THE CYBER TOOLCHAIN

NO. 039

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED AUGUST 8, 2026 · EVERY WEEKDAY

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

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // **CONTENTS** SHOW 35 tools

OFFENSIVE SECURITY

◆ Exploitation & C2

Bugcrowd ⓘ 1 RELEASE · 2026-08-08 CHANGELOG ↗

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

> Bugcrowd added the ability to generate a DNS Records Report within the platform.

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

01 **DNS Records Report generation** NEW 30

Bugcrowd now supports generating a DNS Records Report within the platform.

— Names the report but no detail on fields or access path. changelog-20260808-e39b22ba

WAS THIS USEFUL? 👍 👎

◆ Reconnaissance & OSINT

Nmap ⓘ changes since 2026-07-09 CODE ↗

Nmap - the Network Mapper. Github mirror of official SVN repository.

Nmap 7.991 adds SSL service name matching for portrules and fixes Ncat HTTP proxy hang, libssh2 CVEs, and DNS case matching.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-09
https://github.com/nmap/nmap.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-09
```

› Adds `'s'` service names (e.g. `amt-soap-https`) to NSE portrules so scripts that run on standard unencrypted services also run on their SSL-protected equivalents.

› Ncat now propagates termination signals (`SIGINT`, `SIGTERM`, `SIGHUP`) to child processes when using `-k` keep-open mode, preventing orphaned children on shutdown.

WAS THIS USEFUL? 👍 👎

◆ AI-Powered Offensive Security

Strix

1 RELEASE · 2026-08-07 NOTES ↗

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

> Strix v1.5.1 broadens its vulnerability scanner's coverage of Firebase Storage rules misconfigurations.

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

01 **Expanded Firebase Storage rules coverage** IMPROVED 30

The vulnerability scanner now expands coverage of Firebase Storage rules for misconfiguration detection.

— No detail on which rules or mechanism added. v1.5.1

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Secrets Detection & Management

GitGuardian Detection Engine ⓘ 1 RELEASE · 2026-08-08

CHANGELOG ↗

GitGuardian Detection Engine scans code repositories to identify and alert on exposed secrets, credentials, and sensitive data in real-time.

> GitGuardian's Detection Engine added four new secret detectors (OpenCode, Lovable, Supabase), tightened false-positive filtering across five existing detectors for proxy-config and presigned-URL patterns, and retired or merged three detectors including the Amadeus and GitHub OAuth token checks.

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

01 **False-positive filtering across generic and AWS detectors**

IMPROVED 75

Extends `generic_password` to ignore passwords in public proxy-subscription configs (shadowsocks/vmess/vless/trojan/hysteria URIs, Telegram MTProto links, clash/mihomo node lists); extends `generic_high_entropy_secret` and `base64_generic_high_entropy_secret` to ignore secrets whose decoded context is a public proxy-subscription config or a URL-signing token (S3 presigned parameters, HLS stream tokens, booking/redirect tokens); extends `aws_iam` to ignore presigned-URL parameters and GitHub avatar IDs; and extends `generic_encryption_key` to ignore Talend encrypted passwords (`enc:system|routine.encryption.key.vN:...` format).

— Names exact detectors and false-positive patterns suppressed.

changelog-20260808-28734af6

02 **Detector removals and merge: Amadeus, GitHub OAuth, PostgreSQL variant**

70

BREAKING

The Amadeus OAuth credentials detector and checker have been disabled because the Amadeus self-service API has been discontinued; the `github_oauth_token` detector and checker have been removed since the underlying service is no longer available; and the attached-port PostgreSQL variant detector has been merged into the Generic Database Assignment detector.

— Explains reason and scope for each removal or merge. changelog-20260808-28734af6

THINNER COVERAGE BELOW

03 **New detectors for OpenCode, Lovable and Supabase** NEW

58

Adds `OpenCode API Key` detector for OpenCode API keys, `Lovable Access Token` detector for Lovable project access tokens, `Lovable Git Token` detector for Lovable git tokens, and a new `Supabase Service Role JWT` checker for Supabase `service_role` JWTs.

— Names each detector but gives no usage detail beyond that. `changelog-20260808-28734af6`

---

04 **Azure Event Hub checker hardened** IMPROVED 25

Hardens the Azure Event Hub checker with the same Atom-feed validation requirement.

— Thin one-line note; validation mechanism referenced but unexplained.

`changelog-20260808-28734af6`

└──▷ BREAKING ON UPGRADE

! The Amadeus OAuth credentials detector and checker have been disabled; the Amadeus self-service API has been discontinued.

! The `github_oauth_token` detector and checker have been removed; the service for this detector is no longer available.

! The attached-port PostgreSQL variant detector has been merged into the Generic Database Assignment detector.

WAS THIS USEFUL? 👍 👎

infisical

1 RELEASE · 2026-08-07 NOTES ↗

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

> Infisical v0.162.18 extends gateway routing to Chef app connections and adds an inline role-change action for project members.

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

01 **Gateway support for Chef app connection** NEW 45

The Chef app connection now supports routing through Infisical gateways, enabling Chef integrations to be reached via a gateway rather than direct network access.

— Describes capability but no config or flag names given v0.162.18

02 **Modify Role action on project member rows** NEW 45

A new Modify Role action was added directly on project member role rows in the UI, letting admins change a member's role inline without navigating to a separate page.

— Names UI location but no exact menu path or API v0.162.18

WAS THIS USEFUL? 👍 👎

◆ SAST / SCA / Code Security

Codacy

3 RELEASES · seen 2026-08-08 NOTES ↗

Codacy is an automated code analysis platform that identifies bugs, security vulnerabilities, and code quality issues across repositories.

> Codacy added two new analysis engines — GolangCI-Lint for Go and PHP CS Fixer for PHP code style — and extended repository config-file support to Checkov and ShellCheck so their settings can be version-controlled instead of managed only through the UI.

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

01 Repository config file support for Checkov and ShellCheck NEW

81

Codacy can now read analysis settings directly from repository config files instead of requiring UI-only configuration: enabling "Use Checkov configuration file found in the repository" under Repository > Code Patterns > Checkov picks up a `.checkov.yaml` or `.checkov.yml` file, and enabling "Use Shellcheck configuration file found in the repository" under Repository > Code Patterns > Shellcheck picks up a `.shellcheckrc` file in the repository root.

Let Codacy pick up your existing Checkov config instead of managing every check through the UI.

> 📍 Repository › Code Patterns › Checkov › enable "Use Checkov configuration file found in the repository"

Commit a `.shellcheckrc` to your repo and enable file-based analysis so ShellCheck rules are version-controlled alongside your code.

> 📍 Repository › Code Patterns › Shellcheck › enable "Use Shellcheck configuration file found in the repository"

— Exact file names and UI toggle paths for both engines.

Adding PHP CS Fixer and support to Check… Support for Shellcheck configuration fil…

02 PHP CS Fixer engine for PHP code style NEW

78

Codacy adds PHP CS Fixer as a new supported engine for PHP code-style analysis, configurable via the Codacy UI or via your repository's own `.php-cs-fixer.dist.php` or `.php-cs-fixer.php` file. When both exist, Codacy-selected patterns in the UI take precedence over the repository config file.

— Names exact config files and the precedence rule between them.

Adding PHP CS Fixer and support to Check…

03 **GolangCI-Lint support for Go analysis** NEW 60

Codacy adds GolangCI-Lint as a client-side analysis tool for Go projects: run it locally or in CI and upload results to the Codacy dashboard for reporting. It is not enabled by default on existing repositories and must be turned on manually in repository settings to avoid a large volume of new issues appearing unexpectedly.

— Mentions repository settings but no exact path or command.

Adding GolangCI-Lint as new supported to…

WAS THIS USEFUL? 👍 👎

StackHawk

1 RELEASE · seen 2026-08-08 NOTES ↗

StackHawk is an application security testing platform that performs dynamic security scanning to find vulnerabilities in web applications during development and CI/CD pipelines.

> StackHawk's HawkScan 6.3.0 release adds mid-scan crash detection and a much faster replay-based rescan, alongside new seedPaths validation and a consolidated CLI credential model, while the platform API also changed the audit endpoint's response schema.

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

01

Crash detection for scans NEW

91

A new `hawk.scan.crashDetection` config block detects when a target application crashes or becomes unresponsive mid-scan by passively counting new timeouts and connection failures, then triggering an active liveness probe (using `app.waitForAppTarget` as the probe path) before taking action. The trigger threshold is set via `hawk.scan.crashDetection.maxTimeouts`, which defaults to 5x `concurrentRequests` when unset or zero, and the response is controlled by `hawk.scan.crashDetection.action`, supporting `ABORT` (default, fails the scan with error `scan.target-crashed`), `WARN` (logs and continues), or `DISABLED`.

Abort a scan immediately when the target goes down, using a tighter failure threshold than the default to catch crashes faster.

yaml

```yaml
hawk:
  scan:
    crashDetection:
      maxTimeouts: 10
      action: ABORT
```

Abort a scan immediately and surface a clear error when the target application stops responding mid-scan, rather than burning the full scan window against a dead host.

yaml

```yaml
hawk:
  scan:
    crashDetection:
      action: ABORT
      maxTimeouts: 10
```

— Full mechanism, exact config keys, defaults, and runnable YAML example.

snapshot-20260808

02

Faster rescans via request replay BREAKING

79

`hawk rescan` now replays the exact requests that produced findings — fetched from the platform — instead of re-crawling the app, skipping all discovery (OpenAPI, SOAP, GraphQL, HAR imports) for substantially faster, tightly scoped retests; it falls back to a normal spider automatically if replay cannot run. Because discovery is now skipped entirely, workflows that relied on rescan to re-crawl the application must switch to a full scan instead.

— Clear before/after mechanism and named command, no runnable example given.

snapshot-20260808

03 **seedPaths validation at parse time** NEW 78

`hawk.spider.seedPaths` entries are now validated at configuration parse time by both `hawk scan` and `hawk validate config`: entries over 200 characters or with more than 12 path segments produce a warning, while entries over 1000 characters fail validation outright.

— Exact config key and thresholds named, but no example run shown. `snapshot-20260808`

04 **Unified CLI credential sourcing** BREAKING 75

Commands embedded in the HawkScan CLI (`hawk op`) now read API key, API host, and organization from the same sources as `hawk` itself — environment variables and `~/.hawk/hawk.properties` — eliminating the need for a separate `~/.hawkop/config.yaml`, which is no longer used.

— Names exact config file paths and the removed file; clear migration note.

`snapshot-20260808`

THINNER COVERAGE BELOW

05 **Audit endpoint response schema changed** BREAKING 33

The `GET /api/v1/org/{orgId}/audit` endpoint had its response schema changed.

— Endpoint named but no detail on what changed in the schema. `0.0.1`

└──▷ BREAKING ON UPGRADE

! `hawk rescan` no longer re-runs discovery — all discovery is skipped, including OpenAPI, SOAP, GraphQL, and HAR imports. Workflows that relied on rescan to re-crawl the application will need to use a full scan instead.

! The separate `~/.hawkop/config.yaml` credentials file is no longer used; credentials must now be in environment variables or `~/.hawk/hawk.properties`.

! The separate `~/.hawkop/config.yaml` credential file is no longer used by `hawk op`; credentials are now read exclusively from the environment and `~/.hawk/hawk.properties`.

WAS THIS USEFUL? 👍 👎

Vercel deepsec

changes since 2026-07-09

CODE ↗

Deepsec is a security harness for finding vulnerabilities in your codebase powered by coding agents

deepsec gains `--thinking-level`, one-shot initialization, stable finding IDs, and a broadened SSRF matcher.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-09 https://github.com/vercel-labs/deepsec.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-09
```

- Adds `--thinking-level` CLI flag to control agent reasoning depth during scans.
- Ships one-shot Deepsec initialization (PR #139) to bootstrap a workspace in a single command.
- Introduces stable finding IDs and permissive verdict reconciliation in the revalidate workflow, backed by `ensureFindingIds` and `reconcileVerdicts`.
- Broadens the SSRF matcher to cover a wider range of server-side request forgery patterns.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

ElcomSoft

1 RELEASE · 2026-08-08 CHANGELOG ↗

ElcomSoft develops digital forensics, password recovery, and mobile device extraction tools for law enforcement and security professionals.

> ElcomSoft shipped EIFT Firewall, a free macOS companion app that replaces its 2023 firewall script for safely sideloading the iOS Forensic Toolkit's low-level extraction agent.

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

01 **EIFT Firewall macOS companion app** NEW 58

New free macOS application EIFT Firewall creates a temporary, controlled USB network path between a Mac and a connected evidentiary device to safely sideload the Elcomsoft iOS Forensic Toolkit low-level extraction agent, replacing the macOS firewall script published in 2023.

— Names mechanism and predecessor but no install/usage steps. changelog-20260808-e8228529

WAS THIS USEFUL? 👍 👎

◆ Malware Analysis

quark-engine

1 RELEASE · 2026-08-08 NOTES ↗

Quark-Engine is a static analysis tool for Android malware detection that identifies suspicious behaviors and potential security threats in APK files.

> Quark-Engine v26.8.1 adds AI-assisted analysis skills for Claude Code, a new DexTrace backend for DEX analysis, and fresh malware analysis reports.

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

01 **Claude Code skills for malware analysis and rule generation** NEW

60

Adds Claude Code skills `/quark` and `/quark-rule-gen`, enabling AI-assisted malware analysis and rule generation directly within Claude Code.

— Names two slash commands but no mechanism detail. v26.8.1

THINNER COVERAGE BELOW

02 **Malware analysis reports for AhRat, AndroRAT, and Sova** NEW

40

Adds Quark malware analysis reports for AhRat, AndroRAT, and Sova.

— Names three malware families but no report contents or access path. v26.8.1

03 **DexTrace core-library backend for DEX analysis** NEW 25

Adds a new `DexTrace` core-library backend for DEX analysis.

— Bare name, no mechanism or usage given. v26.8.1

WAS THIS USEFUL? 👍 👎

CAPEv2 ⓘ changes since 2026-07-20 CODE ↗

Malware Configuration And Payload Extraction

CAPEv2 adds Remus malware detection & config extraction, a ThreatFox/Malpedia threat-intel module, and PCI vendor ID spoofing to defeat QEMU fingerprinting

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-20
https://github.com/kevoreilly/CAPEv2.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-20
```

- New `modules/processing/threatintelligence.py` processing module (runs at `order=22`) enriches analysis reports with two phases: an indicator phase that queries ThreatFox to annotate contacted IPs/domains with threat tags, and a family phase that fetches descriptive cards (name, aliases, description, references, attribution) from Malpedia for detected malware families; results are written to `results["threatintelligence"]` and surfaced in a collapsible 'Threat Intelligence' report-page section.
- Adds Remus malware detection and dynamic config extraction.
- Adds PCI vendor ID spoofing to bypass Al-Khaser QEMU detection checks.
- New `require_task_delete` / `can_delete_task` authorization guards in the multi-tenant API enforce that only the submitter or a box-admin can delete a task (`public` policy).
- The API v2 batch-delete endpoint now reports partial failures per task (individual `try/except` per task in batch loops) and returns a `partial_error` field instead of a bodiless 500, with SQL delete committed before irreversible folder/Mongo deletes for rollback safety.

└──▷ BREAKING ON UPGRADE

! The `suspended` parameter has been removed from execute() methods in `process.py` and `abstracts.py` ; callers passing this argument will break.

WAS THIS USEFUL? 👍 👎

◆ Reverse Engineering

Radare2 ⓘ 1 RELEASE · 2026-08-07 NOTES ↗ CODE ↗

UNIX-like reverse engineering framework and command-line toolset

> Radare2 6.2.0 ships a new generic binary resource extraction API alongside broader calling-convention and register analysis for ppc64 ELFv2 and MIPS, an initial Rust v0 demangler, WebAssembly type support, and several breaking changes including a Capstone v4 minimum and renamed resource commands.

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

01 Binary resource extraction and generic resource API BREAKING

95

Adds `iuo` to open binary resources inline and `ix` to extract sections, segments, and resources; adds `iU` (with `iUx` extraction) and `rabin2 -xU` for listing/extracting resources. The old `iR` command is renamed to `iU`, and the previous `iU` command moves to `iu`; the `rabin2 -u` flag behavior moves to an environment variable. A new generic binary resource API adds support for GLib Resources in ELF, Macho code signatures, PE/CLR resources, PE resources from MDMP, OS/2 LE/LX resources, QNX LMF resources, and Pebble resources.

Extract all embedded resources from a PE binary directly from the command line.

```
$ rabin2 -xU malware.exe
```

— Names every command, flag, and rename with runnable example. 6.2.0

02 Calling convention and register analysis improvements NEW

85

Supports the ppc64 ELFv2 calling convention, selected automatically from the ELF ABI, and selects the MIPS calling convention automatically from the ELF o32/n32 ABI. Adds callconv definitions (clobber and preserve registers) for all supported architectures including PPC and ARM, PPC64 TOC pointer base-register tracking gated by `RArchOpMask.STATEFUL` and wired into `aar`, new `r_anal_cc_argslot`, `r_anal_cc_argval`, and `r_anal_cc_isclobber` APIs, `RArchOpMask.STATEFUL` and `RArchSession.reset` for stateful architecture decoding in plugins, and reguse analysis hints with dyncc register role APIs.

— Lists APIs and mechanisms but no direct usage command. 6.2.0

03 New configuration keys for analysis and UI NEW

80

Adds `anal.gpseed` to seed a PIC base register during ESIL analysis, `asm.cmt.reguse` to emit register-use comments in disassembly, `dir.binsdb` for PE/MDMP ordinal lookups from binpath/dirprefix, `types.fields` to gate propagation of callee types into struct and union fields during `aaft`, and `scr.drag` to enable click-and-drag scrolling in disasm, hex, and graph views.

Seed a PIC base register before ESIL-based analysis so position-independent code resolves correctly.

```
$ r2 /bin/firmware.elf
> e anal.gpseed=0x10000
> aaa
```

— Names every key and its file/effect; one has runnable example. `6.2.0`

---

04 **Struct type synthesis via afts command** NEW — 80

Adds struct type synthesis from pointer-argument access patterns via the `afts` command, letting analysts recover unknown data structures after analyzing a function.

Synthesize struct types from pointer-argument access patterns after analyzing a function to recover unknown data structures.

```
$ r2 /bin/target
> aaa
> s sym.process_packet
> afts
```

— Named command with runnable example demonstrating use. `6.2.0`

---

05 **r2agent HTTP auth protection** BREAKING — 60

Adds auth protection to `r2agent` web session spawns and protects `r2agent` with an HTTP auth file; empty lines in the auth file now block all access rather than allowing any password, reversing prior behavior.

— Names the auth file mechanism and behavior change. `6.2.0`

THINNER COVERAGE BELOW

06 **Short-circuit && and || chaining in r2 shell** NEW — 50

Adds `&&` and `||` short-circuit command chaining support in the r2 shell, including after literal commands in the shell pipeline.

— Names the operators but no worked example. `6.2.0`

---

07 **Rust v0 demangler and pluggable demangler framework** NEW

45

Adds an initial Rust v0 symbol demangler, making Rust binaries easier to analyze, and makes RBin demanglers pluggable and selectable.

— Describes capability but no command shown. 6.2.0

---

08 **Lcc command to list registered plugins** NEW

40

Adds `Lcc` command to list all plugins registered with the new command API.

— Named command, minimal further detail. 6.2.0

---

09 **Capstone v4 required, v3 removed** BREAKING

40

Removes Capstone v3 support; Capstone v4 is now the minimum required version for disassembly.

— Clear before/after but no migration steps. 6.2.0

---

10 **New build targets: MiKNix and jslinux** NEW

35

Adds MiKNix platform support in the build system and a `dist/jslinux` target using temu/rv64 for JavaScript-based Linux builds.

— Names targets but no build instructions. 6.2.0

---

11 **RBinPluginAny renamed and LibMagic dropped** BREAKING

35

`RBinPluginAny` is renamed to `"null"` and no longer uses LibMagic.

— Names the rename but no impact guidance. 6.2.0

---

12 **WebAssembly v128, ref and reftype support** NEW

30

Adds `v128`, `ref`, and `reftype` support for WebAssembly binaries.

— Names the types but no mechanism or example. 6.2.0

└──▷ BREAKING ON UPGRADE

! The `iR` command is renamed to `iU`; the old `iU` command is moved to `iu`; the `rabin2 -u` flag behavior is moved to an environment variable.

! Capstone v3 support is removed; Capstone v4 is now the minimum supported version.

! `RBinPluginAny` is renamed to `"null"` and no longer uses LibMagic.

! Empty lines in the r2agent HTTP auth file now deny all access (previously they permitted any password through).

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Fastly ⓘ 1 RELEASE · 2026-08-07 BLOG ↗

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

> Fastly CLI v15.6.0 adds a new command group for managing audit log event mappings.

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

01 **Audit log event-mapping command group in CLI** NEW 60

The Fastly CLI adds the `audit-log event-mapping` command group for managing audit log event mappings.

— Names exact command but not its subcommands or behavior. CLI v15.6.0 (2026-08-07)

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

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

Plerion is a cloud security platform that enables continuous vulnerability discovery and remediation across cloud infrastructure and applications.

> Plerion significantly deepened its cloud security posture coverage for Google Cloud AI services, adding dozens of new checks, guardrails, and expanded asset discovery across the GCP AI stack.

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

01 **Expanded GCP AI security posture coverage** NEW 63

Plerion adds 46 new posture checks spanning Vertex AI, Vertex AI Workbench, Colab Enterprise, Agent Engine, Vertex AI Search, Gemini Enterprise, Model Armor, Dialogflow CX, Cloud TPU, and Gemini Code Assist. It also adds project-level org-policy guardrails covering Vertex AI Workbench access and Model Garden model use, and extends asset discovery to 16 new GCP resource types behind these newly supported AI services.

— Names all services and checks but no config or command to act on.

GCP AI coverage widens further:

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

NVIDIA OpenShell ⓘ 1 RELEASE · 2026-08-07 NOTES ↗

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

> OpenShell v0.0.101 focused on build system plumbing, adding Bazel build targets for two internal drivers.

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

01 **Bazel build targets for credential and VM drivers** NEW

35

Adds Bazel build targets for the credential driver, enabling it to be compiled as part of the standard build, and adds a VM driver built from GitHub sources via Bazel.

— Names build system only; no usage or config detail given. v0.0.101

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Sysdig Secure ⓘ 1 RELEASE BLOG ↗

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

> Sysdig Secure improved vulnerability policy scoping, allowing the same attribute to be used in multiple AND-joined conditions for Admission Control and Runtime stages.

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

01 **Multi-condition attribute scoping in vulnerability policies**

IMPROVED 55

Vulnerability policy scopes for the Admission Control and Runtime stages now support using the same attribute in more than one condition, joined by AND, enabling include and exclude patterns on a single attribute (e.g., all non-production clusters except one).

— Explains mechanism and gives example but no config path Secure SaaS - August 3, 2026

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ SIEM

Google SecOps

1 RELEASE · 2026-08-08 CHANGELOG ↗

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

> Google SecOps expanded its log ingestion catalog with 25 new supported sources spanning security, network, and SaaS vendors, and added self-service license key retrieval for Bindplane Enterprise collection agents.

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

01 Self-service Bindplane Enterprise license key download NEW

75

Enterprise Plus and Google Unified Security (GUS) customers can now download the Bindplane Enterprise (Google Edition) license key directly from `SIEM Settings > Collection Agents` in the console, without opening a support ticket. This capability is in Preview.

Retrieve your Bindplane Enterprise license key from the console so you can activate a self-managed collection agent without opening a support ticket.

> 📍 In the Google SecOps console, go to SIEM Settings › Collection Agents and download the Bindplane Enterprise (Google Edition) license key.

— Names exact UI path and preview status, gives steps. changelog-20260808-95ed47a5

02 26 new log ingestion source integrations NEW

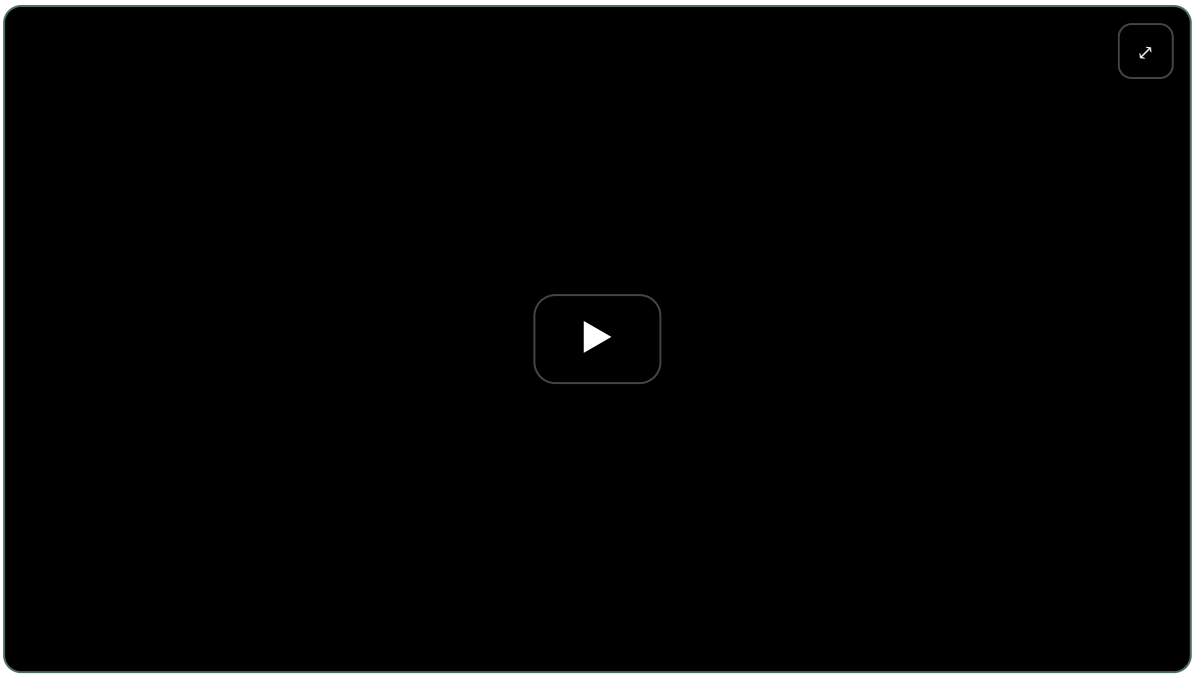
60

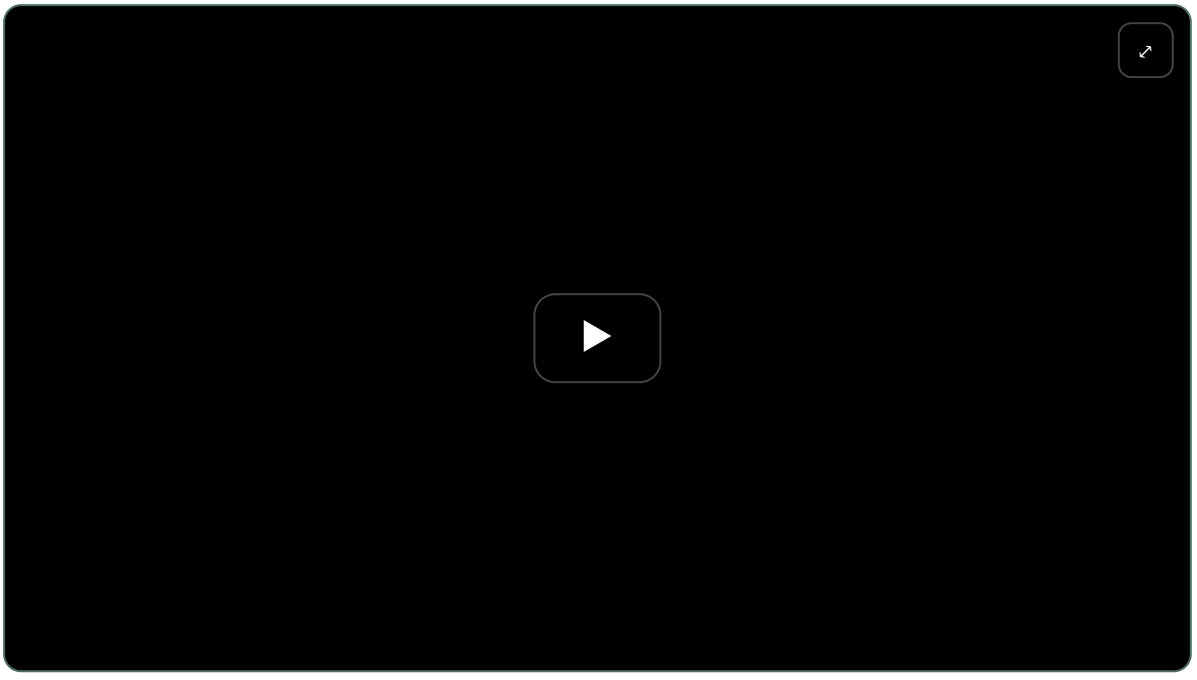
Google SecOps added support for ingesting logs from a broad set of new sources: `ARMIS_ALERTS` (Armis Alerts), `GCP_SECURITYCENTER_EXTERNAL_EXPOSURE` (Security Command Center External Exposure), `ADOBE_EXPERIENCE_PLATFORM` (Adobe Experience Platform), `AUDIOCODES_SBC` (AudioCodes Session Border Controller), `AZURE_GATEWAY_CONTAINERS` (Azure Application Gateway for Containers), `AZURE_LOGIC_APPS` (Azure Logic Apps), `AZURE_NATGW_FLOW` (Azure NAT Gateway Flow), `BROADCOM_DX_NETOPS_SPECTRUM` (Broadcom DX NetOps Spectrum), `CARTO_ACTIVITY` (Carto Activity), `CLAUDE_CODE_OBSERVABILITY` (Claude Code Observability), `CYBLE_ASM` (Cyble Attack Surface Management), `CYBLE_BIP` (Cyble Brand Intelligence & Protection), `DARKWEB_IQ` (Darkweb IQ), `ELLIO_THREAT_INTEL` (Ellio Threat Intelligence), `EXEON_NDR` (Exeon NDR), `GRAVITEE` (Gravitee), `KASPERSKY_ANTI_TARGETED_ATTACK` (Kaspersky Anti Targeted Attack), `MICROSOFT_COPILOT_INTERACTION` (Microsoft Copilot Interaction), `OSTTRA_MARKITWIRE` (OSTTRA MarkitWire), `PROOFPOINT_ADAPTIVE_EMAIL_SECURITY` (Proofpoint Adaptive Email Security), `SECOMEA_GATEMANAGER` (Secomea GateManager),

`TRENDMICRO_VISION_ONE_RISK_EVENT` (Trend Micro Vision One Risk Event), `TXONE_EDGEIPS` (TXOne EdgeIPS), `VECTRA_RUX` (Vectra Respond UX), and `ZOHO_CRM` (Zoho CRM).

— Names every source ID but no ingestion mechanism detail. `changelog-20260808-95ed47a5`

└──▷ ALSO FROM THESE RELEASES





WAS THIS USEFUL? 👍 👎

Elastic Security Labs

1 RELEASE · 2026-08-07 BLOG ↗

Elastic Security Labs publishes security research, threat analysis, and vulnerability findings to help organizations understand emerging cyber threats.

> Elastic Security Labs shipped a purpose-built Elastic Agent integration that watches for silent removal of npm's cooldown protection (`min-release-age`) by snapshotting `.npmrc` files across hosts every six hours.

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

01 **CEL integration detecting npm cooldown removal via .npmrc snapshots**

NEW 90

A roughly 40-line Common Expression Language (CEL) integration in Elastic Agent snapshots every `.npmrc` on a 6-hour heartbeat and sets `cooldown.absent = true` in the ingest pipeline when the `min-release-age` setting disappears between snapshots. It uses the file() and dir() CEL functions to enumerate user home directories (`/Users` , `/home`) and machine-global npm config paths (`/opt/homebrew/etc/npmrc` , `/usr/local/etc/npmrc` , `/etc/npmrc` , `/usr/lib/node_modules/npm/.npmrc`) covering both macOS and Linux from one script, with global paths and poll interval supplied via the "Custom request cursor" YAML field in the Fleet UI, and `resource.url: file:///dev/null` used to satisfy the integration's required URL field without issuing an HTTP request. Registry auth token lines (`//registry.npmjs.org/:_authToken=...`) are filtered out in the ingest pipeline before leaving the host, and per-file `file.hash.sha256` values are emitted alongside file content so in-place rewrites that strip `min-release-age` without changing the file path can still be detected. The write-up also documents why the Custom Logs Filestream integration was unsuitable: files under 1024 bytes are silently held back under the default fingerprint identity strategy, `filestream` only emits on append (not on content removal), and its `clean_inactive` option requires `ignore_older` to also be set, which suits log rotation but not static config files that may sit unchanged for months.

— Names exact paths, YAML field, pipeline keys and fallback rationale.

The security signal log tailing can't se…

WAS THIS USEFUL? 👍 👎

◆ Endpoint Security (EDR/XDR)

Iru ⓘ 1 RELEASE · seen 2026-08-08 NOTES ↗

> Iru's August 7th update adds a new library item assignment criterion based on Apple beta program enrollment and refreshes the Iru AI chat experience.

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

01 **Apple beta enrollment as assignment criterion** NEW 40

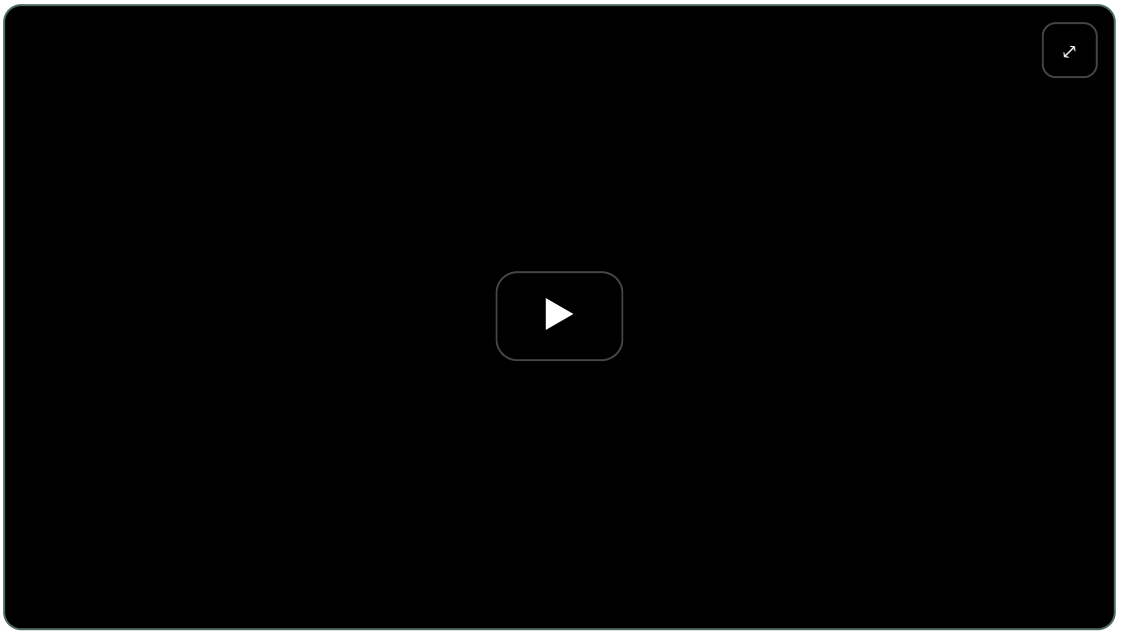
Library items can now be assigned based on a device's Apple beta program enrollment status, adding this as a new criterion for library item assignment.

— Names the criterion but no config path or mechanism given.

What's New in Iru | August 7th

02 **Refreshed Iru AI chat experience** IMPROVED 20

Iru's AI chat experience has been refreshed to provide faster, smarter responses.



— Vague marketing phrasing with no specifics on mechanism or scope.

What's New in Iru | August 7th

WAS THIS USEFUL? 👍 👎

Huntress

1 RELEASE · seen 2026-08-08 NOTES ↗

Huntress is a managed security platform that provides threat detection, incident response, and endpoint protection for small and mid-sized businesses.

> Huntress rolled out a redesigned Managed ITDR Dashboard in Early Access, adding rapid identity triage, failed login characterization, and a built-in SIEM search surface.

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

01 **Quick SIEM Search in ITDR dashboard** NEW 73

A new Quick SIEM Search surface within the ITDR dashboard provides pre-populated queries against the full Entra Unified Audit Log, stored in Huntress SIEM at no cost for ITDR customers for up to one year.

– Names retention and cost terms but no query syntax given

Now Live: Managed ITDR Dashboard in Earl…

02 **Rapid Identity Triage panel** NEW 72

A new panel lets an analyst enter a user's email to get an immediate overview of all activity over the past 24 hours, with options to export a timeline, revoke sessions, or disable the account.

– Names concrete actions but no exact UI path or command

Now Live: Managed ITDR Dashboard in Earl…

THINNER COVERAGE BELOW

03 **Failed Login Characterization view** NEW 50

A new view shows failed login attempts globally, with categorization by residential proxies, VPNs, and datacenters.

– Describes categories shown but no interaction detail

Now Live: Managed ITDR Dashboard in Earl…

04 **Redesigned ITDR dashboard navigation** NEW 50

The redesigned ITDR dashboard is accessible via the ITDR icon in the left-hand navigation, and the existing dashboard remains available in parallel during Early Access.

– Gives exact navigation path but no further mechanism

Now Live: Managed ITDR Dashboard in Earl…

WAS THIS USEFUL? 👍 👎

Trend Vision One ⓘ 2 RELEASES · 2026-08-04 → 2026-08-06 BLOG ↗

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

> Trend Vision One's Agentless Vulnerability and Threat Detection expands scan coverage to Azure Function Apps and adds customer-configurable resource tagging for OCI deployments.

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

01 Customer-configured OCI resource tagging via Terraform NEW

65

Agentless Vulnerability & Threat Detection v1.96.3a0922b adds support for customer-configured tags — both OCI defined tags and freeform tags — applied to all Terraform-deployed OCI resources, including the scanner VM, disk snapshots, and cloned scan disks.

— Names resource scope but no exact config key or command.

Agentless Vulnerability & Threat Detecti…

THINNER COVERAGE BELOW

02 Azure Function Apps scanning in AVTD NEW

55

Agentless Vulnerability and Threat Detection now scans Azure Function Apps for vulnerabilities and malware, accessible via Cyber Risk Exposure Management → Continuous Risk Management → Threat and Exposure Management.

— Names the UI path but no mechanism or limits given.

Agentless Vulnerability and Threat Detec…

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

OpenCTI

1 RELEASE · 2026-08-07 NOTES ↗

Open Cyber Threat Intelligence Platform

> OpenCTI 7.260807.0 introduces shareable saved searches and dashboard filters, direct XTMHub connectivity, advanced Lucene search operators, and several UI and ingestion improvements.

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

01 **Lucene wildcard and fuzzy search operators** NEW 60

Adds support for Lucene advanced search syntax including wildcard prefix (*) and fuzzy search (~) operators.

— Names exact operators reader can use directly in search. 7.260807.0

THINNER COVERAGE BELOW

02 **Admin UI for custom fields configuration** NEW 45

Adds an Admin UI (Settings) for custom fields configuration, described as Chunk 2 of the custom-fields feature.

— Names the settings location but lacks specific fields or flags. 7.260807.0

03 **Shareable saved searches and reusable filters** NEW 35

Adds shareable and reusable saved searches across the platform, plus reusable saved filters inside dashboards and custom views.

— Describes what's shared but no UI path or config detail. 7.260807.0

04 **XTMHub connectivity from OpenCTI interface** NEW 35

Adds XTMHub connectivity, allowing users to connect to Hub directly from the OpenCTI interface.

— Names the integration but not how it's configured. 7.260807.0

05 **Label casing preserved in UI** IMPROVED 35

Preserves user-defined label casing, removing the automatic `text-transform: capitalize` styling previously applied to labels.

— Names the CSS rule removed but minor cosmetic change. 7.260807.0

06 **TAXII feed tabs in ingestion UI** NEW 30

Adds TAXII feed tabs to the ingestion UI for improved feed management.

— Bare description of a UI addition without mechanism. 7.260807.0

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

Tenzir ⓘ 1 RELEASE · 2026-08-07 BLOG ↗

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

> Tenzir Platform v1.38.0 makes telemetry collection opt-in and lets users search pipelines by ID.

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

01 **Pipeline search by ID** IMPROVED 50

The pipeline list search box now matches on pipeline ID in addition to name and definition, letting you paste a pipeline ID to jump straight to it.

— Names the UI mechanism but no exact navigation path

Tenzir Platform v1.38.0: Opt-in telemetr…

02 **Opt-in telemetry consent on Account page** BREAKING 50

Telemetry collection is now gated behind explicit user opt-in, configurable through the Account page, and is now also enabled on `app.tenzir.com` for opted-in users.

— Names the Account page setting but no exact steps

Tenzir Platform v1.38.0: Opt-in telemetr…

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Secure Access (VPN / SASE / ZTNA)

Twingate

1 RELEASE · 2026-08-08 | CHANGELOG ↗

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

> Twingate's Admin Console now lets admins resize table columns, with sizing preferences saved per user.

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

01 **Resizable table columns in Admin Console** IMPROVED 55

Admins can resize table column widths in the Admin Console by hovering over a column edge and dragging. Column sizing is saved per-user, and a reset-to-default option is available.

— Names UI mechanism and behavior but no config/API surface `changelog-20260808-d9a735cb`

WAS THIS USEFUL? 👍 👎

Cohesive VNS3 Controller

1 RELEASE · seen 2026-08-08

NOTES ↗

VNS3 Controller is a network security appliance that creates encrypted virtual networks and manages multi-cloud connectivity with advanced firewall capabilities.

> VNS3 7.2.1 upgrades the underlying kernel to address a privilege escalation issue and adds platform-awareness to the Plugin Catalog.

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

01 **Kernel upgraded to Linux 7.0.0** IMPROVED 50

VNS3 upgraded from the Linux 6.17.x kernel line to 7.0.0 in response to reported local privilege escalation findings against the 6.17 line.

— Names versions and reason but no reader action snapshot-20260808

02 **Plugin Catalog platform and version awareness** IMPROVED 40

The Plugin Catalog now differentiates between ARM and Intel capable plugins, reflecting potential release timeline differences between platforms, and supports more intelligent version management with listings enhanced to show multiple versions and platform compatibility information.

— Describes what changed but no UI path or exact fields snapshot-20260808

WAS THIS USEFUL? 👍 👎

Sophos ZTNA

2 RELEASES · seen 2026-08-08 NOTES ↗

Sophos ZTNA is a zero-trust network access solution that verifies user and device identity before granting access to applications and resources.

> Sophos ZTNA 2025.02 added Let's Encrypt certificate automation and expanded to all Sophos Central regions, while 2025.23 improved on-premises connectivity with traffic bypass and domain controller support.

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

01 **Let's Encrypt certificate generation via Sophos Central** NEW

55

ZTNA gateways can now generate Let's Encrypt certificates directly through Sophos Central, with automatic management and renewal handled by Sophos.

– Names mechanism and scope but no config steps 2025.02

02 **On-premises traffic bypass for the ZTNA agent** NEW

55

A new configuration option prevents the ZTNA agent from intercepting on-premises traffic, avoiding tunnel hairpinning for local resources.

– Explains behaviour and benefit but no exact setting name 2025.23

03 **Domain controller as agent-based resource** NEW

45

ZTNA now supports adding a domain controller as an agent-based resource, providing seamless connectivity to resources behind it.

– Describes new resource type without configuration steps 2025.23

04 **ZTNA expanded to all Sophos Central regions** IMPROVED

25

Sophos ZTNA availability is expanded to all Sophos Central regions.

– Bare statement of expanded availability, no detail 2025.02

WAS THIS USEFUL? 👍 👎

◆ Identity & Access Management

Auth0 ⓘ 1 RELEASE · 2026-08-07 BLOG ↗

Auth0 is a cloud-based identity and access management platform providing authentication, authorization, and user management solutions for applications.

> Auth0 shipped two identity-management additions this window: new My Account API endpoints for managing a user's own authentication methods, and an Early Access Session Delegation capability built on Custom Token Exchange.

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

01 Authentication methods endpoints in My Account API NEW

90

The My Account API gains four endpoints for managing a user's enrolled authentication methods: `GET https://{yourDomain}/me/v1/authentication-methods` to list them, `POST https://{yourDomain}/me/v1/authentication-methods` to create/enroll a new one, `POST https://{yourDomain}/me/v1/authentication-methods/{id}/verify` to verify a specific method by ID, and `DELETE https://{yourDomain}/me/v1/authentication-methods/{id}` to remove one.

List all authentication methods enrolled for the currently authenticated user.

```
$ curl -X GET 'https://{yourDomain}/me/v1/authentication-methods' \
  -H 'Authorization: Bearer <access_token>'
```

Verify a specific authentication method (e.g., after enrollment) to confirm it is active.

```
$ curl -X POST 'https://{yourDomain}/me/v1/authentication-methods/{id}/verify' \
  -H 'Authorization: Bearer <access_token>'
```

— Four full endpoints with methods and paths, plus runnable examples. product docs

02 Session Delegation via Custom Token Exchange NEW

81

In Open Early Access, requesting a Custom Token Exchange token with a `session_transfer` audience returns a Session Transfer Token (instead of an access token), authorized by the same Action logic already controlling delegation; delegated sessions are redeemed through the standard `/authorize` endpoint with no extra login

step for the target user. Both identities are preserved via the `sub` claim (the session owner) and the `act` claim per RFC 8693 (the acting party, e.g. a support agent), and dedicated tenant log events record delegated logins separately from regular logins for compliance and traceability.

.
.
.
.

— Mechanism and claims fully explained but no runnable example given.

Custom Token Exchange - Session Delegati…

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

OpenAI Codex CLI

2 RELEASES · 2026-08-07 → 2026-08-08 NOTES ↗

CODE ↗

Lightweight coding agent that runs in your terminal

> OpenAI Codex CLI's latest alpha releases add session forking and a migrate-rollouts command for thread history, Markdown conversation export and archiving in the TUI, a Guardian circuit-breaker for cyber-specialty models, and expanded MCP hook/event support alongside agent-identity endpoint overrides and several smaller diagnostics and model-list additions.

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

01 Session forking in codex exec NEW

93

Adds a `fork` subcommand to `codex exec` that branches a previous session into a new one, accepting a `SESSION_ID`, optional `--image` / `-i` file attachments, and an optional `PROMPT` (or `-` to read from stdin).

Branch off a previous session to explore an alternative approach without losing the original thread.

```
$ codex exec fork session-123 "try a different refactor approach"
```

Fork a session and attach a screenshot as context for the follow-up prompt.

```
$ codex exec fork session-123 --image screenshot.png "fix the layout issue shown in the image"
```

— Full flag list plus two runnable command examples. `rust-v0.148.0-alpha.5` `rust-v0.148.0-alpha.2`

02 Agent identity endpoint overrides NEW

80

Honors `CODEX_AGENT_IDENTITY_AUTHAPI_BASE_URL` and `CODEX_AGENT_IDENTITY_JWKS_BASE_URL` environment variables to override agent identity registration and verification endpoints.

— Names exact environment variables to set. `rust-v0.148.0-alpha.2`

03 MCP tool hook handler configuration NEW

78

Recognizes `mcp_tool` hook handler configurations (with server, tool, JSON input, and optional timeout and status message) and returns them from `configRequirements/read`.

— Names config fields and the read endpoint. `rust-v0.148.0-alpha.2`

04 **Markdown conversation export in TUI** NEW 75

Adds a `/export` command in the TUI that exports the full conversation history as structured Markdown, supporting clipboard and file destinations, default filename prompting, and an explicit path argument.

— Names command and destinations but no worked example. `rust-v0.148.0-alpha.5` `rust-v0.148.0-alpha.2`

05 **wait_agent timeout clamping** IMPROVED 70

Clamps `wait_agent` timeout requests below `min_wait_timeout_ms` to the configured minimum instead of rejecting them, and reports the adjustment in the tool output.

— Named parameters but no user-facing command. `rust-v0.148.0-alpha.2`

THINNER COVERAGE BELOW

06 **Session archiving and restore in resume picker** NEW 40

Adds session archiving to the resume picker in the TUI, allowing sessions to be archived and restored.

— Describes behaviour but no command or flag named. `rust-v0.148.0-alpha.5` `rust-v0.148.0-alpha.2`

07 **Diagnostics: app-server API and process snapshots** NEW 35

Exposes app-server diagnostics through the experimental API and adds process diagnostics snapshots.

— Groups two thin diagnostics additions without endpoint detail. `rust-v0.148.0-alpha.5`

08 **Multi-agent versions in model/list** IMPROVED 35

Exposes multi-agent versions in `model/list`.

— Names endpoint field but no further detail. `rust-v0.148.0-alpha.5`

09 **Call IDs in MCP requests** IMPROVED 21

Includes call IDs in MCP requests.

— Bare statement without field name or usage. `rust-v0.148.0-alpha.5`

10 **MCP event discovery and subscriptions** NEW 18

Adds MCP event discovery and subscriptions.

— Bare description with no mechanism or endpoint. `rust-v0.148.0-alpha.5`

---

11 **Resilient response streams on connection failure** IMPROVED

18

Keeps response streams alive through connection failures.

— States behaviour without mechanism or scope. `rust-v0.148.0-alpha.5`

---

12 **Unified image budget across agents** NEW

16

Adds a unified image budget across agents.

— No limit, mechanism, or config named. `rust-v0.148.0-alpha.5`

---

13 **Default code-mode exec yield timeout** IMPROVED

16

Configures the default code-mode exec yield timeout.

— No value or config key given. `rust-v0.148.0-alpha.5`

WAS THIS USEFUL? 👍 👎

Claude Code

1 RELEASE · 2026-08-08 | 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.

> v2.1.225 improves gateway spend-limit warnings and expands Remote Control session messaging, alongside a new workspace trust prompt for `claude agents`.

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

01 Remote Control session messaging improvements

IMPROVED

70

`SendMessage` can now initiate a conversation with Remote Control sessions on other machines by name, with `ListAgents` surfacing them as `name [ref]`, instead of only being able to reply after they message you first. `SendMessage` also now locks in a Remote Control recipient you have already confirmed, preventing it from being silently swapped for a same-named session on the local machine when the remote list could not be checked. Separately, photos attached from the Claude app are now passed directly to Claude rather than being read from disk via a separate tool call.

— Names APIs and format string, explains before/after behavior v2.1.225

THINNER COVERAGE BELOW

02 Gateway spend-limit warnings

IMPROVED

55

Claude Code's usage warning now supports gateway spend limits — when the limit is reached, the message names the cap, its reset time, and the operator's message. This requires the gateway to be on version 2.1.225.

— Names cap/reset/message fields but no config surface shown v2.1.225

03 Workspace trust prompt in claude agents

IMPROVED

40

`claude agents` now shows a workspace trust prompt for untrusted directories, matching the existing behavior of `claude`.

— Names the command but no detail on prompt options v2.1.225

WAS THIS USEFUL? 👍 👎

OpenHands

1 RELEASE · 2026-08-07 NOTES ↗

OpenHands: AI-Driven Development

> OpenHands v1.11.0 adds per-run LLM cost visibility, automation tag filtering, and a typed action for launching child conversations, alongside several UI refinements.

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

01 **Typed action for child conversations** NEW 55

Adds a typed agent action for launching local or Cloud child conversations, enabling programmatic multi-conversation workflows.

– Explains purpose and scope but no exact API name. v1.11.0

02 **Per-run LLM cost visibility** NEW 50

Shows per-run LLM cost in the Activity Log and in exports, giving teams immediate spend visibility for each agent run.

– Names Activity Log and exports but no metrics or limits. v1.11.0

03 **Tag metadata in conversations list** IMPROVED 45

Adds tag chips, overflow handling, and hovercard labels to the conversations list for richer conversation metadata display.

– Names three specific UI elements added to the list. v1.11.0

04 **UI polish across Automations, Customize, and Agent Canvas**

IMPROVED 40

Keeps the Automations pane visible even when empty to reduce navigation confusion, reorders the Customize navigation section for improved discoverability, and polishes the Agent Canvas version-update UI in Settings.

– Names three UI areas but each change described thinly. v1.11.0

05 **Automation tag filtering and recognition** NEW 35

Adds an automation tag filter and recognition so automations can be organized and surfaced by tag.

– States the capability but no mechanism or UI path detail. v1.11.0

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

LangGraph ⓘ 1 RELEASE · 2026-08-07 NOTES ↗

Build resilient agents.

> LangGraph's checkpoint package gains an opt-in flag to filter out expired state rows when reading checkpoints.

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

01 **Opt-in omit_expired flag for checkpoint reads** NEW 70

The `omit_expired` parameter can be passed to checkpoint reads, supported in the `checkpoint` and `checkpoint-postgres` packages, to skip expired rows and prevent stale state from surfacing in long-running workflows.

— Names exact flag and packages, but no usage example. `checkpoint==4.2.0`

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Daytona

1 RELEASE · seen 2026-08-08 NOTES ↗

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

> Daytona 0.199.0 adds time-to-live control for sandboxes and raises the minimum supported Python version for its SDK.

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

01 **Sandbox TTL support in SDK and CLI** NEW 45

Adds sandbox time-to-live (TTL) support to both the SDK and CLI, giving users control over how long a sandbox lives before it expires.

— Names SDK and CLI but no flag or parameter names given. snapshot-20260808

02 **Python 3.10 minimum requirement** BREAKING 40

The Python SDK now requires Python >=3.10; installations running older Python versions will break on upgrade.

— Clear version floor but no migration steps given. snapshot-20260808

└──▷ BREAKING ON UPGRADE

! Python SDK now requires Python >=3.10; installations running older Python versions will break on upgrade.

WAS THIS USEFUL? 👍 👎

NoLabs nono

1 RELEASE · 2026-08-07 | NOTES ↗ | CODE ↗

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

> nono v0.72.0 tightens sandbox diagnostics and network enforcement on Linux, while removing two built-in profiles and changing how the external sandbox-policy mode interacts with Landlock.

WHAT SHIPPED · 4 FEATURES

most completely described first — ? what's the number?

01 External sandbox-policy mode no longer bypasses Landlock

70

BREAKING

On Linux, the `external` `--sandbox-policy` mode no longer disables all nono sandboxing: Landlock still handles filesystem/process policy, and only TCP enforcement is delegated externally. Configurations that relied on `external` bypassing Landlock entirely will behave differently.

— Names the flag and describes the exact behavior change. v0.72.0

02 Deny-path reporting in `why` subcommand IMPROVED

65

The `why` subcommand now reports explicit deny paths from the sandbox policy, making it easier to diagnose why a network or filesystem access was blocked.

— Names the exact subcommand and its new diagnostic behavior. v0.72.0

03 Removal of openclaw and swival built-in profiles BREAKING

60

Removes `openclaw` and `swival` as built-in profiles; users relying on them must pull the equivalent packages from the registry, and any setup using them will break on upgrade until they do so.

— Names both removed profiles and the required workaround. v0.72.0

THINNER COVERAGE BELOW

04 Static seccomp network baseline on Linux NEW

25

Adds a static seccomp network baseline on Linux as part of the sandbox policy enforcement.

— Named only in summary, no mechanism or usage detail given. v0.72.0

BREAKING ON UPGRADE

! The `openclaw` and `swival` built-in profiles have been removed; any setup using them will break on upgrade until the equivalent profiles are pulled from the registry.

! The `external` `--sandbox-policy` mode on Linux no longer disables all nono sandboxing: Landlock still handles filesystem/process policy, and only TCP enforcement is delegated externally. Configurations that relied on `external` bypassing Landlock will behave differently.

WAS THIS USEFUL? 👍 👎

SecureAgentics Adrian

changes since 2026-07-09 CODE ↗

Open-source runtime AI agent security tool - monitors and controls AI agents, catching malicious tool use, prompt injection, and policy drift in real time, before the agent acts.

Adrian gains a Claude Code plugin with HITL verdict forwarding, three execution modes, and PII filtering for agent conversations.

└──▷ GET THIS VERSION

```
$ git clone --branch commits-2026-07-09
https://github.com/secureagentics/Adrian.git
# already have the repo? check out this version:
$ git checkout commits-2026-07-09
```

- New Claude Code plugin (`integrations/claude-code/adrian_cc/agent.py`) connects Claude Code hooks directly to the Adrian backend via WebSocket, sending `SessionLogin` + `PairedEvent` protobuf messages and returning allow/deny/advisory verdicts.
- Supports three server-driven execution modes read from `PolicySnapshot` on `LoginAck` : `MODE_ALERT` (fire-and-forget, never blocks), `MODE_BLOCK` (waits for verdict, blocks per policy), and `MODE_HITL` (waits for human approval on the dashboard).
- Adds HITL verdict forwarding: human approval decisions made on the Adrian dashboard are now forwarded inline to Claude Code integrations tagged with the CC source.
- Adds PII filtering for the Claude Code plugin (`integrations/claude-code/adrian_cc/pii/_patterns.py`), detecting and redacting `EMAIL` , `PHONE` , `SSN` , `CREDIT_CARD` , `IP_ADDRESS` , `DATE_OF_BIRTH` , `IBAN` , `PASSPORT` , `STREET_ADDRESS` , `POSTAL_CODE` , `DRIVER_LICENSE` , and `AWS_KEY` categories from agent conversation content.
- Adds a transcript parser (`integrations/claude-code/adrian_cc/transcript.py`) that reads Claude Code's JSONL transcript to extract thinking blocks, user prompts, system prompt context, and invocation metadata, with mtime-based caching to avoid re-parsing on every hook call.

WAS THIS USEFUL? 👍 👎

INFRASTRUCTURE

◆ Internet Services

Cloudflare

1 RELEASE · 2026-08-08 CHANGELOG ↗

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

> Cloudflare's biggest change this window unifies Workers AI and AI Gateway billing and raises frontier-model rate limits, alongside a new low-overhead Kitesurf browser for agentic workloads, expanded Radar BGP APIs and AI-driven Radar tooling, and new AI Search crawling and public-endpoint options.

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

01 **Kitesurf browser option for Browser Run** NEW 85

Adds a `browser=kitesurf` parameter to Browser Run `CDP` and `Quick Action` endpoints to opt into Kitesurf, a new stateless browser using 3–7x less CPU and memory than Chromium, aimed at agentic workloads.

Take a screenshot via Browser Run using Kitesurf to reduce CPU/memory overhead in an AI agent pipeline.

```
$ curl -X GET
'https://api.cloudflare.com/client/v4/accounts/<ACCOUNT_ID>/brow
ser-run/screenshot?browser=kitesurf' \
  -H 'Authorization: Bearer <API_TOKEN>' \
  -H 'Content-Type: application/json' \
  -d '{"url": "https://example.com"}' \
  --output screenshot.png
```

— Exact param, endpoints and a runnable curl example `changelog-20260808-37b7380e`

02 **Discover crawl mode for AI Search web sources** NEW 85

Adds `parse_type: discover` for AI Search web-crawler data sources, with `discover_options` fields (`source` , `limit` , `depth`) to crawl pages beyond a sitemap by following links.

Index a site without a complete sitemap by using the `discover` parse type to follow links up to a set depth and page limit.

```
$ curl -X POST
'https://api.cloudflare.com/client/v4/accounts/<ACCOUNT_ID>/ai-
search/instances' \
  -H 'Authorization: Bearer <API_TOKEN>' \
  -H 'Content-Type: application/json' \
  -d '{
    "id": "my-ai-search",
    "type": "web-crawler",
    "source": "example.com",
    "source_params": {
      "web_crawler": {
        "parse_type": "discover",
        "discover_options": { "source": "links", "limit": 5000,
"depth": 3 }
      }
    }
  }'
```

— Named fields plus a full runnable API example `changelog-20260808-37b7380e`

03 **Unified Workers AI and AI Gateway billing and routing** NEW

80

Introduces unified billing between Workers AI and AI Gateway: setting the gateway's Workers AI billing setting to `Unified Billing` lets prepaid AI Gateway credits cover Workers AI inference in real time. Frontier model rate limits rise to 50 requests per minute per account per model when billed through AI Gateway credits, up from 20 via standard Workers AI billing. env.AI.run() requests can be routed through an AI Gateway by specifying `gateway: { id: "<gateway-id>" }`, with `"default"` auto-creating a gateway on first authenticated request.

— Names config keys and exact limits but no runnable example `changelog-20260808-37b7380e`

04 **Public endpoints for AI Search on custom domains and namespaces**

NEW

75

Enables AI Search public endpoints on custom domains (e.g. `https://search.example.com/search`) with optional Cloudflare Access authentication in front of them, and adds namespace-level public endpoints (`https://ns-<NAMESPACE_ENDPOINT_ID>.search.ai.cloudflare.com/search` and `/chat/completions`) that fan out across multiple AI Search instances via `ai_search_options.instance_ids`.

— Exact endpoint URLs and field named, no example call `changelog-20260808-37b7380e`

05 **Radar Researcher AI assistant** NEW

60

Introduces Radar Researcher, a beta AI assistant on Cloudflare Radar with voice/text queries, interactive charts, searchable conversation history, shareable links, and an `Explain with AI` per-chart option.

— Describes capabilities but only points to the Radar UI `changelog-20260808-37b7380e`

THINNER COVERAGE BELOW

06 **AS-level connectivity graph and upstream widget on Radar** NEW

55

Adds an AS-level connectivity graph and Upstream providers widget to Cloudflare Radar AS pages (e.g. AS13335), with a `Show full paths` toggle and an IP version selector for IPv4/IPv6.

— UI feature with a starting page but no exact steps `changelog-20260808-37b7380e`

07 **cloudflare/mesh Docker image** NEW

55

Publishes the `cloudflare/mesh` Docker image on Docker Hub supporting `amd64` and `arm64`, enabling Cloudflare Mesh nodes to run in Docker Compose, Kubernetes StatefulSet, Kubernetes sidecar, and CI/CD pipelines.

— Names image and platforms but no command shown `changelog-20260808-37b7380e`

08 **WebMCP support in Cloudflare Radar** NEW

45

Adds WebMCP support to Cloudflare Radar, enabling browser-based AI agents to navigate Radar, search data, and use tools such as URL scanning and domain lookup.

— Prose description without concrete steps or interface names `changelog-20260808-37b7380e`

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · August 8, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe