THE CYBER TOOLCHAIN

NO. 019

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 18, 2026 · EVERY WEEKDAY

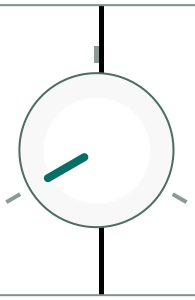
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 10 tools

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Capital One VulnHunter ⓘ v0.1.0 NOTES ↗

Agentic AI security tool that applies proactive, attacker-first analysis directly to source code.

Capital One open-sources VulnHunter: agentic AI SAST with attacker-first analysis, falsification engine, and automated fix-verify loop.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.1.0
https://github.com/capitalone/VulnHunter.git
# already have the repo? check out this version:
$ git checkout v0.1.0
```

└──▷ TRY IT

Run the full attacker-first vulnerability scan against a codebase inside a Claude Code session.

```
$ claude --model opus --add-dir ~/.claude/skills/vulnhunt --add-dir
~/.claude/skills/vulnhunt/phases
# Then inside the Claude Code session:
/vulnhunt
```

Independently verify that a prior fix actually blocked the exploit before merging, using the read-only verifier agent.

```
$ claude --model opus --add-dir ~/.claude/skills/vulnhunt-fix-verify \
       --add-dir ~/.claude/skills/vulnhunt-fix-verify/phases
# Then inside the Claude Code session:
/vulnhunt-fix-verify repo=<abs_path> report=<abs_path> fixed=VULN-
001,... out=<abs_path>
```

Batch-scan a list of repositories and collect all findings for centralized review, skipping repos already processed.

```
$ cd harness
python -m local_harness.batch.run scan --resume
python -m local_harness.batch.run collect
```

- Adds `/vulnhunt` Claude Code skill: entry-point-forward scanner that maps attacker paths to dangerous sinks and filters findings through a multi-stage falsification pipeline (Recon → Parallel Hunt → Adversarial Disprove → Capability Filter).
- Adds `/vulnhunter-fix` skill: test-driven remediation loop that writes an exploit demo, creates a failing security test (RED), implements the fix (GREEN), verifies the exploit is blocked, and opens a reviewable PR.
- Adds `/vulnhunt-fix-verify` skill: independent, read-only agent that validates whether a finding was successfully remediated and emits a per-finding verdict.
- Adds `vulnhunter-agent/` headless runtime: clones targets, runs `/vulnhunt`, publishes results, and opens GitHub issues for confirmed vulnerabilities — suitable for CI/CD pipelines.
- Adds `harness/` batch-scanning and benchmarking tooling: manage a repo list, resume interrupted scans, collect findings centrally, and evaluate detection accuracy against a known-vulnerable corpus with LLM-judge scoring.

WAS THIS USEFUL? 👍 👎

◆ Web Application Security

Cloudflare ⓘ

Cloudflare WAF protects WordPress applications from two high-severity vulnerabilities (2026-07-17)

2 RELEASES · 2026-07-17 | BLOG ↗ | LATEST | STABLE

Cloudflare is a cloud-based platform that provides website security, performance optimization, DDoS protection, CDN services, and DNS management for businesses and developers.

Cloudflare WAF deploys two new managed rules blocking critical WordPress RCE (CVE-2026-63030) and SQL injection (CVE-2026-60137) for all plan tiers.

└──▷ HOW TO FIND IT

Verify Managed Rules are enabled on a Pro/Business/Enterprise zone so the new WordPress rules are active with the Block action rather than Log.

📍 In the Cloudflare dashboard, go to Security › WAF › Managed Rules, confirm the Cloudflare Managed Ruleset is enabled, then check Security Events and filter by the two new rule IDs to monitor matching traffic while you patch WordPress.

› New WAF managed rule detects and blocks SQL injection attempts targeting WordPress 6.8+ (CVE-2026-60137) before crafted parameter values reach the application.

› New WAF managed rule detects and blocks unauthenticated RCE attempts via the WordPress REST API batch endpoint (CVE-2026-63030) affecting WordPress 6.9+.

› Both rules are automatically active for free-plan customers via the Free Ruleset and deployed with the default Block action for Pro, Business, and Enterprise Managed Ruleset customers.

▾ 1 more release in this issue · 2026-07-17

WAF - WAF Release - 2026-07-17 - Emergency (2026-07-17) | BLOG ↗ | STABLE

Emergency WAF release adds new managed rules blocking unauthenticated RCE and SQLi exploits in popular web frameworks.

› New detection rule (7dfb2bd4708d4b88b9911dc0550664b6) added to Cloudflare Managed Ruleset to block unauthenticated RCE via malicious path sequences at Block action.

› New detection rule (1c060d3a371549219ee290d7ed933fcc) added to Cloudflare Managed Ruleset to block SQLi via unsanitized request parameters at

Block action.

- New detection rule (ebd3f2df15c74ddcbf6220c9b5ec246a) added to Cloudflare Free Ruleset to block unauthenticated RCE at Block action.
- New detection rule (db003b39b7774859a8d588ce33697a1a) added to Cloudflare Free Ruleset to block SQLi at Block action.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Container & Image Security

Sigstore cosign ⓘ v3.1.2 NOTES ↗

Code signing and transparency for containers and binaries

cosign v3.1.2 adds a `bundle inspect` command and a `--base-config` flag for signing-config overrides.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.1.2 https://github.com/sigstore/cosign.git
# already have the repo? check out this version:
$ git checkout v3.1.2
```

└──▷ TRY IT

Inspect a local cosign bundle file to examine its contents and structure.

```
$ cosign bundle inspect <bundle-file>
```

Override specific services from a base signing config when configuring cosign for a custom environment.

```
$ cosign signing-config --base-config <base-config-file>
```

› Adds `cosign bundle inspect` command to examine cosign bundle files.

› Adds `--base-config` flag to the signing-config command to override services from a base config.

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Endpoint Security (EDR/XDR)

Microsoft Defender ⓘ changelog · Jul 17, 2026 BLOG ↗

Microsoft Defender is a built-in security tool that protects Windows devices from malware, viruses, ransomware, and other cyber threats in real time.

- New Microsoft Defender Experts Threat Intelligence service delivers continuous, curated, expert-led intelligence tailored to an organization.
- Microsoft Defender Experts MDR extended with third-party and multicloud coverage.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Immuta ⓘ snapshot-20260718 NOTES ↗

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

└──▷ HOW TO FIND IT

Access in-product documentation and ask questions in natural language without leaving Immuta.

> 📍 In the Immuta UI, click the Documentation icon in the top menu, then type your question into the AI-powered documentation chatbot.

Switch the Alation integration to OAuth 2.0 to eliminate recurring 90-day API key rotation.

> 📍 In the Immuta console, go to App Settings › Integrations › Alation, select OAuth 2.0 as the authentication method, and enter your OAuth credentials.

- Supports automated sensitive data discovery (regex, dictionary, and column-name regex identifiers) for the Amazon Redshift viewless integration.
- Embeds Immuta documentation and an AI-powered chatbot directly in the Immuta UI for in-product search.

└──▷ BREAKING ON UPGRADE

! When a data source is disabled, Immuta will revoke data owners' Immuta-managed grants in the data platform (opt-in 6/30–7/31, on by default 8/1, enforced for all 9/1) — workflows that depend on owners retaining access after disabling a source will break.

WAS THIS USEFUL? 👍 👎

◆ Secure Access (VPN / SASE / ZTNA)

Zscaler ⓘ changelog · Jul 17, 2026 BLOG ↗

Zscaler is a cloud-based security platform that protects users, devices, and applications by routing traffic through its global network to enforce security policies.

Zscaler launches Zero Trust Enterprise Browser and Browser Detection and Response (BDR) extension to bring Zero Trust into the browser itself.

- Introduces Zero Trust Enterprise Browser: a purpose-built, fully managed Chromium browser with Zero Trust security integrated directly into the browsing experience.
- Adds Browser Detection and Response (BDR) via the Zero Trust Browser Extension, delivering browser-native threat detection and response inside existing browsers for threats that evade upstream network and endpoint controls.
- Delivers three Zero Trust Browser deployment form factors — Cloud Browser Isolation, Browser Extension, and Enterprise Browser — allowing organizations to mix and match based on user, device, and workflow requirements.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Compliance Automation & Trust Management

Vanta ⓘ July 2026 NOTES ↗

Vanta is a trust management platform that automates security compliance, monitors risks, and helps companies achieve certifications like SOC 2, ISO 27001, and HIPAA.

Vanta July 2026 adds internal auditor roles, control assessments, Dropbox deprovisioning, MCP support, and 8 more capabilities.

└──▷ **TRY IT**

Use the Vanta Agent's `/triage-issues` shortcut to immediately surface the highest-priority compliance gaps without manually scanning every control.

```
$ In the Vanta Agent chat, type: /triage-issues
```

- Adds version-specific policy selection when attaching policies as evidence to information requests, giving auditors precise control over which policy version is submitted.
- Enables auditors to formally assess controls directly in Vanta for internal and external IRL audits, with framework-specific statuses (Conforming/Non-conformity for ISO; Satisfied/Not Satisfied for FedRAMP) and required justification comments.
- Introduces an Internal Auditor Role, allowing an internal employee to be designated as auditor on an internal audit engagement with a scoped IRL, request tracker, and population views without requiring an external audit firm.
- Adds Custom Commitment Types for Customer Commitments, letting users define their own contract commitment types with a name, definition, and example clauses — Vanta automatically extracts them from future and previously uploaded contracts.
- Enables the Vanta Agent to assign control owners inline during agentic Control Program Onboarding, with validation of missing descriptions, ineligible owners, and ambiguous entries before bulk-review commit.

+7 more

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Cursor ⓘ Improvements to Cursor in Slack NOTES ↗

Cursor is an AI-powered code editor that helps developers write, edit, and understand code faster using built-in AI assistance.

Cursor in Slack adds upfront planning, multi-repo environment support, and cross-channel read/write workflows.

└──▷ HOW TO FIND IT

Kick off a task that spans multiple repos so Cursor automatically targets the right multi-repo environment instead of a single default repo.

> 📍 In Slack, mention @Cursor with your request (e.g., '@Cursor update the API contract in the backend and reflect it in the frontend'). Cursor reads the request and selects the named multi-repo environment that covers both repos. If it needs an additional repo mid-task, click the 'Switch repository' button it surfaces, choose the repo or environment, and it resumes.

Leverage cross-channel context so Cursor can pull information from a separate channel and post results back without manual copy-paste.

> 📍 In Slack, mention @Cursor in your thread (e.g., '@Cursor summarize the decisions from #architecture-review and draft an implementation plan here'). Cursor reads the target channel and posts the output back in your original thread.

- Cursor now shares a plan before starting a task in Slack, letting you redirect it before work begins.
- Live status updates show each step as Cursor works through a task in Slack.
- Supports named multi-repo environments from Slack, letting Cursor access frontend, backend, and shared code across separate repos in a single session.
- Mid-task repo switching via a 'Switch repository' button lets Cursor resume work after changing to a different repo or environment.
- Cross-channel workflows enable Cursor to read context from other Slack channels and threads and post updates back to the relevant channel or thread.

└──▷ BREAKING ON UPGRADE

! In-message buttons are removed and replaced by compact footer links, which will change existing interaction patterns for users who relied on those buttons.

WAS THIS USEFUL? 👍 👎

Claude Code ⓘ v2.1.214 NOTES ↗ CODE ↗

Claude Code is an agentic coding tool that lives in your terminal, understands your codebase, and helps you code faster by executing routine tasks, explaining complex code, and handling git workflows - all through natural language commands.

Claude Code v2.1.214 adds EndConversation tool, OTel message correlation attributes, docker daemon-redirect permission prompts, and more.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v2.1.214 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.214
```

└──▷ **TRY IT**

Cap OTel content attribute size below the default 60 KB to stay within your observability backend's event-size limits.

```
$ CLAUDE_CODE_OTEL_CONTENT_MAX_LENGTH=16384 claude
```

- › Adds the EndConversation tool so Claude can autonomously terminate sessions involving highly abusive users or jailbreak attempts.
- › Adds a periodic progress heartbeat for long-running tool calls that previously went silent.
- › Adds an ISO `modified` timestamp to memory file frontmatter.
- › Adds `message.uuid`, `client_request_id`, and `tool_source` attributes to OpenTelemetry log events for message-level correlation and tool provenance.
- › Adds `CLAUDE_CODE_OTEL_CONTENT_MAX_LENGTH` environment variable to configure the 60 KB truncation limit on OpenTelemetry content attributes.

+2 more

└──▷ **BREAKING ON UPGRADE**

- ! Single-segment `dir/**` hook `if:` conditions now match only `<cwd>/dir`; write `**/dir/**` to match at any depth.
- ! `file` commands using `-m` / `--magic-file` or `-f` / `--files-from` now require permission instead of being auto-allowed as read-only.

! SessionStart hooks now report source `"fork"` when a session begins as a fork instead of `"resume"`.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI ⓘ `rust-v0.145.0-alpha.22` NOTES ↗ CODE ↗

PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI alpha.22 stabilizes the memories feature, adds managed Bedrock login, and imports plugins from known marketplaces.

└──▷ **GET THIS VERSION**

```
$ git clone --branch rust-v0.145.0-alpha.22
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.22
```

- Stabilizes the memories feature flag, making it generally available in the CLI.
- Adds a managed Bedrock login API for authenticating with AWS Bedrock as a provider.
- Imports enabled plugins automatically from known marketplaces.
- Makes GPT-5.6 Sol the default model for Bedrock sessions.
- Adds dedicated storage for paginated thread history, enabling larger conversation histories.

+7 more

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

▷ Subscribe