THE CYBER TOOLCHAIN

NO. 022

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED JULY 22, 2026 · EVERY DAY

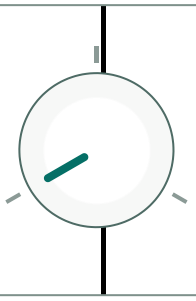
The daily firehose — everything the toolchain shipped today, already filtered.

▸ // HOW THIS ISSUE IS MADE SHOW

VIEW

ISSUE VIEW

scroll



SCROLL | EXPAND IN PLACE | PREVIEW PANE

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 30 tools

OFFENSIVE SECURITY

vim ? help

◆ AI-Powered Offensive Security

Strix

Feedback ⓘ v1.2.0 NOTES ↗

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

Strix v1.2.0 adds recon asset-discovery skill, AD/Kerberos testing, Docker resource limits, shared-network port resolution, and a local run viewer with email reports.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v1.2.0 https://github.com/usestrix/strix.git
# already have the repo? check out this version:
$ git checkout v1.2.0
```

Adds a recon asset-discovery skill for automated attack surface enumeration.

Adds an Active Directory / Kerberos domain testing skill for enterprise network assessments.

Adds a Grafana/Prometheus observability pivot skill to identify exposed monitoring infrastructure.

Enables opt-in resource limits (CPU/memory) for Docker sandbox containers to cap agent resource usage.

Introduces a local run viewer with email report delivery, run history, and access to the full platform suite.

WAS THIS USEFUL? 👍 👎

APPLICATION SECURITY

◆ SAST / SCA / Code Security

Promptfoo

Feedback ⓘ snapshot-20260722 NOTES ↗

Promptfoo is an open-source tool for testing, evaluating, and red-teaming LLM applications to detect vulnerabilities, improve reliability, and compare model outputs.

Promptfoo snapshot-20260722 adds adaptive rate limiting, Transformers.js local inference, telecom red team plugins, and video generation providers.

└──▷ TRY IT

View Promptfoo service logs directly from the CLI without navigating the filesystem.

```
$ promptfoo logs
```

Load multiple environment files (e.g., shared secrets plus environment-specific overrides) in a single eval run.

```
$ promptfoo eval --env-file .env.shared --env-file .env.staging
```

Add a description to a red team generation run for traceability across scans.

```
$ promptfoo redteam generate -d "Sprint 42 telecom compliance scan"
```

Adds Transformers.js provider to run Hugging Face models locally in Node.js or the browser.

Adds Vercel AI Gateway and Cloudflare AI Gateway providers for request routing.

Adds AWS Bedrock Video (Nova Reel and Luma Ray 2) and Azure AI Foundry Video (Sora) providers for video generation.

Adds native session endpoint support to the HTTP provider for stateful conversations.

Adds `apiBaseUrl`, `websocketUrl` override, and function call support to the xAI Voice provider.

+20 more

WAS THIS USEFUL? 👍 👎

StepSecurity

Feedback ⓘ

`Workflow Run Policies: allow-list enforcement, runner constraints, and richer policy ins`

2 RELEASES · 2026-07-17 → 2026-07-21 | BLOG ↗ | LATEST | STABLE

StepSecurity is a platform that secures CI/CD pipelines by detecting and preventing supply chain attacks, hardening GitHub Actions workflows, and monitoring runtime behavior.

StepSecurity Workflow Run Policies gain allow-list enforcement, container blocking, bulk-secret detection, and richer per-run evaluation visibility.

└──▷ **HOW TO FIND IT**

Block bulk secret harvesting (e.g. toJSON(secrets)) across all branches while allowing normal targeted secret access — configure in the Secret Exfiltration Policy.

> 📍 In the StepSecurity console, go to Workflow Run Policies › Secret Exfiltration Policy, enable 'Block only when a workflow accesses all secrets at once' (recommended), then enable `Analyze default branch runs` to extend coverage to the default branch.

Restrict CI jobs to an approved set of runner labels and RunsOn structured constraints, preventing use of arbitrary or untrusted runner configurations.

> 📍 In the StepSecurity console, go to Workflow Run Policies › Runner Label Policy and switch to `Allowed labels & constraints` mode. Add the permitted labels and structured runs-on keys (e.g. family, cpu, image) to the allow list; each configured dimension is enforced independently.

Ensure every targeted job sources its runtime policy from the StepSecurity policy store, preventing jobs from bypassing centrally managed policies.

> 📍 In the StepSecurity console, go to Workflow Run Policies › Harden-Runner Policy and enable `Require policy store`. Targeted jobs must then include `use-policy-store: true` on their Harden-Runner step or the run will be blocked.

Adds `Currently known compromised actions` display in the Compromised Actions Policy editor, showing affected versions, with the list snapshotted per evaluation for auditability.

New 'Block only when a workflow accesses all secrets at once' mode in Secret Exfiltration Policy blocks bulk patterns like `${{ toJSON(secrets) }}` while allowing targeted references, reducing false positives.

New `Analyze default branch runs` option in Secret Exfiltration Policy extends coverage to the default branch.

Secret Exfiltration violation details now include the specific secret reference and its line number in the workflow file.

New Allowed labels & constraints allow-list mode in Runner Label Policy enforces that only configured labels and constraint values may be used, with each dimension enforced independently.

+7 more

▾ **1 more release in this issue** · 2026-07-17 → 2026-07-21

Agent Skills Inventory for Developer Machines (2026-07-17) BLOG ↗ STABLE

Dev Machine Guard adds fleet-wide AI agent skills inventory with executable content flagging and provenance tracking.

└──▷ **HOW TO FIND IT**

Review which AI agent skills on your fleet ship executable content or lifecycle hooks before they run on a developer's machine.

> 📍 In the Dev Machine Guard console, go to Agent Skills, then filter by Flags → hooks or shell to see every skill that registers lifecycle hooks or injects shell commands, and drill into per-device installation details.

Spot developer machines running locally edited or outdated skill copies that differ from the rest of the fleet.

> 📍 In the Dev Machine Guard console, go to Agent Skills and look for skills highlighted under version drift detection (SKILL.md content hash mismatch), then use Export CSV to capture the full list for remediation tracking.

New Agent Skills page in the Dev Machine Guard dashboard inventories every AI agent skill (Claude Code, Codex, GitHub Copilot) detected across your fleet, with per-device installation details and scope (Global, Project, or System).

Flags skills that contain executable script files, declare lifecycle hooks, or inject shell commands that run with developer privileges — surfacing supply chain risk from AI coding agent skill folders.

Traces provenance of skills installed via the skills.sh CLI back to their source GitHub repository, plugin, and upstream folder hash; standalone skills are marked as authored on device.

Detects version drift by comparing SKILL.md content hashes across devices and highlights skills that differ between machines.

Supports filtering by source (skills.sh managed or local), agent, scope, and executable-content flags, plus full inventory export as CSV.

WAS THIS USEFUL? 👍 👎

Veracode

Feedback ⓘ Static Analysis updates (2026-07-21) BLOG ↗

Veracode is a cloud-based application security platform that helps developers identify and remediate vulnerabilities in their code through static, dynamic, and software composition analysis.

Veracode Static Analysis July 2026 engine adds Android 17, .NET 10 enhancements, Python httplib SSL detection, and pnpm module support.

Adds support for Android 17 in the static analysis engine.

Enhances support for .NET 10, including improved analysis of applications using Saturn.dll.

Adds support for Python httplib to detect CWE-295 and CWE-297 (SSL/TLS certificate validation) flaws.

Expands pnpm module support to improve JavaScript dependency resolution.

Improves detection of CWE-80 flaws in React Native (JavaScript) applications.

+3 more

WAS THIS USEFUL? 👍 👎

◆ Secrets Detection & Management

GitGuardian Detection Engine Feedback ⓘ

Detection Engine Updates Version 2.167 (2026-07-07) BLOG ↗

GitGuardian Detection Engine scans code repositories and developer environments to automatically detect secrets, credentials, and sensitive data exposures in real time.

GitGuardian v2.167 adds 11 new secret detectors spanning Azure, AI/LLM platforms, and developer tools, plus expanded validity checking.

New detector for LiteLLM API Key (without host capture).

New detector for Azure SQL Credentials (Assignment).

New detector for Azure Machine Learning keys.

New detector for Azure Video Indexer keys.

New detector for Apify Proxy Token.

+11 more

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Reverse Engineering

Capstone

Feedback ⓘ 6.0.0-Alpha10 NOTES ↗ PRE-RELEASE

Capstone disassembly/disassembler framework for ARM, ARM64 (ARMv8), Alpha, BPF, Ethereum VM, HPPA, LoongArch, M68K, M680X, Mips, MOS65XX, PPC, RISC-V(rv32G/rv64G), SH, Sparc, SystemZ, TMS320C64X, TriCore, Webassembly, XCore and X86.

Capstone 6.0.0-Alpha10 adds AArch64 explicit wide-immediate printing, RISC-V FP rounding mode details, and a RISC-V v5 compat macro.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 6.0.0-Alpha10 https://github.com/capstone-engine/capstone.git
# already have the repo? check out this version:
$ git checkout 6.0.0-Alpha10
```

└──▷ USE IT

Maintain a single codebase that builds against both Capstone v5 and v6 RISC-V headers without renaming `CS_MODE_RISCVC`.

c

```c
#define CAPSTONE_RISCV_COMPAT_HEADER
#include <capstone/capstone.h>

// CS_MODE_RISCVC now works as an alias for CS_MODE_RISCV_C
cs_open(CS_ARCH_RISCV, CS_MODE_RISCV32 | CS_MODE_RISCVC, &handle);
```

Adds `CS_OPT_SYNTAX_AARCH64_EXPLICIT_WIDE_IMM` flag to print `MOVN` / `MOVZ` instructions in explicit shifted form instead of `MOV` aliases.

RISC-V instruction details now expose the FP rounding mode via `cs_riscv.rounding_mode`.

New `CAPSTONE_RISCV_COMPAT_HEADER` define restores the `CS_MODE_RISCVC` alias for codebases targeting the v5 API.

Adds ColdFire `EMAC_B` dual-accumulator MAC support to the M68K module.

└──▷ BREAKING ON UPGRADE

! Alpha architecture enum names are renamed from `Alpha_` `-prefixed` mixed-case to `ALPHA_` `-prefixed` uppercase (e.g., `Alpha_REG_Foo` → `ALPHA_REG_FOO`); any code referencing the old names will fail to compile.

! `CS_MODE_RISCVC` was renamed to `CS_MODE_RISCV_C`; builds not using the new `CAPSTONE_RISCV_COMPAT_HEADER` workaround will break.

! `aarch64_op_mem.disp` changed from `int32_t` to `int64_t`; code that stores or passes this field as a 32-bit integer may break.

WAS THIS USEFUL? 👍 👎

NETWORK SECURITY

◆ Firewalls & Network Defense

Vectra AI

Feedback ⓘ v9.13 - Observe (2026-07-23)

2 RELEASES · 2026-07-23 BLOG ↗ LATEST STABLE

Vectra AI is a cybersecurity platform that uses AI-driven threat detection and response to identify and prioritize attacks across hybrid cloud, network, and identity environments.

Vectra AI v9.13 expands detections for Windows post-exploitation, HTTP attack tools, covert tunneling, and credential/data exfiltration techniques.

New detections for SMB-based lateral movement, malicious named pipe activity, and Cobalt Strike SMB beaconing across Windows environments.

Expanded HTTP-based behavioral signatures covering malicious file downloads/uploads, certificate abuse, and offensive framework activity using trusted Windows utilities.

New detection coverage for web-based tunnels, public tunneling services, and custom BackConnect protocols used to bypass perimeter controls.

New detections for Active Directory Certificate Services abuse, suspicious AD database file transfers, and cloud storage exfiltration via sync tools.

▾ 1 more release in this issue · 2026-07-23

> v9.13 - Platform (2026-07-23) BLOG ↗ STABLE
>
> Vectra AI v9.13 expands ERSPAN support to 8-core+ virtual Sensors and 16 IPs per capture interface.
>
> Supports 8-core and higher virtual Sensors for ERSPAN-based monitoring deployments.
>
> Supports up to 16 IPs per capture interface in ERSPAN environments.
>
> Increases processing efficiency for encapsulated traffic workloads in ERSPAN deployments.

WAS THIS USEFUL? 👍 👎

Rspamd

Feedback ⓘ 4.1.2 NOTES ↗

Rapid spam filtering system.

Rspamd 4.1.2 adds per-symbol Lua composites, a static embedding Neural provider, new Settings policies, and external Rspamadm command discovery.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch 4.1.2 https://github.com/rspamd/rspamd.git
# already have the repo? check out this version:
$ git checkout 4.1.2
```

└──▷ **USE IT**

Prevent automatic Neural network updates while still allowing manual forced-learn during a model evaluation period.

ini

```
neural {
  train {
    frozen = true;
  }
}
```

Load Rspamadm command modules from a custom directory to deploy organisation-specific admin tooling without modifying core config.

```
$ RSPAMADM_COMMAND_PATH=/opt/rspamd/custom-commands rspamadm <command>
```

Adds per-symbol Lua conditions and explicit dependencies in Composites rules.

Adds `static_embed` Neural provider with a WordPiece subword tokenizer over a static embedding matrix and per-token vector access.

Adds Neural `train.frozen` control to halt automatic training and a forced-learn minimal-scan fast path.

Adds `implicit_allow` policy option in Settings `apply` with additive `symbols_enabled`, plus strict `explicit_enable` gating.

Enables discovery of Rspamadm command modules from external directories via `rspamadm.d` or `RSPAMADM_COMMAND_PATH`.

+3 more

└──▷ **BREAKING ON UPGRADE**

- ! In `lua_task`, the `orig` address flavour now returns only the wire (transmitted) addresses when a rewrite preserved the originals — code expecting rewritten addresses under `orig` will receive different values.
- ! WebUI: FooTable has been removed and all tables are now Tabulator-based; any customizations or integrations targeting FooTable will break.
- ! Maps: glob map patterns are now anchored to match the whole subject; entries relying on substring matching must be updated to use explicit wildcards or they will stop matching.

WAS THIS USEFUL? 👍 👎

Fastly Feedback ⓘ

`Added virtual patch for CVE-2026-63030 (wp2shell, also covers CVE-2026-60137) (2026-07-21`

BLOG ↗

Fastly is a cloud edge platform that accelerates, secures, and delivers websites, apps, and APIs by caching content closer to users globally.

Fastly adds virtual patch for WordPress RCE chain CVE-2026-63030/CVE-2026-60137 (wp2shell) in WAF and control panel.

└──▷ HOW TO FIND IT

Enable the wp2shell virtual patch via the Fastly control panel to immediately protect WordPress services from the RCE chain.

> 📍 In the Fastly control panel, go to Virtual Patches › search for CVE-2026-63030 › click the pencil icon › set Status to Enabled › (optional) choose Block requests or Log requests › click Update virtual patch.

Enable the wp2shell templated rule in the Next-Gen WAF control panel to block and alert on CVE-2026-63030 signals.

> 📍 In the Next-Gen WAF control panel, go to Rules › Templated Rules › search CVE-2026-63030 › click View › Configure › Add trigger › select 'Block requests from an IP immediately if the CVE-2026-63030 signal is observed' › Update rule. Then go to the Signals tab › search CVE-2026-63030 › View › Detections tab › Add detection (Enabled) › Create detection. Then go to the Alerts tab › Add alert › set Status to Enabled › Save alert.

WAS THIS USEFUL? 👍 👎

CLOUD & CONTAINER SECURITY

◆ CNAPP

Wiz Cloud

Feedback ⓘ changelog · Jul 22, 2026 BLOG ↗

Wiz Cloud is a cloud security platform that provides full-stack visibility, risk assessment, and threat detection across multi-cloud environments.

Wiz adds a Threat Actors page listing historical threat actors that have targeted Fortinet products.

└──▷ HOW TO FIND IT

Review which threat actors have historically targeted Fortinet products to prioritize investigation and response.

> 📍 In the Wiz console, navigate to the Threat Actors page to view the comprehensive list of threat actors that have targeted Fortinet products.

WAS THIS USEFUL? 👍 👎

◆ Runtime Security

Wiz Defend

Feedback ⓘ changelog · Jul 22, 2026 BLOG ↗

Wiz Defend is a cloud detection and response platform that identifies, investigates, and remediates real-time threats across cloud environments.

Wiz Agentless Workload Detection gains EDR-like visibility into FortiGate appliances, Redis, npm supply chain, and more virtual appliance threats.

└──▷ HOW TO FIND IT

Hunt for rogue admin account creation on FortiGate devices during active CVE-2026-24858 exploitation — search the Cloud Event page for admin accounts named `helpdesk`.

> 📍 In the Wiz console, go to Cloud Events and run the pre-built query to search for creation of an admin named `helpdesk` on FortiGate devices.

Detect dangerously permissive FortiGate admin accounts (trusthost 0.0.0.0/0) that allow management access from any IP — a key persistence indicator in documented FortiGate campaigns.

> 📍 In the Wiz console, go to Cloud Events and run the pre-built event query to search for creation of admin accounts with trusthost set to 0.0.0.0/0 on FortiGate devices.

Extends Agentless Workload Detection to FortiGate firewalls, collecting and parsing local FortiOS logs not typically streamed to expose attacker activity including super-user account creation and configuration theft.

Adds Cloud Event page visibility for FortiGate authentication logs, enabling real-time monitoring of login activity and admin account changes for agentless EDR customers.

Adds built-in threat detection rules for Redis replication abuse: `Redis suspicious replication target configured`, `Redis module loaded from unusual path`, and `Attack Detected and Prevented by Redis`.

Enables agentless monitoring of npm package installation logs to detect supply chain compromises such as malicious versions of @asyncapi/specs in real time.

Provides agentless visibility into HTTP/HTTPS access logs (Nginx, Apache, IIS) to detect webshell deployments, unusual executable access, and exploitation of hosted applications such as WordPress via wp2shell.

+4 more

WAS THIS USEFUL? 👍 👎

◆ Threat Intelligence

Watcher

Feedback ⓘ v3.5.2 NOTES ↗

Watcher - Open Source AI-powered Cyber Threat Intelligence & Hunting Platform. Developed with Django & React JS.

Watcher v3.5.2 adds a Connectors dashboard for centralizing integrations, Fernet-encrypted credentials, and SSO auto-provisioning control.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.5.2 https://github.com/thalesgroup-cert/Watcher.git
# already have the repo? check out this version:
$ git checkout v3.5.2
```

└──▷ TRY IT

Prevent SSO from auto-creating accounts for unknown users — only pre-approved admin-created accounts can log in via OIDC.

```
$ OIDC_CREATE_USER=False
```

Verify all external integrations are reachable after a network change, without waiting for the Monday scheduled job.

📍 Navigate to `/connectors` (as superuser), then click Test next to each connector to see Healthy / Unhealthy status in real time.

New `/connectors` dashboard (superuser only) centralizes configuration, live overrides, and health-check testing for all external integrations (SMTP, Slack, TheHive, MISP, LDAP, OIDC, MISP, CertStream, SearxNG, CyberWatch, WIPO UDRP, MySQL).

Per-connector live connectivity tests report Healthy / Unhealthy / Not tested yet status alongside Configured / Incomplete / Disabled state.

Scheduled job automatically tests every connector each Monday at 06:00 and raises a Pending Action alert only for configured connectors that fail, without duplicating alerts on repeated failures.

Sensitive connector fields (passwords, tokens, API keys) are now encrypted at rest with Fernet; existing stored values are re-encrypted automatically on migration.

New `OIDC_CREATE_USER` setting lets admins restrict SSO login to pre-existing accounts, disabling auto-provisioning of new users on first login.

+3 more

└─▷ **BREAKING ON UPGRADE**

! This release includes database schema changes — run `python manage.py migrate` after updating or connector credential storage and new settings will not function correctly.

WAS THIS USEFUL? 👍 👎

OpenCTI

Feedback ⓘ 7.260722.0 NOTES ↗

Open Cyber Threat Intelligence Platform

OpenCTI 7.260722.0 adds a faceted connector marketplace, workflow reset, webhook deny-lists, admin-forced password changes, and auto-reloading public dashboards.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch 7.260722.0 https://github.com/OpenCTI-Platform/opencti.git
# already have the repo? check out this version:
$ git checkout 7.260722.0
```

└──▷ HOW TO FIND IT

Block webhooks from calling internal or sensitive URIs by configuring a deny-list at the platform level.

> 📍 In the console, go to Settings › Customization › Webhook and add URIs to the deny list to prevent outbound webhook calls to those destinations.

Force a specific user to reset their password on next login — useful after an account compromise or routine credential rotation.

> 📍 In the console, go to Settings › Security › Users, select the target user, and use the Force password change action.

Enable auto-reload on a public dashboard so external stakeholders always see up-to-date intelligence without manually refreshing.

> 📍 In the console, open a dashboard, go to its settings, and set a refresh_interval value to enable auto-reload for the public link.

Redesigns the connector catalog into a faceted marketplace experience for easier discovery and filtering.

Adds ability to fully reset a workflow to its initial state.

Adds a configurable deny-list of URIs for webhooks to block outbound calls to restricted endpoints.

Enables admins to force a password change for any user account.

Adds auto-reload support for public dashboards so shared views stay current without manual refresh.

+3 more

WAS THIS USEFUL? 👍 👎

Feedly Threat Intelligence

Feedback ⓘ

Enrich and triage Apple security releases in the Vulnerability Intel Agent (2026-07-21)

BLOG ↗

Feedly Threat Intelligence is a security tool that aggregates and analyzes threat feeds, helping security teams track emerging threats, vulnerabilities, and adversary activity in real time.

Feedly's Vulnerability Intel Agent now ingests Apple security releases within minutes, with enrichment and filtering by exploitation, product, and patch policy.

.
.
.

└──▷ **HOW TO FIND IT**

Focus remediation efforts on only the Apple CVEs in your macOS fleet that are already being exploited or have a published PoC.

> 📍 In the Vulnerability Intel Agent, select `Latest Apple Security Release` as the filter and set the date range. Add filters for `exploited` and `proof of concept`. Hide iOS and iPadOS rows using the product-level filter to scope results to macOS only.

Automatically flag Apple CVEs that breach your patch-now SLA thresholds (e.g., CVSS ≥ 8 and EPSS ≥ 80%) so they can be routed to immediate remediation.

> 📍 In the Vulnerability Intel Agent, select the target Apple security release, then set CVSS threshold filter to ≥ 8 and EPSS threshold filter to ≥ 80%. Deploy the Intel Agent as an alert to get notified when new CVEs matching these thresholds appear in future releases.

Adds Apple security release ingestion to the Vulnerability Intel Agent, with each CVE appearing in the output table within minutes of publication.

Supports filtering CVEs by exploitation activity (exploited, PoC, weaponized, CISA KEV) to surface the highest-priority items from a release.

Enables product-level filtering so practitioners can scope a release to only the Apple products they actually run (e.g., macOS only).

Supports threshold-based patch policy filters on CVSS score and EPSS probability to automatically surface CVEs that breach SLA triggers.

Adds a `Latest Apple Security Release` filter option that always points to the most recent Apple advisory.

+1 more

WAS THIS USEFUL? 👍 👎

◆ Security Data Pipeline

Cribl Feedback ⓘ

Cribl MCP Server: Interact with Cribl using the AI tools you already use (2026-07-22)

BLOG ↗

Cribl is an observability pipeline platform that helps organizations collect, process, route, and transform machine data from any source to any destination.

Cribl 4.19 embeds an MCP Server natively in the Leader, letting AI clients like Claude Code, Cursor, and OpenAI Codex CLI query configs, metrics, and Search Notebooks directly.

└──▷ HOW TO FIND IT

Enable the embedded MCP Server and connect an AI client to your Cribl Leader for the first time.

> 📍 In the Cribl UI, go to Settings › AI Settings and enable MCP Server. Copy the generated bearer token, then add the Cribl MCP endpoint and token to your client configuration (e.g., Claude Code, Cursor, or OpenAI Codex CLI). Restart the client and confirm the Cribl tools appear and the server shows as connected.

Adds a native, embedded MCP Server in the Cribl Leader — no separate container or service required to connect MCP-compatible AI clients.

Supports MCP-compatible clients including Claude Code, Cursor, and the OpenAI Codex CLI via remote HTTP, with no additional proxy to install.

Enables AI clients to retrieve Sources, Destinations, Worker Groups, Pipelines, metrics, and alerts directly from the Cribl environment.

Allows AI clients to connect to Cribl Search to generate and refine queries, retrieve results, and create or update Search Notebooks.

Supports building AI-powered runbooks and security playbooks as Notebooks that combine queries, Markdown summaries, findings, and next steps.

+3 more

WAS THIS USEFUL? 👍 👎

◆ SIEM

Sumo Logic

Feedback ⓘ

Searchable Time Support for Monitors (Alerts) (2026-07-22) BLOG ↗

Sumo Logic is a cloud-native platform that collects, analyzes, and visualizes machine data and logs for security, observability, and operational intelligence.

Sumo Logic Monitors now support Searchable Time, letting alerts evaluate logs based on when they became searchable.

WAS THIS USEFUL? 👍 👎

◆ SOAR

Tines

Feedback ⓘ JSON in records (2026-07-21) BLOG ↗

Tines is a no-code automation platform that helps security and operations teams build workflows to automate repetitive tasks and incident response processes.

Tines Records now supports a JSON field type for storing structured data alongside other record fields.

Adds a JSON field type to Records, enabling structured data to be stored directly within a record.

WAS THIS USEFUL? 👍 👎

Detection Engineering

Tenzir

Feedback ⓘ Tenzir Node v6.8.0: Splunk search input (2026-07-21)

BLOG ↗

Tenzir is an open-source security data pipeline platform that helps teams collect, route, shape, and analyze large-scale security telemetry efficiently.

Tenzir v6.8.0 adds a Splunk search input operator and improves ClickHouse throughput with JSON column support.

New `from_splunk` operator lets nodes run bounded searches against a Splunk Search Head and emit results as events.

Upgraded `to_clickhouse` operator re-batches events before insert for higher throughput.

Adds support in `to_clickhouse` for writing directly into JSON columns.

Nodes now quarantine corrupt or truncated partitions instead of crashing repeatedly, improving resilience.

WAS THIS USEFUL? 👍 👎

Endpoint Security (EDR/XDR)

Trend Vision One

Feedback ⓘ

Integrate AI Secure Access with a third-party proxy (2026-07-20) BLOG ↗

Trend Vision One is a unified cybersecurity platform that integrates extended detection and response (XDR) across endpoints, networks, email, and cloud environments.

Trend Vision One AI Secure Access can now chain with third-party secure web gateways (e.g., Zscaler, Netskope) for AI-specific inspection.

WAS THIS USEFUL? 👍 👎

CrowdStrike

Feedback ⓘ undated-6a01ac0f BLOG ↗

CrowdStrike is a cybersecurity platform that provides cloud-native endpoint protection, threat intelligence, and real-time detection to defend organizations against cyberattacks and breaches.

CrowdStrike adds detections for `SANDWORM_MODE`, a multi-stage npm supply chain worm targeting AI coding assistants and LLM toolchains.

Adds indicators-of-attack (IOA) coverage for `SANDWORM_MODE`, detecting a multi-stage npm supply chain worm that targets AI-augmented CI/CD pipelines including GitHub Actions, Cursor, Claude Code, and Windsurf.

Deploys process-tree-ancestry-based detections for Node.js runtime behaviors characteristic of `SANDWORM_MODE`, including obfuscated loader execution, in-memory payload staging via `/dev` /shm, and rogue MCP server registration.

Extends detection coverage to AI toolchain compromise behaviors: rogue MCP server injection into Claude Desktop, Cursor, VSCode, and Windsurf configuration files, and LLM API key harvesting from environment variables and .env files.

Adds detection for `SANDWORM_MODE` persistence mechanisms, including malicious git hook injection via ~/.git-templates/hooks/ and manipulation of git config `--global init.templateDir`.

Introduces detection for multi-channel exfiltration behaviors including HTTPS POST to attacker infrastructure, GitHub private repository mirroring, and DNS tunneling with Base32 encoding and domain generation algorithm fallback.

WAS THIS USEFUL? 👍 👎

CrowdStrike Falcon Insight XDR

Feedback ⓘ

undated-6cd80cac BLOG ↗

CrowdStrike Falcon Insight XDR is an extended detection and response platform that unifies endpoint and cross-domain telemetry to detect, investigate, and stop advanced threats.

Falcon Insight XDR adds `SANDWORM_MODE` detection surfaced in the Endpoint Detections UI.

WAS THIS USEFUL? 👍 👎

EXPOSURE & VULNERABILITY MANAGEMENT

◆ Vulnerability Management

Qualys VMDR

Feedback ⓘ changelog · Jul 22, 2026 BLOG ↗

Qualys VMDR is a cloud-based vulnerability management platform that continuously discovers, assesses, prioritizes, and remediates security vulnerabilities across IT assets.

Qualys VMDR adds QID coverage for Oracle's July 2026 Critical Patch Update across 10+ Oracle products.

Adds QID detections for Oracle July 2026 Critical Patch Update (CPUJUL2026) covering JDeveloper, Java SE, Coherence, VirtualBox, HTTP Server, Solaris 11.4, WebLogic Server, E-Business Suite, and MySQL Server.

WAS THIS USEFUL? 👍 👎

IDENTITY, ACCESS & DATA PROTECTION

◆ Identity & Access Management

Nudge Security

Feedback ⓘ

Export a complete log of admin activity · 3 RELEASES · SEEN 2026-07-22 · NOTES ↗

LATEST · STABLE

Nudge Security discovers and secures all SaaS apps used across an organization, helping IT teams manage shadow IT and enforce security policies through automated nudges.

Nudge Security adds exportable admin audit log capturing OAuth revocations, role changes, and app connect/disconnect events.

└──▷ HOW TO FIND IT

Pull a CSV of all admin actions for a SOC 2 audit or incident investigation.

> 📍 In the console, go to Settings › Audit event, review the log of administrative actions, then click Export to download as CSV.

▾ 2 more releases in this issue · seen 2026-07-22

Know your real vendor risk with inherent and residual risk scores, now in open beta

NOTES ↗ · STABLE

Nudge Security adds continuously updated inherent and residual risk scores for every SaaS and AI vendor, now in open beta.

└──▷ HOW TO FIND IT

Review the full risk breakdown and recommended remediations for a specific app to decide what to fix next.

> 📍 In the Nudge Security console, open your App Inventory, select any SaaS or AI app, then click the Risk tab to view its inherent and residual risk scores, contributing risk factors, and recommended actions ranked by projected residual risk reduction.

Filter your entire app inventory by risk level to quickly surface Critical or High residual-risk vendors across the organization.

> 📍 In the Nudge Security console, go to App Inventory and use the risk score filter to narrow the view to Critical or High residual risk apps.

Adds inherent and residual risk scores for every SaaS and AI vendor, mapped to Low/Medium/High/Critical levels and recalculated automatically as risk factors change.

Surfaces a transparent breakdown of risk drivers: vendor security posture, internal deployment/usage factors (e.g., critical downstream connections, AI agents), and compensating controls.

Provides recommended actions ranked by projected residual risk reduction so teams know what to remediate first.

Scores every discovered app on Day One without requiring vendor cooperation, covering the long tail of shadow IT and AI apps.

Continuously monitors for new risk signals — including new MCP server connections and vendor breach disclosures — and updates scores in real time.

+1 more

Prioritize your SaaS and AI inventory automatically with data types and business criti

NOTES ↗ STABLE

Nudge Security adds automatic business criticality tiers and 27 data-type filters for SaaS/AI inventory, now in open beta.

└──▷ **HOW TO FIND IT**

Customize which data types are considered high-sensitivity to align automatic criticality tiers with your organization's governance model.

📍 In the console, go to Settings › Risk and adjust the data sensitivity tier assignments to match your organization's data governance model.

Automatically assigns each app a business criticality tier (Low, Medium, High, or Critical) based on inferred data types the vendor typically handles.

Filters SaaS and AI inventory across 27 distinct data types with no manual research required.

Enables customization of data sensitivity tiers to match your organization's data governance model via Settings > Risk.

Surfaces security posture findings selectively for Critical and High tier apps, reducing noise for lower-risk apps.

Displays business criticality tiers and data types within the app inventory view and individual app overviews.

└──▷ **BREAKING ON UPGRADE**

! Previously tagged data types are moved to a `Data Types (Legacy)` field, separate from the new automatically populated `Data Types` field.

WAS THIS USEFUL? 👍 👎

Auth0 Feedback ⓘ

`Self-Service Enterprise Configuration: Customer admins can now configure third-party appl`

BLOG ↗

Auth0 is a cloud-based identity platform that provides authentication and authorization services, enabling developers to securely add login and access control to applications.

Auth0 Self-Service Enterprise Configuration lets customer admins control third-party app SSO access during their own setup wizard.

└──▷ **HOW TO FIND IT**

When creating a Self-Service Enterprise Configuration ticket, delegate the third-party access decision to the customer admin instead of setting it yourself — useful when you want customers to self-govern whether third-party apps can use their SSO connection.

> 📍 In the Auth0 console, go to Authentication › Enterprise › Self-Service Enterprise Configuration, create or edit a ticket, and enable third-party access delegation on the ticket instead of setting domain connection access directly.

WAS THIS USEFUL? 👍 👎

GOVERNANCE, RISK & COMPLIANCE

◆ Integrated & IT Risk Management

CISO Assistant

Feedback ⓘ v3.20.0 NOTES ↗

CISO Assistant is a one-stop-shop GRC platform for Risk Management, AppSec, Compliance & Audit, TPRM, BIA, Privacy, and Reporting. It supports 150+ global frameworks with automatic control mapping, including ISO 27001, NIST CSF, SOC 2, CIS, PCI DSS, NIS2, DORA, GDPR, HIPAA, CMMC, and more.

CISO Assistant v3.20.0 adds technical posture management, SSO single logout, MCP maturity scores, domain export, and a library builder.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch v3.20.0 https://github.com/intuitem/ciso-assistant-community.git
# already have the repo? check out this version:
$ git checkout v3.20.0
```

New technical posture management capability for tracking and assessing technical security posture.

Implements Single Logout (SLO) for both OIDC and SAML authentication providers.

MCP integration now exposes maturity scores, enabling AI-driven access to assessment data.

New CLI domain export capability for extracting domain data.

New in-product library builder with requirement weight support, enabling custom framework authoring.

+4 more

WAS THIS USEFUL? 👍 👎

BUILD

◆ AI Coding Agents

Kiro

Feedback ⓘ Global Hooks, Session Search, and Reliability Fixes

NOTES ↗

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 global hooks, searchable session history, proxy reliability, and dynamic MCP tool refresh.

Adds user-level global hooks that apply across all projects without per-project configuration.

New searchable session history panel lets practitioners quickly locate and revisit past agent sessions.

Agent now recovers automatically from interrupted responses in corporate proxy environments.

MCP tools refresh dynamically without requiring a restart.

WAS THIS USEFUL? 👍 👎

OpenAI Codex CLI

Feedback ⓘ rust-v0.146.0-alpha.2

4 RELEASES · 2026-07-21 → 2026-07-22 NOTES ↗ CODE ↗ PRE-RELEASE

Lightweight coding agent that runs in your terminal

Codex CLI adds Codex Auto Review model metadata, session headers for realtime conversations, and routes LM Studio traffic through the shared HTTP client.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.146.0-alpha.2
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.146.0-alpha.2
```

Configures Codex Auto Review with an action-risk review prompt template, default generic security policy, approval rules, and permission messages for `danger_full_access`, `workspace_write`, and `read_only` modes.

Includes `session-id` and `thread-id` headers when starting realtime conversations over WebSocket or WebRTC.

Routes LM Studio server requests through the shared route-aware HTTP client pool, adding connection-timeout support scoped to connection establishment.

▾ 3 more releases in this issue · 2026-07-21 → 2026-07-22

rust-v0.146.0-alpha.1 NOTES ↗ CODE ↗ PRE-RELEASE

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.146.0-alpha.1
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.146.0-alpha.1
```

Adds `app/read` app-server API endpoint for fetching metadata for up to 100 app IDs with optional tool summaries, deduplication, and caching.

Adds audio input support forwarded to the Responses API, gated by model input modalities.

Adds API for reading installed app runtime state.

rust-v0.145.0-alpha.30 NOTES ↗ CODE ↗ PRE-RELEASE

Codex CLI alpha.30 adds Esc-key session interruption and parallel stdin writes across unified-exec sessions.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch rust-v0.145.0-alpha.30
https://github.com/openai/codex.git
# already have the repo? check out this version:
$ git checkout rust-v0.145.0-alpha.30
```

Runs write_stdin in parallel across unified-exec sessions for faster multi-session I/O.

Updates bundled model catalog with GPT-5.6 context windows set to 272,000 tokens and new personality instruction variables.

rust-v0.145.0 NOTES ↗ CODE ↗ LATEST STABLE

Codex CLI v0.145.0 adds paginated thread history, Amazon Bedrock support, audio inputs, multi-agent V2 GA, and `/import` for Cursor/Claude settings.

└──▷ GET THIS VERSION ▾

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

Adds experimental paginated thread history with resume, search, persisted thread names, sub-agent support, and memories.

Expands `/import` to migrate settings, MCP servers, plugins, sessions, commands, and project-scoped memories from Cursor and Claude Code.

Adds experimental Amazon Bedrock login with custom endpoint and authentication support; sets GPT-5.6 Sol as the default Bedrock model.

Supports audio inputs and tool outputs (common local audio formats) and introduces streaming realtime V3 conversations.

Adds secure, clickable inline visualization links in the terminal UI.

WAS THIS USEFUL? 👍 👎

Cline

Feedback ⓘ desktop-v0.0.2 NOTES ↗

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

Cline Code desktop app launches for macOS with signed/notarized builds for Apple Silicon and Intel, plus automatic background updates.

└──▷ GET THIS VERSION ▾

```
$ git clone --branch desktop-v0.0.2 https://github.com/cline/cline.git
# already have the repo? check out this version:
$ git checkout desktop-v0.0.2
```

New native macOS desktop app for running and inspecting Cline agent sessions, signed and notarized for both Apple Silicon and Intel.

Adds automatic update mechanism: checks on launch and every 2 hours, downloads new versions in the background, and prompts for a one-click restart.

WAS THIS USEFUL? 👍 👎

Claude Code

Feedback ⓘ v2.1.217 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 v2.1.217 adds emoji autocomplete, subagent concurrency controls, and AWS gateway deployment assets.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v2.1.217 https://github.com/anthropics/claude-code.git
# already have the repo? check out this version:
$ git checkout v2.1.217
```

└──▷ **TRY IT**

Cap parallel subagents to a lower value in resource-constrained CI environments to prevent runaway agent fan-out.

```
$ CLAUDE_CODE_MAX_CONCURRENT_SUBAGENTS=5 claude
```

Allow nested subagent spawning up to two levels deep for multi-stage agentic workflows that require it.

```
$ CLAUDE_CODE_MAX_SUBAGENT_SPAWN_DEPTH=2 claude
```

Adds emoji shortcode autocomplete in the prompt input (type `:heart:` to insert or `:hea` for suggestions); disable with the `emojiCompletionEnabled` setting.

Adds visible warnings when transcript writes are failing (e.g. disk full) or when session saving is disabled via an inherited environment variable, preventing silent transcript loss.

Adds a cap on concurrently-running subagents (default 20), overridable with `CLAUDE_CODE_MAX_CONCURRENT_SUBAGENTS`, preventing unbounded agent fan-out from a single message.

Adds `CLAUDE_CODE_MAX_SUBAGENT_SPAWN_DEPTH` to control nested subagent depth; nested spawning is now disabled by default.

Makes footer PR badge links clickable hyperlinks in environments where terminal support can't be detected (e.g. over SSH/tmux); opt out with `FORCE_HYPERLINK=0`.

+1 more

└──▷ **BREAKING ON UPGRADE**

! Subagents no longer spawn nested subagents by default; existing workflows that relied on nested subagent spawning will stop working unless `CLAUDE_CODE_MAX_SUBAGENT_SPAWN_DEPTH` is set to permit deeper nesting.

WAS THIS USEFUL? 👍 👎

Pi

Feedback ⓘ v0.81.1 NOTES ↗

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

Pi v0.81.1 adds verifiable, checksummed release source archives for reproducible standalone binary builds.

└──▷ **GET THIS VERSION** ▾

```
$ git clone --branch v0.81.1 https://github.com/earendil-works/pi.git
# already have the repo? check out this version:
$ git checkout v0.81.1
```

Adds deterministic, checksummed source archives to GitHub releases with documented instructions for rebuilding standalone binaries, supporting supply-chain verification.

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · July 22, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe Feedback on this issue