THE CYBER TOOLCHAIN

NO. 034

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 3, 2026 · EVERY WEEKDAY

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 18 tools

OFFENSIVE SECURITY

◆ AI-Powered Offensive Security

Hacktron

2 RELEASES · seen 2026-08-03 NOTES ↗

> Hacktron expanded integrations with Azure DevOps and Bitbucket PR scanning, and added bulk triage, live AI-assisted runs, interactive taint trace navigation, and findings export.

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

01 **Interactive taint trace navigation** NEW 35

Lets users walk a taint trace interactively to follow data flows from source to sink during code analysis.

– Explains source-to-sink concept but not UI mechanics

Walk a taint trace, export your findings…

02 **Bitbucket PR scanning** NEW 35

Scans Bitbucket pull requests to surface vulnerabilities before code is merged.

– Clear purpose but no config or trigger details Walk a taint trace, export your findings…

03 **Azure DevOps integration** NEW 30

Connects Hacktron findings and workflows directly with Azure DevOps projects.

– Names integration target but no setup steps or scope

Azure DevOps is here, plus bulk triage a…

04 **Bulk triage of findings** NEW 25

Allows practitioners to triage multiple findings simultaneously rather than one at a time.

– States capability but no UI path or limits given

Azure DevOps is here, plus bulk triage a…

05 **Live AI-assisted runs** NEW 15

Enables real-time AI-assisted analysis during assessments.

– Vague on mechanism, trigger, or where it appears

Azure DevOps is here, plus bulk triage a…

06 **Findings export** NEW 15

Export findings for reporting or downstream workflows.

– No format, API, or export destination named Walk a taint trace, export your findings…

WAS THIS USEFUL? 👍 👎

◆ Exploitation & C2

shellfire ⓘ 1 RELEASE · 2026-08-03 NOTES ↗

An exploitation shell focusing on exploiting command injection vulnerabilities, eg., LFI, RFI, SSTI, etc.

> shellfire v0.16 adds a configurable request timeout to keep GET, POST, and form requests from hanging indefinitely against unresponsive targets.

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

01 **Configurable request timeout for HTTP requests** NEW 65

Adds a configurable request timeout (cfg.timeout, default 30 seconds) applied to GET, POST, and form requests so shellfire no longer blocks indefinitely against a hung target.

— Names config key and default but no usage example. v0.16

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Snyk Code

1 RELEASE · 2026-08-03 BLOG ↗

Snyk Code is a developer security tool that identifies and fixes vulnerabilities in source code during development.

> Snyk Code introduces an Ignore Approval Workflow that turns developer-submitted code ignores into reviewer-gated requests, backed by a dedicated review UI, granular permissions, and audit-log visibility.

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

01 **Ignore Approval Workflow for ignores** NEW 65

Developer-submitted ignores now become Pending requests that require explicit reviewer approval or rejection before they take effect. Requests can be submitted from the CLI, IDE, or API, and a reason is required on every submission.

– Names submission surfaces and reason requirement, no exact command shown 2026-08-03-f4bb32cd

02 **Granular ignore permissions and roles** NEW 65

New granular permissions for ignores (read, create, edit, review, delete) let admins control who can act on ignores; Collaborators default to request-level access and Admins default to review-level access, with custom roles also supported.

– Lists exact permission types and default role mapping 2026-08-03-f4bb32cd

THINNER COVERAGE BELOW

03 **Ignore Requests page in web UI** NEW 50

A new Ignore Requests page in the web UI lets reviewers approve or reject pending ignore requests in one central place.

– Clear UI surface named but no navigation path given 2026-08-03-f4bb32cd

04 **Ignore-request data in audit-logs API** NEW 50

Ignore-request activity is now exposed via the audit-logs API, enabling full auditability and export of the approval workflow.

– Names the API but no schema or query example given 2026-08-03-f4bb32cd

05 **Reviewer auto-approval for self-created ignores** NEW 40

Reviewers can have their own submitted ignores auto-approved, bypassing the triage queue entirely.

— States behaviour but no config detail on enabling it `2026-08-03-f4bb32cd`

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Dalfox

1 RELEASE · 2026-08-02 NOTES ↗ CODE ↗

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

> Dalfox's docs and v3.2.0 release focus on making large-scale, CI-driven XSS scanning practical: baseline diffing, URL deduplication, resumable state files, and mid-scan session-loss detection headline the window, alongside new DOM-XSS sinks, an expanded payload library with a query subcommand, and MCP/Caido integrations for agent- and proxy-driven testing.

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

01 **Mid-scan authenticated session-loss detection** NEW 93

Dalfox fingerprints the authenticated response before scanning and re-probes it after each target's injection stage, detecting session loss via status-code changes, redirect-to-login, or appearance of a password field, rather than silently returning zero findings. `--session-check` supplies a regex the authenticated body must keep matching (overriding built-in heuristics) and `--session-check-url` designates a lightweight probe endpoint (e.g. `/api/me`) instead of the scan target; `--on-session-loss abort` halts the affected target and skips same-host targets (exiting non-zero if nothing was found), while `continue` keeps scanning and preserves the exit code. Affected targets are reported with `"status": "incomplete"` and `"error_code": "SESSION_LOST"` plus an `error_message` and a top-level `incomplete` flag.

Abort a scan automatically if the authenticated session is lost mid-run

yaml

```yaml
# dalfox.yaml
session_check: "Welcome, "
session_check_url: "https://app.example.com/api/me"
on_session_loss: "abort"
```

Scan an authenticated app and detect session expiry using an exact string marker and a cheap dedicated probe endpoint, aborting the host group if the session is lost.

```
$ dalfox scan https://app.example.com/dashboard?q=1 --cookies
sid=$SESSION --session-check 'Signed in as' --session-check-url
https://app.example.com/api/me --on-session-loss abort
```

— Detection mechanism, flags, and exact status/exit codes specified v3.2.0

02 **Baseline diffing against prior scan reports** NEW 91

A `--baseline` / `baseline` config key diffs a new scan against a previous JSON/JSONL report so only net-new findings are surfaced. `--baseline-mode filter` drops known findings from counts, output and exit code, while `annotate` keeps all findings and tags each with `new: true` / `false` for dashboards; structured output gains a `meta.baseline` block reporting path, mode, enabled state, and total/new/known counts. Fingerprinting matches on host+path, parameter name, parameter location, injection context, CWE, finding tier and evidence family — deliberately excluding payload text and AST line numbers so reruns and bundler rebuilds don't resurrect triaged findings.

In CI, report only findings that are new since the last scan by diffing against the previous run's JSON report.

```
$ dalfox url https://example.com -f json -o current.json --
baseline previous.json --baseline-mode filter
```

Record a full-backlog baseline on first run, then gate CI on new findings only — the build goes red only when something new appears.

```
$ # Record baseline once
dalfox scan scope.txt -f json -o .dalfox/baseline.json

# Every CI run after — exits non-zero only on new findings
dalfox scan scope.txt \
  --baseline .dalfox/baseline.json \
  --only-poc v \
  -f json -o dalfox.json --silence
```

Use annotate mode to keep the full finding set while marking novelty, useful for dashboards that need historical and new findings together.

```
$ dalfox scan scope.txt \
  --baseline .dalfox/baseline.json \
  --baseline-mode annotate \
  -f json -o dalfox-annotated.json
```

Refresh the baseline after triage without accidentally overwriting it with a diff-only report — omit `--baseline` so the output holds the full set.

```
$ dalfox scan scope.txt -f json -o .dalfox/baseline.json
git commit -am "chore: refresh dalfox baseline"
```

— Full mechanism, flags, fingerprinting fields, and CI commands given `v3.2.0`

---

03 **URL deduplication before scanning** NEW

83

`--dedup-urls` (or `dedup_urls` config key) collapses duplicate targets before scanning with `exact` (identical URL+method), `signature` (URLs differing only in parameter values, matched by method/scheme/host/port/path/sorted parameter names), or `off` `/no-dedup` modes. The report envelope adds `dedup_mode` and `targets_deduplicated` fields showing how many targets were collapsed.

Collapse thousands of harvested value-variant URLs into one scan per unique endpoint so your `--max-targets-per-host` budget covers different endpoints rather than one endpoint's parameter values.

```
$ gau target.app | dalfox scan --dedup-urls signature
```

— Exact matching criteria, report fields, and pipeline example given `v3.2.0`

---

04 **Payload query subcommand and expanded payload categories** NEW

83

A new `dalfox payload` subcommand queries built-in payload libraries by named selector — `special-chars`, `functions`, `awesome-alert`, `dom-clobbering`, `mxss`, and `blind` — printing one payload per line so output composes with tools like `grep` and `wc`. Generation is CSP-aware: when a `Content-Security-Policy` header is present, output is tailored to DOM script-gadgets and nonce-based policy-callback payloads. New categories cover `dom-clobbering` vectors, `mxss` mutation-XSS/sanitizer-bypass payloads for testing sanitizer pipelines, `blind` blind-XSS callback skeletons with a `{}` OOB URL placeholder, and `awesome-alert` screenshot-ready PoCs.

Quickly count how many special characters Dalfox considers relevant — useful before manual reflection testing to scope your character-probe list.

```
$ dalfox payload special-chars | wc -l
```

Filter the confirmable-sink list down to prompt-based payloads when alert is blocked by the target's CSP or JS environment.

```
$ dalfox payload functions | grep -i prompt
```

— Subcommand, named selectors, and runnable shell examples given v3.2.0

05 **Resumable mass scans via state file** NEW 65

`--state-file` (or `state_file` config key) records completed targets during a bulk scan and lets a later run resume without restarting from scratch, skipping already-completed targets while retrying ones that were cancelled or errored.

Resume a large scan that was interrupted by Ctrl-C or a dropped SSH session without re-scanning already-completed targets.

```
$ dalfox scan --input-type file urls.txt --state-file scan.state
```

Resume a large mass scan that was interrupted — targets already scanned are skipped automatically.

```
$ dalfox scan urls.txt --state-file scan-state.json
```

Resume a large scan without re-testing already-completed targets, and only report findings not seen in a previous run

```yaml
# dalfox.yaml
state_file: "scan.state"
baseline: "baseline.json"
baseline_mode: "filter"
dedup_urls: "signature"
```

— Flag/config key and resume commands shown, mechanism brief v3.2.0

THINNER COVERAGE BELOW

06 **MCP server for AI agent workflows** NEW 55

Exposes Dalfox as a tool via an MCP stdio server, making it available to any MCP-compatible client including Claude and Cursor, with drop-in instructions for skill-aware agents (Claude Code, Cursor, OpenCode, Codex). Gains a `max_payloads_per_param` control and a synchronous wait mode for agent workflows.

— Named clients and a control but no config example shown v3.2.0

07 **Six new DOM-XSS sink detections** NEW 50

Adds detection of six additional DOM-XSS sinks: drag-drop events, the async clipboard API, `FileReader`, `setAttributeNS`, indirect `eval`, and `DOMParser`.

— Six sinks named but no usage or example given `v3.2.0`

08 **Per-target timeout and request-rate capping** NEW 35

Adds a per-target wall-clock timeout as a hard upper bound in seconds to abort stalled endpoints, and a global request-rate cap intended to be friendly to shared IPs and edge WAF thresholds; HTTP 429 responses are always retried regardless of the rate-cap setting.

— Behavior described but no flag names or config keys given `product docs`

09 **Interactsh OOB integration for blind XSS** NEW 35

Adds Interactsh integration to enable out-of-band/OAST detection of blind XSS, with support for optional custom server domains.

— Integration named but no usage shown `product docs`

10 **Caido Active Workflow integration** NEW 35

Integrates with Caido's Active Workflows and Findings to automate real-time XSS testing from within a proxy session.

— Integration named, no setup steps given `product docs`

11 **HTTP QUERY method support** NEW 35

Adds support for scanning targets that use the HTTP `QUERY` method (RFC 10008).

— Method and RFC named, no usage example `v3.2.0`

12 **Finding confidence, detection method and impact metadata** NEW 25

Findings now carry three separate metadata axes: confidence, detection method, and impact.

— Axes named but no schema, values or example shown `v3.2.0`

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft ⓘ

EIFT Firewall: a new free tool for controlled agent sideloading (2026-08-03)

BLOG ↗

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

ElcomSoft releases EIFT Firewall, a free macOS app for controlled USB-network agent sideloading to iOS devices.

- New free macOS companion app EIFT Firewall replaces the 2023 firewall script for sideloading the iOS Forensic Toolkit low-level extraction agent.
- Creates a temporary, controlled USB network path between a Mac and a connected evidentiary iOS device to safely deliver the extraction agent.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

HyperDbg ⓘ 1 RELEASE · 2026-08-02 NOTES ↗

State-of-the-art native debugging tools

> HyperDbg v0.23 adds a user-mode CPUID command with SDK support, introduces floating-point scripting, and continues progress on its Linux port.

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

01 ucpuid command for user-mode CPUID execution NEW 70

A new `ucpuid` command lets users execute CPUID in user-mode context directly from the debugger, and a corresponding SDK function enables programmatic invocation via the HyperDbg API.

Inspect CPUID leaves from user-mode context without switching privilege levels manually.

```
$ ucpuid
```

— Names command and SDK, includes runnable example. v0.23

THINNER COVERAGE BELOW

02 Linux port advances IMPROVED 50

Progress on the Linux port includes completion of the script engine Linux build, addition of a Linux kernel module build file (Kbuild), and inclusion of a Linux implementation for PlatformCPU.

— Names specific components changed but no usage steps. v0.23

03 Floating-point support in script engine NEW 30

The script engine now supports floating-point (float) data types and operators, enabling FP-aware scripting.

— States capability but no usage example or syntax. v0.23

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Plerion ⓘ 1 RELEASE · seen 2026-08-03 NOTES ↗

Plerion is a cloud security platform that detects and responds to infrastructure vulnerabilities and misconfigurations across multi-cloud environments.

> Plerion added CIS Google Cloud Platform Foundation Benchmark v5.0.0 mapping, enabling control-by-control compliance tracking for Google Cloud posture.

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

01 **CIS Google Cloud Foundation Benchmark mapping** NEW 55

Plerion maps Google Cloud posture to the CIS Google Cloud Platform Foundation Benchmark v5.0.0, enabling control-by-control compliance tracking from compliance settings.

— Names exact benchmark version and location but no steps or scope detail.

`CIS benchmark coverage for Google Cloud:`

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Managed Detection & Response (MDR)

Arctic Wolf ⓘ 1 RELEASE · 2026-08-03 PRESS ↗

Arctic Wolf is a managed security service provider offering 24/7 threat detection, response, and managed SOC services for enterprises.

> Arctic Wolf launched Cyber Resilience, a bundled offering combining MDR, exposure management, endpoint defense, security awareness, and incident response, backed by a new industry warranty of up to $3 million.

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

01 **Cyber Resilience bundled offering** NEW 55

A single package combining Aurora MDR, Attack Surface Management, Vulnerability Management with Resolve Patch Management, Managed Endpoint Defense, Managed Security Awareness, and Incident Response 360 (IR360).

— Names bundled components but no pricing or setup path launch-20260803-84fbc5d4

02 **Security Operations Warranty up to $3M** NEW 45

Arctic Wolf's Security Operations Warranty offers up to $3 million in financial coverage when a security incident occurs while using the Cyber Resilience offering, described as the largest such warranty in the industry to date.

— Specific dollar figure given but no claim process detailed launch-20260803-84fbc5d4

03 **Aurora Incident Response 360 (IR360)** NEW 40

A component of the Cyber Resilience offering providing immediate threat containment and full recovery coverage for any incident type.

— States purpose but not how containment or recovery works launch-20260803-84fbc5d4

WAS THIS USEFUL? 👍 👎

◆ SIEM

Sumo Logic ⓘ 1 RELEASE · 2026-08-03 BLOG ↗

Sumo Logic is a cloud-native platform that collects, analyzes, and visualizes machine data and logs for security, observability, and operational intelligence.

> Sumo Logic updated credit tracking in its Manage module so mid-subscription credit purchases are handled as separate blocks rather than repricing existing credits.

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

01 **Separate credit blocks for mid-subscription purchases** IMPROVED

45

Sumo Logic now tracks additional credits purchased mid-subscription as a distinct credit block, so originally purchased credits continue burning at their original rate instead of being repriced.

— Explains mechanism but no UI path or numbers given.

Credit Block Burn Rates (Manage) (2026-0…

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

OSSEC HIDS

1 RELEASE · 2026-08-02 NOTES ↗

OSSEC is an Open Source Host-based Intrusion Detection System that performs log analysis, file integrity checking, policy monitoring, rootkit detection, real-time alerting and active response.

> OSSEC 4.2.0 overhauls its manager-side processing with a multi-threaded analysisd pipeline and pthread-based daemons, while enriching rootcheck alerts with file hashes and simplifying Windows agent deployment.

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

01 **Rootcheck alerts enriched with file hashes** NEW 55

Appends MD5, SHA1, and SHA256 hashes of detected rootcheck files to alerts, enriching forensic context.

– Names exact hash algorithms but no alert format shown. 4.2.0

02 **IPv6 key validation for authd** NEW 55

Adds IPv6 key validation via `OS_I` sValidIP for authd `-i` keys.

– Names function and flag, but no usage example. 4.2.0

03 **Self-contained Windows agent installer** IMPROVED 50

Bundles required MinGW/OpenSSL runtime DLLs in the Windows agent installer, enabling self-contained deployment.

– Names dependencies bundled; no install steps given. 4.2.0

04 **Multi-threaded analysisd event pipeline** NEW 40

Adds an always-on multi-threaded analysisd event pipeline for Linux servers, improving analysisd scalability.

– Names analysisd and platform but no config or metrics. 4.2.0

05 **Manager daemons migrated to pthreads** IMPROVED 35

Migrates manager daemons to pthreads and adds shared thread pool support.

– States mechanism change, no behavioural detail or numbers. 4.2.0

06 **Truncated diff indicator in alerts** IMPROVED 35

Indicates in alerts when report_changes diffs are truncated.

– Names report_changes feature; lacks detail on threshold or format. 4.2.0

WAS THIS USEFUL? 👍 👎

Trend Vision One

2 RELEASES · 2026-07-15 BLOG ↗

Trend Vision One is a unified cybersecurity platform that integrates extended detection and response (XDR) across endpoints, networks, email, and cloud environments.

> Trend Vision One's Service Gateway firmware and Forward Proxy Service both added FQDN support for the Indonesia (ID) site, alongside smaller reliability and IoT connectivity improvements for Service Gateway appliances.

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

01 **Disk space allocation for extended drives** IMPROVED 45

When a disk is extended, Service Gateway now allocates the remaining disk space to the `/data` partition, improving storage utilization on expanded appliances.

— Clear behavior change named but no config steps given

Service Gateway firmware version 3.0.31 …

02 **FQDN support for Indonesia (ID) site** NEW 35

Service Gateway firmware and the Forward Proxy Service both gained FQDN support for the Indonesia (ID) site, enabling Service Gateway deployments in that region.

— Names region and component but no setup steps or mechanism

Service Gateway firmware version 3.0.31 … Service Gateway: Forward Proxy Service v…

03 **Service installation waits for CNI readiness** IMPROVED 35

Improves service installation reliability by waiting for CNI network readiness before proceeding.

— Names CNI dependency but no further detail on impact

Service Gateway firmware version 3.0.31 …

04 **IoT connection recovery via independent channel** IMPROVED 25

Improves IoT connection recovery through an independent channel.

— States mechanism briefly but lacks detail on behavior

Service Gateway firmware version 3.0.31 …

05 **IoT device configuration update handling** IMPROVED 15

Enhances IoT device configuration update handling for better connectivity management.

– Vague wording with no mechanism or scope specified

Service Gateway firmware version 3.0.31 …

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Cloudmersive ⓘ 2 RELEASES · 2026-08-02 BLOG ↗

Cloudmersive provides cloud-based APIs for document conversion, validation, virus scanning, and data processing tasks.

> Cloudmersive shipped two incremental API updates: added batch job overrides for the DLP AI API and expanded virus scanning to recognize additional Python model file formats.

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

01 **Batch job API overrides for DLP AI** IMPROVED 25

The Cloud DLP AI API adds additional batch job API overrides.

– No detail on what overrides do or how to use them.

Cloudmersive Cloud DLP AI API v2-8-0 Rel…

02 **Recognition of additional Python model file formats** IMPROVED 25

The Cloud Virus API now recognizes additional Python model file formats during virus scanning.

– No specifics on which file formats or scanning mechanism.

Cloudmersive Cloud Virus API v11-10-10 (…

WAS THIS USEFUL? 👍 👎

◆ Data Security & DLP

Immuta

1 RELEASE · seen 2026-08-03 NOTES ↗

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

> Immuta's single release this window centers on the Connections experience — per-connection health states, granular policy sync controls, and bulk tagging — alongside dual request forms for data products, Redshift viewless auto-discovery, and new system-managed groups for policy authoring; it also ships several breaking changes around Snowflake policy roles and grant revocation.

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

01 **Configurable 'New' tag behavior** IMPROVED 60

Makes the 'New' tag behavior configurable in App Settings, enabled for columns and disabled for data sources by default.

— Names exact setting, location, and default state snapshot-20260803

02 **Automated sensitive data discovery for Redshift viewless integration**

NEW 60

Adds automated sensitive data discovery for Immuta's Amazon Redshift viewless integration, supporting regex, dictionary, and column name regex identifier types.

— Names identifier types and integration but no config path snapshot-20260803

03 **Tag ingestion for Databricks Unity Catalog foreign catalogs** NEW 60

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

— Names exact mechanism and integration surface snapshot-20260803

04 **Grant revocation on data source disable** BREAKING 60

When a data source is disabled, Immuta now revokes data owners' Immuta-managed grants in the data platform; this behavior change is on by default starting 8/1 and enabled for all users by 9/1.

— Includes exact rollout dates and behavior change snapshot-20260803

THINNER COVERAGE BELOW

05 **Per-object policy sync toggle** NEW — 55

Adds a per-object policy sync toggle in Connections, letting admins enable or disable policy push to the platform for individual data objects, with visibility into child object configuration.

— Clear mechanism and location, no exact UI steps `snapshot-20260803`

06 **Bulk tag assignment at database/schema level** IMPROVED — 55

Supports bulk or individual tag assignment and removal at the database or schema level, with tags inherited by all attached data sources.

— Explains scope and inheritance mechanism clearly `snapshot-20260803`

07 **Dual request forms per data product** NEW — 55

Supports two distinct request forms per data product — one for data access requests and one for masking exception requests — each with independent questions and data use agreements.

— Describes the two forms and independence but no setup steps `snapshot-20260803`

08 **Built-in system-managed groups for policies** NEW — 55

Introduces three built-in system-managed groups — All Agents, All Users, and All Identities — available for use in subscription and data policies without manual membership management.

— Names groups and use case, lacks usage steps `snapshot-20260803`

09 **Per-connection integration health states** NEW — 50

Adds integration health states (ACTIVE, INACTIVE, UNHEALTHY) displayed per connection for Databricks Unity Catalog and Snowflake, giving admins at-a-glance integration status.

— Names states and platforms but no navigation path given `snapshot-20260803`

10 **Snowflake policy role renaming** BREAKING — 50

Immuta policy role names used to grant users access to Snowflake data sources will be renamed as part of the Snowflake integration architecture update rolling out from July 23, 2026; existing policy roles will be recreated under new names.

— Gives rollout date and mechanism, lacks migration steps `snapshot-20260803`

11 **Removal of Snowflake integration error banners** BREAKING — 45

Integration error banners for Snowflake integration validation failures are removed; these messages are now only available through the new Connections experience.

– Explains removal and where messages now live `snapshot-20260803`

12 **Expanded Snowflake manual policy resync** IMPROVED 40

Expands Snowflake manual policy resync to cover both data policies and subscription policies simultaneously.

– States scope expansion but no mechanism detail `snapshot-20260803`

13 **Force-deletion of unhealthy connections** NEW 35

Enables force-deletion of connections that are in an UNHEALTHY state.

– Minimal detail beyond the capability name `snapshot-20260803`

14 **Case-insensitive Starburst principal matching** IMPROVED 35

Enables case-insensitive Starburst (Trino) principal name matching for more reliable user and service account identification.

– Names the integration and behavior, no further detail `snapshot-20260803`

15 **Detail pages for leaf data objects** NEW 30

Adds detail pages for leaf data objects in Connections, including a direct link to the corresponding data source.

– Bare description of a new page with minimal detail `snapshot-20260803`

╰──▷ BREAKING ON UPGRADE

! Immuta policy role names used to grant users access to Snowflake data sources will be renamed as part of the Snowflake integration architecture update rolling out from July 23, 2026; existing policy roles will be recreated under new names.

! Integration error banners for Snowflake integration validation failures are removed; these messages are now only available through the new Connections experience.

! When a data source is disabled, Immuta now revokes data owners' Immuta-managed grants in the data platform (behavior change release; on by default 8/1, enabled for all 9/1).

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Zscaler

1 RELEASE · 2026-08-03 BLOG ↗

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

> Zscaler brought Zero Trust Branch into FedRAMP Moderate and High environments and added a bundle of branch security capabilities — per-device network segmentation, OT/IoT device isolation, a ransomware killswitch, one-arm connectivity mode, and zero-touch appliance provisioning.

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

01 **Per-device network segmentation via DHCP** NEW 65

Provides device segmentation using a /32 network per device (with variable subnet length support), with Zscaler acting as a DHCP server to discover and isolate every device at a branch site.

— Concrete mechanism named, but no setup walkthrough. 2026-08-03-b73e1ef3

02 **One-arm mode with outbound-only tunnels** NEW 60

Supports one-arm mode alongside existing SD-WAN or gateway mode to terminate multiple internet links and load balance traffic, using outbound-only tunnels to the Zero Trust Exchange that eliminate publicly exposed routes.

— Mechanism detailed, lacks configuration instructions. 2026-08-03-b73e1ef3

03 **Zero-touch provisioning via Branch Appliance** NEW 60

Offers zero-touch provisioning through the integrated Branch Appliance: define a site, activate the appliance, and it automatically establishes secure outbound connectivity.

— Gives concrete setup steps, lacks command-level detail. 2026-08-03-b73e1ef3

THINNER COVERAGE BELOW

04 **TIC 3.0 Branch Office local breakout support** NEW 55

Supports the CISA TIC 3.0 Branch Office Use Case with a local breakout architecture, allowing branch users to access the web and agency-sanctioned CSPs directly without backhauling traffic to a central location.

— Names standard and architecture but no configuration steps. 2026-08-03-b73e1ef3

05 **OT/IoT device discovery and isolation** NEW 55

Enables OT and IoT device discovery, tagging, and least-privilege policy enforcement to isolate cameras, badge systems, building management systems, and other hard-to-

patch devices from user networks.

— Lists device types and policy type, no setup steps. 2026-08-03-b73e1ef3

06 **Ransomware Killswitch with color-coded policies** NEW

45

Introduces a Ransomware Killswitch concept using color-coded policies that let teams rapidly tighten enforcement during suspicious activity to limit lateral spread.

— Concept named but policy definitions and triggers unspecified. 2026-08-03-b73e1ef3

07 **Zero Trust Branch in FedRAMP Moderate/High** NEW

40

Zscaler Zero Trust Branch is now available in FedRAMP Moderate and High authorized environments, giving civilian federal agencies and public sector organizations a compliant path to deploy zero trust branch security.

— States availability and audience, no deployment mechanism given. 2026-08-03-b73e1ef3

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Kiro ⓘ 1 RELEASE · seen 2026-08-03 CHANGELOG ↗

Kiro is an AI-powered IDE by AWS that uses "specs" to translate requirements into structured plans, then automatically generates and implements code.

> Kiro 2.16 introduces `/tangent` side-conversations with a visual tree picker for exploring off-thread questions without losing context, alongside per-tool token visibility and several workflow refinements to specs, plan mode, and hooks.

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

01 **/tangent side-conversations with tree picker** NEW 88

The new `/tangent` command branches into a side-conversation that inherits full conversation history, letting you explore freely and return to the original thread. Supports named tangents via `/tangent <name>`, a visual conversation-tree picker via `/tangent ls`, and one-step return to the root via `/tangent root`.





Investigate a tangential question mid-task without losing your main thread — then jump straight back.

```
$ /tangent investigate-auth-edge-case
```

Browse your full conversation tree to navigate between active tangents and the main thread.

```
$ /tangent ls
```

– Names exact commands and behaviour; runnable examples included

changelog-20260803-9deba171

02 **Per-tool token breakdown in /context** IMPROVED 65

The `/context` command now shows a per-tool token breakdown so you can see exactly what is consuming your context window.

Check which tools are eating your context budget before starting a large refactor.

```
$ /context
```

– Clear command given, but mechanism of breakdown not detailed changelog-20260803-9deba171

THINNER COVERAGE BELOW

03 **Guided spec creation with description prompt** IMPROVED 53

`/spec new` now prompts for a description before drafting requirements, using your input as ground truth instead of inferring from the spec name alone.

– Names command and behaviour change, no example given changelog-20260803-9deba171

04 **Plan mode auto-executes approved plans** IMPROVED 40

Plan mode now auto-executes an approved plan immediately, removing the manual mode-switch step after plan review.

– Explains before/after but no navigation or example changelog-20260803-9deba171

05 **Kiro submenu in editor context menu** NEW 40

A Kiro submenu was added to the editor context menu, including an 'Ask Kiro to Fix' quick fix on errors and warnings.

– Names the UI surface but no exact steps to use it changelog-20260803-9deba171

06 **Guided form for creating hooks** NEW 20

A guided form was added for creating hooks.

– Bare mention with no mechanism or location detail changelog-20260803-9deba171

07 **Replies in IDE display language** IMPROVED 20

Kiro now replies in the IDE display language.

— Bare statement, no detail on supported languages or config changelog-20260803-9deba171

---

08 **Chat history persists across workspace changes** IMPROVED

20

Chat history persists across workspace folder changes.

— Bare statement with no mechanism or scope detail changelog-20260803-9deba171

WAS THIS USEFUL? 👍 👎

Cline

1 RELEASE · 2026-07-29 NOTES ↗

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

> Cline's desktop app v0.0.7 adds a system tray status icon, paginated and favoritable session history, visibility into subagent/teammate runs, several UI refinements, and support for free-tier models with rate-limit messaging.

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

01 **UI polish for tool disclosures and messages** IMPROVED 45

Adds tool-specific icons on tool disclosures, elapsed thinking time display, restyled reasoning sections, aligned timestamps, and a layout-stable message actions bar with a viewport-anchored scroll.

– Lists several concrete UI changes, no usage steps `desktop-v0.0.7`

02 **Free-tier model support with rate-limit messaging** NEW 45

Free Cline models are now supported and labeled "(free)" in model pickers, with a clear rate-limit message that includes the reset time.

– Names label and message content, points to model picker `desktop-v0.0.7`

03 **Paginated session history** IMPROVED 40

Session history is now paginated in ten-session pages, loading older history on demand rather than all at once.

– Gives exact page size and mechanism, no UI path `desktop-v0.0.7`

04 **Session favorites and recency ordering** NEW 35

Enables favoriting sessions and orders all sessions by most recent activity, with consistent status dot colors applied across views.

– Describes behavior but not how to favorite a session `desktop-v0.0.7`

05 **System tray status icon with session count** NEW 30

Adds a system tray icon that displays live app status and the count of currently running agent sessions.

– Names the surface but no interaction detail `desktop-v0.0.7`

06 **Subagent and teammate run visibility** NEW 30

Subagent and teammate runs launched from a session now appear in the app along with their status and results.

— States what appears but not where or how `desktop-v0.0.7`

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ AI Model & Data Infrastructure

Hugging Face ⓘ 1 RELEASE · seen 2026-08-03 NOTES ↗

Hugging Face is a platform providing pre-trained machine learning models, datasets, and tools for natural language processing and computer vision tasks.

> Hugging Face added label-based filtering to Jobs, letting users narrow job lists via clickable label chips or free-form `key=value` queries.

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

01 **Filter Jobs by label** NEW 63

Hugging Face Jobs now supports filtering by label through clickable chips that show the most-used labels along with job counts, available on both user and organization jobs pages. A free-form `key=value` input also lets users filter by any label, including ones not surfaced as chips.

— Names UI mechanism and input format but no limits or API details. `Filter Jobs by Label`

WAS THIS USEFUL? 👍 👎

◆ Local LLM Runtimes

vMLX

2 RELEASES · 2026-08-02 → 2026-08-03 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 DeepSeek V4 Flash (DSV4) runtime support with dedicated cache handling and settings, then extended it with a memory-scaling cache ceiling that delivers roughly 11% faster decode throughput for DSV4.

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

01 **DSV4 cache handling with memory-scaled ceiling** IMPROVED

80

Native DSV4 cache handling preserves short append checkpoints, lossless L2 writes, valid eviction ancestry, partial-tail replay, SSD-only operation, and RAM-to-SSD refault. The MLX cache ceiling now scales dynamically with installed memory (16 GB → 4 GB, 64 GB → 11.5 GB, 128 GB → 23 GB, capped at 24 GB), delivering ~11% faster decode throughput for DeepSeek V4 Flash at every context length by eliminating redundant buffer re-allocation.

— Concrete numbers and mechanism given, but no config key to set it. v1.6.22 v1.6.20

02 **Native DeepSeek V4 Flash session support** NEW

70

Adds native DSV4 sessions with native reasoning, DSML tool support, sampling guidance, cache-tier capabilities, and pool-quantization behavior. Server and Chat settings now distinguish native DSV4 compiled decode and pooled-cache state from generic TurboQuant and unsupported whole-model cache modes, with explicit Activation-QAT control and effective-state reporting.

— Names DSML, TurboQuant, Activation-QAT but no exact UI path or command. v1.6.20

THINNER COVERAGE BELOW

03 **No-more-tools continuation pass after tool use** NEW

53

Explicit visible-final instructions after tool use now enter a no-more-tools continuation pass, with terminal abort cleanup drained before the runtime returns to idle.

— Describes mechanism but not how to trigger or observe it. v1.6.20

04 **Electron settings validate DSV4 top-p values** IMPROVED

48

Electron settings warn when a DSV4 top-p value differs from bundle guidance and preserve remote model tool contracts.

— Names the warning behavior but no exact setting path. v1.6.20

05 **Release packaging attests DSV4 encoder and build provenance**

IMPROVED 48

Release packaging now attests the native Python DSV4 encoder, nested affine defaults, mixed module quantization metadata, and the clean JANG runtime source, with separate Sequoia and Tahoe packaging/notarization contracts.

— Lists named artifacts but no action a user can take. v1.6.20

06 **Health output distinguishes served vs loaded model** IMPROVED 28

Health output now distinguishes the served model from the currently loaded model.

— Bare statement of change with no detail on where or how viewed. v1.6.20

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 3, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe