THE CYBER TOOLCHAIN

NO. 002

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 2, 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 20 tools

APPLICATION SECURITY

vim ? help

◆ Web Application Security

ProjectDiscovery nuclei v3.10.0

Nuclei is a fast, customizable vulnerability scanner powered by the global security community and built on a simple YAML-based DSL, enabling collaboration to tackle trending vulnerabilities on the internet. It helps you find vulnerabilities in your applications, APIs, networks, DNS, and cloud configurations.

Nuclei v3.10.0 adds per-host HTTP client pooling for improved scan throughput and connection efficiency.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.10.0
https://github.com/projectdiscovery/nuclei.git
# already have the repo? check out this version:
$ git checkout v3.10.0
```

› Adds per-host HTTP client pooling, reducing connection overhead and improving parallelism when scanning many distinct hosts.

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

Truffle Security trufflehog v3.95.8

trufflehog discovers, classifies, and validates leaked credentials

TruffleHog v3.95.8 adds Octopus Deploy detection, Prometheus engine metrics, and a flag to drop unverified JWT results.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.95.8
https://github.com/trufflesecurity/trufflehog.git
# already have the repo? check out this version:
$ git checkout v3.95.8
```

└──▷ TRY IT

Suppress noisy unverified JWT hits when scanning a repo where JWTs are common but you only care about confirmed leaks.

```
$ trufflehog git https://github.com/example/myrepo --
results=verified,unknown --drop-unverified-jwt-results
```

- Adds Prometheus metrics exposing engine channel depths/capacities, worker counts, and bytes/chunks/secrets found — enabling real-time observability of scan progress.
- New `--drop-unverified-jwt-results` flag suppresses JWT detector results that are unverified with no verification errors, reducing noise in high-volume scans.
- Azure SAS token detector now matches tokens regardless of query-parameter order and handles URL-encoded timestamps, catching tokens emitted by tools like Azure Storage Explorer that were previously missed.
- Grafana detector updated to treat HTTP 401 responses as unverified (instead of verified), tightening classification of rotated/revoked tokens.

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Malware Analysis

quark-engine v26.7.1

Quark Engine is an Android malware analysis tool that uses a threat scoring system to detect malicious behaviors by analyzing APK files.

v26.7.1 adds NGate & Cerberus detection rules and Phase 5 dynamic string resolution via DexTrace.

└──▷ GET THIS VERSION

```
$ git clone --branch v26.7.1 https://github.com/ev-flow/quark-engine.git
# already have the repo? check out this version:
$ git checkout v26.7.1
```

› Enables Phase 5 dynamic string resolution via DexTrace, extending analysis to runtime-resolved obfuscated strings.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

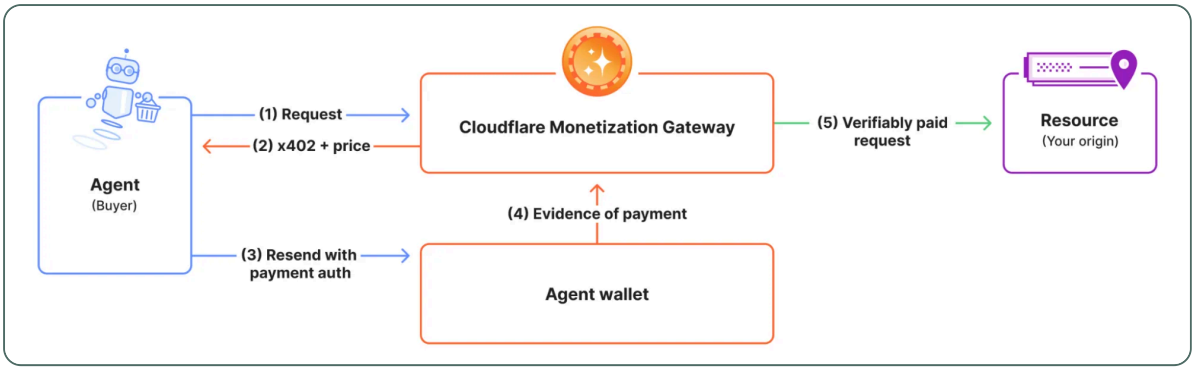
Cloudflare

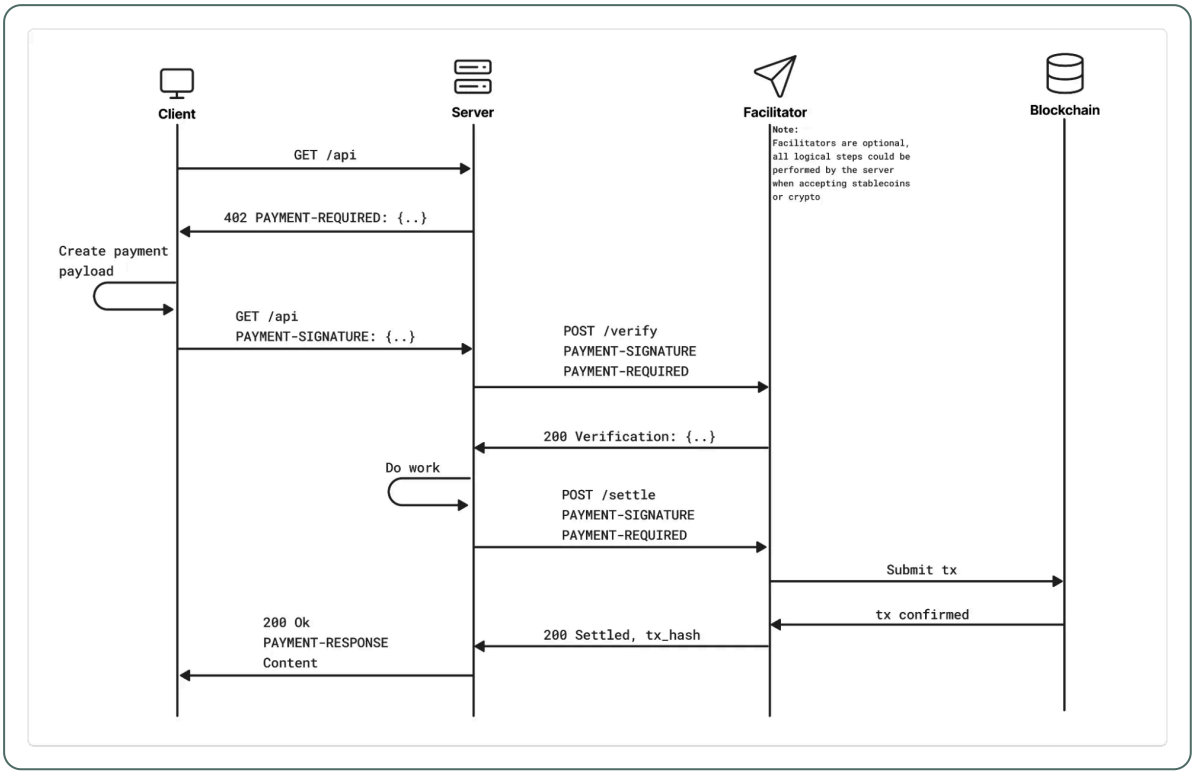
Announcing the Monetization Gateway: charge for any resource behind Cloudflare via x40

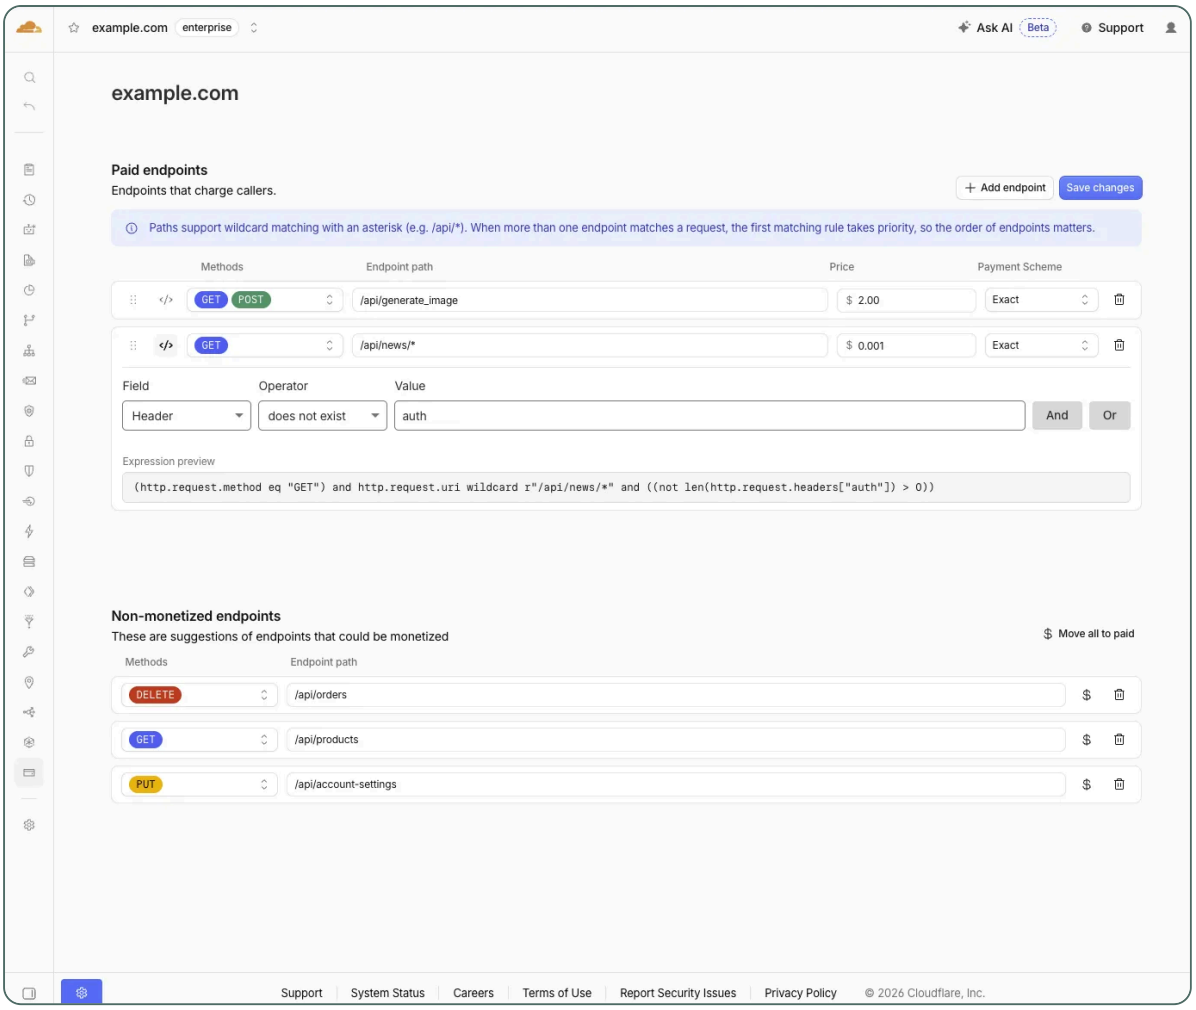
4 RELEASES · 2026-07-01 LATEST STABLE

Cloudflare Blog is a publication sharing updates, technical insights, and announcements about Cloudflare's internet security, performance, and developer infrastructure products.

Cloudflare Monetization Gateway lets you charge any caller for any resource via x402 stablecoin payments at the edge.







└──▷ HOW TO FIND IT

Gate a premium API route so any agent caller — with no prior account — must pay $0.01 per request before your origin is reached.

> 📍 In the Cloudflare dashboard, go to Monetization Gateway › Payment Rules › Create Rule. Set the route expression to `http.request.uri.path matches "/api/premium/*"`, verb to GET/POST, price to $0.01 USD (stablecoin), then Save and Deploy.

Monetize unauthenticated traffic by converting 401 responses into 402 Payment Required — no origin code changes needed.

> 📍 In the Cloudflare dashboard, go to Monetization Gateway › Payment Rules › Create Rule. Enable `Intercept HTTP 401 responses` and set the 402 price and stablecoin asset. Callers that would have received 401 will instead receive a 402 with payment instructions.

› New Cloudflare Monetization Gateway enables usage-based payment enforcement at the edge for web pages, datasets, APIs, and MCP tool calls — no origin billing system required.

- Supports the x402 open protocol (HTTP 402 Payment Required), allowing agentic callers to discover price, pay in stablecoins, and receive access in a single HTTP round-trip with no prior account or API key.
- Flexible payment rules API lets you gate specific REST verbs, routes (e.g., `/api/premium/*`), or unauthenticated callers with per-request pricing as low as fractions of a cent.
- Variable pricing support enables charging different amounts based on task complexity — for example, up to $2 for compute-intensive requests like image generation.
- Interceptor mode converts HTTP 401 Unauthorized responses from your origin into 402 Payment Required with pricing instructions, monetizing unauthenticated traffic without origin changes.

+4 more

▾ **3 more releases in this issue** · 2026-07-01

Making AI search smarter (2026-07-01) STABLE

Cloudflare launches AI search quality research program and Pay Per Use experiments to compensate content creators.

└──▷ HOW TO FIND IT

Express interest in the Ceramic.ai pay-per-query or You.com pay-per-content programs to start earning compensation when your content appears in AI search results.

> 📍 Navigate to the Cloudflare interest form linked in the blog post (`blog.cloudflare.com`/making-ai-search-smarter/) and submit your details to opt in to the Ceramic or You.com compensation programs.

If you operate an answer engine and want to integrate Cloudflare's freshness/traffic signals to crawl smarter and reduce redundant fetches.

> 📍 Contact aeo@cloudflare.com to join the crawl-quality research program as an answer engine partner.

- Launches research program to share content-freshness and traffic signals with answer engines, enabling smarter crawling that skips unchanged pages and surfaces higher-quality content.
- Introduces Pay Per Use model (evolving from Pay Per Crawl), letting AI companies pay content owners per query or per result rather than per crawl.
- New partnership with Ceramic.ai enables a pay-per-query model where publishers are compensated each time their content appears in Ceramic search results.
- New partnership with You.com enables on-demand, per-content payments by agents

for specific premium content without upfront commitment.

› Adds Answer Engine Optimization (AEO) reporting for participating publishers, showing top queries driving content appearances, specific snippets surfaced, average ranking position, and more.

`Your site, your rules: new AI traffic options for all customers (2026-07-01)`

STABLE

Cloudflare launches granular AI traffic controls for all tiers, BotBase bot directory for Enterprise, and a new content-use signal in robots.txt.

⋮

└──▷ USE IT

Express a strict content-use preference in robots.txt so crawlers know they may index and link but must not store or reproduce content beyond excerpts.

```ini
# robots.txt — manually authored or via Cloudflare managed
robots.txt
User-agent: *
Content-Signal: search=yes,ai-train=no,use=reference
Allow: /
```

Allow only Search, SEO, and Ads Verification bots up to the reference content-use level — useful for a publisher that wants discoverability but no content reproduction.

📍 In the Cloudflare dashboard, go to Security › Bots › Bot Management configuration card › BotBase, filter by classification (Search, SEO, Ads Verification), then set content use level to `reference` and block all other classifications.

Find a specific bot's detection ID in BotBase to target it precisely in a Security rule — useful when you want to allow one verified agent while blocking the rest of its category.

📍 In the Cloudflare dashboard, go to Security › Bots › Bot Management configuration card › BotBase, search for the bot by name, copy its Detection ID, then reference that ID in a custom Security rule.

› New per-category AI traffic controls — Search, Agent, and Training — available to all customers including Free tier, replacing the single `Block AI Bots` toggle.

› New default enforcement (September 15, 2026): Training and Agent crawlers blocked by default on ad-serving pages for new domains; multi-purpose crawlers (e.g., Googlebot, BingBot) subject to the most restrictive applicable rule.

- Introduces BotBase, a searchable database of all known bots and agents with taxonomy classification, detection IDs, and Security-rule integration — available to Enterprise Bot Management customers via the Bot Management configuration card.
- New `content use` signal (use=) extends Content Signals in robots.txt with three levels — immediate, reference, and full — letting owners express how crawlers may store and reshare content.
- Managed robots.txt now automatically includes `use=reference` alongside existing `search=yes`, `ai-train=no` Content Signals for all customers who have enabled it.

+3 more

`Unmasking the crawls with Attribution Business Insights (2026-07-01)`

STABLE

Cloudflare launches Attribution Business Insights dashboard for Bot Management customers to analyze AI crawler impact.

.
.
.

└──▷ HOW TO FIND IT

Check your site's crawl-to-referral ratio over the past 30 days to quantify how extractive a specific AI crawler is before entering a licensing conversation.

📍 In the Cloudflare dashboard, go to Security › Bots › Attribution Business Insights, set the time range to 30 days, and review the crawl-to-referral ratio column in the per-bot-operator breakdown.

- New Attribution Business Insights dashboard gives Bot Management customers a dedicated analytics hub for AI crawler and bot traffic, with no manual filtering required.
- Adds site-wide and per-bot-operator crawl-to-referral ratio metrics, viewable across 24-hour, 7-day, and 30-day windows.
- Introduces updated bot taxonomy that classifies AI crawlers as Training, Search, or Agent — going beyond a generic `AI Crawler` label.
- Top bots breakdown surfaces each bot's country of origin, bandwidth consumption, and current allow/block status in a single view.
- Per-operator breakdown consolidates all bots owned by a single AI company, enabling data-driven content licensing and business negotiations.

WAS THIS USEFUL? 👍 👎

◆ Container & Image Security

Aqua Security Trivy v0.72.0

Find vulnerabilities, misconfigurations, secrets, SBOM in containers, Kubernetes, code repositories, clouds and more

Trivy v0.72.0 adds Bottlerocket scanning, OpenAI/GitHub App secret rules, .NET self-contained runtime detection, and JAR license extraction.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.72.0
https://github.com/aquasecurity/trivy.git
# already have the repo? check out this version:
$ git checkout v0.72.0
```

└──▷ TRY IT

Scan a Bottlerocket-based container image for OS-level CVEs — now matched against the Bottlerocket advisory database.

```
$ trivy image --scanners vuln <bottlerocket-based-image>
```

Audit a self-contained .NET binary or publish directory for vulnerabilities in the bundled runtime, not just your app dependencies.

```
$ trivy fs --scanners vuln <path/to/self-contained-dotnet-publish/>
```

Scan a codebase for leaked OpenAI API keys and the new stateless GitHub App installation token format alongside other secrets.

```
$ trivy fs --scanners secret <your-repo-path/>
```

- Adds vulnerability matching for Bottlerocket OS images.
- Detects bundled .NET runtime in self-contained deployments, enabling CVE scanning of apps that ship their own runtime.
- Adds OpenAI secret detection rules to the secrets scanner.
- Adds detection of the new stateless format for GitHub App installation tokens in secrets scanning.

› Extracts JAR licenses from embedded pom.xml files for more complete Java license inventory.

+2 more

WAS THIS USEFUL? 👍 👎

◆ Cloud Security Posture

Wiz

Build AI Security Agents with Wiz MCP (2026-07-02)

2 RELEASES · 2026-06-30 → 2026-07-02 LATEST STABLE

Wiz Blog is a cloud security publication that shares research, threat intelligence, best practices, and insights on vulnerabilities, compliance, and cloud infrastructure protection.

Wiz MCP is now generally available, connecting AI agents and assistants to the Wiz Security Graph, AI Agents, and reusable AI Skills.

└──▷ HOW TO FIND IT

Let an AI agent surface only exploitable vulnerabilities an attacker could chain to reach sensitive data — ranked by remediation priority rather than raw CVE severity.

> 📍 Surface only the open vulnerabilities an attacker could realistically chain to reach sensitive data, explain the attack path for each, and rank them by what to remediate first.

Operationalize a penetration test report by importing it into Wiz, mapping findings to cloud assets, and creating tracked Penetration Test Findings — eliminating the static-report problem.

> 📍 Import this penetration test report into Wiz, map findings to my environment, and create Penetration Test Findings.

Give a security leader an auditable compliance gap analysis by comparing an internal framework against existing Wiz policies and publishing results to a documentation tool.

> 📍 Compare our internal security framework against existing Wiz policy coverage, identify the gaps, recommend custom rules to close them, and produce an auditable coverage summary.

- Introduces Wiz MCP (general availability): a secure integration layer that connects AI assistants, custom agents, and AI-powered applications to the Wiz platform.
- Enables AI to query and reason over the Wiz Security Graph — including findings, issues, vulnerabilities, attack paths, identities, cloud infrastructure, code, data, and runtime relationships.
- Exposes Wiz AI Agents (Red, Blue, and Green) through MCP so external AI can consume agent analyses and recommendations without recreating the same reasoning.
- Provides access to Wiz AI Skills — predefined, reusable agentic playbooks for vulnerability triage, remediation, threat investigation, compliance mapping, and more.
- Supports an Import Security Reports Skill that extracts findings from external penetration test reports, maps them to cloud resources, and creates Penetration Test

Findings in Wiz.

+3 more

▾ **1 more release in this issue** · 2026-06-30 → 2026-07-02

Start Secure in the AI Era: Accelerating AI Threat Readiness with WizOS (2026-06-30)

STABLE

Wiz adds Secure Architecture Opportunities, WizOS Migration Skill, and Green Agent WizOS remediation to accelerate container hardening.

└──▷ HOW TO FIND IT

Find which container image types carry the most CVE risk and get a prioritized migration plan before touching any Dockerfiles.

> 📍 In the Wiz console, go to Findings › Secure Architecture Opportunities, filter by image type (e.g., Python, nginx), and review the associated vulnerabilities, security issues, and estimated effort score to pick your highest-impact migration target.

Patch a critical container CVE without leaving your AI coding agent — let the WizOS Migration Skill swap the base image automatically.

> 📍 In your AI coding agent, invoke the WizOS Migration Skill against your Dockerfile; the skill identifies the appropriate WizOS replacement image and applies the base-image swap directly in the file.

Accelerate MTTR on a critical container finding by routing a WizOS fix straight to an engineering issue with full context.

> 📍 In a Wiz Green Agent remediation recommendation for a container image CVE, click Send to Coding Agent — Wiz opens a repository issue, tags your AI coding agent, and passes the specific WizOS image and migration context for automated resolution.

- New Secure Architecture Opportunities page (public preview) in Wiz Findings groups WizOS migration candidates by image type (e.g., Python, nginx) with associated CVEs, related issues, and estimated migration effort for risk-ordered prioritization.
- New WizOS Migration Skill for AI coding agents surfaces the correct WizOS replacement image from Dockerfile context and can apply the base-image swap directly within the development environment.
- Wiz Green Agent now includes WizOS image recommendations in remediation guidance for container CVEs, with a Send to Coding Agent action that automatically opens a repo issue and passes full context to the coding agent for hands-off migration.

WAS THIS USEFUL? 👍 👎

Weaveworks eksctl v0.229.0

The official CLI for Amazon EKS

eksctl v0.229.0 adds cluster version rollback support and EC2 instance store for EKS on Outposts.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.229.0 https://github.com/eksctl-io/eksctl.git
# already have the repo? check out this version:
$ git checkout v0.229.0
```

└──▷ **TRY IT**

Roll back a cluster to a previous Kubernetes version after a failed upgrade by declaring rollback behaviour in your ClusterConfig.

```
$ eksctl upgrade cluster -f cluster-rollback.yaml
```

- Adds cluster version rollback support via `rollbackConfig` in the ClusterConfig spec.
- Adds EC2 instance store support for EKS on Outposts node groups.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

CrowdStrike

Falcon Cloud Security June 2026 Release: Updates for Azure and Google Cloud

CrowdStrike Blog is a cybersecurity publication sharing threat intelligence, research, and insights on adversaries, vulnerabilities, and endpoint protection best practices.

Falcon Cloud Security June 2026 expands real-time CSPM, Cloud Risks, DSPM, and CIEM to Azure and Google Cloud.

- Extends real-time CSPM to Google Cloud, surfacing new assets, misconfigurations, and cloud risks within minutes of introduction — matching existing AWS and Azure coverage.
- Expands correlated Cloud Risks breach-path detection to Azure and Google Cloud, previously limited to AWS environments.
- Adds adversary-informed risk prioritization to Azure and Google Cloud, mapping known adversary activity to cloud attack paths across all three major cloud providers.
- Introduces DSPM for Google Cloud Storage, enabling agentless discovery, sensitive data classification, and exposure identification alongside existing AWS and Azure DSPM coverage.
- Adds CIEM for Azure, allowing security teams to analyze permissions, understand effective access, and identify excessive entitlements across AWS and Azure using a unified workflow.

+1 more

WAS THIS USEFUL? 👍 👎

Microsoft Security

What's new in Microsoft Security: June 2026 (2026-06-30)

Microsoft Security provides comprehensive cybersecurity solutions, tools, and threat intelligence to help organizations protect identities, data, endpoints, and cloud environments from cyberattacks.

Microsoft Security June 2026: AI agent protection, multicloud DB coverage, Entra backup GA, and unified identity risk scoring.

- Introduces Codename MDASH (private preview): a multi-model agentic scanning system that discovers, validates, and routes software vulnerabilities into Defender workflows and engineering pipelines for closed-loop remediation.
- Expands Microsoft Defender endpoint protection to discover 25+ types of local AI agents and Model Context Protocol (MCP) servers on managed Windows and macOS devices, with runtime prompt-injection detection and blocking for coding agents such as GitHub Copilot CLI and Claude Code.
- Adds Advanced Hunting support for investigating AI agent exposure across an environment (preview).
- Extends Microsoft Defender for Cloud database threat protection (GA) to open-source relational databases on AWS RDS, adding anomalous-access and brute-force detection alongside automated sensitive data discovery.
- Ships Microsoft Purview customizable reports (GA) in Data Security Posture Management (DSPM), enabling role-specific, exportable report views with out-of-the-box and custom report creation for data security trend analysis.

+2 more

WAS THIS USEFUL? 👍 👎

◆ Detection Engineering

SigmaHQ sigma-cli v3.0.3

The Sigma command line interface based on pySigma

sigma-cli v3.0.3 adds Homebrew installation support for macOS users.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.0.3 https://github.com/SigmaHQ/sigma-cli.git
# already have the repo? check out this version:
$ git checkout v3.0.3
```

└──▷ TRY IT

Install sigma-cli on macOS without Python tooling — useful for quick setup on analyst workstations.

```
$ brew install sigma-cli
```

› Supports Homebrew installation on macOS, enabling `brew install sigma-cli` as a new install path.

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Claude Code v2.1.198

3 RELEASES · 2026-06-29 → 2026-07-01 LATEST

STABLE

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 v2.1.198 adds Chrome GA, automated draft PRs from background agents, `/dataviz` skill, and AWS gateway support.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.1.198 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.198
```

└──▷ TRY IT

Get notified in your hook handler when a background agent needs input or completes, so you can automate follow-up actions.

```
$ # In your Notification hook handler, inspect the event type:
# agent_needs_input — agent is waiting for a response
# agent_completed  — agent finished its work
claude agents
```

Use the new `/dataviz` skill inside a session to get chart design guidance and validate color palettes for dashboards.

```
$ /dataviz
```

- Launches Claude in Chrome as generally available.
- Adds `agent_needs_input` and `agent_completed` Notification hook events in `claude agents` so background sessions fire alerts when they need input or finish.
- New `/dataviz` skill provides chart and dashboard design guidance with a runnable color-palette validator.
- Adds Claude Platform on AWS (`anthropicAws`) as a gateway upstream provider; model-not-found responses now advance the failover chain instead of failing hard.

› Background agents launched from `claude agents` now automatically commit, push, and open a draft PR when they finish code work in a worktree.

+4 more

└──▷ **BREAKING ON UPGRADE**

! The `--bg` flag combined with `--print` / `-p` is now rejected up front instead of silently creating an unattachable session.

! The `/agents` wizard has been removed; subagents must now be created by asking Claude or by editing `.claude/agents/` directly.

▾ **2 more releases in this issue** · 2026-06-29 → 2026-07-01

v2.1.197 STABLE

Claude Code now defaults to Claude Sonnet 5 with a native 1M-token context window and promotional pricing through August 31.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.1.197 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.197
```

› Switches default model to Claude Sonnet 5, featuring a native 1M-token context window for working across very large codebases in a single session.

› Introduces promotional pricing of $2/$10 per million input/output tokens for Claude Sonnet 5 through August 31.

└──▷ **BREAKING ON UPGRADE**

! Claude Sonnet 5 is now the default model in Claude Code, replacing the previously default model

v2.1.196 STABLE

Claude Code v2.1.196 adds org-default models, clickable file attachments, MCP security hardening, and an always-on stream watchdog.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.1.196 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.196
```

└──▷ TRY IT

Inspect which MCP servers are approved before trusting a cloned repo — untrusted servers now show as pending instead of silently starting.

```
$ claude mcp list
```

- Adds organization default model support — admins configure it in the org console and users see `Org default` or `Role default` in `/model` when no personal selection is made.
- Adds readable auto-generated names for sessions at startup, making sessions easier to identify and reference.
- Adds clickable file attachments in chat — Cmd/Ctrl-click a file to reveal it in Finder/Explorer.
- Security hardening: `claude mcp list` / `get` no longer auto-spawn MCP servers that a repo self-approved via a committed `.claude/settings.json`; untrusted workspaces now show a `⏸ Pending approval` indicator instead.
- Enables the streaming idle watchdog by default for all providers — aborts and retries any response stream that produces no events for 5 minutes; set `CLAUDE_ENABLE_STREAM_WATCHDOG=0` to opt out.

+4 more

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI rust-v0.143.0-alpha.33

3 RELEASES · 2026-06-29 → 2026-07-02 PRE-RELEASE

Lightweight coding agent that runs in your terminal

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.143.0-alpha.33
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.143.0-alpha.33
```

- Surfaces rollout/session budget exhaustion to the user instead of silently aborting.
- Enables image generation when using provider-based authentication.
- Adds MCP tool call error metrics for tracking MCP tool failures.
- Traces MCP startup latency in OTEL for diagnosing slow MCP server initialization.
- Instruments rollout persistence bytes for observability into context-window compaction.

▾ 2 more releases in this issue · 2026-06-29 → 2026-07-02

rust-v0.143.0-alpha.32 PRE-RELEASE

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.143.0-alpha.32
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.143.0-alpha.32
```

- Enables image generation when using provider-based authentication (not just API-key flows).
- Adds MCP tool call error metrics for observability into MCP tool failures.
- Traces MCP startup latency via OTEL for diagnosing slow MCP server initialization.

rust-v0.143.0-alpha.31 PRE-RELEASE

Remote plugins enabled by default and `max` reasoning effort promoted to first-class in the TUI picker.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.143.0-alpha.31
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.143.0-alpha.31
```

- Enables remote plugins by default — no configuration required; set `features.remote_plugin` to false to opt out.
- Promotes `max` to a first-class reasoning effort level, rendering it as a productized `Max` label in the TUI reasoning picker (relevant for AWS Bedrock GPT-5.6 catalog users).
- Uses model metadata (`include_skills_usage_instructions`) to drive skills usage instructions, enabling per-model skill rendering control starting with `gpt-5.5`.

WAS THIS USEFUL? 👍 👎

Cline v4.0.5

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

Cline v4.0.5 adds Claude Sonnet 5 support across Anthropic, Bedrock, Vertex, Claude Code, SAP AI Core, OpenRouter, and Vercel AI Gateway.

└──▷ GET THIS VERSION

```
$ git clone --branch v4.0.5 https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout v4.0.5
```

- Adds Claude Sonnet 5 (including 1M-context variant) to the model picker across Anthropic, Bedrock, Vertex, Claude Code, SAP AI Core, OpenRouter, and Vercel AI Gateway providers.
- Sets Claude Sonnet 5 as the new default/recommended model for providers where Sonnet is the default path (Anthropic, Claude Code, Bedrock, OpenRouter/Cline defaults).

WAS THIS USEFUL? 👍 👎

Goose v1.40.0

an open source, extensible AI agent that goes beyond code suggestions - install, execute, edit, and test with any LLM

Goose v1.40.0 adds 13+ new AI providers, hooks system with tool-use denial, local code review, new CLI slash commands, and prompt injection mitigation.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.40.0 https://github.com/aaif-goose/goose.git
# already have the repo? check out this version:
$ git checkout v1.40.0
```

└──▷ TRY IT

Cap how much output a noisy tool can return so long responses don't blow the context window.

```
$ GOOSE_MAX_TOOL_RESPONSE_SIZE=32768 goose run --recipe audit.yaml
```

Run an AI-powered local code review on the current working directory before opening a pull request.

```
$ goose review
```

- New `/status` slash command in the CLI to check session state at a glance.
- New `goose review` subcommand for running local AI-powered code review.
- Adds a hooks system for extensibility, including `PreToolUse` denial to block tool calls before execution and blocking `Stop` hook decisions.
- Prompt injection mitigation enabled by default.
- Supports 13+ new model providers: Together AI, EmpirioLabs, OrcaRouter, NEAR AI Cloud, Perplexity, Alibaba/Qwen via DashScope, Databricks AI Gateway, Scaleway, Vercel AI Gateway, xAI SuperGrok OAuth, OpenAI GPT via AWS Bedrock mantle endpoint, GLM-5.2, and MLX models for local inference.

+11 more

WAS THIS USEFUL? 👍 👎

OpenCode

v1.17.13 · 2 RELEASES · 2026-06-30 → 2026-07-01 · LATEST

STABLE

The open source coding agent.

Desktop app gains per-window tab scoping, searchable model picker, session tab hover previews, and streamlined WSL setup.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.17.13
https://github.com/anomalyco/opencode.git
# already have the repo? check out this version:
$ git checkout v1.17.13
```

- Adds searchable v2 model picker and model management flow in the composer for easier model switching.
- Adds hover preview for session tabs showing the project, path, branch, and server at a glance.
- Streamlines WSL server setup with clearer distro checks and guided install flows.

▾ 1 more release in this issue · 2026-06-30 → 2026-07-01

v1.17.12

STABLE

OpenCode v1.17.12 adds yolo auto-approve mode, live SDK event streams, paged session history, and MCP resource autocomplete.

└──▷ GET THIS VERSION

```
$ git clone --branch v1.17.12
https://github.com/anomalyco/opencode.git
# already have the repo? check out this version:
$ git checkout v1.17.12
```

- Adds yolo mode to the TUI that auto-approves all permission prompts without manual confirmation.
- Adds a live event subscription stream to the SDK for real-time session monitoring.
- Adds SDK access to active sessions.
- Adds paged durable session history via a new API endpoint.

> › Adds session permission request create and fetch endpoints to the SDK/API.
>
> +4 more

WAS THIS USEFUL? 👍 👎

Cursor MCPs and Organizations in Team Marketplaces

Cursor is an AI-powered code editor that helps developers write, edit, and understand code faster using integrated large language model assistance.

Cursor team marketplaces now support centrally managed Team MCP servers and org-group access controls.

.
.

└──▷ **HOW TO FIND IT**

Centrally publish approved MCP servers so team members can install them locally without configuring servers themselves.

> 📍 In the Cursor dashboard, go to Dashboard -> Integrations & MCP, configure your Team MCP servers, then enable them in the team marketplace so members can install from the marketplace.

Restrict a team marketplace to a specific organization group rather than the entire team or a SCIM directory group.

> 📍 In the Cursor dashboard, go to Dashboard -> Plugins -> Team Marketplaces, select the marketplace, and restrict access to the desired organization group.

- › Enables admins to configure Team MCP servers once and distribute them to cloud agents, the agents window, IDE, and CLI via a team marketplace — members install approved integrations locally without manual server setup.
- › Adds organization group support for team marketplace access control, letting admins restrict marketplace visibility to specific org groups (in addition to existing SCIM directory groups).

WAS THIS USEFUL? 👍 👎

Windsurf snapshot-20260702

Windsurf is an AI-powered IDE and coding assistant that helps developers write, edit, and debug code faster using intelligent suggestions and automation.

Windsurf snapshot-20260702 covers v3.3.18: ACU usage display, configurable subagent models, MCP registry pre-warming, and commit attribution control.

└──▷ USE IT

Suppress Devin attribution in commit messages — useful when org policy requires clean, tool-agnostic commit history.

```yaml
# In your Devin Local config file
attribution: false
```

- Enables configuring subagents with a default model, giving teams control over model selection per agent role.

WAS THIS USEFUL? 👍 👎

Kiro

IAM Roles and Authorize Powers for Third-Party Services

2 RELEASES · SEEN 2026-07-02 LATEST STABLE

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

Kiro adds AWS IAM role assumption for sandboxes and a standardized OAuth flow for third-party service Powers.

└──▷ HOW TO FIND IT

Grant Kiro's sandbox scoped AWS access so agent tasks and MCP servers can call AWS APIs without long-lived keys.

> 📍 In the console, go to Settings › Agent › Sandbox › AWS Credentials and configure the IAM role Kiro Web should assume.

Authorize a third-party service (e.g., Stripe or Figma) once so the agent can use it across all future tasks without re-prompting.

> 📍 In the console, go to Settings › Agent › Manage Powers, select the desired Power, and complete the authorization flow.

- New IAM Role for Sandbox: configure a role that Kiro Web assumes to supply the agent, CLI tools, and MCP servers with auto-refreshing, short-lived AWS credentials during task execution.
- New Authorize Powers flow: connect third-party services (Figma, Stripe, Supabase, etc.) once via a standard authorization flow in Agent settings; tokens are encrypted at rest and never exposed to the sandbox or agent.
- Instant revocation: disconnecting a Power immediately revokes access to the linked third-party service.

▾ **1 more release in this issue** · seen 2026-07-02

> Claude Sonnet 5 Now Available STABLE
>
> Claude Sonnet 5 is now available in Kiro IDE, CLI, and Web with 1M context and stronger agentic reasoning.



└──▷ HOW TO FIND IT

Switch to Claude Sonnet 5 in the Kiro IDE or CLI to leverage its stronger agentic coding and longer autonomous runs.

> 📍 Restart your IDE or CLI, then open the model selector and choose Claude Sonnet 5.

- Adds Claude Sonnet 5 to Kiro IDE, CLI, and Web — Anthropic's most agentic Sonnet model with stronger reasoning, tool use, and coding performance approaching Opus 4.8.
- Supports full 1M token context window at a 1.3× credit multiplier, matching Sonnet 4.6 pricing.
- Available to Kiro Pro, Pro+, Pro Max, and Power customers in AWS US-East-1 (N. Virginia) and AWS Europe (Frankfurt) with cross-region inference support.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

CrewAI 1.15.2a1 PRE-RELEASE

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

CrewAI 1.15.2a1 adds inline skill definitions, a stream frame protocol for Flows, and new tool/app types in CrewDefinition.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.15.2a1 https://github.com/crewAIInc/crewAI.git
# already have the repo? check out this version:
$ git checkout 1.15.2a1
```

- Supports inline skill definitions, enabling skills to be declared directly within agent or crew configurations without separate files.
- Introduces a stream frame protocol for Flows, defining a structured message format for streaming Flow execution events.
- Adds `tool` and `app` type options to `CrewDefinition`, expanding the kinds of crew artifacts that can be declared.
- Adds a generated Flow Definition authoring skill for AI coding agents, guiding them when scaffolding or editing Flow definitions.

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

Ollama v0.31.1

Get up and running with Kimi-K2.6, GLM-5.1, MiniMax, DeepSeek, gpt-oss, Qwen, Gemma and other models.

Gemma 4 on Apple Silicon gets ~90% faster token generation via automatic multi-token prediction (MTP), no config required.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.31.1 https://github.com/ollama/ollama.git
# already have the repo? check out this version:
$ git checkout v0.31.1
```

- Enables automatic multi-token prediction (MTP) for Gemma 4 on Apple Silicon, delivering ~90% faster token generation on coding-agent workloads with no configuration or output changes.
- Adds minimum-version enforcement for `ollama launch hermes-desktop`, automatically upgrading older Hermes CLI installs (v0.16.0+ required) instead of failing silently.
- Improves streamed tool-call parsing so closing braces or brackets inside JSON string argument values no longer cause tool calls to be silently dropped as plain text.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 2, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe