET /v1/otCollectors/otCollectorsByName`, `GET /v1/otCollectors/{id}`, `GET /v1/plan/pendingUpdateRequest`, and `POST /v1/otCollectors`. A new `dryRun` query parameter was added to `POST /v1/sourceTemplate` and `POST /v1/sourceTemplates`.

— Names exact endpoints and a new parameter, but not what the schema change is. 1.0.0

THINNER COVERAGE BELOW

02 **Archive Destinations API for AWS S3** NEW 55

A new public API enables creating and managing AWS S3 archive destinations for Installed Collectors programmatically, complementing the existing UI flow for at-scale automation.

— Describes purpose and scope but no endpoint path or method given.

Create and Manage Archive Destinations …

03 **Mobot dashboard panel creation from prompts** NEW 55

A new Mobot skill lets practitioners describe a desired panel in plain language, after which Mobot drafts the query, selects a visualization, and adds the panel to a new or existing dashboard.

— Clear mechanism described but no exact UI path or command given.

Create Dashboard Panels with Mobot (Das…

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Attack Surface & Exposure Management

Intruder

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

Intruder is a cloud-based vulnerability scanner that automatically discovers and tests systems for security weaknesses.

> Intruder redesigned its side navigation, adding pinnable favourites, activity indicators and single-click access to every page.

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

01 **Reorganised side navigation** IMPROVED 55

The sidebar now groups pages into expandable theme sections, replacing the previous nested-tab structure, and every page is reachable in a single click. A pinnable Favourites shortcut bar lets users surface their most-visited pages at the top of the nav, and icon and pill indicators on nav items surface new data or activity at a glance.

— Describes UI changes and mechanism but no config/API names.

Reorganised side navigation (2026-08-13)

WAS THIS USEFUL? 👍 👎

runZero

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

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

> runZero's 5.1 release adds AWS Inspector vulnerability findings to its AWS integration and overhauls the reports page with pagination, search, and delete, while also refining how the AWS integration handles closed linked accounts.

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

01 **Paginated, searchable reports page** IMPROVED 48

The reports page now displays a paginated list of all analysis task-based reports and supports searching and deleting reports directly from that page.

— Describes UI capabilities with a clear navigation starting point

runZero 5.1.260813.0 (2026-08-13)

02 **AWS Inspector findings for EC2 and Lambda** NEW 43

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

— Names the feature and asset types but no mechanism detail

runZero 5.1.260813.0 (2026-08-13)

03 **AWS integration skips closed linked accounts** IMPROVED 27

The AWS integration now skips linked organizational accounts that are closed, avoiding scans against inactive accounts.

— States the behavior change but no further mechanism or config detail

runZero 5.1.260813.0 (2026-08-13)

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Twingate

Feedback ⓘ 1 RELEASE · SEEN 2026-08-14 NOTES ↗

Twingate is a zero-trust network access platform that securely connects remote users and devices to private resources without VPN.

> Twingate rolled out granular role-based access control for its Admin Console, introducing six scoped admin roles with visible role badges and enforced page-level restrictions.

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

01 **Scoped admin roles in Admin Console** NEW 79

Twingate introduces six admin roles — DevOps, Helpdesk, Support, Access Reviewer, Billing, and full admin — each scoped to specific Admin Console pages and actions. DevOps admins are limited to the Network tab (Resources, Connectors, Remote Networks, Group-to-Resource assignments, Access Requests); Support admins get all pages except `Internet Security > Secure DNS` and `Internet Security > Client Configuration`; Billing admins are restricted to the `Settings > Billing` page for plan changes and payment details. The Helpdesk role gives read-only access to selected areas plus targeted actions — resetting MFA, verifying/unverifying devices, and managing trusted serial numbers — so support staff can troubleshoot without full admin rights. DevOps, Helpdesk, Support, Access Reviewer, and Billing admins see a role badge next to the Twingate logo (hover reveals access details), and any attempt to edit settings outside a role's scope is blocked with an inline error message.

. . . . .

— Names exact console pages and role behaviors but only UI navigation, no API/CLI.

Helpdesk Admin Role

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Didit

Feedback ⓘ 1 RELEASE · SEEN 2026-08-14 NOTES ↗

Didit is a decentralized identity verification platform enabling secure, privacy-preserving credential management and verification across applications.

> Didit's biggest changes this window are a new endpoint for skipping or re-requiring frozen KYB Key People parties and dedicated 502 error codes for KYB registry and phone-verification outages, alongside workflow branching on marital status, richer session API responses (barcode data, liveness method), capture-flow, country-selection, React Native SDK, and prepaid-credit pricing updates.

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

01 Skip or re-require frozen KYB Key People parties NEW

93

New `PATCH /v3/session/{sessionId}/kyb/parties/{partyUuid}/` endpoint lets operators set `is_skipped: true` to exempt a frozen KYB Key People party from verification — immediately recomputing `KYBKeyPeopleCheck` and the parent KYB session status — or `is_skipped: false` to restore it to scope; requires the `write:businesses` permission via API key or console user token. The endpoint also surfaces new party-level status signals `kyc_session_deleted` and `kyb_sub_session_deleted` to explain why a party is frozen.

Unblock a KYB session stalled in `Awaiting User` because a UBO's individual KYC session was deleted — skip that party so the business session can settle.

```
$ curl -X PATCH
'https://api.didit.me/v3/session/sess_abc123/kyb/parties/0a3f1e2c
-7c2f-4a1b-9d55-2f8d2a5d6b41/' \
  -H 'Authorization: Bearer <api_key>' \
  -H 'Content-Type: application/json' \
  -d '{"is_skipped": true}'
```

Restore a previously skipped Key People party to scope after the individual re-submits their KYC session.

```
$ curl -X PATCH
'https://api.didit.me/v3/session/sess_abc123/kyb/parties/0a3f1e2c
-7c2f-4a1b-9d55-2f8d2a5d6b41/' \
  -H 'Authorization: Bearer <api_key>' \
  -H 'Content-Type: application/json' \
  -d '{"is_skipped": false}'
```

— Full endpoint, fields, permission and runnable curl examples given `3.0.0`

02

502 error codes for KYB registry and phone-provider outages

NEW

83

New `kyb_registry_provider_unavailable` error code (with `code`, `summary`, and `detail` fields) is returned as HTTP `502` on `POST /v3/kyb/search/` when the KYB registry provider is temporarily unavailable — the search is not billed, and Didit's hosted flow plus Web, iOS, and Android verification flows automatically offer manual company entry instead. A parallel `phone_provider_unavailable` error code is returned as HTTP `502` on `POST /v3/phone/send/` and `POST /v3/phone/check/` when Prelude is unavailable; nothing is billed on send and the attempt budget is only consumed by completed checks.

Handle a KYB registry outage gracefully by detecting the new error code and prompting the user to enter company details manually.

```python
if response.status_code == 502 and response.json().get('code') ==
'kyb_registry_provider_unavailable':
    prompt_manual_company_entry()
```

— Named endpoints, error codes and billing behaviour with a code sample `3.0.0` `snapshot-20260814`

03

Barcode data and liveness method in session responses NEW

73

Retrieve-session responses now include `id_verifications[].barcodes`, with `type`, `position`, `data`, `data_raw`, and `side` fields per entry (empty array if the document has no barcode output), and face-step payloads now always include `face_liveness_method`, defaulting to `PASSIVE` when a workflow doesn't set it explicitly.

Inspect barcode data extracted from a scanned ID document after a verification session completes.

```
$ GET /sessions/{session_id}

// Response excerpt:
// {
//   "id_verifications": [{
//     "barcodes": [
//       { "type": "PDF417", "side": "back", "data": "...",
"data_raw": "...", "position": {...} }
//     ]
//   }]
// }
```

— Names every response field with an example response payload `snapshot-20260814`

04 **Marital-status branching in workflow builder** NEW 60

`kyc.marital_status` is now a usable OCR field in conditional branches and custom status rules within the workflow builder; when the source document omits marital status the field is empty and the fallback branch applies.

Branch a workflow on marital status to apply different KYC rules for specific applicant profiles.

```json
{
  "condition": "kyc.marital_status == 'married'",
  "then": "enhanced_due_diligence_step"
}
```

— Named field and fallback behaviour with a sample branch condition `snapshot-20260814`

THINNER COVERAGE BELOW

05 **Copilot-driven Database Validation step setup** NEW 45

The workflow copilot can create and configure a Database Validation step in a single action, resolving the supported validation service per country from a list you specify and displaying the saved configuration on the canvas.

— Describes mechanism but no exact UI path or command `snapshot-20260814`

06 **React Native SDK codegen and dSYM support** IMPROVED 45

The React Native TurboModule specification now supports React Native 0.77 code generation, and the React Native SDK (4.5.4) adds iOS dSYM artifacts.

– Names versions and artifact type but no usage step snapshot-20260814

07 **Prepaid credit discounts replace bonus credits** BREAKING

40

Prepaid credits now use charge discounts instead of bonus credits: 4% off at $250 and 7% off at $1,000.

– Exact thresholds and percentages given, no action needed snapshot-20260814

08 **Document capture improvements on Web, iOS, Android** NEW

35

French 2021 ID cards can now be read and validated via the ANTS 2D-Doc Data Matrix on the back of the document across Web, iOS, and Android, and capture flows on all three platforms can now continue with a still photo when video recording is unavailable.

– Describes two capabilities but no exact trigger or config snapshot-20260814

09 **Continent grouping for compliance country selection** IMPROVED

25

Compliance country questions in the console now group countries by continent and support selecting or clearing an entire continent at once.

– Bare UI description with no navigation detail snapshot-20260814

WAS THIS USEFUL?  

Nudge Security

Feedback

2 RELEASES · SEEN 2026-08-14

NOTES ↗

Nudge Security monitors SaaS application usage and security risks across organizations to reduce cloud infrastructure threats.

> Nudge Security shipped browser-level protection against employees sharing sensitive data with AI tools, plus a new dashboard for tracking passkey adoption and storage across the organization.

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

01 Passkey usage dashboard NEW

83

A new `Dashboards > Passkey usage` dashboard shows org-wide passkey adoption rate, top apps used with passkeys, and a per-provider storage breakdown with a toggle between number of passkeys and number of users. It also surfaces device-loss survivability and login-verification status for each detected passkey, and adds per-employee passkey usage detail on each user's user details page.

Audit where your org's passkeys are stored and whether they survive device loss — useful before a hardware refresh or incident review.

📍 In the Nudge Security console, go to Dashboards › Passkey usage. Use the provider breakdown toggle to switch between number of passkeys and number of users, then review the device-loss survivability and user-verification columns.

— Names exact dashboard path, toggle, and data fields shown

Track passkey adoption and storage in t…

02 In-browser sensitive data nudges for AI conversations NEW

80

The browser extension now detects PII, secrets, or financial info typed into AI conversations and prompts employees to redact in one click or justify sending, configurable at `Settings > In-browser nudges`. An optional enforcement mode, configured at the same settings path, makes redaction the only way to send, preventing unredacted sensitive data from leaving the browser.

Enable sensitive data nudges with mandatory redaction so employees can never paste unredacted PII or secrets into AI tools.

📍 In the console, go to Settings › In-browser nudges, enable sensitive data nudges, then turn on enforcement so redaction is required before sending.

— Names exact settings path and both modes; missing detection scope limits

Nudge employees in the browser before t…

WAS THIS USEFUL?  

BeyondTrust

Feedback

1 RELEASE · 2026-08-13

BLOG ↗

BeyondTrust provides privileged access management and endpoint security solutions for controlling and monitoring user access to critical systems.

> BeyondTrust's Endpoint Credential Manager 1.6.2606 adds configuration self-heal and environment-variable support for ConfigCLI, alongside a cluster of security hardening changes — per-installation credential-store encryption, enforced certificate revocation checking, and fail-closed log appenders — several of which require action before upgrading.

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

01 **Per-installation credential store encryption** BREAKING

85

Switches the credential store to per-installation encryption entropy, replacing a previously shared value; a new `ECSM.entropy` file is created in the install directory and should be backed up alongside `ECSM.settings`.

— Names exact new file and backup requirement before upgrading

Endpoint Credential Manager 1.6.2606 re…

02 **Environment-variable resolution in ConfigCLI** NEW

75

Adds `env:` prefix support in `ConfigCLI` arguments so sensitive values can be resolved from environment variables instead of the command line, keeping them out of process creation logs.

— Names exact `env:` prefix and mechanism; usable in ConfigCLI today

Endpoint Credential Manager 1.6.2606 re…

03 **Enforced certificate revocation checking on appliance connection**

75

BREAKING

Enables certificate revocation checking on the appliance connection and enforces size limits on protocol and HTTP reads; CRL/OCSP endpoints must be reachable from the ECM host before upgrading.

— Gives concrete pre-upgrade check but no specific limit values

Endpoint Credential Manager 1.6.2606 re…

04 **Configuration self-heal on startup and install** NEW

70

The ECM service and Configurator automatically detect and repair a missing or stale configuration file at startup and during installation, preserving `appSettings`, proxy, and log configuration while refreshing only version-driven sections.

— Clear mechanism and scope, but behavior is automatic, not user-triggered

Endpoint Credential Manager 1.6.2606 re…

05 **Log appenders fail closed on credential values** BREAKING 70

Log appenders now fail closed rather than falling back to an unprotected credential value; plaintext credential values in log appender configuration must be re-encrypted before upgrading.

— Clear before/after and required pre-upgrade migration step

Endpoint Credential Manager 1.6.2606 re…

06 **Plugin loading diagnostics and hardening** IMPROVED 65

Plugin load diagnostics now log assembly load failures with full detail, replacing the opaque `Found 0 plugins` message seen when tools such as WDAC, AppLocker, or EDR solutions block plugin loading; plugin loading is also hardened to load from the install directory rather than the process working directory.

— Names affected tools and message change; no direct reader action

Endpoint Credential Manager 1.6.2606 re…

THINNER COVERAGE BELOW

07 **log4net dependency upgrade** IMPROVED 35

Upgrades the `log4net` dependency from 3.1.0 to 3.3.1.

— Exact version numbers given but no behavioral detail

Endpoint Credential Manager 1.6.2606 re…

08 **Restricted configuration deserialization** IMPROVED 15

Restricts configuration deserialization to known types.

— One-line hardening note with no named types or scope

Endpoint Credential Manager 1.6.2606 re…

└──▷ BREAKING ON UPGRADE

! Certificate revocation checking is now enforced on the appliance connection — CRL/OCSP endpoints must be reachable from the ECM host before upgrading.

! Plaintext credential values in log appender configuration must be re-encrypted before upgrading, as log appenders now fail closed rather than falling back to an unprotected credential value.

! The credential store now uses per-installation encryption entropy; a new `ECSM.entropy` file is created in the install directory and must be backed up alongside `ECSM.settings`.

WAS THIS USEFUL?  

Okta Identity Engine

Feedback ⓘ 1 RELEASE · 2026-08-12

BLOG ↗

Okta Identity Engine is a cloud-based identity platform providing authentication, authorization, and user management for applications.

> Okta Identity Engine's biggest addition this window is multi-resource refresh tokens for custom authorization servers, alongside new granular admin permissions for user behavior detection, expanded dynamic zone IP categories, and GA promotions for a list-endpoint field filter and an identity verification integration.

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

01 Multi-resource refresh tokens on custom authorization servers

NEW 95

`/authorize` endpoints on custom authorization servers now accept multiple `resource` parameters, letting a single refresh token cover multiple resources; each `grant_type=refresh_token` exchange then mints a precisely scoped, single-`aud` access token for one resource, eliminating the need for one refresh token per resource. This is GA in Preview and requires an Okta for AI Agents subscription with agent-to-agent connections enabled.

Request a refresh token scoped to multiple resources at authorization time, then exchange it for a precisely scoped access token for one resource — eliminating the need for one refresh token per resource.

```
$ # Step 1: Authorize with multiple resource parameters
GET /oauth2/{authorizationServerId}/v1/authorize?client_id=
<client_id>&response_type=code&scope=openid
offline_access&resource=https://api.example.com/service-
a&resource=https://api.example.com/service-b&redirect_uri=
<redirect_uri>

# Step 2: Exchange the refresh token for a single-audience access
token
POST /oauth2/{authorizationServerId}/v1/token
  grant_type=refresh_token
  &refresh_token=<refresh_token>
  &resource=https://api.example.com/service-a
  &client_id=<client_id>
```

— Full mechanism, exact endpoints, params, and a runnable example. product docs

02 **Granular user behavior admin permissions** NEW 60

Adds `okta.users.behaviors.manage`, `okta.users.behaviors.read`, and `okta.users.behaviors.reset` custom admin permissions for fine-grained control over user behavior detection.

— Names all three permission strings but no usage example. product docs

THINNER COVERAGE BELOW

03 **`fields` query parameter for list endpoints** IMPROVED 40

New `fields` query parameter for Groups, Realms, and Devices list endpoints is now GA in Production.

— Names endpoints and parameter but gives no usage example. product docs

04 **New IP service categories for dynamic zones** IMPROVED 30

Adds new IP service categories for enhanced dynamic zones, expanding the criteria and classification options available when building network-based policies and defining dynamic network zones.

— Repeated across three releases but never names the actual categories.

Weekly release 2026.08.1 (2026-08-12)

05 **Submit Identity Verification integration GA** IMPROVED 23

Submit Identity Verification integration is now GA in Production.

— Only a bare GA announcement with no detail on mechanism. product docs

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

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

RegScale is a compliance and regulatory management platform that automates governance, risk, and compliance (GRC) workflows for organizations.

> RegScale's CLI 6.40.0 release adds a configurable artifact output directory and an optional Wiz data-retention reconciliation feature that auto-closes stale issues.

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

01 **Wiz data-retention reconciliation** NEW 69

Adds optional reconciliation that closes RegScale issues whose findings have aged out of Wiz's retention window, enabled via `wizRetentionReconciliation` and tuned with additional `w…` settings.

Enable Wiz retention reconciliation to automatically close RegScale issues whose Wiz findings have aged out of Wiz's retention window.

yaml

```yaml
wizRetentionReconciliation: true
```

— Explains behavior and config key but additional tuning settings are unnamed.

CLI (2026-08-13)

02 **Configurable artifact output directory** NEW 65

Adds the `artifactsDir` setting to configure the base directory where scanner and integration output files are written, replacing the default location.

Set a custom output directory so all scanner and integration artifacts land in a dedicated path instead of the default location.

yaml

```yaml
artifactsDir: /var/regscale/artifacts
```

— Names exact config key and example value, no further mechanism given. CLI (2026-08-13)

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Daytona

Feedback ⓘ 1 RELEASE · SEEN 2026-08-14 NOTES ↗

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

> Daytona expanded its API with organization-level SSO/identity-provider management, new audit and runner-snapshot endpoints, and force/filter query parameters, alongside a broad wave of schema refinements across sandbox, organization, admin, and API-key endpoints; a separate 0.200.0 release enforced TLS verification and made SDK/CLI security tighter.

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

01 **SSO identity provider and pending-link management endpoints**

NEW 88

Adds organization-level SSO/identity-provider CRUD and testing via `POST /organizations/{organizationId}/identity-providers`, `GET /organizations/{organizationId}/identity-providers`, `GET /organizations/{organizationId}/identity-providers/{id}`, `PATCH /organizations/{organizationId}/identity-providers/{id}`, `DELETE /organizations/{organizationId}/identity-providers/{id}`, and `POST /organizations/{organizationId}/identity-providers/test-connection`, plus `POST /organizations/{organizationId}/sso-enabled` to toggle SSO. User-facing pending SSO link handling is added via `GET /users/me/pending-sso-links`, `DELETE /users/me/pending-sso-links/{id}`, and `POST /users/me/pending-sso-links/{id}/confirm`.

— Ten concrete endpoints named with clear CRUD purpose 1.0

02 **Force-delete and snapshot source filtering query parameters**

NEW 80

Adds a `force` query parameter to `DELETE /admin/runners/{id}` and `DELETE /runners/{id}`, and a `sourceSandboxId` query parameter to `GET /snapshots` for filtering snapshots by their originating sandbox.

— Exact params and endpoints named, directly usable 1.0

THINNER COVERAGE BELOW

03 **Configurable HTTP timeout in Go SDK** NEW 45

Makes the HTTP timeout configurable in the Go SDK to honor long execution timeouts.

— States purpose but no config key or method name given snapshot-20260814

04 **Restricted permissions on CLI config files** IMPROVED

45

Creates CLI config files with `0600` permissions for tighter access control.

— Concrete permission value named but no file path given `snapshot-20260814`

05 **Audit scenarios endpoint** NEW

41

New `GET /audit/scenarios` endpoint added to the API.

— Endpoint named but purpose and response undescribed `1.0`

06 **Runner-scoped snapshots listing endpoint** NEW

41

New `GET /runners/me/snapshots` endpoint added to list snapshots for the calling runner.

— Endpoint named but response detail unspecified `1.0`

07 **Schema updates across sandbox, organization, admin, and API-key endpoints**

IMPROVED

34

Request or response schemas changed with no described content beyond the change itself for a large set of endpoints. Sandbox lifecycle: `DELETE /sandbox/{sandboxIdOrName}`, `DELETE /sandbox/{sandboxIdOrName}/ssh-access`, `GET /sandbox`, `GET /sandbox/for-runner`, `GET /sandbox/paginated`, `GET /sandbox/{sandboxIdOrName}`, `GET /sandbox/{sandboxIdOrName}/ancestors`, `GET /sandbox/{sandboxIdOrName}/forks`, `GET /sandbox/{sandboxIdOrName}/parent`, `GET /sandbox/{sandboxId}/organization`, `GET /sandbox/{sandboxId}/region-quota`, `POST /sandbox` (request body and response schema changed), `POST /sandbox/{sandboxIdOrName}/archive`, `POST /sandbox/{sandboxIdOrName}/autoarchive/{interval}`, `POST /sandbox/{sandboxIdOrName}/autodelete/{interval}`, `POST /sandbox/{sandboxIdOrName}/autopause/{interval}`, `POST /sandbox/{sandboxIdOrName}/autostop/{interval}`, `POST /sandbox/{sandboxIdOrName}/backup`, `POST /sandbox/{sandboxIdOrName}/fork`, `POST /sandbox/{sandboxIdOrName}/network-settings`, `POST /sandbox/{sandboxIdOrName}/pause`, `POST /sandbox/{sandboxIdOrName}/public/{isPublic}`, `POST /sandbox/{sandboxIdOrName}/recover`, `POST /sandbox/{sandboxIdOrName}/resize`, `POST /sandbox/{sandboxIdOrName}/snapshot`, `POST /sandbox/{sandboxIdOrName}/start`, `POST /sandbox/{sandboxIdOrName}/stop`, `POST`

`/sandbox/{sandboxIdOrName}/ttl/{ttlMinutes}`, `PUT /sandbox/{sandboxIdOrName}/secrets`. Organization and admin: `GET /admin/organizations/{organizationId}/quota/{regionId}/{sandboxClass}`, `GET /admin/users/{id}`, `PATCH /admin/organizations/{organizationId}/quota/{regionId}` (request body changed), `PATCH /admin/snapshots/{id}/general`, `POST /admin/organizations`, `POST /admin/organizations/{organizationId}/quota/{regionId}` (request body and response schema changed), `POST /admin/sandbox/{sandboxId}/recover`, `POST /admin/users` (request body changed), `GET /organizations`, `GET /organizations/invitations`, `GET /organizations/{organizationId}`, `GET /organizations/{organizationId}/invitations`, `GET /organizations/{organizationId}/roles`, `GET /organizations/{organizationId}/users`, `PATCH /organizations/{organizationId}/quota/{regionId}` (request body changed), `POST /organizations`, `POST /organizations/invitations/{invitationId}/accept`, `POST /organizations/{organizationId}/invitations`, `POST /organizations/{organizationId}/roles` (request body changed), `POST /organizations/{organizationId}/users/{userId}/access`, `PUT /organizations/{organizationId}/invitations/{invitationId}`, `PUT /organizations/{organizationId}/roles/{roleId}` (request body changed). API keys and snapshots: `GET /api-keys`, `GET /api-keys/current`, `GET /api-keys/{name}`, `POST /api-keys` (request body changed), `GET /config`, `GET /users/me`, `GET /snapshots/{id}`, `POST /snapshots`, `POST /snapshots/{id}/activate`.

— Endpoints named but actual schema change content unknown `1.0`

---

08 **Enforced TLS certificate verification** BREAKING 21

Daytona 0.200.0 enforces TLS verification, per the release summary; no further mechanism or affected surfaces are described.

— Only a one-line summary, no detail on scope or migration `snapshot-20260814`

WAS THIS USEFUL? 👍 👎

Kiro

Feedback

1 RELEASE · SEEN 2026-08-14

NOTES ↗

> Kiro's 1.0.309 release adds support for nested `AGENTS.md` files for directory-scoped instructions, alongside efficiency improvements for long sessions and background resource usage.

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

01 Nested AGENTS.md files for directory-scoped instructions NEW

55

Kiro now supports nested `AGENTS.md` files placed throughout the workspace, giving Kiro instructions that are scoped to each directory tree rather than a single workspace-wide file.

— Names the exact file and scoping behavior, no example given

Nested AGENTS.md and Long-Session Effic…

02 Long-session context retention and idle efficiency IMPROVED

25

Kiro preserves more context during long sessions, reducing loss of earlier conversation state, and reduces idle resource usage for improved background efficiency.

— Describes benefit but no mechanism or metrics given

Nested AGENTS.md and Long-Session Effic…

03 Cloud configuration, chat, hooks and MCP improvements

IMPROVED

15

Kiro improves Cloud configuration, chat, hooks, and MCP connections, though the source does not specify what changed in each area.

— Bare list of areas with no detail on what changed

Nested AGENTS.md and Long-Session Effic…

WAS THIS USEFUL? 👍 👎

Cursor

Feedback ⓘ 1 RELEASE · SEEN 2026-08-14 NOTES ↗

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

> Cursor shipped shareable deeplinks for prompts, commands, and rules, brought Bugbot PR review to Azure DevOps, and gave Cloud Agents pre-built environments that boot up to 3x faster.

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

01 Bugbot PR review for Azure DevOps NEW 95

Bugbot now reviews Azure DevOps Services pull requests, posting review comments, inline comments, and build statuses under the `cursor-bugbot/review` context. On-demand reviews can be triggered with `cursor review` or `bugbot run` comments, restricted to users whose Azure DevOps sign-in address matches a Cursor account on the team. Repository settings let you disable Run for All Authors and configure 'Only Review PRs by...' or 'Skip PRs by...' using sign-in addresses (e.g. `taylor@contoso.com`) found under Organization settings → Users. A Microsoft Entra service principal is provisioned per tenant for all Bugbot actions so reviews survive role changes; service hooks install automatically when a repo is enabled under Bugbot in Automations and are removed when disabled or disconnected. Cloud Agents can also clone, push, commit, and open PRs on Azure DevOps using the connected identity. The integration is limited to team-scope repositories (personal-scope unsupported), and Auto-Enable for New Repositories, Automatically Learn Rules, Personal Bugbot settings, and Autofix are not supported — repo-level settings apply to all authors instead.

Trigger an on-demand Bugbot review on an open Azure DevOps pull request by posting a comment.

```
$ bugbot run
```

Enable Bugbot on specific Azure DevOps repositories and restrict reviews to a subset of authors by sign-in address.

> 📍 In Cursor, open Bugbot in Automations, find your Azure DevOps repository, turn Bugbot on, disable Run for All Authors, and add sign-in addresses like `taylor@contoso.com` under 'Only Review PRs by...'.

— Named endpoints, settings, identity mechanism, and exact steps. product docs

02 Deeplinks for prompts, commands, and rules NEW 85

Adds shareable deeplinks using the `cursor://anysphere.cursor-deeplink/` protocol or the `cursor.com/link/` web base URL: prompt deeplinks via `cursor://anysphere.cursor-deeplink/prompt?text=...` (also `cursor.com/link/prompt?text=...`) share pre-filled chat prompts the recipient must confirm before execution; command deeplinks save shared custom commands into the recipient's `.cursor/commands/` directory pending review; rule deeplinks save shared custom rules into `.cursor/rules/` pending review. All deeplink URLs support a maximum encoded length of 8,000 characters.

Share a reusable chat prompt with a teammate so they can open it pre-filled in Cursor and confirm before running.

```
$ cursor://anysphere.cursor-deeplink/prompt?
text=Research%20and%20find%20one%20bug%20in%20this%20codebase
```

— Names exact protocols, paths, and the 8,000-char limit. product docs

---

03 **Pre-built environments for Cloud Agents** NEW 85

Cloud Agents can now boot into Builds — pre-baked copies of the dev environment (repos cloned, dependencies installed, install script run) — instead of setting up from scratch each session, delivering 10x faster internal environment boot and 3x faster time to first token at no additional cost. A Builds tab in the Cloud Agents dashboard shows build status, logs, commit SHAs, and which build each agent run used; builds can be triggered manually, debugged with an agent, and automatically fall back to the last successful build when a bad commit or dependency update breaks the environment. A configurable threshold controls how stale builds get refreshed, and agents can inspect and manage builds using built-in tools.

Enable Builds on an existing Cloud Agents environment to get faster agent starts without reconfiguring from scratch.

> 📍 In the Cloud Agents dashboard, open your environment › go to the Builds tab › click Enable Builds. Optionally click Run setup agent first to preview any proposed config changes before enabling.

— Gives UI steps and metrics but no named config key.

Cloud Agents Start 3x Faster with Builds

WAS THIS USEFUL? 👍 👎

Cline

Feedback ⓘ 2 RELEASES · 2026-08-14 NOTES ↗ CODE ↗

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

> Cline's SDK and CLI releases in this window add a provider-executed `web_search` model tool with persisted search history, a dedicated Cline gateway provider, and safer Hub session handling around updates and build swaps, alongside a refreshed model catalog with Crusoe support and streaming markdown fixes.

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

01 `cline doctor fix` honest process reporting IMPROVED

85

`cline doctor fix` now reports honestly: processes that survived a kill are separated from ones that appeared during the fix, a live parent respawning a daemon is named, and a startup lock held by a running process is reported as held rather than leaked.

— Names the exact command and each corrected behavior sdk/sdk/v0.0.75

02 Provider-executed web search tool NEW

80

Adds a `web_search` model tool, off by default and enabled in settings, letting supported models search the web during a turn; search calls and their results are persisted in session history for replay on reload.

Enable provider-executed web search so a supported model can look up live information mid-turn during an agent session.

> 📍 In the Cline settings panel, locate the model tools section and enable the `web_search` tool, then start or resume a session — search calls and their results will be saved in session history and replayed on reload.

— Names exact tool, settings path, and usage steps sdk/sdk/v0.0.75

03 Safer auto-update sequencing for Hub sessions IMPROVED

80

Auto-updates no longer install while a CLI is attached to the Hub; the update is recorded at startup and installed on exit, preventing a background update from killing a live session (previously surfacing as Hub connection closed with `code=1006`). `cline update` still installs immediately and now reports that the update applies on next start.

— Names command and exact error code it fixes cli-v3.0.55

04 Dedicated Cline gateway provider NEW

70

Adds a dedicated Cline provider for the Cline gateway, replacing the generic OpenAI-compatible path, so extended thinking budgets and other gateway options now apply to both `cline` and `cline-pass` models.

— Names affected models and provider path change sdk/sdk/v0.0.75

THINNER COVERAGE BELOW

05 **Streaming markdown rendering stability** IMPROVED 55

Streaming assistant markdown no longer flashes back to raw text — settled headings, links, and code stay rendered as new chunks arrive instead of the whole message being rebuilt on every chunk.

— Explains behavior change with no config surface cli-v3.0.55

06 **Hub build handoff without disrupting live sessions** IMPROVED 50

Newer builds no longer replace a Hub serving live sessions — they attach over the compatible wire protocol and the swap happens once the Hub is idle. Development builds now run their own Hub daemon per build ID instead of contending for a single record.

— Describes mechanism but has no user-facing command sdk/sdk/v0.0.75

07 **Crusoe provider and model catalog refresh** NEW 40

The model catalog now includes Crusoe as a provider, with model lists and per-provider default models updated across all providers.

— Names new provider but no config detail given sdk/sdk/v0.0.75 cli-v3.0.55

08 **Idle plugin sandbox process cleanup** IMPROVED 30

Idle plugin sandbox processes are now reclaimed instead of lingering for the life of the session.

— Brief description with no mechanism or command sdk/sdk/v0.0.75 cli-v3.0.55

WAS THIS USEFUL?  

OpenAI Codex CLI

Feedback ⓘ 2 RELEASES · 2026-08-14 NOTES ↗

CODE ↗

Lightweight coding agent that runs in your terminal

> Across two alpha releases, Codex CLI added an Amazon Bedrock Runtime model provider, skill-level model delegation to Luna, and per-server MCP OAuth callback ports, alongside session, Guardian, and telemetry improvements.

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

01 **Per-server MCP OAuth callback port** NEW 90

Adds `oauth.callback_port` to MCP server configuration so each server can use its own OAuth callback port instead of the global `mcp_oauth_callback_port`; also accepts `oauth.callbackPort` from plugin MCP declarations and skill dependency metadata.

Configure a per-server OAuth callback port so that an MCP server uses port 9005 instead of the global default during `codex mcp login`.

yaml

```yaml
oauth:
  callback_port: 9005
```

— Named config key plus runnable YAML example rust-v0.148.0-alpha.14

02 **Skill-level model delegation to Luna** NEW 80

Skill frontmatter now supports an optional `model` field (e.g. `model: luna`) letting a skill declare a preferred model; unrecognized values are ignored without blocking the skill. A new `SkillModelDelegationInstruction` routes skills requesting Luna while running on Sol or Terra to Luna when available in the current provider namespace, with bounded and validated model identifiers and instruction size limits.

— Concrete frontmatter example and named mechanism, but no full config path

rust-v0.148.0-alpha.15

03 **Amazon Bedrock Runtime model provider** NEW 75

Adds the built-in `amazon-bedrock-runtime` provider for regional `bedrock-runtime` OpenAI-compatible endpoints, with SigV4 auth, per-provider AWS profile, region, and transport overrides, global and US cross-region GPT-5.6 model variants; web search is disabled since it is unsupported on Bedrock.

— Names provider id and auth/config options but no exact config key shown

`rust-v0.148.0-alpha.15`

04 **Preserved developer messages across compaction** IMPROVED 65

Client-authored developer messages are now preserved across context compaction when `retain_client_developer_messages` is enabled.

— Named flag but no file location or default given `rust-v0.148.0-alpha.14`

THINNER COVERAGE BELOW

05 **Guardian review evidence and context improvements** IMPROVED 57

Guardian V2 now receives full tool action context — the original `ToolPayload` including tool name and arguments plus a bounded conversation transcript — for richer risk classification, and Guardian review evidence now includes `node_repl` images.

— Names payload and evidence types but no config surface to invoke `rust-v0.148.0-alpha.14`

06 **Running-task exit choices on Ctrl-C** NEW 55

Pressing Ctrl-C while a task is running in local daemon sessions now shows a menu to cancel the task and stay in Codex, exit Codex while leaving the task running, or stop the task and exit Codex.

— Describes UI behavior with a clear trigger but no flag/command `rust-v0.148.0-alpha.14`

07 **Current-time reminder handling in model context** IMPROVED 50

Injected current-time reminders are now wrapped in `<current_time_reminder>` tags in model context, and full-history subagents get a refreshed reminder by excluding parent reminders and injecting a single fresh reminder into the child request.

— Named tag format but no user-facing control `rust-v0.148.0-alpha.14`

08 **Experimental thread queue APIs** NEW 23

Adds experimental thread queue APIs to the app server.

— Bare mention with no endpoint names or usage detail `rust-v0.148.0-alpha.14`

09 **Structured telemetry for response retries** NEW 16

Adds structured telemetry for response retries.

— One-line mention with no schema or access point `rust-v0.148.0-alpha.14`

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ 1 RELEASE · 2026-08-13 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 v2.1.232 introduces cross-session @`-mentions` and per-session messaging controls, turns on subagent forking by default, adds GitLab secret redaction and full GitLab plugin marketplace support, and tightens gateway and sandbox settings with new boot-time validation.

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

01 **GitLab support and settings aliases for plugin marketplaces**

NEW 85

Bare `gitlab.com` repo URLs (including nested subgroups) now clone like `github.com` URLs for plugin marketplaces. Settings now accept `additionalMarketplaces` and `allowedMarketplaces` as friendlier aliases for `extraKnownMarketplaces` and `strictKnownMarketplaces`. An enterprise-policy url-typed `blockedMarketplaces` entry for a bare repo URL continues blocking that URL even when the CLI classifies it as a git clone.

Block a private GitLab plugin marketplace from being used, even when the CLI identifies the URL as a git clone.

```json
{
  "blockedMarketplaces": [
    { "type": "url", "value": "gitlab.com/my-org/internal-plugins" }
  ]
}
```

Use the friendlier `allowedMarketplaces` alias in settings instead of the legacy `strictKnownMarketplaces` key to restrict which marketplaces Claude Code may use.

```json
{
  "allowedMarketplaces": [
    "gitlab.com/my-org/approved-plugins"
  ]
}
```

— Config keys and behavior both named with runnable examples. v2.1.232

02 **Cross-session @-mentions and messaging controls** NEW

83

Typing `@` in the prompt now mentions another Claude session by name, using `SendMessage` to reach that session directly; `SendMessage` also delivers to a bare name that exactly matches one live session without a confirmation step. New `/config` rows for Dialog expiry and Messages from your other sessions let you set per-session inbound messaging to accept, hold, or refuse. Interactive sessions on one machine now keep unique names — renaming to a name already in use generates a `name-word-word` variant and notifies the user.

Configure cross-session messaging policy so inbound messages from your other sessions are held for review rather than accepted automatically.

📍 `/config`

— Names exact commands and config rows, mechanism explained. v2.1.232

03 **Gateway config validation at boot** BREAKING

80

The `desktop:` overlay now accepts every released Desktop setting (previously limited to 11 hand-listed keys), validated at boot against Desktop's own schema, with unknown or invalid keys now failing boot. Empty `managed.policies[].match.groups` / `admin.admin_groups` entries and malformed `email_domain` values (empty, or containing `@`, whitespace, or commas) now fail at boot instead of silently matching no one or granting admin access.

— Names exact config keys and boot-failure conditions. v2.1.232

04 **GitLab token redaction** NEW

75

Adds secret redaction for GitLab token families `glrt-`, `gloas-`, `glptt-`, `glagent-`, `glimt-`, `glsoat-`, `glcbt-`, `glft-`, and `glffct-`; fully redacts routable `glpat-` and `gldt-` tokens; extends the same sandbox and credential-path protection to the `glab` CLI config store.

— Every token prefix named but no user action required. v2.1.232

05 **Managed settings approval dialog improvements** IMPROVED

65

The managed settings approval dialog now shows endpoint URLs, uses clearer wording for telemetry-only changes, skips routine OpenTelemetry options, and requires explicit approval for server-managed sandbox binary overrides (`sandbox.bwrapPath`, `sandbox.socatPath`, `sandbox.ripgrep`).

— Names the specific override keys the dialog now gates. v2.1.232

06 **Sandbox ripgrep binary setting restricted to trusted scopes** 60

BREAKING

`sandbox.ripgrep` is now honored only from user, managed, and `--settings` settings; project settings can no longer override the sandbox's ripgrep binary.

— Names the exact key and scopes but no example given. v2.1.232

07 **Fable 5 advisor restored** IMPROVED 60

Fable 5 is offered again as an advisor in `/advisor` for organizations with Fable access, with usage-credits consent set up through `/model fable`.

— Gives exact commands to access the restored advisor. v2.1.232

08 **Marketplace auto-refresh on plugin install** IMPROVED 60

`/plugin install plugin@marketplace` now refreshes the marketplace first, so newly published plugins install without a manual marketplace update.

— Names the exact command whose behavior changed. v2.1.232

THINNER COVERAGE BELOW

09 **Subagent forking enabled by default** IMPROVED 55

A `subagent_type: "fork"` subagent now inherits the full conversation and prompt cache by default, and non-teammate agent spawns in interactive sessions now run in the background by default.

— Explains new default behavior, no config surface given. v2.1.232

10 **/feedback and /bug open immediately** IMPROVED 45

`/feedback` and `/bug` now open immediately when invoked while Claude is responding, instead of waiting for the turn to finish.

— Names exact commands but change is minor UX timing. v2.1.232

11 **/code-review high-effort runs as background agent** IMPROVED 45

`/code-review` at high, xhigh, and max effort now runs in a background agent, consistent with other effort levels.

— Names the command and effort levels affected. v2.1.232

12 **Permission checks for bash input redirection** IMPROVED 40

Bash input redirections (`< file`) are now permission-checked like their argument spellings, on all platforms.

— Brief statement of behavior change, no further detail. v2.1.232

13 **Remote Control reconnection reliability** IMPROVED 40

Remote Control now reconnects for approximately 30 minutes after a network blip and no longer drops after a few blips spread across an hour.

— Gives a concrete duration but no config or command. v2.1.232

└──▷ **BREAKING ON UPGRADE**

! `sandbox.ripgrep` set in project settings is no longer honored; it must now be placed in user, managed, or `--settings` settings.

! Gateway: unknown or invalid keys under the `desktop:` overlay now fail boot (previously accepted silently).

! Gateway: empty `managed.policies[].match.groups` / `admin.admin_groups` entries and malformed `email_domain` values now fail at boot instead of silently matching no one or granting admin access.

WAS THIS USEFUL?  

AGENT

◆ AI Agent Frameworks

CrewAI

Feedback ⓘ 1 RELEASE · 2026-08-14 NOTES ↗

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

> CrewAI 1.15.16 focuses on execution observability, adding UUID-based tracking for flow and agent runs plus richer AMP deployment telemetry.

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

01 **Flow execution tracking and exception recording** NEW

38

Adds execution context management with UUID support for tracking individual flow and agent executions, and records the type of exception that ended a flow so observability pipelines get structured failure context.

— Describes what changed but no concrete API or field names given. 1.15.16

02 **AMP deployment telemetry** IMPROVED

35

Records when a trace batch is shared with AMP for auditable telemetry handoff visibility, and counts deployments from any origin while recording where they started, supporting multi-origin deployment tracking in AMP.

— Explains behaviour but lacks specific endpoints or config to act on. 1.15.16

WAS THIS USEFUL? 👍 👎

Hermes

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

The agent that grows with you

> Hermes shipped a full community plugin management CLI (install, search, doctor, and pack commands) alongside new desktop ws RPCs for profiles and image generation, an inline MCP setup consent tool, a `concise` messaging mode, an installer flag to skip bundling `cua-driver`, and a relaxed approval-gated policy for editing `~/.ssh/config`.

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

01 **Plugin management CLI for community plugins** NEW 100

New CLI surface for installing and vetting community plugins: `install <identifier> [--force] [--ref COMMIT_SHA]` installs from a Git URL or bare community index name, with `--ref` accepting a full 40-character commit SHA to pin an exact immutable revision; `search [term] [--json] [--capability CAP] [--refresh]` fuzzy-searches the community index (fetched from `plugins.index_url`, cached under `~/.hermes/cache/` for 24 hours with offline fallback); `doctor [path-or-id] [--ci]` validates a plugin through the real manifest parser, loader, and registration path, with `--ci` exiting code 1 on errors for CI pipelines; `pack install <path-or-url> [--force]` installs a `hermes-pack.yaml` plugin pack whose plugins are each pinned to an exact 40-character commit SHA, behind a mandatory review screen and per-plugin capability consent, exiting non-zero on any plugin failure; `pack export [--enabled-only] [--name NAME]` emits a `hermes-pack.yaml` to stdout recording repo and exact SHA per git-installed plugin while stripping secrets, capability grants, and `allow_*` gates; and `pack show <path-or-url>` dry-runs a pack without installing. Git installs record source, revision, and pin status in the `plugins/.install-metadata.json` sidecar (never secrets or grants); unpinned plugins update via `pull`, while pinned plugins require `--force --ref <new-commit>` to move revision.

Pin a community plugin to an exact immutable commit SHA so your environment never drifts on upgrade.

```
$ hermes install my-plugin --ref
a3f1c2e4b5d6789012345678901234567890abcd
```

Find plugins that expose a specific capability before committing to an install.

```
$ hermes search --capability network_access --json
```

Validate a plugin in CI and fail the build immediately if the manifest or loader registration has errors.

```
$ hermes doctor ./my-plugin --ci
```

— Every subcommand, flag, and file path is named with runnable examples. product docs

02 **profiles.list and profiles.create ws RPCs** NEW 95

`profiles.list` and `profiles.create` ws JSON-RPC methods, implemented in `tui_gateway/methods_profiles.py`, let desktop plugins enumerate profiles (name, path, model/provider, description, skill_count, and an optional `last_session` preview via `include_sessions`) and create new profiles without going through the dashboard REST router.

List all profiles from a desktop plugin to build a roster UI, including a last-session preview per profile.

```
$ { "method": "profiles.list", "params": { "include_sessions": true } }
```

— Named RPC methods, fields, and file, with a runnable example. v2026.8.13

03 **Approval gating for `~/.ssh/config` edits** IMPROVED 83

`~/.ssh/config` is promoted from hard-blocked to approval-gated: edits to SSH client config such as host aliases, `ProxyJump`, and VS Code Remote-SSH targets now go through the same approve-once/session/always prompt used for other `~/.ssh` writes instead of a flat refusal. Non-interactive callers (the ACP file bridge, background jobs with no human channel) still fail closed on `~/.ssh/config` writes, and private keys, `authorized_keys`, and all other material under `~/.ssh` remain hard-blocked regardless of approval state.

— Names exact paths and exceptions; behavioral not a runnable command. product docs

04 **Installer skip-Computer-Use flag** NEW 80

The installer adds a `--skip-computer-use` flag to opt out of pre-installing `cua-driver`; without it, `cua-driver` is bundled automatically so the Computer Use toolset is ready as soon as it's enabled.

— Exact flag and default behavior given, mechanism is minimal. product docs

05 **Concise messaging verbosity mode** NEW 65

Adds a `concise` messaging mode (the new default) that surfaces a one-line status message on completion, with a short output tail appended on failure, joining the existing `all`, `result`, `error`, and `off` modes.

— Names all mode values but not the setting used to switch modes. product docs

THINNER COVERAGE BELOW

06 **Inline MCP setup consent tool** NEW 55

Adds a `setup_mcp` tool that renders an inline MCP consent card in the desktop transcript using a clarify-style blocking bridge, with composer draft suggestions surfaced as brand pills.

— Mechanism described but no command or config to invoke it. v2026.8.13

07 **Desktop composer suggestion improvements** IMPROVED 55

Labels the agent's recommended choice on clarify surfaces across the desktop UI; adds skill-match, connection-repair, and recurrence-to-cron suggestion providers to the desktop composer; and suppresses composer suggestions the user has repeatedly ignored.

— Names three suggestion providers but no config surface to tune them. v2026.8.13

08 **image.generate ws RPC** NEW 35

Adds an `image.generate` ws RPC for plugin surfaces.

— Bare method name with no parameters or usage shown. v2026.8.13

09 **Cron and blueprint recipes in sidebar nav** NEW 30

Adds cron and blueprint recipes to the sidebar nav rail.

— Names the UI location but no further detail on the recipes. v2026.8.13

10 **Per-route webhook toolset overrides** NEW 25

Adds per-route toolset overrides for webhook agent runs.

— No config key, route, or example given. v2026.8.13

11 **Model drift indicator in Desktop cron view** NEW 25

Surfaces model drift impact in the Desktop cron view.

— One-line description with no mechanism or metric detail. v2026.8.13

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

vMLX

Feedback ⓘ 1 RELEASE · 2026-08-14 NOTES ↗

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

> vMLX added native support for the Qwen3.6-27B model bundle with automatic multi-token prediction, measuring a 33% throughput gain with no configuration required.

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

01 **Qwen3.6-27B bundle support with auto multi-token prediction**

NEW 76

vMLX now supports the Qwen3.6-27B bundle line (and the upcoming Qwen 3.8 line): the `qwen3_coder` tool-parser name resolves to the XML-function parser, capability-only stamps route through the stock loader instead of the JANG codec, and the bundled multi-token prediction (MTP) head is automatically constructed, quantized per the stamp's per-module overrides, and engaged at the stamp's trained speculative depth — no flags required. Measured throughput on the 4-bit Qwen3.6-27B bundle went from 23.3 to 31 tokens/sec (+33%), with greedy output byte-identical whether the MTP head is on or off, and image/video probes answering correctly through both API and app.

— Names parser, loader routing, quantization mechanism and exact benchmark numbers.

v1.6.28

WAS THIS USEFUL? 👍 👎

◆ Agentic Infrastructure

Ollama

Feedback ⓘ 1 RELEASE · 2026-08-14 NOTES ↗ CODE ↗

Get up and running with Kimi-K2.6, GLM-5.2, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.

> Ollama v0.32.11 adds one-command launchers for two external agent tools, extends prompt template support to Nemotron 3.5, and brings web search to the OpenAI Responses API.

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

01 Agent harness launchers for DeepSeek and Meta tools NEW

70

Adds `ollama launch dsh` to integrate DeepSeek Harness, DeepSeek's open-source agent harness, and `ollama launch muse` to integrate Muse Code, Meta's agentic coding CLI, letting either be run locally backed by an Ollama model.

Launch DeepSeek's open-source agent harness locally through Ollama.

```
$ ollama launch dsh
```

Launch Meta's Muse Code agentic coding CLI backed by a local Ollama model.

```
$ ollama launch muse
```

— Names exact commands but no further mechanism detail v0.32.11

THINNER COVERAGE BELOW

02 Nemotron 3.5 prompt layout support NEW

35

Adds support for the Nemotron 3.5 prompt layout, selecting the correct parser and renderer from its checkpoint template.

— Explains mechanism but no usage example given v0.32.11

03 Web search in OpenAI Responses API NEW

25

Brings web search to the OpenAI Responses API.

— Only mentioned in summary, no mechanism or config detail v0.32.11

WAS THIS USEFUL?  

emisar

Feedback

1 RELEASE · 2026-08-13

NOTES ↗

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

> emisar v0.40.0 adds self-updating runners with verified binary swaps, pagination for large MCP runbook results, and expands its integration catalog to 100 packs while tightening scope enforcement across discovery surfaces.

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

01 **Self-updating runners via `emisar update`** NEW 88

The new `emisar update` subcommand lets installer-managed runners update themselves, performing automatic release-checksum and GitHub build-attestation verification before executing stop, swap, restart, and rollback steps.

Update an installer-managed runner in place — the runner verifies checksum and attestation before swapping binaries.

```
$ emisar update
```

— Runnable command with clear verification and rollback mechanism described. v0.40.0

02 **JSON output for pack updates** IMPROVED 77

The `--json` flag on `emisar pack update` now emits its partial report before returning a post-update validation error, preserving machine-readable output for automation pipelines even on failure.

Capture machine-readable pack-update output even when post-update validation fails, so CI pipelines can parse the partial report.

```
$ emisar pack update --json
```

— Exact flag and behavior on failure named, runnable example given. v0.40.0

03 **Paginated results for wait_for_run** IMPROVED 60

`wait_for_run` now returns paginated 64 KiB continuation pages when a terminal runbook result exceeds one MCP response, keeping the summary in the first response.

— Mechanism and size limit named but no direct usage example. v0.40.0

THINNER COVERAGE BELOW

04 **Catalog expands to 100 packs, 1,682 actions** NEW 54

The catalog now covers 100 packs and 1,682 actions, adding JFrog Artifactory, Databricks, Sentry, Symbolicator, and NTPsec, along with larger Cassandra and Cloudflare coverage.

— Concrete counts and integrations named, but no usage detail given. v0.40.0

05 **Scope enforcement for discovery across console and MCP**

IMPROVED 46

Runner, pack, action, approval, and audit discovery now enforces the member's current runner and pack scope consistently across both the console and MCP.

— States what changed but no mechanism or configuration path given. v0.40.0

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NVIDIA OpenShell

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

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

> OpenShell added lifecycle control for sandboxes, letting users stop and start them via the CLI.

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

01 **Stop and start operations for sandboxes** NEW 60

The `openshell sandbox` subcommand gains `stop` and `start` operations, enabling lifecycle control of running sandboxes.

— Names exact subcommand and operations but no further mechanism. v0.0.105

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 14, 2026

Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue