THE CYBER TOOLCHAIN

NO. 009

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 9, 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 17 tools

OFFENSIVE SECURITY

◆ Post-Exploitation

SpecterOps BloodHound ⓘ v9.4.0 NOTES ↗

Six Degrees of Domain Admin

BloodHound v9.4.0 adds ADCS ESC14 detection, new graph/relationship APIs, webhook creation, auditor role expansions, and ingest metrics.

└──▷ GET THIS VERSION

```
$ git clone --branch v9.4.0
https://github.com/SpecterOps/BloodHound.git
# already have the repo? check out this version:
$ git checkout v9.4.0
```

- Adds ADCS ESC14 Scenario A attack path coverage for certificate-based privilege escalation detection.
- New GET `api/v2/relationships/{relationship_id}` endpoint for retrieving individual graph relationship details.
- New GET node endpoint (`api/v2/...`) for retrieving individual graph node details.
- Introduces webhook creation capability for event-driven alerting integrations.
- Adds `data_quality_stats` table and supporting APIs for tracking data quality statistics.

+12 more

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ Web Application Security

humble ⓘ 1.63 NOTES ↗

A humble and fast security-oriented HTTP headers analyzer.

humble 1.63 adds HAR file analysis via '`-if`' and a new insecure header check.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.63 https://github.com/rfc-st/humble.git
# already have the repo? check out this version:
$ git checkout 1.63
```



└──▷ TRY IT

Analyze a captured HAR file from a browser session to audit its HTTP response headers without making live requests.

```
$ python humble.py -if <path/to/capture.har>
```

- Supports analyzing HAR (HTTP Archive) files as input via the '`-if`' option, enabling offline inspection of captured browser traffic.
- Adds 1 new insecure header check, expanding detection coverage.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ Runtime Security

Wiz Defend ⓘ changelog · Jul 8, 2026 BLOG ↗

Wiz Defend is a cloud detection and response platform that identifies and helps remediate real-time threats across cloud environments.

Wiz Defend adds AI coding assistant context-aware detections for symlink and SSH key attacks at clone and write time.

- Detects malicious symlinks pointing to sensitive files at git clone time, catching GhostApproval-style attacks before AI coding assistant involvement.
- Monitors writes to sensitive files (e.g., SSH authorized_keys) at write time, triggering alerts regardless of whether the write originates via a symlink.
- Adds AI agent context awareness: alerts auto-elevate severity when malicious activity originates from an AI coding assistant process operating on untrusted repositories.
- Leverages Wiz Runtime Sensor telemetry to identify active endpoints and their exposure based on inbound network traffic.

WAS THIS USEFUL? 👍 👎

◆ CNAPP

Wiz Cloud ⓘ changelog · Jul 8, 2026 BLOG ↗

Wiz Cloud is a cloud security platform that scans and visualizes risks across cloud environments, helping teams identify vulnerabilities, misconfigurations, and threats.

Wiz ASM adds automated subdomain recon, API endpoint discovery, runtime sensor telemetry, and custom target input to map attack surface.

└──▷ HOW TO FIND IT

Extend ASM coverage to legacy or air-gapped systems not reached by automated discovery by manually specifying their addresses.

> 📍 In the Wiz console, go to ASM › Custom Targets and add the IP addresses, domains, or IP ranges of legacy or isolated systems you want analyzed.

- Adds automated subdomain reconnaissance seeded from owned root domains to surface shadow exposures.
- Discovers REST API endpoints via cloud-native API gateways, Wiz Runtime Sensor telemetry, and Attack Surface Scanner-detected API specifications.
- Leverages Wiz Runtime Sensor inbound network traffic telemetry to identify and assess active endpoint exposure.
- Supports manual input of IP addresses, domains, and IP ranges as Custom Targets to cover legacy or isolated systems outside automated discovery.
- Maps cloud network configurations to calculate effective exposure for every cloud object.

WAS THIS USEFUL? 👍 👎

Zscaler Posture Control

changelog · Jun 26, 2026

BLOG ↗

Zscaler Posture Control is a cloud-native application protection platform (CNAPP) that identifies and remediates security risks, misconfigurations, and compliance violations across multi-cloud environments.

Zscaler Posture Control adds AI-SPM with shadow AI discovery, drift detection, and automated red teaming mapped to MITRE ATLAS, NIST AI RMF, and OWASP LLM Top 10.

- New AI Security Posture Management (AI-SPM) provides continuous discovery and posture assessment of AI assets including shadow AI, unsanctioned models, embedded SaaS AI, developer-built agents, and MCP servers.
- New AI Bill of Materials (AI-BOM) maps model lineage across discovered AI assets for full supply-chain visibility.
- Adds posture assessment for AI-specific risks: misconfigurations, excessive permissions, vulnerable RAG frameworks, and exposed data pipelines.
- New AI posture drift detection continuously compares permission scope, configuration state, and data access patterns against established baselines to surface slow-erosion changes point-in-time assessments miss.
- Automated red teaming runs continuous, high-scale tests across the AI lifecycle and maps findings to MITRE ATLAS, NIST AI RMF, OWASP LLM Top 10, and the EU AI Act.

+1 more

WAS THIS USEFUL? 👍 👎

Palo Alto Prisma Cloud ⓘ changelog · Jul 7, 2026 BLOG ↗

Palo Alto Prisma Cloud is a cloud-native security platform that protects applications, data, and infrastructure across multi-cloud environments throughout the development lifecycle.

Prisma AIRS 3.0 adds Japan local cloud region, AI Agent SSPM, AI Red Teaming, and expanded model/runtime scanning.

- New local cloud region in Japan for Prisma AIRS, enabling data residency and low-latency processing for Japanese enterprises.
- New AI Agent SSPM (SaaS Security Posture Management) detects and governs agentic threats including identity impersonation, memory manipulation, and tool misuse across autonomous AI assistants.
- New AI Red Teaming capability performs automated, adaptive penetration tests against AI deployments using an agent that learns and evolves like a real attacker.
- New AI Model Security scanning checks third-party models for vulnerabilities, malicious scripts, tampering, and deserialization attacks before deployment.
- New AI Runtime Security protection for LLM-powered apps covers prompt injection, sensitive data leaks, and Indirect Prompt Injection (IDPI) embedded in websites.

+1 more

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Detection Engineering

SigmaHQ sigma-cli ⓘ v3.1.0 NOTES ↗

The Sigma command line interface based on pySigma

sigma-cli v3.1.0 adds JUnit XML export and per-rule file output with directory structure preservation.

└──▷ GET THIS VERSION

```
$ git clone --branch v3.1.0 https://github.com/SigmaHQ/sigma-cli.git
# already have the repo? check out this version:
$ git checkout v3.1.0
```

└──▷ TRY IT

Output each converted rule as its own file mirroring the source directory layout — useful for version-controlled rule repos or selective SIEM deployment.

```
$ sigma convert -t esql -p ecs_windows --output-dir translated/ --
output-filename-template "{path}/{stem}.esql" rules/
```

Flatten all converted rules into a single directory with per-rule files, handling rules that emit multiple queries by appending an index suffix.

```
$ sigma convert -t esql -p ecs_windows --output-dir translated/ --
output-filename-template "{stem}-{index}.esql" rules/
```

- Adds JUnit XML export for sigma convert results, enabling CI pipeline integration with standard test reporting tools.
- Adds `--output-dir` and `--output-filename-template` flags to write each converted rule to a separate file, with optional preservation of the source directory hierarchy.

└──▷ BREAKING ON UPGRADE

! Python 3.9 is no longer supported

WAS THIS USEFUL? 👍 👎

◆ Vulnerability Management

Qualys Patch Management ⓘ changelog · Jul 8, 2026 BLOG ↗

Qualys Patch Management is a cloud-based tool that automates the discovery, prioritization, and deployment of patches across IT assets to reduce vulnerabilities.

Qualys Patch Management adds autonomous zero-day remediation via TruRisk Eliminate and Windows 11 upgrade orchestration through QQL and CAR scripts.

.
.
.

└──▷ HOW TO FIND IT

Identify and deploy Windows 11 enablement packages at scale to upgrade endpoints already running recent Windows 11 builds with minimal operational impact.

> 📍 In TruRisk Eliminate Patch Catalog, run QQL: patch.title:"Enablement" — select the matching patch, create a patch job, and target the relevant asset groups.

Find ISO-based Windows feature update patches for environments where enablement packages are unavailable (e.g., 23H2 → 24H2 transition).

> 📍 In TruRisk Eliminate Patch Catalog, run QQL: `patch.isFeatureUpdate` :true AND patch.vendor: Microsoft — enable the required ISO package, then create and deploy a patch job to the target asset groups.

Directly upgrade Windows 10 endpoints to the latest Windows 11 build without intermediate steps using a pre-built CAR library script.

> 📍 In TruRisk Eliminate, navigate to Script Library → System Updates and Upgrades → Windows OS Upgrades, select the Windows 10 to Windows 11 upgrade script, import it in approved mode, and execute it on the required asset groups.

- New autonomous remediation workflows via Qualys TruRisk Eliminate enable machine-speed patching, mitigation, container image replacement, workload isolation, and firewall policy updates triggered directly from confirmed vulnerability findings.
- Supports patchless/compensating mitigation actions when immediate patching is not feasible, reducing exposure window while permanent fixes are prepared.
- Enables enterprise-wide Windows 11 upgrade orchestration through TruRisk Eliminate, covering three paths: enablement packages, ISO-based feature updates, and direct Windows 10-to-Windows 11 upgrades.
- New QQL filters for patch targeting: `patch.title:"Enablement"` for enablement packages and `patch.isFeatureUpdate:true AND patch.vendor:\` Microsoft\ for

ISO-based feature updates.

- Windows 10-to-Windows 11 upgrade scripts now available in the Custom Assessment and Remediation (CAR) library under Script Library → System Updates and Upgrades → Windows OS Upgrades.

WAS THIS USEFUL? 👍 👎

◆ Attack Surface & Exposure Management

Qualys CyberSecurity Asset Management ⓘ

changelog · Jul 7, 2026 BLOG ↗

Qualys CyberSecurity Asset Management provides continuous visibility into all IT assets, identifying security gaps, risks, and unauthorized devices across hybrid environments.

Qualys CSAM integrates with Cisco Cloud Control, exposing asset inventory, TruRisk scores, and remediation workflows to Cisco AI agents.

└──▷ HOW TO FIND IT

Use Cisco AI Canvas to query Qualys asset and risk context for a specific host directly from the AI agent workflow.

> 📍 In Cisco AI Canvas, navigate to the AI Agents panel, select the Qualys integration tools (Unified Asset Inventory, TruRisk Prioritized Findings), and invoke a tool call for a target host to retrieve vulnerability detections, TruRisk score, exploit context, and available remediation workflows.

- New Cisco Cloud Control integration connects Qualys Unified Asset Inventory, TruRisk Prioritized Findings, and remediation workflows into Cisco's AI Canvas environment.
- Cisco AI agents can now query Qualys at runtime — pulling asset inventory, vulnerability detections per host/workload, TruRisk scores, exploit context, and triggering remediation or mitigation workflows as typed tool calls.
- Enables cross-domain asset and risk visibility spanning cloud, network, identity, datacenter, and security domains without manual tool-switching.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Data Security & DLP

Zscaler Data Protection ⓘ changelog · Jul 8, 2026 BLOG ↗

Zscaler Data Protection secures sensitive data across cloud applications, endpoints, and networks by enforcing DLP policies and preventing unauthorized data exfiltration.

Zscaler Workflow Automation adds end-user DLP triage, self-service email quarantine release, and automated exception management.

└──▷ HOW TO FIND IT

Allow end users to self-release quarantined emails after justification, eliminating helpdesk tickets for routine outbound email DLP hits.

> 📍 In the Zscaler console, go to Advanced Account Settings and enable 'Enable Email Quarantine Release for End Users'.

Manually release a quarantined email to specific recipients from the incident details view when admin review is required.

> 📍 In the Zscaler console, open the incident in the advanced incident details interface, then select the 'Release Email Quarantine' action and choose specific recipients to deliver the message to.

- Adds end-user justification workflows that engage users directly via Slack, Microsoft Teams, or email when an inline DLP policy triggers, routing responses to auto-remediation, manager escalation, or analyst enrichment.
- Enables manager-approved, time-bounded DLP policy exceptions created automatically through workflow responses — no manual policy edits required.
- Introduces 'Release Email Quarantine' action in the advanced incident details interface, letting admins deliver quarantined messages to all or specific recipients directly from the platform.
- New 'Enable Email Quarantine Release for End Users' setting in Advanced Account Settings allows end users to release their own quarantined emails after providing a business justification or obtaining manager approval.
- Supports native integration with ServiceNow and Jira to auto-create enriched incident tickets with user justification and forensic context attached.

+2 more

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Claude Code ⓘ v2.1.205 2 RELEASES · 2026-07-07 → 2026-07-08 NOTES ↗

CODE ↗ 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 v2.1.205 upgrades `/doctor` into a full setup checkup, hardens auto mode safety rules, and improves the agents view.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.1.205 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.205
```

└──▷ TRY IT

Run a full environment health check and let Claude Code automatically fix any issues it finds.

```
$ /doctor
```

Review the enriched agent list — with colored state labels, classifier headlines, and PR links — for all background sessions.

```
$ claude agents
```

- Adds auto mode rule that blocks tampering with session transcript files.
- Improves auto mode to prompt for confirmation before running `rm -rf` on an unresolvable variable, reducing accidental data loss.
- Upgrades `/doctor` into a full setup checkup that can both diagnose and fix issues; adds `/checkup` as an alias.
- Enhances `claude agents` view: PR-linked sessions (edit, merge, comment, push) now surface the PR link directly in the list.
- Enhances `claude agents` view: rows show a colored state word and classifier-written headline instead of raw tool call text; peek view shows full status and the exact ask for blocked sessions.

+3 more

└──▷ BREAKING ON UPGRADE

! User-configured MCP servers can no longer register under the names "Claude Browser" or "Claude Preview"

▾ **1 more release in this issue** · 2026-07-07 → 2026-07-08

v2.1.203 NOTES ↗ CODE ↗ STABLE

Claude Code v2.1.203 adds login-expiry warnings, manual-mode visibility, MCP roots updates, and a VSCode remote-control toggle.

└──▷ GET THIS VERSION

```
$ git clone --branch v2.1.203 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.203
```

› Adds a pre-expiry warning when your login session is about to expire, letting you re-authenticate before background sessions are interrupted.

› Adds a grey ⏸ badge in the footer when manual permission mode is active, so the current permission mode is always visible.

› Exposes the session's additional working directories in MCP `roots/list`, and sends `notifications/roots/list_changed` whenever that set changes — enabling MCP servers to stay in sync with multi-root workspaces.

› Changes the Escape key (instead of left arrow) to close background tasks, diff, and workflow detail views — practitioners must update muscle memory.

› Changes the empty `claude agents` view to always display organized sections (Needs Input / Working / Completed) with descriptions, even when no agents are running.

+2 more

└──▷ BREAKING ON UPGRADE

! The left arrow key no longer closes the background tasks, diff, and workflow detail views — press Esc instead

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI ⓘ rust-v0.144.0-alpha.2

4 RELEASES · 2026-07-07 → 2026-07-08 NOTES ↗ CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Adds Ultra concurrency warning, proxy-aware WebSocket connector, MCP tool snapshot reuse, and Windows sandbox runtime access.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.144.0-alpha.2
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.144.0-alpha.2
```

- Adds a TUI warning when Ultra reasoning is selected with high multi-agent concurrency (≥8 threads), displaying configured thread and subagent limits so users understand cost implications.
- Adds a proxy-aware WebSocket connector that respects outbound proxy policy and custom-CA trust, matching the same effective HTTP proxy configuration used for REST calls.
- Reuses a single MCP tool snapshot within each sampling request, reducing latency and ensuring consistent tool state between Apps World State and tool-router construction.
- Grants Windows sandbox sessions read/execute access to the managed primary runtime cache (`%USERPROFILE%\.cache\codex-runtimes`), enabling bundled Python, Node, and native tools in elevated sandbox contexts.

▾ 3 more releases in this issue · 2026-07-07 → 2026-07-08

rust-v0.144.0-alpha.1 NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI v0.144.0-alpha.1 adds pnpm install detection, faster branch picker, and hook execution tracing.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.144.0-alpha.1
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.144.0-alpha.1
```

└──▷ TRY IT

Verify that a pnpm-managed Codex install is recognized correctly in your environment.

```
$ codex doctor
```

- Detects Codex installs managed by pnpm, surfacing pnpm status in `codex doctor` output alongside npm and bun.
- Speeds up the `/review` branch picker by switching branch enumeration to `git for-each-ref` scoped to `refs/heads`, preventing timeouts in large repos.
- Adds trace spans around configured hook command execution, making slow installed hooks visible directly in traces.
- Supports externally provided Codex auth snapshots with explicit runtime capabilities, installable via the existing ExternalAuth provider path.
- Switches code-mode to hosted mode by default.

rust-v0.143.0 NOTES ↗ CODE ↗ LATEST STABLE

Codex CLI v0.143.0 adds remote plugins by default, system proxy support, daemon pairing codes, new Bedrock models, and richer app-server APIs.

└──▷ GET THIS VERSION

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

└──▷ TRY IT

Generate a manual pairing code from a running Codex daemon to connect a remote-control client without automatic discovery.

```
$ codex remote-control pair
```

- Enables remote plugins by default, with richer catalog rows, npm marketplace sources, and visible remote/local version info.
- Supports routing authentication and Responses API traffic through macOS and Windows system proxies, including PAC and WPAD configurations.
- Adds `codex remote-control pair` subcommand for generating manual pairing codes from a running daemon.

› Adds Amazon Bedrock GPT-5.6 Sol, Terra, and Luna models with first-class support for `max` reasoning effort.

› Enables MCP tool search by default; ChatGPT-hosted MCP servers can now explicitly use session authentication.

+1 more

`rust-v0.143.0-alpha.39` NOTES ↗ CODE ↗ PRE-RELEASE

Adds 'writes' approval mode, serialized MCP OAuth credential stores, canonical turn item emission, and improved Windows URL safety detection.

└──▷ GET THIS VERSION

```
$ git clone --branch rust-v0.143.0-alpha.39
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.143.0-alpha.39
```

└──▷ TRY IT

Set the new 'writes' approval mode on an MCP tool so the agent auto-approves read operations but prompts only for writes.

```
$ codex mcp get
```

› Adds `writes` approval mode for MCP app tool approval, giving practitioners a middle ground between `auto` and `prompt` that gates only write operations

› Emits canonical `TurnItem` lifecycle events for command execution, dynamic tool calls, sub-agent activity, collab tool calls, and collab wait items, enabling more structured agent event streams

› Enables auth elicitation by default, so authentication prompts are surfaced automatically without manual configuration

› Supports hosted success redirects in the login flow

› Shows reset credit details (title, expiry, idempotency key) in the TUI redemption picker so users can choose exactly which credit to consume

+2 more

WAS THIS USEFUL? 👍 👎

OpenHands

cloud-1.43.0 · 4 RELEASES · 2026-07-07 → 2026-07-08 · NOTES ↗

LATEST · STABLE

OpenHands: AI-Driven Development

OpenHands cloud-1.43.0 adds SMTP email service, user login tracking, and a renamed Usage & Monitoring dashboard.

└──▷ GET THIS VERSION

```
$ git clone --branch cloud-1.43.0
https://github.com/OpenHands/OpenHands.git
# already have the repo? check out this version:
$ git checkout cloud-1.43.0
```

└──▷ HOW TO FIND IT

Check whether email delivery (SMTP or Resend) is enabled for your OpenHands Cloud instance — useful before configuring notification workflows.

> 📍 In the console, go to Usage & Monitoring and look for the email-enabled status indicator showing SMTP or Resend as active.

- Adds SMTP email service integration, enabling transactional email delivery alongside the existing Resend provider.
- New email-enabled status surface shows whether SMTP or Resend email is active in the platform configuration.
- Tracks user login timestamps, giving admins visibility into last-active data per user.
- Renames the Admin Dashboard to 'Usage & Monitoring' for clearer intent.

▾ 3 more releases in this issue · 2026-07-07 → 2026-07-08

1.10.0 · NOTES ↗ · STABLE

OpenHands 1.10.0 adds SMTP email service, user login tracking, and a renamed Usage & Monitoring dashboard.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.10.0
https://github.com/OpenHands/OpenHands.git
# already have the repo? check out this version:
$ git checkout 1.10.0
```

└──▷ HOW TO FIND IT

Check whether email delivery is active for your deployment — useful when onboarding users who need invite or notification emails.

> 📍 In the console, go to Settings › Usage & Monitoring and check the Email Enabled indicator for your configured SMTP or Resend provider.

- Adds SMTP email service integration, enabling email delivery alongside the existing Resend provider.
- Surfaces email-enabled status in the UI for both SMTP and Resend configurations.
- Tracks user login timestamps, giving admins visibility into account activity.
- Renames the admin dashboard to 'Usage & Monitoring' for clearer intent.

cloud-1.42.0 NOTES ↗ STABLE

OpenHands cloud-1.42.0 adds a super-admin management API for granting, revoking, and listing admin privileges.

└──▷ GET THIS VERSION

```
$ git clone --branch cloud-1.42.0
https://github.com/OpenHands/OpenHands.git
# already have the repo? check out this version:
$ git checkout cloud-1.42.0
```

- New enterprise super-admin management endpoint enables granting, revoking, and listing super-admin roles via API.

1.9.1 NOTES ↗ STABLE

OpenHands 1.9.1 adds a super-admin management API for granting, revoking, and listing admin privileges.

└──▷ GET THIS VERSION

```
$ git clone --branch 1.9.1
https://github.com/OpenHands/OpenHands.git
# already have the repo? check out this version:
$ git checkout 1.9.1
```

› New enterprise/auth super-admin endpoint enables granting, revoking, and listing super-admin roles via API.

WAS THIS USEFUL? 👍 👎

Crush ⓘ v0.83.0 2 RELEASES · 2026-07-07 → 2026-07-08 NOTES ↗ CODE ↗

LATEST STABLE

Glamourous agentic coding for all

Crush v0.83.0 adds GPT-5.6 support, routing it through the Responses API.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.83.0
https://github.com/charmbracelet/crush.git
# already have the repo? check out this version:
$ git checkout v0.83.0
```

› Adds GPT-5.6 model support, routed through the OpenAI Responses API.

▾ 1 more release in this issue · 2026-07-07 → 2026-07-08

v0.82.0 NOTES ↗ CODE ↗ STABLE

Crush v0.82.0 adds preserved newlines in expanded tool output and Alibaba US region support.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.82.0
https://github.com/charmbracelet/crush.git
# already have the repo? check out this version:
$ git checkout v0.82.0
```

└──▷ TRY IT

Use Alibaba's US-region endpoint instead of Singapore for lower-latency access from North America.

```
$ export ALIBABA_US_API_KEY=<your-key>
crush
```

› Preserves newlines in expanded tool content in the chat UI, improving readability of multi-line tool outputs.

› Adds Alibaba (United States) as a supported provider region, configurable via `ALIBABA_US_API_KEY`.

WAS THIS USEFUL? 👍 👎

gemini-cli ⓘ v0.50.0 NOTES ↗

An open-source AI agent that brings the power of Gemini directly into your terminal.

Gemini CLI v0.50.0 adds tool registry discovery for MCP integrations.

└──▷ **GET THIS VERSION**

```
$ git clone --branch v0.50.0 https://github.com/google-gemini/gemini-cli.git
# already have the repo? check out this version:
$ git checkout v0.50.0
```

› Adds tool registry discovery, enabling automatic detection and registration of available tools/MCP integrations.

WAS THIS USEFUL? 👍 👎

AGENT

◆ AI Agent Frameworks

CrewAI ⓘ 1.15.2 NOTES ↗

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

CrewAI 1.15.2 adds inline skill definitions, templated Flow action inputs, stream frame protocol, and AgentExecutor feedback handling.

└──▷ GET THIS VERSION

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

- Supports inline skill definitions, eliminating the need for separate skill files in crew/flow configurations.
- Adds templated Flow action inputs, enabling dynamic parameterization of Flow actions at runtime.
- Defines a stream frame protocol for Flows, enabling structured streaming output from Flow executions.
- Adds a generated Flow Definition authoring skill to guide AI coding agents when scaffolding Flows.
- Implements message setup and feedback handling in AgentExecutor for richer agent interaction loops.

+4 more

WAS THIS USEFUL? 👍 👎

DEPLOY

◆ Local LLM Runtimes

Ollama ⓘ v0.31.2 NOTES ↗ CODE ↗

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

Ollama v0.31.2 adds flash attention on Pascal GPUs, iGPU vision-model offload, and privacy-hardened Claude Code launch.

└──▷ GET THIS VERSION

```
$ git clone --branch v0.31.2 https://github.com/ollama/ollama.git
# already have the repo? check out this version:
$ git checkout v0.31.2
```

└──▷ TRY IT

Launch Claude Code backed by a local model with telemetry and feedback reporting disabled by default (no extra flags needed — privacy-off is now the default).

```
$ ollama launch claude
```

- Enables flash attention on older NVIDIA Pascal GPUs (compute capability 6.x), unlocking faster inference on GTX 10-series and similar cards.
- Enables iGPU to offload vision models with padding to fit available memory, broadening multimodal support on integrated graphics.
- Disables Claude Code telemetry, error reporting, feedback commands, and non-essential traffic by default when using `ollama launch claude`.
- Updates the MLX and llama.cpp engines to their latest upstream versions.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 9, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe