THE CYBER TOOLCHAIN

NO. 046

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 15, 2026 · EVERY DAY

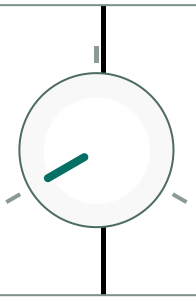
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 36 tools

OFFENSIVE SECURITY

vim ? help

◆ Exploitation & C2

Bugcrowd Researcher Platform

Feedback ⓘ

1 RELEASE · 2026-08-15 CHANGELOG ↗

Bugcrowd Researcher Platform connects security researchers with organizations to identify and report vulnerabilities through crowdsourced bug bounty programs.

> Bugcrowd Researcher Platform now supports researcher participation in Bugcrowd Gov (US) engagements under FedRAMP compliance.

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

01 **FedRAMP support for Bugcrowd Gov engagements** NEW 35

The platform now supports researcher participation in Bugcrowd Gov (US) engagements under FedRAMP compliance, letting researchers hack on government programs subject to that standard.

— States compliance scope but no mechanism or how to access engagements.

changelog-20260815-67db9996

WAS THIS USEFUL? 👍 👎

Bugcrowd

Feedback ⓘ 1 RELEASE · 2026-08-15 CHANGELOG ↗

Bugcrowd is a crowdsourced security testing platform connecting organizations with ethical hackers to identify and report vulnerabilities.

> Bugcrowd shipped FedRAMP support for compliance-sensitive customers alongside two new reporting surfaces: an on-demand Vulnerability Detections dashboard and a DNS Records CSV export for auditors and leadership.

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

01 DNS Records CSV export report NEW

75

A new DNS records (csv) report type under `Organization > Reports` exports a point-in-time inventory of discovered assets — names, types, and discovery source — filterable by creation date range. Generated reports are emailed when ready and remain accessible to all Owner roles from the Reports page as timestamped snapshots.

Generate a DNS records CSV for a quarter to hand auditors a defensible record of attack surface growth without waiting on engineering.

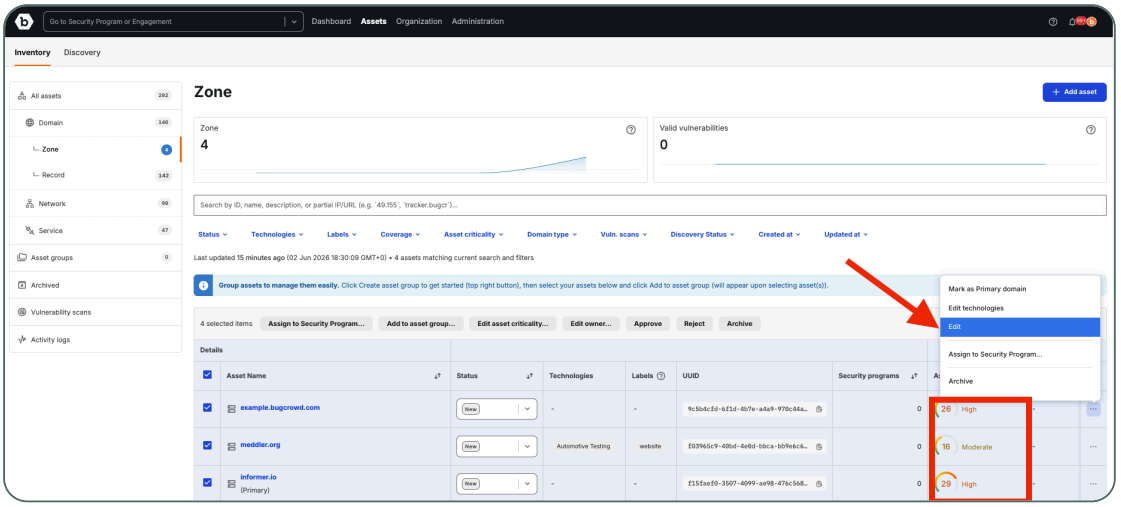
> 📍 Go to Organization › Reports, click Generate report, select DNS records (csv) from the dropdown, provide a Report Title and Date Range, then click Generate report. Download the artifact from the Reports page once the email notification arrives.

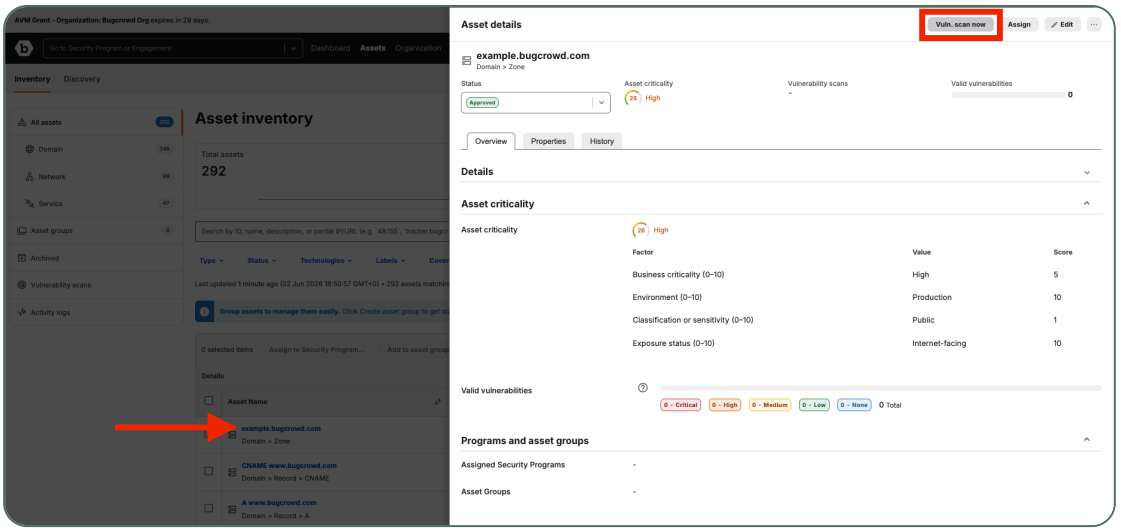
— Names exact report type, fields, and full generation steps product docs

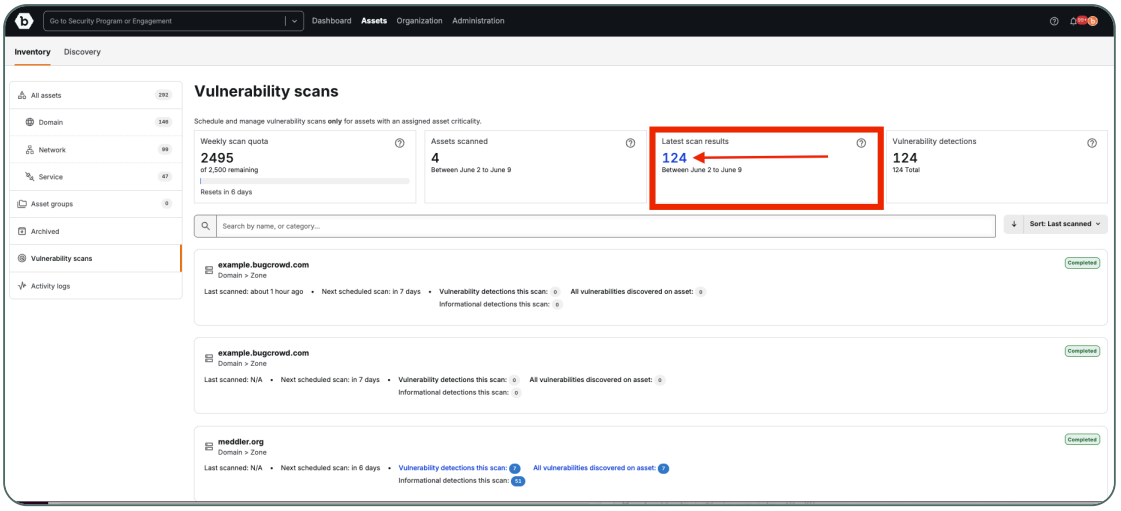
02 Vulnerability Detections dashboard in AI Analytics NEW

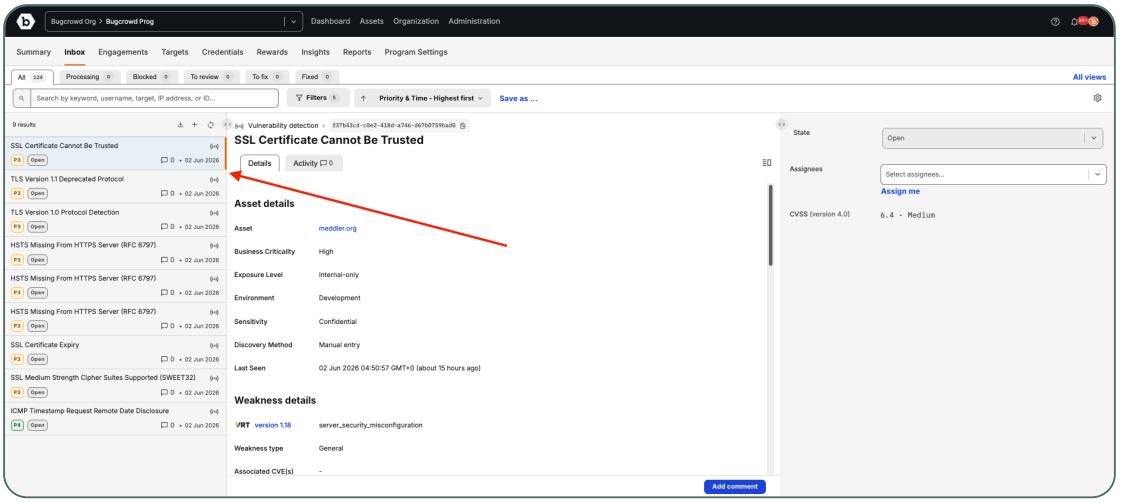
70

A new Vulnerability Detections tab under `Organization > AI Analytics` lets Owners filter by date range, vulnerability priority, asset group, and asset owner to scope posture views to a specific business unit or asset owner before a leadership review.









Scope the Vulnerability Detections dashboard to a specific business unit before a leadership review.

> 📍 Go to Organization › AI Analytics, click the Vulnerability Detections tab, then apply filters for date range, vulnerability priority, asset group, or asset owner to focus on the relevant business unit.

— Names exact tab and filters with navigation steps, but no mechanism detail `product docs`

THINNER COVERAGE BELOW

03 **FedRAMP support for compliance environments** NEW 25

Bugcrowd added FedRAMP support, enabling government and compliance-sensitive organizations to use the platform within FedRAMP-authorized environments.

— Bare statement of support with no scope, timeline, or setup detail

changelog-20260815-94c2ea16

WAS THIS USEFUL? 👍 👎

NetSPI OCInferno

Feedback ⓘ changes since 2026-07-16

CODE ↗

A pentesting tool for enumeration/download/graphical analysis of OCI content. Includes an OpenGraph generator for Bloodhound-style analysis.

OCInferno gains exploit modules, MFA stripping, delegation auth, container image tokens, IDD self-service, and new enumeration flags.

└──▷ **GET THIS VERSION**

```
$ git clone --branch commits-2026-07-16
https://github.com/NetSPI/OCInferno.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-16
```

└──▷ **TRY IT**

Enumerate OCI identity resources and save discovered domain URLs to the database for later use in IDD attacks.

```
$ ocinferno enum_identity --save-domains
```

Scope an enum_all run to specific regions to avoid timeouts on unsupported OCI endpoints.

```
$ ocinferno enum_all --regions us-ashburn-1,us-phoenix-1
```

Strip MFA factors from a target OCI account after gaining initial access, using the new exploit module.

```
$ ocinferno exploit_remove_authentication_factors
```

Adds `--save-domains` flag to `enum_identity` to persist discovered domain URLs to the database.

Adds `--regions` flag to enumeration modules to scope scans to specific OCI regions.

Logs `[TIMEOUT]` in `enum_all` output when a module hits a connect or read timeout, making stalled regions visible.

Adds `exploit_remove_authentication_factors` module to strip MFA (classic and

IDD) from OCI accounts.

Adds exploit modules for DevOps repositories, including download capability for DevOps repos.

+7 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

Acunetix 360

Feedback ⓘ 1 RELEASE · 2026-08-15 CHANGELOG ↗

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

> Acunetix 360 adds partial-update support to its scan profile and policy APIs, extends scan data retention to two years, and upgrades core services to .NET 10.

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

01 **PATCH support for scan profile and policy APIs** IMPROVED

70

Adds `PATCH` semantics to scan profile and policy API endpoints, allowing partial updates that preserve existing values without resubmitting the full configuration object.

– Names exact HTTP method and endpoints affected but no path given.

changelog-20260815-6f9ff8e2

THINNER COVERAGE BELOW

02 **.NET 10 upgrade for Auth Verifier and scan agent** IMPROVED

35

Upgrades the Auth Verifier Service and scan agent to .NET 10 for long-term support, improved performance, and enhanced stability.

– Names components and version but no user-facing action. changelog-20260815-6f9ff8e2

WAS THIS USEFUL? 👍 👎

c/side

Feedback

1 RELEASE · SEEN 2026-08-15

NOTES ↗

C/Side is a cloud-based IDE for Dynamics 365 Business Central development, enabling creation and modification of extensions and customizations.

> c/side introduced CSP Builder, a tool that turns observed script hosts into a deployable Content-Security-Policy with a review and report-only rollout workflow.

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

01 **CSP Builder for policy generation and rollout** NEW

75

CSP Builder generates a suggested Content-Security-Policy from hosts c/side has already observed on your site, covering seven directives: `script-src`, `connect-src`, `img-src`, `style-src`, `frame-src`, `font-src`, and `media-src`. You can review the suggested policy, toggle off unexpected hosts, and add hosts not yet observed before deploying, and the workflow supports deploying a report-only header first, monitoring violations in the dashboard, then switching to enforce mode once the policy is clean.

— Names directives and workflow steps but no exact header/config syntax.

Introducing CSP Builder

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

Truffle Security trufflehog

Feedback ⓘ

1 RELEASE · 2026-08-14 NOTES ↗ CODE ↗

trufflehog discovers, classifies, and validates leaked credentials

> TruffleHog's headline additions this window are Slack scanning for Enterprise Edition, AWS credential analysis via `aws analyze`, and a cross-source Activity Log with a read API; the open-source CLI gains SARIF output (`--sarif`), a large set of custom regex detector tuning options, S3 custom endpoints, and eight new detectors, while Enterprise also picks up source archiving, SharePoint historical scanning, Google Drive domain-wide delegation, and 16 more detectors.

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

01 Custom regex detector configuration options IMPROVED

100

Custom regex detectors (beta), wired in via the `detectors` field of the TruffleHog config YAML with `keywords`, `name`, `regex`, and `verify` (with `endpoint`, `headers`, `unsafe`) sub-keys, gain a set of tuning parameters: `successranges` under `verify` to specify HTTP status codes or inclusive ranges (e.g. `'200-202'`) that mark a secret as live/verified; `rotatedranges` under `verify` to specify codes or ranges (e.g. `'401'`, `'403'`) that mark a secret as definitively rotated/unverified; `entropy` to filter detected strings by randomness score and cut false positives; `exclude regexes capture` to strip sub-portions of a detected secret from the reported result; `exclude regexes match` to suppress entire matches from being reported; `exclude words` to ignore strings containing specified words; and `primary regex name` to designate which named regex pattern determines the reported line number when multiple regexes are defined under the `regex` field, with capture groups for extracting specific secret segments and a requirement that all named regexes match for a detection to succeed.

Distinguish live secrets from rotated ones during verification so your triage queue only surfaces actionable findings.

```yaml
detectors:
  - name: hogtokendetector
    keywords:
      - hog
    regex:
      token: '[^a-zA-Z0-9+\/]{0,1}([a-zA-Z0-9+\/]{40})[^a-zA-Z0-9+\/]{0,1}'
    verify:
      - endpoint: http://localhost:8000/
        unsafe: true
        headers:
          - "authorization: super secret authorization header"
        successranges:
          - "200"
        rotatedranges:
          - "401"
          - "403"
```

— Names every config key with a runnable example YAML config. product docs

02 **Activity Log for scan events (Enterprise)** NEW 90

New Activity Log gives a top-level view of scan lifecycle events — `scan started`, `scan completed`, `scan failed`, `scan scheduled`, and `scan anomaly detected` — across every source, filterable by `date range`, `event`, `status`, `source`, and `actor`, with filters stored in the page URL for shareable views; events are append-only with 90-day retention. Scan events carry one of six statuses — `success`, `partial`, `failure`, `in progress`, `scheduled`, or `warning` — shown as color chips, and anomaly-detection heuristics flag completed scans that scanned no data or finished suspiciously fast. Each row opens a detail drawer with full message text and a direct, shareable URL. A read API is available so scan failures can be routed into external monitoring and alerting tools, with API documentation shipping alongside it; source IDs are exposed via the API only. Archiving a source preserves its Activity Log history, while deleting a source removes its events along with its findings.

— Full statuses, filters, retention and API named; exact endpoint path omitted.

product docs

03 **SARIF output for scan results** NEW 85

Adds `--sarif` flag for SARIF-format output compatible with GitHub code scanning's `github/codeql-action/upload-sarif` action, tracking findings as new/fixed across scans.

Upload TruffleHog results to GitHub code scanning in a CI workflow — SARIF lets GitHub inline findings on PR diffs and deduplicate across runs.

```
$ trufflehog git https://github.com/your-org/your-repo --
results=verified,unknown --sarif > results.sarif
```

— Named flag with a runnable CI command. v3.97.0

04 **Source archiving in Enterprise** NEW 80

A source can be archived to stop scanning while retaining all findings with triage history intact, accessible from the three-dot menu on a source row in Integrations; `show archived` reveals archived sources, and archiving requires admin or editor role. Archived sources are now filtered server-side in the sources list, keeping the integration list responsive for large deployments.

— Exact UI navigation, filter and role requirement given. product docs

05 **Slack scanning in Enterprise Edition** NEW 75

TruffleHog Enterprise adds Slack as a listed app in the Slack Marketplace, scanning messages, threads, and attached files for credentials and verifying each finding as live or not. Continuous scanning covers public channels, private channels the app is invited to, DMs, and group DMs, with edited messages rescanned and the app auto-joining new public channels. Historical scanning covers messages, threads, and attachments including archives, base64-encoded data, binaries, and Microsoft Office documents, with include/exclude lists to scope channel coverage and resumable scans after interruption.

— Full mechanism and scope described, but no exact setup steps given. product docs

06 **AWS credential analysis (`aws analyze`)** NEW 75

Adds `aws analyze` for customers on the Analyze SKU — resolves effective permissions, discovers IAM role chaining paths, and warns on IAM conditions and policy variables.

— Names the exact command and SKU with concrete analysis behaviour. product docs

07 **New detectors and Cloudflare format support in Enterprise scanning**

NEW 60

Adds 16 new detectors to Enterprise scanning: Braintrust, Cloudflare API Token, Cloudflare Global API Key, Cloudinary, Duffel, Duo, GitLab OAuth, HashiCorp Vault Token, HashiCorp Vault Batch Token, New Relic Insights Insert Key, Octopus Deploy, pganalyze Read Key, Pinecone, Red Hat Pyxis, Shippo, and SonarQube Cloud — the SonarQube Cloud scoped organization token, Duo API secret key, and Cloudflare detectors are new in both Enterprise and open source. Also adds recognition of the new

2026 prefixed Cloudflare credential formats for both the Cloudflare API Token and Global API Key detectors in Enterprise and open source.

— Names every detector and format but gives no config or flag to invoke. product docs

08 **Verification freezing for archived sources** NEW 60

Scanning stops immediately when a source is archived, including in-flight scans. Verification only freezes for a finding once every source behind it is archived; frozen findings retain their last verification status and are excluded from MTTR calculations.

— Explains exact freeze condition and MTTR effect. product docs

THINNER COVERAGE BELOW

09 **Encrypted private keys surfaced as findings** IMPROVED 55

Encrypted (passphrase-protected) private keys are now reported as findings even when the passphrase cannot be recovered, surfaced as unverified findings carrying `encrypted:true` metadata instead of being silently dropped.

— Names the metadata field but not a flag to control it. v3.97.0

10 **Domain-wide delegation for Google Drive integration** NEW 50

The add integration wizard for Google Drive now supports domain-wide delegation, enabling scanning with the access of all domain users in a single integration, for both hosted and self-hosted scanners.

— Names the exact wizard and scanner types but no setup steps. product docs

11 **Per-clone subdirectories for `--clone-path`** IMPROVED 50

Each clone under `--clone-path` now gets its own subdirectory, preventing collisions when scanning multiple repositories concurrently.

— Names the flag and the before/after behaviour. v3.97.0

12 **New detectors in v3.97.0** NEW 45

Adds New Relic License Key, New Relic Browser Key, New Relic User Key, New Relic Insights Query Key, New Relic Mobile App Token, Microsoft Teams Webhook V2, Figma Personal Access Token v3, and SolarWinds Observability detectors.

— Names all eight detectors but no usage detail. v3.97.0

13 **SharePoint historical version scanning** NEW 40

Scans historical versions of SharePoint items, surfacing credentials that were removed from the current version but remain present in earlier versions of items in generic lists

and search-driven document libraries.

— Explains scope but no configuration detail is given. product docs

14 **Android APK detection by content inspection** NEW 40

Detects Android package (APK) files by inspecting file contents rather than relying on the `.apk` extension, so packages surfaced through byte-stream enumeration are scanned.

— States the mechanism but no config or example. product docs

15 **Jira auto-close on remediation** NEW 40

Jira issues can now be configured, in integration settings, to close automatically when their underlying secret is remediated.

— Names the settings location but no field-level detail. product docs

16 **Location type filter on secrets list page** IMPROVED 40

Adds a location type filter update on the secrets list page, consolidated into a single popover view that explains what is being filtered.

— Names the exact UI location but little mechanism. product docs

17 **S3 source custom endpoint support** NEW 40

Adds custom endpoint support to the S3 source, enabling scans against S3-compatible storage services.

— Names the source and capability but no config example. v3.97.0

18 **Revoke access to a single shared secret** IMPROVED 35

Adds the ability to revoke access to a single shared secret rather than revoking the entire share.

— Thin bullet with no UI or API path named. product docs

19 **Expired verified JWTs excluded from findings** IMPROVED 35

Verified but expired JWTs are now excluded from findings at ingestion time, reducing noise from tokens that carry no remaining risk.

— Explains behaviour but no configuration or opt-out mentioned. product docs

20 **Editor-level notification triggering** IMPROVED 35

Editors can now trigger notifications directly, removing the previous dependency on admin-level access for routine notification workflows.

— States the permission change but no workflow detail. product docs

21 **Paginated user audit log** IMPROVED 35

The user audit log under Settings › Users now paginates, allowing navigation through long user histories.

— Exact UI path named but no further mechanism. product docs

22 **AWS Secrets Manager metadata type** NEW 25

Adds an AWS Secrets Manager metadata type for richer analysis output.

— Bare mention with no mechanism or example. v3.97.0

WAS THIS USEFUL?  

GitGuardian SaaS

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

GitGuardian SaaS is a cloud-based security platform that detects and prevents secret leaks in Git repositories and source code.

> GitGuardian shipped an API to set a custom validity status on secrets GitGuardian can't reach itself, an automated-detector option for custom source configuration, and a new Azure DevOps extension for one-click ggshield pipeline scanning.

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

01 **Custom validity status API for secrets** NEW 80

A new API endpoint lets practitioners set a custom validity status (`valid` or `invalid`) on secrets from internal or private systems, requiring a personal access token with the `secrets:write` scope. The status is displayed as Custom validity with the date it was set on every incident tied to that secret, surfaces in filters, severity and downstream incident workflows, and setting or removing it is recorded in both the incident activity log and the audit log.

– Names scope and behavior but not exact endpoint path

Set your own validity status on a secre…

THINNER COVERAGE BELOW

02 **Azure DevOps ggshield pipeline extension** NEW 55

A new Azure DevOps extension lets admins add a `ggshield` secret scan job to every pipeline in their organization via a single pipeline decorator applied through Azure Pipelines.

– Names tool and mechanism but no setup steps given product docs

03 **Detector-validators filter for custom sources** NEW 45

The public API now supports an Only use detectors with validators option when creating or updating custom sources, enabling fully automated and consistent source configuration.

– Names the option but no field name or example call

Set your own validity status on a secre…

└──▷ **ALSO FROM THESE RELEASES**

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

CAPEv2

Feedback ⓘ changes since 2026-08-09 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 integrates Google Magika deep-learning content-type detection alongside libmagic for headerless and obfuscated file identification.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-08-09
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-08-09
```

└──▷ USE IT

Enable Magika in your CAPE processing config to get deep-learning content-type labels on dropped files and decoded buffers that libmagic identifies only as `data`.

custom/conf/processing.conf

```
[magika]
enabled = yes
```

Adds `[magika]` section to `processing.conf` (via `processing.conf.default`) to enable/disable Google Magika content-type detection and tune confidence thresholds.

Adds get_magika() method to the `FileObject` class (`lib/cuckoo/common/objects.py`), returning a dict with `label`, `description`, `mime_type`, `group`, and `score` fields; runs below get_type() and never modifies the libmagic verdict.

Adds `magika` block to file info in CAPE's processing pipeline (`modules/processing/CAPE.py`), backfilling results for files retrieved from the MongoDB cache that predate Magika being enabled.

Adds clear_magika_cache() call at task boundaries in `CAPE.py` to prevent stale per-path predictions from persisting across tasks in long-lived workers.

Adds `magika` as a searchable hash/label type in the web search interface (`search.html`) and file info template.

+1 more

WAS THIS USEFUL?  

◆ Digital Forensics & Incident Response

ALEAPP

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

Android Logs Events And Protobuf Parser

> ALEAPP v2026.3.0 adds SQLCipher decryption for Signal, deep Telegram artifact parsing, and more than a dozen new app and drone parsers, alongside expanded Samsung and Google system database support and a hardening pass on HTML report generation.

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

01 **Stale-allowlist detection in claim language checker** NEW

70

Adds stale-allowlist detection and always-on `NOT CHECKED` reporting to `admin/scripts/check_claim_language.py`, closing a coverage-audit gap present in ALEAPP but not the other LEAPP cores.

— Names exact script path, a reader could run it directly v2026.3.0

02 **Telegram for Android artifact parsing** NEW

65

Adds Telegram for Android artifacts covering accounts, chats, peers, group/channel members, VoIP call logs, and service message actions parsed from `cache4.db` and shared preferences.

— Names source database and artifact scope, no runnable step v2026.3.0

03 **Snapchat arroyo.db conversation recovery** NEW

65

Adds Snapchat `arroyo.db` conversation parsing, including recovery of WAL-only rows with row-provenance columns.

— Names database and recovery mechanism, no usage steps given v2026.3.0

04 **Signal for Android decryption and message linking** NEW

60

Adds Signal for Android parsing with full SQLCipher decryption, attachment and thumbnail decryption, and message-linking.

— Names SQLCipher and mechanism but no exact command or path v2026.3.0

THINNER COVERAGE BELOW

05 **New app and drone artifact parsers across a dozen apps** NEW

55

Adds new parsers for Session (account parsing from shared preferences), Kik, BeReal (from app caches), JusTalk and JusTalk Kids (from the Realm store and MMKV stores for account identity and app state), Weibo (including bidirectional image-to-post linking), Xiaohongshu (RED), Slack, AllTrails, BRD (BreadWallet), Dropbox, what3words (via a vendored Realm parser), DJI drone flight tracks, Claude, MeWe (posts, comments, media, polls, reactions, and groups), and Smart Sidebar File Dock artifacts for OxygenOS/ColorOS.

— Each addition described in one line, no depth beyond names v2026.3.0

06 **Google Play services database parsing** NEW 55

Adds Google Play services database parsing covering constellation, account history, icing, and ODLH databases.

— Named databases but no parsing detail or usage step v2026.3.0

07 **Samsung system database parsing** NEW 55

Adds Samsung system database parsing covering Privacy Dashboard, WiFi, Sleep Detection, CMH tagging, and media provider.

— Named sub-areas but no mechanism or usage detail v2026.3.0

08 **HTML report injection and remote destination hardening** IMPROVED 55

Adds a CI guard for HTML report injection and remote destinations; evidence names are now escaped in report media tags and all remote CDN/URL destinations removed from HTML reports.

— Explains the fix mechanism, no config key or command given v2026.3.0

09 **ProtonVPN offline hostname resolution** IMPROVED 45

ProtonVPN hostname resolution no longer touches the network at parse time.

— States the behavior change but no further mechanism v2026.3.0

10 **Modernized UsageStats and Recent Tasks parsing** IMPROVED 25

Adds modernized UsageStats and Recent Tasks artifact parsing.

— Bare description, no explanation of what changed v2026.3.0

11 **Google Dialer and Calendar artifact parsing** NEW 25

Adds Google Dialer and Google Calendar artifact parsing.

— Bare name of two apps, no further detail v2026.3.0

WAS THIS USEFUL?  

iLEAPP

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

iOS Logs, Events, And Plist Parser

> iLEAPP v2026.3.0 adds Apple Unified Logs support and a rotated PowerLog archive parser that extends recoverable history by weeks, alongside new parsers for Telegram, Wickr, Session/Oxen, and seven other apps.

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

01 **Telegram artifact parsing for iOS** NEW 80

Adds Telegram artifacts for iOS: Accounts (from `accounts-metadata/atomic-state` and Postbox `t0`), Contacts & Peers (full `t2` peer dump with avatar images and Spotlight cache), Chats (Postbox chat list), Cached Peer Details (Postbox `t18`), Device Contacts (Postbox `t54`), Peer Presence, and Message Tags.

— Names every underlying database table and source path. v2026.3.0

02 **Rotated PowerLog archive (PLSQL.gz) parsing** NEW 75

Adds parsing of `*.PLSQL.gz` rotated PowerLog archives, decompressed to a session temp directory and reused across all PowerLog artifacts, alongside the existing `CurrentPowerlog.PLSQL`, with per-archive time-offset correction that extends recoverable history by weeks on most images.

— Names exact file formats and correction mechanism, no CLI step given. v2026.3.0

03 **Wickr artifact and identity recovery** NEW 65

Adds Wickr artifacts parsed from `wickrLocal.sqlite` and app logs, plus Wickr account identity recovered from the keychain.

— Names source database and keychain recovery, no exact fields. v2026.3.0

THINNER COVERAGE BELOW

04 **New app parsers for DuckDuckGo, MEGA, Onion Browser, what3words, Dubox/Terabox, Claude, Apple Intelligence Report, and Lidl Plus**

NEW 55

Adds new app parsers for DuckDuckGo, MEGA, Onion Browser (including Realm-based artifacts via a vendored Realm parser), what3words, Dubox/Terabox, Claude, Apple Intelligence Report, and Lidl Plus (SQLCipher + Keychain).

— Names each app and two decryption mechanisms, no field-level detail. v2026.3.0

05 **Apple Unified Logs support** NEW 45

Adds Apple Unified Logs support with the parser bundled into all builds, surfacing 16+ new artifacts identified from a predicate survey.

— States scope and count but not which artifacts. v2026.3.0

06 **CI hardening and expanded test coverage** NEW 45

Adds a CI tripwire guarding against unsupported claim language in artifact documentation, a CI guard for HTML report injection and remote destinations in generated reports, and CI test coverage for the Python 3.13 and 3.14 runtime contract plus Windows/Linux/ARM64 test builds.

— Describes CI scope but not user-facing usage. v2026.3.0

07 **Biome/SEGB coverage for edge selection, notifications, privacy accounting, and contact feedback**

NEW 40

Adds Biome/SEGB coverage for Networking Edge Selection, DuetExpertCenter notifications, Privacy Accounting streams, and Contact Autocomplete Feedback.

— Names four artifact areas without mechanism or format. v2026.3.0

08 **Session/Oxen artifact parsing** NEW 30

Adds Session/Oxen artifacts including decrypted messages, contacts, and attachments.

— Thin description, no source file or mechanism named. v2026.3.0

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Finite State Platform

Feedback ⓘ 1 RELEASE · 2026-08-12

BLOG ↗

Finite State Platform is a software supply chain security tool that analyzes firmware and software components for vulnerabilities and compliance risks.

> Finite State added an in-product What's New drawer for release highlights and made targeted reliability and usability fixes to license search and the License Policy page.

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

01 License search and policy page improvements IMPROVED

60

Improves license search so prefix matches appear before substring matches in the Add Licenses and Component license pickers, and updates the License Policy page to source the live license catalog and display a retryable error state when the catalog cannot load.

— Names exact UI surfaces and behavior change, no config keys

August 12, 2026 - What's New Drawer and…

THINNER COVERAGE BELOW

02 In-product What's New drawer NEW

40

Adds an in-product What's New drawer that surfaces release highlights from the documentation feed whenever new customer-facing updates are published.

— Describes what it does but no navigation detail given

August 12, 2026 - What's New Drawer and…

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Network Analysis & Monitoring

tcpreplay

Feedback ⓘ 1 RELEASE · 2026-08-15 NOTES ↗ CODE ↗

Edit and replay captured network traffic at arbitrary speeds — a suite of pcap tools for *NIX and Windows.

> tcpreplay 4.6.1 focuses on testing infrastructure — continuous OSS-Fuzz fuzzing, UBSan hardening, and a new unit test suite — while removing the `--xdp-batch-size` flag.

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

01 **Continuous OSS-Fuzz fuzzing pipeline** NEW 73

Adds three libFuzzer targets under `test/fuzz/` — `fuzz_pcap`, `fuzz_fragroute`, and `fuzz_services` — built against ASan/UBSan and submitted to OSS-Fuzz for continuous fuzzing of the pcap parser, the fragroute rules engine, and the tcpprep `--services` parser.

— Names exact fuzz targets and paths but no run instructions v4.6.1

02 **Dependency-free unit test suite** NEW 65

Adds a dependency-free unit test suite under `test/unit/`, runnable via `make check` or `ctest`.

— Gives exact commands to run the new suite v4.6.1

THINNER COVERAGE BELOW

03 **Removal of --xdp-batch-size flag** BREAKING 50

The `--xdp-batch-size` flag has been removed; any script or config that passes it will fail on upgrade.

— Names exact flag and upgrade impact, no migration path given v4.6.1

04 **UBSan checks and packed wire-format structs** IMPROVED 48

UBSan now runs alongside ASan in CI to catch misaligned-access bugs, and wire-format header structs are now marked `packed`, hardening packet parsers against a class of defects ASan alone misses.

— Explains mechanism but is internal CI/build change v4.6.1

└──▷ BREAKING ON UPGRADE

! The `--xdp-batch-size` flag has been removed; any script or config that passes it will fail on upgrade.

WAS THIS USEFUL?

◆ Firewalls & Network Defense

Sublime Security

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

Sublime Security is an email security platform that detects and prevents advanced threats like phishing, malware, and business email compromise attacks.

> Sublime Security introduced Tailored Phishing Simulations in Beta, a new capability for running environment-aware phishing campaigns with staggered delivery and cross-campaign tracking.

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

01 **Tailored Phishing Simulations (Beta)** NEW 65

New beta feature lets teams run phishing campaigns prioritized by the attack patterns Sublime is already detecting in the customer's environment. Campaigns support sending multiple templates per campaign with staggered delivery schedules, and a cross-campaign dashboard tracks click rates, report rates, and repeat clickers with recipient-level detail per campaign. Available to all customers during Beta; at GA it will move to a paid Sublime Tailored Phishing Simulations and Training add-on, enabled by contacting Sublime Support.

— Explains mechanism and scope but no config/API names given.

Platform v.1.88 (Tailored Phishing Simu…

WAS THIS USEFUL?

CLOUD & CONTAINER SECURITY

◆ Cloud Security Posture

Weaveworks eksctl

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

The official CLI for Amazon EKS

> eksctl v0.230.0 adds FIPS-compliant Bottlerocket support for managed nodegroups, introduces control plane component configuration, and relaxes required addons for IPv6 clusters.

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

01 **BottlerocketFips AMI family for managed nodegroups** NEW

55

Adds `BottlerocketFips` AMI family support in managed nodegroups, enabling FIPS-compliant Bottlerocket nodes.

— Names the AMI family but no config example given v0.230.0

02 **Optional vpc-cni and kube-proxy for IPv6 clusters** IMPROVED

55

Makes `vpc-cni` and `kube-proxy` optional when creating IPv6 clusters, reducing mandatory addon installation.

— Names the two addons and the cluster type affected v0.230.0

03 **Control plane component configuration** NEW

35

Adds control plane component config support, enabling configuration of individual EKS control plane components.

— No specifics on which components or config keys v0.230.0

WAS THIS USEFUL? 👍 👎

Sysdig Secure

Feedback ⓘ 1 RELEASE BLOG ↗

Sysdig Secure is a container security platform providing runtime monitoring, vulnerability scanning, and compliance enforcement for cloud-native applications.

> Sysdig Secure ships image layer-level vulnerability attribution across the Findings and Resource Details views, plus a new Vulnerability Policies evaluation tab and base image allow/deny list rules for rule bundles.

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

01 Per-layer vulnerability analysis for container images NEW

75

Adds a Layers tab in the Resource Details drawer for container images, listing every layer grouped by base image or Application Layers, with the Dockerfile instruction and per-severity vulnerability counts for each layer; selecting a layer filters the Vulnerabilities tab. Adds a Layer filter on the Findings view to slice the CVE list by a specific image layer directly. Adds Layer Number and Layer Command fields to the Findings Details drawer and findings panel showing which layer and Dockerfile instruction (e.g. `RUN apt-get install -y curl`) introduced the affected component — selecting the Layer Command opens the Layers tab for deeper inspection.

— Names exact tabs, filter and fields but only a UI path, no API.

Secure SaaS - August 13, 2026

02 Base image allow/deny list rules NEW

70

Adds Base Image Allow List and Base Image Deny List rules under Image Configuration in rule bundles, accepting base image pull strings to approve or block entire registries or namespaces in a single entry.

— Named rule types and config location, but no exact syntax given.

Secure SaaS - August 13, 2026

03 Vulnerability Policies tab in Resource Details drawer NEW

65

Adds a Vulnerability Policies tab in the Resource Details drawer, opened from Vulnerability Findings, showing how each resource was evaluated against configured Vulnerability Policies. A details panel per rule surfaces CVEs, affected components, and remediation guidance, and exceptions can be created directly from the drawer.

— Describes mechanism and scope but only points to a UI path.

Secure SaaS - August 13, 2026

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

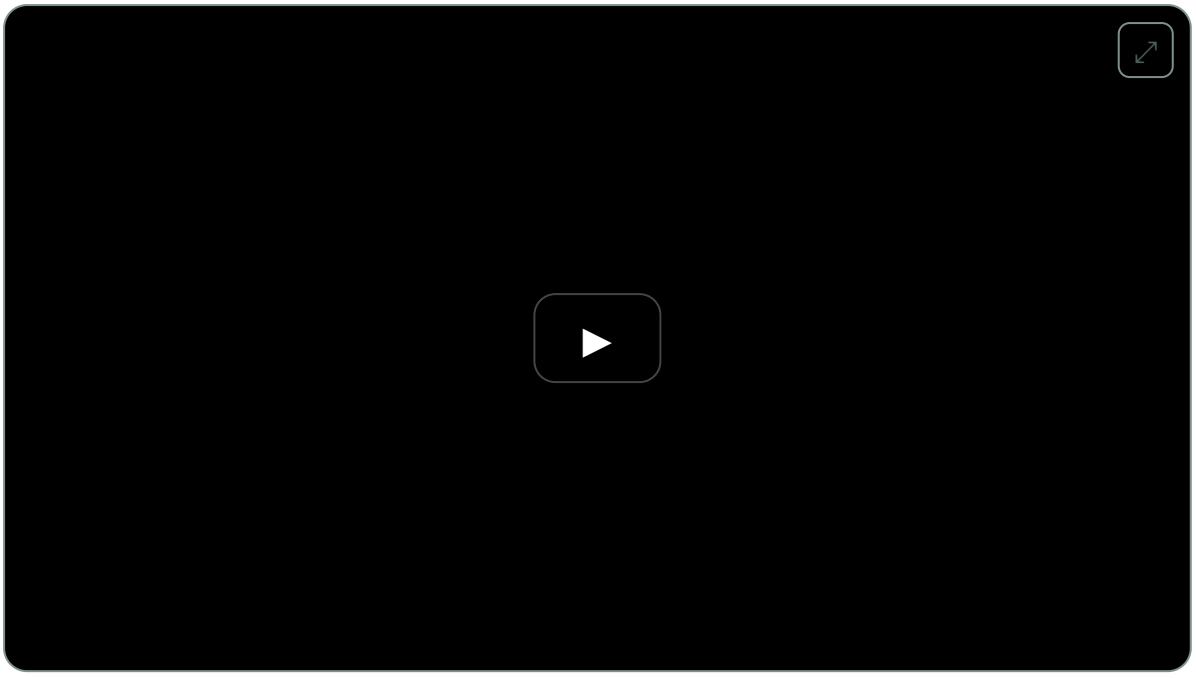
◆ SIEM

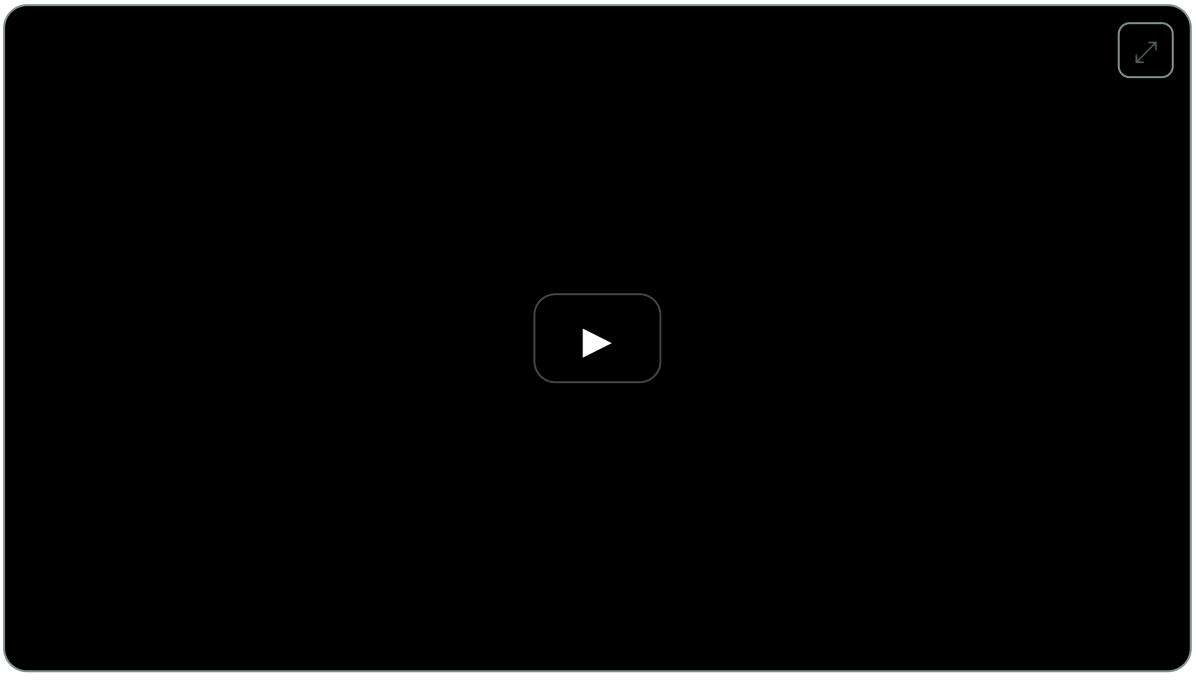
Google SecOps

Feedback ⓘ changelog · Aug 15, 2026 CHANGELOG ↗

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

Google SecOps gains Cloud Logging integration for SIEM ingestion pipeline monitoring and feed debugging (public preview, BYOP required).





└──▷ HOW TO FIND IT

Quickly jump from a failing feed to its pre-filtered Logs Explorer view to diagnose missing or delayed log delivery.

📍 In the Google SecOps console, go to Feed management, locate the feed you want to debug, and select `Debug with logs` to open Logs Explorer pre-filtered for that feed.

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Iru

Feedback ⓘ 1 RELEASE · SEEN 2026-08-15 NOTES ↗

> Iru shipped phased OS rollouts for Apple devices, an authenticated guest mode for SSO, and upfront compliance framework selection in Compliance Automation.

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

01 **Phased rollouts for Apple OS updates** NEW 45

Managed OS for Apple devices now supports phased rollouts, enabling staged OS update deployments across a fleet rather than pushing updates to all devices at once.

– Describes behaviour but no config or rollout controls named

What's New in Iru | August 14th

02 **Authenticated Guest Mode for SSO Library Item** NEW 45

Authenticated Guest Mode is now available on the Single Sign-On Library Item, allowing guest users to authenticate through SSO flows.

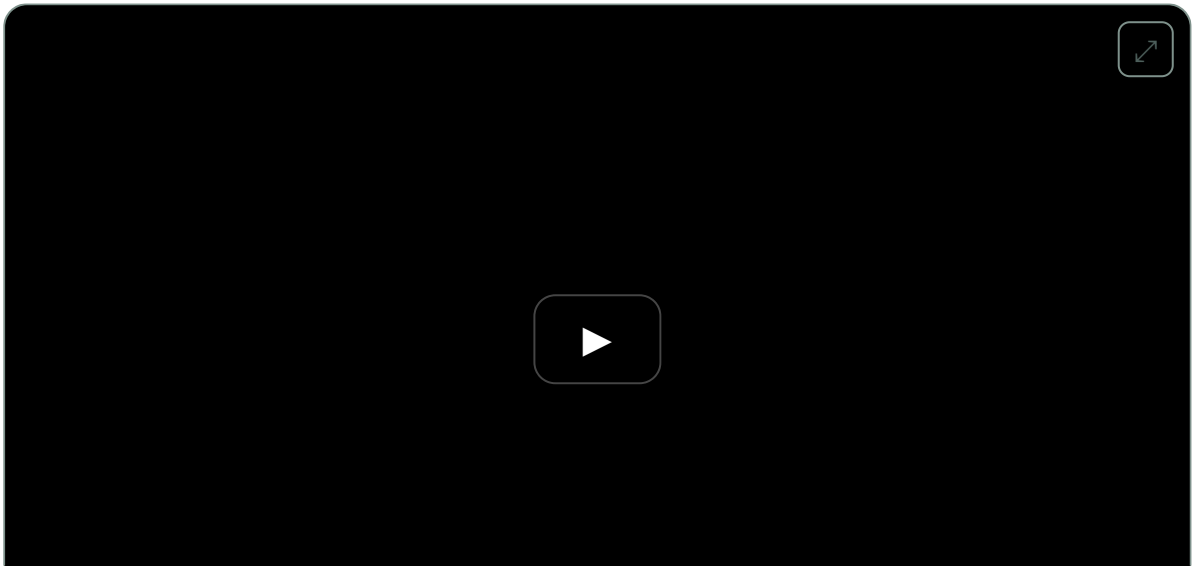
– Names the exact library item but no setup steps What's New in Iru | August 14th

03 **Upfront framework selection in Compliance Automation** IMPROVED 40

Compliance Automation now lets you select the compliance framework upfront when generating new policies, rather than configuring it after creation.

– Clear before/after but no named frameworks or UI path What's New in Iru | August 14th

└──▷ ALSO FROM THESE RELEASES





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 added OpenTelemetry-based instrumentation for Pydantic AI into its LLM Observability pipeline, and shipped a new Privacy 2.1 model family that scans telemetry faster and with less memory in Cribl Guard.

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

01 **Cribl Privacy 2.1 detection model family** IMPROVED 80

Cribl Privacy 2.1 introduces a new AI-based telemetry privacy detection model family for Cribl Guard, delivering 64%–163% higher detection throughput and 13%–24% lower average memory use versus Privacy 2.0, with F1 accuracy held steady or improved and CPU utilization unchanged. A new model pruning and distillation pipeline shrinks feed-forward layers by 75% or more and trains a smaller `student` model against the original frozen `teacher` to recover near-equivalent detection quality at higher throughput. It introduces a three-tier lineup — Fast (maximum throughput), Base (balanced), and Pro (maximum detection accuracy, coming soon) — so operators can pick an accuracy/throughput/resource tradeoff per pipeline, trained on data enriched with feedback from production Guard customers. It also previews `cribl-privacy-3.0-pro`, a forthcoming larger, high-accuracy model for lower-volume, accuracy-sensitive workloads, alongside additional personalization controls for tuning detections per environment.

— Concrete metrics, mechanism, and named model tiers given; no config command shown.

Cribl Privacy 2.1: More telemetry scann…

THINNER COVERAGE BELOW

02 **Pydantic AI instrumentation for LLM Observability** NEW 45

Cribl LLM Observability now supports instrumenting Pydantic AI applications via OpenTelemetry, with telemetry routed to Cribl Stream and Cribl Search.

— Names the integration but no config or setup steps given. product docs

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir

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

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

> Tenzir v6.11.0 ships a new `merge` operator, parallel pipeline execution, and extended windowing, alongside full Sigma v2.1 detection support with OCSF Detection Findings output, new YARA-X detection findings, `summarize` eventstats/streamstats modes, and new geo, math, and string functions.

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

01 Window operator gains new modes IMPROVED

95

Extends `window` operator with processing-time (omit `on`), count-based (unsigned `size`), and trailing windows (`trailing=true`), exposing the triggering event through `$window.event`; adds a `trigger` expression to selectively fire subpipelines and `tolerance` for bounded reordering of late events.

Compute a rolling 5-minute byte sum that travels with each event, enabling streamstats-style enrichment without leaving the pipeline.

```
$ window size=5min, trailing=true, on=ts {
  summarize rolling_bytes = sum(bytes)
  this = { ...$window.event, rolling_bytes: rolling_bytes }
}
```

— Full mechanism plus a runnable pipeline example

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

02 Full Sigma v2.1 rule support in sigma operator NEW

95

Adds `path=` and `rules=` named arguments to the `sigma` operator — `path=` accepts a file, directory, or mixed list (searched recursively) and `rules=` embeds YAML rule strings directly in a pipeline, with positional path now deprecated — and implements the full Sigma v2.1.0 detection rule surface, including `exists`, `cased`, `neq`, `windash`, `fieldref`, UTF-16 encoding modifiers, `re` sub-modifiers `i` / `m` / `s`, time-part modifiers `minute` through `year`, list-form OR conditions, and global filter rules targeting rules by `id` or `name` (or `rules: any` for all compatible sources). Also adds `tenzir.metrics.sigma` events emitted per batch reporting input event count, rule evaluations, and matches, and Sigma matches now output as OCSF Detection Findings.

– Every modifier and arg named but no runnable example shown

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

03 **Sigma engine breaking behavior changes** BREAKING 85

Passing a rule path positionally to the `sigma` operator now emits a deprecation warning (use `path=` instead); several Sigma modifiers that were previously silently misinterpreted now match correctly, changing results for rules relying on the old behavior; Sigma wildcards now treat all regex metacharacters literally and match across newlines (so `a+b` no longer accidentally matches `ab`); field names containing dots now deterministically prefer an exact top-level key before nested traversal; and Sigma rules with custom taxonomies or `expand` placeholder expansion are now rejected with an explicit diagnostic instead of being silently accepted.

– Clear before/after behavior but no migration command given product docs

04 **Parallel pipeline execution** NEW 80

Enables parallel pipeline execution via a `// parallelism: <n>` leading comment or the `--parallelism` flag on the `tenzir` binary, running CPU-bound operators across multiple cores while preserving per-operator semantics.

– Exact flag and comment syntax given, directly usable

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

05 **summarize operator eventstats and streamstats modes** NEW 75

Adds `output: 'events'` option to the `summarize` operator to attach final aggregate values to every event in a finite input (TQL equivalent of Splunk `eventstats`), and adds an incremental `summarize` output for running aggregates on event streams (cumulative event emission for unbounded streamstats-style use cases).

– Names the config key and comparison but no worked example

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

06 **geo_distance function** NEW 70

Adds geo_distance(lon1, lat1, lon2, lat2) function returning surface distance in meters, with optional `spheroid=true` for WGS-84 accuracy.

– Exact signature and optional flag given, no example call

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

07 **exp, log, and pow math functions** NEW 65

Adds exp(), log(), and pow() math functions to TQL; log() defaults to the natural base and accepts an optional second argument for arbitrary logarithm bases.

– Names functions and default base but no example usage

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

08 **Dynamic patterns in replace function** IMPROVED 65

Extends `replace` function to accept dynamic patterns and replacements (e.g. message.replace(user.name, '<redacted>')); literal calls with an empty pattern leave the subject unchanged.

— Includes example call and edge-case behavior Tenzir Node v6.11.0: Sigma 2.1 and YARA…

THINNER COVERAGE BELOW

09 **Merge operator for combining pipeline streams** NEW 55

Adds a `merge` operator that attaches an additional source subpipeline and contributes its events into the main event stream — the dual of `fork`.

— Names operator and its dual but gives no usage example

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

10 **YARA-X detection findings support** NEW 40

Adds support for the YARA-X detection engine, surfacing YARA-X scan results as structured findings.

— Only names the engine, no operator syntax or config given

Tenzir Node v6.11.0: Sigma 2.1 and YARA…

└──▷ **BREAKING ON UPGRADE**

! Passing a rule path positionally to the `sigma` operator now emits a deprecation warning; use `path=` instead.

! Several Sigma modifiers that were previously silently misinterpreted now match correctly, meaning rules relying on the old incorrect behavior will produce different (corrected) results.

! Sigma wildcards now treat all regex metacharacters literally and match across newlines, so patterns such as `a+b` no longer accidentally match `ab`.

! Field names containing dots in Sigma rules now deterministically prefer an exact top-level key before nested traversal, changing resolution behavior for ambiguous field names.

! Sigma rules with custom taxonomies or that use `expand` placeholder expansion are now rejected with an explicit diagnostic rather than silently accepted.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

OpenCVE

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

Vulnerability Intelligence Platform

> OpenCVE v3.1.0 introduces a REST API v2 with OpenAPI documentation and scoped tokens, adds project-level access control with membership roles, and extends organization audit logging.

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

01 **REST API v2 with OpenAPI docs** NEW 60

OpenCVE ships a new REST API v2 with OpenAPI documentation and scoped tokens for integrations. An API Docs link is now visible to all users in the header for easy access to the new documentation.

– Names API version and docs but no endpoint paths given. v3.1.0

02 **Project-level access control and membership roles** NEW 60

OpenCVE adds project-level access control with membership roles and enforced ACLs across CVE views and dashboard widgets, plus a new project members management UI for administering per-project roles.

– Explains scope and enforcement points, includes UI location. v3.1.0

THINNER COVERAGE BELOW

03 **Access mode and scopes for org API tokens** IMPROVED 45

Organization API tokens gain an access mode and fine-grained scopes, enabling tighter permission control on existing org-level tokens, alongside an improved organization API token management UI.

– Describes new controls but no concrete scope names or config keys. v3.1.0

04 **Extended organization audit logs** IMPROVED 35

Organization audit logs now cover automations, CVE comments, and project membership changes.

– Lists new log categories but no mechanism or access path. v3.1.0

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

runZero

Feedback ⓘ runZero 5.1.260814.1 (2026-08-14) BLOG ↗

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

runZero 5.1 expands IPMI vulnerability detection, improves asset geolocation, and re-evaluates service fingerprints via the Refingerprint action.

The `Refingerprint` asset action now also re-evaluates service fingerprints for improved accuracy.

Asset geolocation now uses location data from additional integration-sourced attributes.

The scanner detects a broader range of IPMI vulnerabilities.

Directory and device data are now processed concurrently, improving task performance.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Didit

Feedback

ⓘ 1 RELEASE · 2026-08-15 CHANGELOG ↗

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

> Didit's changelog this window centers on richer Session API data and verification branching, alongside expanded error signaling and React Native SDK compatibility.

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

01 **Barcode fields in Session API** NEW 67

The Session API now returns barcode data extracted from scanned ID documents, exposed via `id_verifications[].barcodes` in the session response returned by `GET /v1/sessions/<session_id>`.

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

```
$ curl -X GET 'https://api.didit.me/v1/sessions/<session_id>' \
  -H 'Authorization: Bearer <token>' \
  | jq '.id_verifications[].barcodes'
```

— Has a runnable example but no further mechanism detail changelog-20260815-e9f20562

THINNER COVERAGE BELOW

02 **Marital status branching in KYC** NEW 40

Verification workflows can now branch based on the `kyc.marital_status` field, letting sessions route logic depending on the applicant's marital status.

— Names the field but no example or config shown changelog-20260815-e9f20562

03 **502 error codes for KYB and phone checks** NEW 38

New 502 error codes have been added for KYB (Know Your Business) and phone verification flows to signal these check failures more clearly.

— Names error code and areas but no endpoint or sample response

changelog-20260815-e9f20562

04 **React Native 0.77 TurboModule support** IMPROVED 38

The Didit SDK now supports the React Native 0.77 TurboModule architecture, updating compatibility for React Native based integrations.

— States version support but no migration or code detail changelog-20260815-e9f20562

WAS THIS USEFUL? 👍 👎

Nudge Security

Feedback ⓘ 1 RELEASE · SEEN 2026-08-15 NOTES ↗

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

> Nudge Security launched a new Risk dashboard (open beta) that consolidates business criticality and residual risk scoring into a single view of SaaS risk posture.

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

01 Risk dashboard with trend chart and risk matrix NEW

75

A new Risk dashboard at `Dashboards > Risk` (open beta) shows tiles for apps assessed, business-critical apps, overall residual risk, and apps missing approval, built automatically from Nudge's business criticality tiers and residual risk scores. It includes a risk trend chart that can be filtered to business-critical apps only over time, a business criticality vs. residual risk matrix where clicking any cell reveals the exact apps behind that risk concentration, and a highest-risk apps ranking table ordering apps by risk score so reviewers can see what needs attention first.

Review your SaaS risk posture and identify apps missing approval or carrying high residual risk.

> 📍 In the Nudge Security console, go to Dashboards › Risk to view tiles for apps assessed, business-critical apps, overall residual risk, and apps missing approval.

— Names exact nav path and dashboard components; open beta scope only caveat

See your risk posture at a glance with …

WAS THIS USEFUL? 👍 👎

◆ Public Key Infrastructure (PKI)

DigiCert CertCentral

Feedback ⓘ 1 RELEASE · 2026-08-15

CHANGELOG ↗

DigiCert CertCentral is a platform for managing SSL/TLS certificates, automating issuance, renewal, and deployment across organizations.

> DigiCert CertCentral's biggest changes this window are compliance-driven: shorter certificate validity and reuse windows, mandatory CT logging, DNSSEC enforcement during DCV/CAA checks, and stricter Multi-Perspective Issuance Corroboration, alongside the retirement of legacy ACME, CertCentral Discovery and Managed Automation, plus smaller usability additions like CSR-driven CN/SANs population and new DCV polling status fields in the Services API.

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

01 **DCV polling status in Services API** NEW 88

Adds `dcv_polling_status` and `dcv_polling_last_checked` fields to the Services API responses for certificate requests and DNS name validations, surfacing whether DigiCert's automatic DCV polling is active and when it last ran.

Check whether DCV polling is active and when it last ran for a pending certificate order, to decide whether to trigger a manual validation check.

```json
// Example Services API response fragment for a pending order
{
  "dcv_method": "dns-txt-token",
  "dcv_token": {
    "token": "{{dcv_token}}",
    "status": "pending",
    "expiration_date": "2022-02-24T16:25:52+00:00"
  },
  "dcv_polling_status": true,
  "dcv_polling_last_checked": "2026-07-09T22:07:24Z"
}
```

— Named API fields with a runnable example response shown changelog-20260815-5e950aeb

02 **Use CN/SANs from CSR on certificate requests** NEW 70

Adds a Use CN/SANs from CSR option on certificate request forms (DV, OV, EV TLS; QWAC and QWAC PSD2; X9 PKI for TLS; Private TLS) to replace existing Common Name and SANs fields with values from the uploaded CSR, with an Undo option to revert.

— Names exact UI controls and certificate types covered changelog-20260815-5e950aeb

03 **Mandatory CT logging for public TLS certificates** BREAKING

70

Removes CT logging opt out options from CertCentral, enforcing CT logging on all public TLS certificates (DV, OV, EV, QWAC, QWAC PSD2) across all DigiCert brands. Certificates that previously needed to be kept out of CT logs can no longer be issued that way from June 1, 2026 onward.

— Names cert types and effective date, no config surface changelog-20260815-5e950aeb

04 **Certificate validity and validation reuse periods reduced**

70

BREAKING

Moves public TLS certificate maximum validity to 199 days and public code signing certificate maximum validity to 459 days, and reduces domain validation and organization validation reuse periods to one day less than the CA/Browser Forum permitted maximums, per CA/Browser Forum requirements. Orders requesting longer validity periods will be rejected or truncated.

— Exact day limits named, though no config key given changelog-20260815-5e950aeb

05 **DNSSEC validation during DCV and CAA checks** BREAKING

65

Begins DNSSEC validation during domain control validation and DNS CAA checks, returning DNSSEC-specific error and status messages in CertCentral order details pages and via the Services API when issues are detected. Domains with misconfigured DNSSEC will block certificate issuance until the configuration is resolved.

— Explains mechanism and effect, no exact error codes named changelog-20260815-5e950aeb

06 **Stricter Multi-Perspective Issuance Corroboration** BREAKING

65

Updates Multi-Perspective Issuance Corroboration (MPIC) to require corroboration from at least four remote network locations across at least two different Regional Internet Registry regions. Certificate issuance fails if more than one perspective does not corroborate.

— Precise numeric requirements named, no reader action available

changelog-20260815-5e950aeb

07 **Legacy CertCentral ACME service retirement**

BREAKING

60

Retires the legacy CertCentral ACME service, with legacy ACME credential warnings shown in the ACME credentials table and requests to legacy ACME URLs failing to issue certificates after the retirement date. Legacy credentials must be replaced with new CertCentral ACME credentials.

— Names exact UI location and required credential replacement changelog-20260815-5e950aeb

THINNER COVERAGE BELOW

08 **CertCentral Discovery and Managed Automation retirement**

DEPRECATED

55

Announces retirement of CertCentral Discovery and Managed Automation services on October 1, 2026, with capabilities moving to DigiCert Trust Lifecycle Manager. Integrations relying on these services will stop working after that date.

— Exact retirement date and successor product named changelog-20260815-5e950aeb

09 **Simplified delegation request workflow** IMPROVED

50

Removes the organization selection step from the delegation request workflow — a single delegation request now applies to all organizations the selected approver represents and all certificate products they are authorized to approve.

— Explains before/after but no UI path or API named changelog-20260815-5e950aeb

10 **Client Authentication EKU removed from public TLS certs**

45

BREAKING

Public TLS certificates will no longer be issued with the Client Authentication EKU; existing workflows that depend on this EKU on newly issued certificates will break.

— Names the EKU but gives no migration guidance changelog-20260815-5e950aeb

11 **Editable Service User names** IMPROVED

35

Enables updating a Service User's first and last names in CertCentral, making it easier to distinguish Service Users and assign them to teams or operational roles.

— Thin description, no exact UI path or field name given changelog-20260815-5e950aeb

12 **TLS issuance moves to X9 PKI hierarchies** IMPROVED

25

Transitions TLS certificate issuance to DigiCert's X9 PKI hierarchies.

— Single sentence, no detail on impact or migration steps changelog-20260815-5e950aeb

└──▷ ALSO FROM THESE RELEASES

└──▷ **BREAKING ON UPGRADE**

! Public TLS certificates will no longer be issued with the Client Authentication EKU; existing workflows that depend on this EKU on newly issued certificates will break.

! CT logging opt-out options are removed from CertCentral — certificates that previously needed to be kept out of CT logs can no longer be issued that way from June 1, 2026 onward.

! Requests to legacy ACME URLs will fail to issue a certificate after the legacy CertCentral ACME service retirement date; legacy credentials must be replaced with new CertCentral ACME credentials.

! CertCentral Discovery and Managed Automation services retire on October 1, 2026 — integrations relying on these services will stop working.

! TLS certificate validity is capped at 199 days and code signing certificate validity at 459 days; orders requesting longer validity periods will be rejected or truncated.

! DNSSEC validation is now enforced during DCV and CAA checks — domains with misconfigured DNSSEC will block certificate issuance until the configuration is resolved.

! MPIC now requires corroboration from at least four remote network locations across at least two RIR regions — certificate issuance fails if more than one perspective does not corroborate.

WAS THIS USEFUL?  

◆ Secure Access (VPN / SASE / ZTNA)

Twingate

Feedback ⓘ 1 RELEASE · SEEN 2026-08-15 NOTES ↗

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

> Twingate extended its Identity Firewall to protect internal HTTP web apps, authenticating users at the Twingate layer before they reach the app.

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

01 **HTTP web app support in Identity Firewall** NEW 70

Twingate Identity Firewall now supports HTTP web apps, enabling protection of internal web apps, data portals, CI/CD infrastructure, and AI-coded business apps. It authenticates users at the Twingate layer using JWT and custom header injection, supporting MFA, JIT access, security policies, and access logging, so users reach internal web apps already authenticated.

— Names JWT and header injection but no config key or command.

Privileged Access for Web Apps

WAS THIS USEFUL? 👍 👎

Firezone

Feedback ⓘ 1 RELEASE · 2026-08-15 BLOG ↗

Firezone is an open-source VPN and secure access platform that replaces traditional corporate VPNs with modern zero-trust networking.

> Firezone's macOS/iOS client gains a command-line interface aimed at unattended or remotely managed Macs.

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

01 **New `firezone-cli` for macOS Client** NEW 45

Adds `firezone-cli`, a command-line interface to the macOS Client intended for Macs that run unattended or are managed remotely.

— Names the CLI but no subcommands or usage shown. macOS / iOS 1.5.19 (2026-08-15)

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Continuous Control Monitoring (CCM)

RegScale

Feedback ⓘ RegScale 6.33.1.0 (2026-08-14) BLOG ↗

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

RegScale 6.33.1.0 expands AI Agent coverage to Lines of Inquiry, CAPs, Risk Treatment, Questionnaire Scoring, and more.

Extends AI Agent capabilities to Lines of Inquiry, Corrective Action Plans, Risk Treatment Mappings, Questionnaire Scoring, Lightning Assessment Tests, and Security Control Test Generation.

Adds RegML Mapping Recommendations (details not fully disclosed in release notes).

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Daytona

Feedback ⓘ 1 RELEASE · SEEN 2026-08-15 NOTES ↗

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

> Daytona 0.201.0 adds direct S3-compatible file transfers via pre-signed sandbox URLs, along with typed SDK error codes and a date-range filter for sandbox listing.

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

01 **Pre-signed sandbox file transfer URLs** NEW 50

Introduces pre-signed sandbox file download and upload URLs, allowing direct S3-compatible file transfers without proxying through the API.

— Explains mechanism but no exact API call or flag given snapshot-20260815

02 **Sandbox list date-range filter** NEW 20

Adds a date-range filter to the sandbox list functionality, per the release summary, with no further detail on parameters or endpoint given.

— No parameter names or endpoint path provided snapshot-20260815

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ 1 RELEASE · 2026-08-14 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 v2.1.233 adds GitLab merge request support, an opt-in apps gateway identity-forwarding setting, a Linux memory limit for Bash tool commands, and a WebFetch cache TTL knob, while also disabling todo/task-tracking tools by default on newer models.

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

01 Todo/task-tracking tools disabled by default on newer models

85

BREAKING

Todo/task-tracking tools (`TaskCreate` , `TaskGet` , `TaskUpdate` , `TaskList` , `TodoWrite`) are no longer available by default on Opus 4.8, Sonnet 5, Fable 5, Mythos 5, and newer models; set `CLAUDE_CODE_ENABLE_TODO_TOOLS=1` to restore them.

Re-enable todo/task-tracking tools on newer models that disable them by default.

```
$ CLAUDE_CODE_ENABLE_TODO_TOOLS=1 claude
```

— Names exact tools, affected models, and restore command. v2.1.233

02 Memory limit for Bash tool commands NEW

82

Adds the `CLAUDE_CODE_TOOL_MEMORY_LIMIT` environment variable (Linux only) to cap memory for Bash tool commands via cgroup support, preventing a runaway build from stalling a session.

Cap memory for Bash tool commands on Linux to prevent a runaway build from consuming all RAM and stalling the session.

```
$ CLAUDE_CODE_TOOL_MEMORY_LIMIT=2G claude
```

— Named env var with runnable example and cgroup mechanism. v2.1.233

03 WebFetch cache TTL control NEW

72

Adds the `CLAUDE_CODE_WEBFETCH_CACHE_TTL_MS` environment variable to control the WebFetch session URL cache TTL, defaulting to 15 minutes.

Shorten the WebFetch URL cache lifetime so repeatedly-fetched pages stay fresh during a long session.

```
$ CLAUDE_CODE_WEBFETCH_CACHE_TTL_MS=300000 claude
```

— Named env var, default value, and runnable example. v2.1.233

04 **GitLab merge request support in worktree and agents view**

IMPROVED 65

Extends the `--worktree` flag and the `claude agents` view to accept GitLab merge request URLs, displaying MRs as `!N`.

— Names flag, view, and display format for MRs. v2.1.233

05 **Skills directory validation in plugin validate** IMPROVED

65

Improves `claude plugin validate` to inspect a bare `.claude/skills` directory and report SKILL.md files with unparseable frontmatter.

— Names command and directory/file targets. v2.1.233

06 **Per-user identity forwarding on apps gateway** NEW 60

Adds an opt-in `forward_user_identity` apps gateway setting on Anthropic upstreams, forwarding the signed-in user's identity as headers so a proxy can attribute spend per user.

— Named setting and purpose, but no exact config path shown. v2.1.233

07 **Unrecognized model diagnostic in print mode** NEW 60

Adds a `[claude-code:unrecognized_model]` diagnostic line to stderr in print mode when a request is sent for an unrecognized model ID; suppress it with `modelOverrides`.

— Exact diagnostic string and suppression key named. v2.1.233

08 **Apps gateway upstream error forwarding** IMPROVED 60

Improves apps gateway error forwarding so 400/413 responses from Vertex, Foundry, and Claude Platform on AWS upstreams now surface the upstream's own error message.

— Names upstreams and status codes but no user action. v2.1.233

THINNER COVERAGE BELOW

09 **Faster self-hosted-runner session start** IMPROVED 55

Improves `claude self-hosted-runner` session start time by creating the session branch without rewriting the working tree and eliminating two blocking server round trips.

— Mechanism described but no numbers or user action. v2.1.233

---

10 **Screen reader mode improvements** IMPROVED 50

The `/effort` selector now renders as a numbered list with a typed-number prompt, and hint and dialog text is no longer clipped in screen reader mode.

— Names the /effort selector as a starting point. v2.1.233

---

11 **GitHub app setup tip suppressed for non-GitHub hosts** IMPROVED 45

Suppresses the GitHub app setup tip in repositories whose origin remote is on gitlab.com or bitbucket.org; the enterprise marketplace tip now covers non-GitHub internal git hosts.

— Named hosts but purely cosmetic, no action needed. v2.1.233

└──▷ **BREAKING ON UPGRADE**

! Todo/task-tracking tools (`TaskCreate` , `TaskGet` , `TaskUpdate` , `TaskList` , `TodoWrite`) are no longer available on Opus 4.8, Sonnet 5, Fable 5, Mythos 5, and newer models; set `CLAUDE_CODE_ENABLE_TODO_TOOLS=1` to restore them.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI

Feedback ⓘ

3 RELEASES · 2026-08-13 → 2026-08-14 NOTES ↗ CODE ↗

Lightweight coding agent that runs in your terminal

> OpenAI Codex CLI's alpha builds this window centered on Guardian V2 risk scoring and its new configuration surface, defaulted unified exec on Windows, expanded gRPC code-mode and MCP protocol handling, and added thread usage monitoring alongside numerous smaller reliability and telemetry changes.

└──▷ **WHAT SHIPPED · 17 FEATURES** most completely described first ? what's the number?

01 **Guardian V2 risk classification and scoring** NEW 88

Guardian V2 introduces risk scoring on threads that flags high-risk actions for automatic review, recording scores as snapshots, initializing per-thread samplers, and preserving transcript boundaries. `features.guardianv2` now accepts either a boolean toggle or an object configuring classifier instructions, review threshold, reasoning effort, action and instruction token limits, transcript source inclusion, per-entry and total token budgets, and the number of recent non-user entries.

— Names exact config key and its full parameter set. rust-v0.148.0-alpha.18 rust-v0.148.0-alpha.16

02 **Unified exec enabled by default on Windows** BREAKING 75

`unified_exec` is now enabled by default on all platforms including Windows, exposing `exec_command` and `write_stdin` on Windows instead of the previous `shell_command`.

— Named exact commands but no usage example given. rust-v0.148.0-alpha.18

03 **Thread usage monitoring in status and terminal title** NEW 73

Adds `thread-credits` and `estimated-thread-cost` as configurable terminal title items for monitoring per-thread quota consumption, shows estimated thread usage in the `/status` command and TUI status surfaces, and adds per-thread usage queries to the backend client.

— Names exact title items and status command. rust-v0.148.0-alpha.11

04 **gRPC code-mode host support in app server** IMPROVED 65

The app server now supports gRPC code-mode hosts alongside existing WebSocket-based hosts, and removes the gRPC code-mode open session limit so a host can register more than `MAX_IN_FLIGHT_REQUESTS` open sessions simultaneously.

— Names the session-limit constant but no config path. `rust-v0.148.0-alpha.18` `rust-v0.148.0-alpha.11`

05 **agent_name in Responses API turn metadata** NEW 65

Adds `agent_name` to Responses API turn metadata, falling back to `/root` when no canonical agent path is set, with the field reserved from client metadata overrides.

— Names field and fallback path, but no API call example. `rust-v0.148.0-alpha.16`

06 **MCP server protocol discovery and header customization**

IMPROVED 63

MCP servers gain dynamic HTTP header helpers enabling per-request header customization, and MCP client protocol discovery is now instrumented with a counter and duration metric tagged `legacy` or `auto` mode and classified as `modern`, `legacy`, or `failure` outcomes.

— Names metric tags but no way to query them directly. `rust-v0.148.0-alpha.18` `rust-v0.148.0-alpha.11`

THINNER COVERAGE BELOW

07 **Strict auto-review for MCP tool calls** BREAKING 50

Codex now enforces strict auto-review for MCP tool calls even when the approval policy, tool annotations, or a remembered session decision would otherwise skip review.

— Describes behavior change without exact policy names. `rust-v0.148.0-alpha.16`

08 **tool_search for plugin app discovery** NEW 50

The model is now instructed to use `tool_search` to discover relevant tools within a selected plugin's apps before falling back to unrelated or built-in tools.

— Names tool but no configuration surface for users. `rust-v0.148.0-alpha.16`

09 **MCP connection reuse during reconciliation** IMPROVED 45

Reuses active, uncancelled pending MCP connections during reconciliation when identity, catalog limit, and protocol mode still match, avoiding redundant startups.

— Explains condition and benefit, no exposed control. `rust-v0.148.0-alpha.16`

10 **Manifest-defined metrics for plugin scripts** NEW 40

Adds manifest-defined metrics and bounded measurement analytics for trusted plugin scripts and shell commands.

— Names capability but no manifest schema or command. `rust-v0.148.0-alpha.11`

11 **Workload identity integration for authentication** NEW

30

Integrates workload identity with Codex authentication for enterprise environments.

— No mechanism or config details given. `rust-v0.148.0-alpha.11`

12 **Model ETag tracking via WebSocket metadata** IMPROVED

30

Reads model ETags from WebSocket metadata events for more efficient model version tracking.

— Mechanism named but no user-facing control. `rust-v0.148.0-alpha.11`

13 **Non-interactive approval policy for delegate sessions**

30

BREAKING

Enforces non-interactive approval policy for Codex delegate sessions.

— Names policy but no config key or example. `rust-v0.148.0-alpha.11`

14 **Implicit skill invocation detection from PowerShell** NEW

25

Detects implicit skill invocations from PowerShell reads.

— No trigger conditions or command surface described. `rust-v0.148.0-alpha.11`

15 **Startup override to skip project configuration** NEW

20

Adds an override to skip project configuration at startup.

— One-line description, no flag name or mechanism given. `rust-v0.148.0-alpha.18`

16 **Conversation history creation timestamps** NEW

15

Stamps conversation history items with creation times.

— Bare description, no format or field name. `rust-v0.148.0-alpha.11`

17 **Durable reverts for paginated threads** NEW

15

Adds durable reverts for paginated threads.

— No mechanism or command given. `rust-v0.148.0-alpha.11`

WAS THIS USEFUL? 👍 👎

Cline

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

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

> Cline added web search for supported models across its extension and desktop app, introduced a font size slider in desktop Settings, and expanded its model catalog with Crusoe as a new provider alongside updated per-provider default models.

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

01 **Web search for supported models** NEW 65

Adds a Web Search toggle in Settings (desktop) and Feature Settings (extension/CLI) that lets supported models search the web during a task; search calls and results are recorded in the session transcript/conversation and persist across reopens and reloads.

— UI toggle path given, but no API or flag named desktop-v0.0.13 v4.1.10

THINNER COVERAGE BELOW

02 **Font size slider in desktop Settings** NEW 55

Adds an app font size slider in Settings that scales the interface and applies the chosen size before the window paints, eliminating the flash-to-old-size on launch.

— UI path clear, no config key or default value desktop-v0.0.13

03 **Crusoe model provider and catalog updates** NEW 45

Adds Crusoe as a new model provider in the catalog, alongside updated model lists and per-provider default models across all existing providers.

— Provider named but no setup or config steps given v4.1.10 desktop-v0.0.13

WAS THIS USEFUL? 👍 👎

Pi

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

AI agent toolkit: unified LLM API, agent loop, TUI, coding agent CLI

> Pi's v0.84.2 release (mirrored across several docs updates) bundles fullscreen transcript search, granular tool-selection controls, per-run theming, SDK model-catalog caching with offline mode, and new extension/gateway APIs.

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

01 Model catalog caching and offline mode in SDK NEW

91

ModelRuntime.create() now accepts `allowModelNetwork` and `modelRefreshTimeoutMs` to opt into a network refresh of the model catalog at creation time and bound how long it may take, plus `modelsStorePath` to set a custom local catalog store path and a `modelsStore` injection point for a custom persistence implementation. A new modelRuntime.refresh({ allowNetwork: true, force: true, signal }) method forces an immediate catalog refresh, bypassing the default four-hour per-provider throttle, and the `PI_OFFLINE` environment variable disables all model network access.

Refresh the model catalog at startup with a 15-second timeout — useful in CI or containerized environments where stale catalogs would cause failures.

javascript

```
const refreshedRuntime = await ModelRuntime.create({
  allowModelNetwork: true,
  modelRefreshTimeoutMs: 15_000
});
```

Force an immediate catalog refresh mid-session, bypassing the four-hour throttle — useful after a provider publishes a new model you need right away.

javascript

```
await modelRuntime.refresh({ allowNetwork: true, force: true,
signal });
```

— Full API surface named with runnable examples product docs

02 Tool selection controls: defaultTools and CLI flags NEW

88

A new `defaultTools` config key selects which built-in tools are enabled at startup (e.g. `["bash", "edit", "write"]`, or an empty array for none), while extension and

SDK custom tools remain enabled regardless. New CLI flags refine this: `--tools` acts as a strict allowlist for all tools (replacing `defaultTools` behavior), `--no-tools` disables all tools, `--no-builtin-tools` disables only built-in defaults, and `--exclude-tools` filters the resulting tool list after other tool flags are applied.

Start Pi with only bash and edit available at launch, keeping any extension tools active — useful for restricting the default toolset in a project.

```json
{"defaultTools": ["bash", "edit"]}
```

Set the default built-in tool selection for all projects globally so every new session starts with only the tools you want.

```json
{
  "defaultTools": ["read", "bash"]
}
```

— Every flag and config key named with example v0.84.2

03 Per-run theme selection with --use-theme flag NEW

80

The `--use-theme` flag starts an interactive run with a specific theme without changing the saved setting, e.g. `pi --use-theme light`, and supports `light/dark` syntax such as `pi --use-theme light/dark` to follow the terminal's appearance for that run.

Start a single session with the light theme without overwriting your saved theme preference.

```
$ pi --use-theme light
```

Follow the terminal's light/dark appearance for one run without changing the saved setting.

```
$ pi --use-theme light/dark
```

Apply a specific theme for a single Pi session without overwriting your saved theme preference.

```
$ pi --use-theme dark/compact
```

Lock in a specific visual theme for a single interactive session without overwriting your saved preferences.

```
$ pi --use-theme dracula interactive
```

Apply a one-off theme for a single session without permanently changing your saved theme preference.

```
$ pi --use-theme dracula
```

— Exact flag and syntax with runnable commands v0.84.2

04 **Fullscreen transcript search and scrolling** NEW 78

Fullscreen mode gains transcript search via `Ctrl+Shift+F` with incremental match highlighting and next/previous navigation using `Enter` / `Ctrl+G` and `Shift+Enter` / `Ctrl+Shift+G`. Search-match highlight colors are customizable via new `searchMatchBg` and `searchMatchText` theme color tokens, with documented fallback behavior when omitted. Fullscreen mode also adds unbound single-line transcript scrolling actions.

Add search-match highlight colors to a custom theme; omit these tokens to use the built-in fallbacks.

```json
"searchMatchBg": "#2d2d30",
"searchMatchText": ""
```

— Named keybindings and theme tokens, but no CLI command v0.84.2

05 **Experimental strict JSON-schema constrained sampling** NEW 78

Under `PI_EXPERIMENTAL=1`, the default `read`, `bash`, `edit`, and `write` tools support experimental strict JSON-schema constrained sampling for schema-conformant outputs.

Enable strict JSON-schema constrained sampling for the built-in read/bash/edit/write tools to get schema-conformant outputs during experimentation.

```
$ PI_EXPERIMENTAL=1 pi interactive
```

Enable experimental strict JSON-schema constrained sampling for built-in tools to improve structured output reliability.

```
$ PI_EXPERIMENTAL=1 pi
```

— Named env var and tools with runnable command v0.84.2

06 **Configurable fullscreen exit output** NEW 68

The `fullscreenExitOutput` config setting controls what prints when exiting fullscreen mode: `"transcript"` prints the final transcript and resume hint, while `"resume-hint"` restores the previous screen and prints only the resume hint; it has no effect in regular TUI mode.

— Exact config key and values, no example shown v0.84.2

THINNER COVERAGE BELOW

07 **expandPromptTemplates option for sendUserMessage** NEW 55

The extension pi.sendUserMessage() API gains an `expandPromptTemplates` option for explicitly dispatching commands and expanding skills and prompt templates.

— Names exact API and option, no example v0.84.2

08 **PI_TUI_ESC_TIMEOUT for Escape-key latency** NEW 55

New `PI_TUI_ESC_TIMEOUT` environment variable tunes Escape-key detection latency for high-latency terminals such as over SSH.

— Named env var, purpose given, no example v0.84.2

09 **Cloudflare AI Gateway binding fetch helper** NEW 54

createGatewayBindingFetch() routes Cloudflare AI Gateway requests through a Workers AI binding without needing an API token.

— Named function, no usage example given v0.84.2

10 **Windows path syntax documented for JSON config** IMPROVED 50

Docs now clarify that Windows paths in JSON config (e.g. `shellPath`) must use forward slashes (`C:/Program Files/Git/bin/bash.exe`) or escaped backslashes (`C:\\Program Files\\Git\\bin\\bash.exe`).

— Names exact config key and path formats `product docs`

11 **AssistantMessage.endTurn field for Codex diagnostics** NEW

48

A new `AssistantMessage.endTurn` field preserves OpenAI Codex's terminal `end_turn` signal for diagnostics.

— Named field, no example or mechanism detail `v0.84.2`

12 **AI_AGENT=pi process marker documented** IMPROVED

38

Docs now document the `AI_AGENT=pi` process marker and how it differs from `PI_CODING_AGENT=true`.

— Names both markers but no usage detail `v0.84.2`

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Agentic Infrastructure

Ollama

Feedback ⓘ 2 RELEASES · 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 added support for the new Qwen3.8 model family, including a 27B variant with an Apple Silicon MLX build, and refined how developer instructions from OpenAI-compatible coding agents are handled for Qwen3.8.

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

01 **Qwen3.8 27B model with MLX variant** NEW 80

Adds the `qwen3.8:27b` model, with performance gains across coding, research, and long-horizon agentic tasks, runnable via `ollama run qwen3.8:27b`. Also adds a `qwen3.8:27b-mlx` variant optimized for maximum performance and output quality on Apple Silicon devices, runnable via `ollama run qwen3.8:27b-mlx`.

Run the Qwen3.8 27B MLX variant on Apple Silicon for high-quality coding agent workloads.

```
$ ollama run qwen3.8:27b-mlx
```

— Names exact model tags and runnable commands. v0.32.12

THINNER COVERAGE BELOW

02 **Qwen3.8 developer-instruction folding** IMPROVED 45

Adds Qwen3.8 support for developer instructions: leading system/developer instruction prefixes sent by OpenAI-compatible coding agents are folded into a single system turn before Qwen3.8 validation, preserving instruction precedence without affecting Qwen3.5 or other renderers.

— Explains mechanism but gives no command or config to try. v0.32.13

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

NVIDIA OpenShell

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

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

> OpenShell v0.0.106 ships official Go and TypeScript SDKs and adds cert-manager and OpenShift Route support to its Helm chart deployment.

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

01 **cert-manager and OpenShift Route support in Helm chart** NEW

40

The Helm chart deployment adds cert-manager external issuer integration and OpenShift passthrough Route support.

— Names the integrations but no config keys or flags given v0.0.106

02 **Go SDK for OpenShell** NEW

35

A complete Go SDK is now available, including domain clients, authentication, and hardening support.

— Names components but no install path or repo link v0.0.106

03 **TypeScript SDK published as @nvidia/openshell-sdk** NEW

35

A TypeScript SDK is now published as `@nvidia/openshell-sdk`.

— Names the exact package but no usage detail v0.0.106

WAS THIS USEFUL? 👍 👎

ToolHive

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

ToolHive is an enterprise-grade platform for running and managing Model Context Protocol (MCP) servers.

> ToolHive v0.43.0 adds plugin lifecycle tracking to lock files, OAuth confidential client and RFC 8693 delegate client support, and makes push signing the default while removing the lock feature's opt-in gate, alongside smaller improvements to certificate provenance enforcement, rate-limiting observability, SSE parsing, and LLM tooling.

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

01 **Plugin lifecycle tracking in lock files** NEW 65

The lock file schema now includes a `plugins` key so plugin entries are tracked in lock provenance, backed by a new `PluginLockService` and a managed install flag for plugin lifecycle management.

— Names schema key and service but no usage example given v0.43.0

THINNER COVERAGE BELOW

02 **Build fingerprint removed from proxy health endpoint** IMPROVED 55

Drops the build fingerprint from the proxy `/health` response to reduce information exposure.

— Names exact endpoint but lacks further detail v0.43.0

03 **LLM config lifecycle and token helper command** IMPROVED 55

Resets the LLM config automatically when the last tool is torn down, and uses a bare `thv` command as the LLM token helper.

— Names the `thv` command but no flags or example v0.43.0

04 **Confidential clients and RFC 8693 delegate clients in OAuth** NEW 50

ToolHive now supports confidential clients in dynamic client registration for OAuth flows, and makes RFC 8693 delegate clients reachable and usable for token-exchange workflows.

— Names standards involved but no endpoint or flag given v0.43.0

05 **Certificate provenance enforcement in lock verification**

IMPROVED

45

Lock provenance now records the certificate ref and runner used for a build, and these values are enforced during skill verification.

— Describes mechanism but no config key or flag named v0.43.0

06 **Default push signing and lock feature GA** BREAKING

45

Pushes are now signed by default, and the lock feature's feature gate has been removed — it previously required explicit opt-in.

— States behavior change but no migration guidance v0.43.0

07 **Rate-limiting observability via OpenTelemetry** NEW

40

Adds rate-limiting observability through metrics and tracing using OpenTelemetry.

— Names OpenTelemetry but no metric names or setup steps v0.43.0

08 **Multi-line Modern SSE event parsing** IMPROVED

35

Parses multi-line Modern SSE events, broadening SSE protocol compatibility.

— Names protocol feature but no config or example v0.43.0

09 **Trust state display in CLI** NEW

20

Displays recorded trust state to the user in the CLI.

— Bare description with no command or flag named v0.43.0

WAS THIS USEFUL?  

INFRASTRUCTURE

◆ Internet Services

Cloudflare

Feedback ⓘ 1 RELEASE · 2026-08-15 CHANGELOG ↗

Cloudflare is a global network platform providing CDN, DDoS protection, DNS, security, and developer services for websites and applications.

> Cloudflare shipped new DeepSeek V4 reasoning models in Workers AI, a full set of Worker-level Access controls (runtime identity context, local dev simulation, per-Worker policy attachment, and an account-wide default-protect mode), plus package-registry-aware Gateway policies, new Realtime SFU DataChannel reliability options, data-residency controls for Artifacts, FIDO2 MFA for infrastructure apps, GA Certificate Transparency Monitoring, and regex-based blocked-content rules in Email Security.

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

01 **Jurisdiction field for Artifacts namespaces** NEW 90

Adds a `jurisdiction` field to the Artifacts namespace API (`POST /api.cloudflare.com/client/v4/accounts/$ACCOUNT_ID/artifacts/namespaces`), supporting `eu` and `us` values for data localization; jurisdictions are immutable after namespace creation.

Configure an Artifacts namespace restricted to EU data residency at creation time.

```
$ curl --request POST \

'https://api.cloudflare.com/client/v4/accounts/$ACCOUNT_ID/artifacts/namespaces' \
  --header 'Authorization: Bearer $CLOUDFLARE_API_TOKEN' \
  --header 'Content-Type: application/json' \
  --data '{"namespace": "my-eu-namespace", "jurisdiction": "eu"}'
```

— Endpoint plus a runnable curl example changelog-20260815-3fbd774e

02 **Local Access simulation in wrangler.jsonc** NEW 88

Adds an `access` block inside the `dev` section of `wrangler.jsonc` to simulate Cloudflare Access identities locally — set `aud` , `email` , and related fields to test Access-gated Worker handlers without deploying; remove the block to simulate unauthenticated requests.

Test Cloudflare Access-gated Worker handlers locally with a simulated authenticated identity, without deploying.

```json
// wrangler.jsonc
{
  "dev": {
    "access": {
      "aud": "my-app",
      "identity": {
        "email": "admin@example.com"
      }
    }
  }
}
```

— Exact config block with a runnable example changelog-20260815-3fbd774e

03 **Reliability fields in Realtime SFU DataChannels API** NEW

85

Adds `ordered`, `maxRetransmits`, and `maxPacketLifeTime` fields to the Realtime SFU DataChannels API (`POST /v1/apps/$APP_ID/sessions/$SESSION_ID/datachannels/new`), enabling unordered and partially-reliable delivery modes.

— Named endpoint and fields, no request body example changelog-20260815-3fbd774e

04 **DeepSeek V4 models in Workers AI** NEW

80

Adds `@cf/deepseek-ai/deepseek-v4-pro-0813` and `@cf/deepseek-ai/deepseek-v4-flash-0731` to Workers AI, both with a full 1,048,576-token context window, thinking-mode reasoning, and function calling; requires a Workers Paid plan or prepaid credits.

— Names exact models and limits but no sample call changelog-20260815-3fbd774e

05 **Package registry selectors in Gateway HTTP policies** NEW

80

Adds `pkg.ecosystem`, `pkg.name`, `pkg.version`, `pkg.namespace`, and `pkg.purl` selectors to Cloudflare Gateway HTTP policies, enabling allow/block rules on software package downloads across npm, PyPI, RubyGems, Cargo, Go, Maven, and NuGet; requires TLS decryption.

— Named selectors and ecosystems but no sample policy changelog-20260815-3fbd774e

06 **Access identity in Worker request context** NEW

70

Adds `ctx.access` and ctx.access.getIdentity() to the Worker request context so authenticated requests expose the user's email, name, and groups without manual JWT

validation.

— Names the API but lacks a full usage example changelog-20260815-3fbd774e

07 **Worker-level Access policy attachment** NEW 65

Enables attaching an Access policy to a Worker itself, rather than to individual routes or domains, so all associated domains, Custom Domains, `workers.dev` URLs, and preview URLs are covered automatically. An Access tab in the Workers & Pages dashboard lets admins view and manage these Worker-level Access policies.

— Clear UI path but no config/API name given changelog-20260815-3fbd774e

08 **Blocked content rules in Email Security** NEW 60

Adds a Blocked content area under Policies & rules in Cloudflare Email Security, allowing administrators to define plaintext or regex rules scoped to message `subject`, `body`, or both; available on Enterprise and Enterprise + PhishGuard plans.

— UI location and scope given, no rule syntax example changelog-20260815-3fbd774e

THINNER COVERAGE BELOW

09 **Global protect-all-Workers default mode** NEW 55

Adds a global 'protect all Workers by default' mode that requires an Access policy for every Worker account-wide, with per-Worker bypass exemptions to make specific Workers publicly accessible.

— Behaviour described but no config key named changelog-20260815-3fbd774e

10 **FIDO2 MFA for infrastructure apps** NEW 55

Adds `ssh_fido2_key` and `piv_key` as supported MFA methods in application-level and policy-level settings for infrastructure applications via independent FIDO2 MFA.

— Named methods but no config path shown changelog-20260815-3fbd774e

11 **Certificate Transparency Monitoring reaches GA** IMPROVED 40

Promotes Certificate Transparency Monitoring to general availability across all Cloudflare plans, with automatic filtering of Cloudflare-issued certificate alerts and structured alert emails.

— GA status noted but no setup detail changelog-20260815-3fbd774e

WAS THIS USEFUL?  

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 15, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe | Feedback on this issue