THE CYBER TOOLCHAIN NO. 103

Tail

THE DAILY RELEASE FIREHOSE

PUBLISHED OCTOBER 11, 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

INFORMATION DENSITY

everything



BRIEFED | **EVERYTHING**

Every tool in the issue, in full.

all 4 changes

Do you prefer **scroll**? 👍 👎

▸ // CONTENTS SHOW 5 tools

▾ // THE BRIEF HIDE

What stands out across today's releases, grouped by what it lets you do. Every tool named links to its entry below.

DETECT

Find the AI traffic leaving your network that nobody approved

vim ? help

Two curated indicator lists of AI infrastructure, pulled through the List API, give you something to match egress logs and proxy policy against. Shadow AI has been hard to quantify because the destinations are ordinary SaaS endpoints; a maintained list turns 'we think people are pasting code into chatbots' into a query you can run and a block you can write.

Recorded Future

IDENTIFY

See what the AI coding agents on your developers' machines can actually run

Inventories agent skills across 18 AI coding agents on both developer endpoints and GitHub repos, flagging executable content, where it came from, and risky configuration. Skill and MCP-style extensions install quietly per-developer and carry arbitrary commands, so until now the only way to know what was on a laptop was to ask the person using it.

StepSecurity

PROTECT

Get a machine-generated patch you can actually review before merging

The open-source Auto Patcher harness for OpenAnt produces AI-written vulnerability fixes with a Trust Report laying out the evidence behind each one. The bottleneck on AI remediation has never been generating a diff, it is a reviewer deciding whether to trust it; shipping the reasoning alongside the patch is what makes that review cheap enough to do at volume.

Knostic

APPLICATION SECURITY

◆ Software Supply Chain Security

StepSecurity ⓘ announced · Oct 8, 2026 BLOG ↗ RANK

StepSecurity hardens CI/CD pipelines by detecting and preventing supply chain attacks through runtime security monitoring and vulnerability detection.

> StepSecurity launched an AI agent skill inventory feature spanning both developer machines and GitHub repositories, surfacing executable content, provenance, and risky configurations across 18 AI coding agents.

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

01 **Agent skill inventory for developer machines** NEW 87

A new Agent Skills page in Dev Machine Guard (under `IDE & AI Agents`) inventories AI agent skills across developer machines, scanning `~/.agents/skills`, `~/.claude/skills`, project-level skill folders, and `skills.sh` lock files across 18 AI coding agents including Claude Code, Codex, GitHub Copilot, Cursor, and Gemini CLI. It adds `code`, `hooks`, and `shell` flags per skill to surface executable content, plus a `command` flag for Claude Code commands, filterable via `Has code`, `Has hooks`, and `Shell execution` filters. Provenance tracing for `skills.sh` `-managed` skills surfaces upstream source, ref, path, upstream folder hash, and the `skills-lock.json` file; the inventory supports CSV export, notifications for new skills detected after a device's initial baseline with per-device alerts, and content-hash comparison across devices to spot installs where a skill's `SKILL.md` content differs.

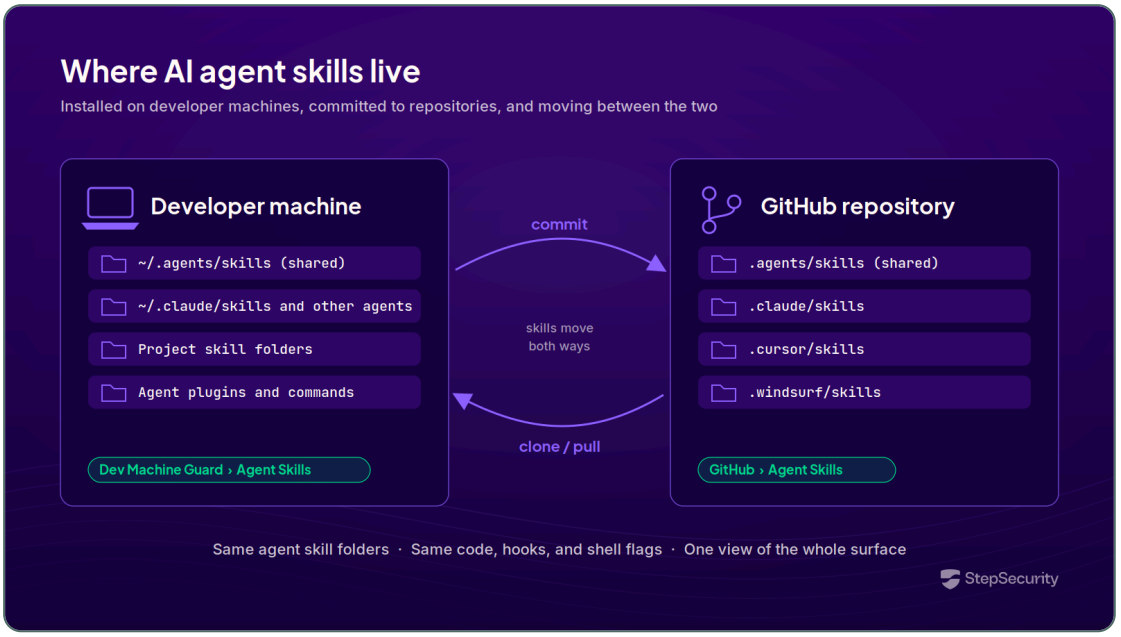

Find developer machines where the same skill has diverging content — a signal that one copy may have been tampered with or is outdated.

📍 In the StepSecurity dashboard, go to IDE & AI Agents › Agent Skills (Dev Machine Guard). Sort or filter by `Content hashes` › 1, then select the skill to open

the per-device panel grouped by content hash and compare variants.

Audit unmanaged (`Local`) skills on developer machines that have no upstream source in `skills.sh` or a plugin marketplace — the class most likely to have been hand-planted.

📍 In the StepSecurity dashboard, go to IDE & AI Agents › Agent Skills (Dev Machine Guard) and filter by Source = `Local` . Export the result as CSV for offline review.

— Names exact paths, flags, filters, and export; UI navigation given in examples.

`launch-20261008-9cfac2cd`

---

02 **Agent skill inventory for GitHub repositories** NEW 75

A new Agent Skills page under GitHub scans the default branch of each non-archived repository in the organization daily, detecting skills in folders such as `.claude/skills` , `.cursor/skills` , `.windsurf/skills` , and `.agents/skills` . A Declared Behavior panel shows frontmatter details including pre-approved tools (e.g., `Read` , `Write` , `Bash`), invocation restrictions (`no model invocation` , `not user invocable`), forked subagent context, and model overrides, and repository skills are labeled `invalid yaml` when frontmatter can't be parsed or `symlink` when the skill folder is reached via a symbolic link.

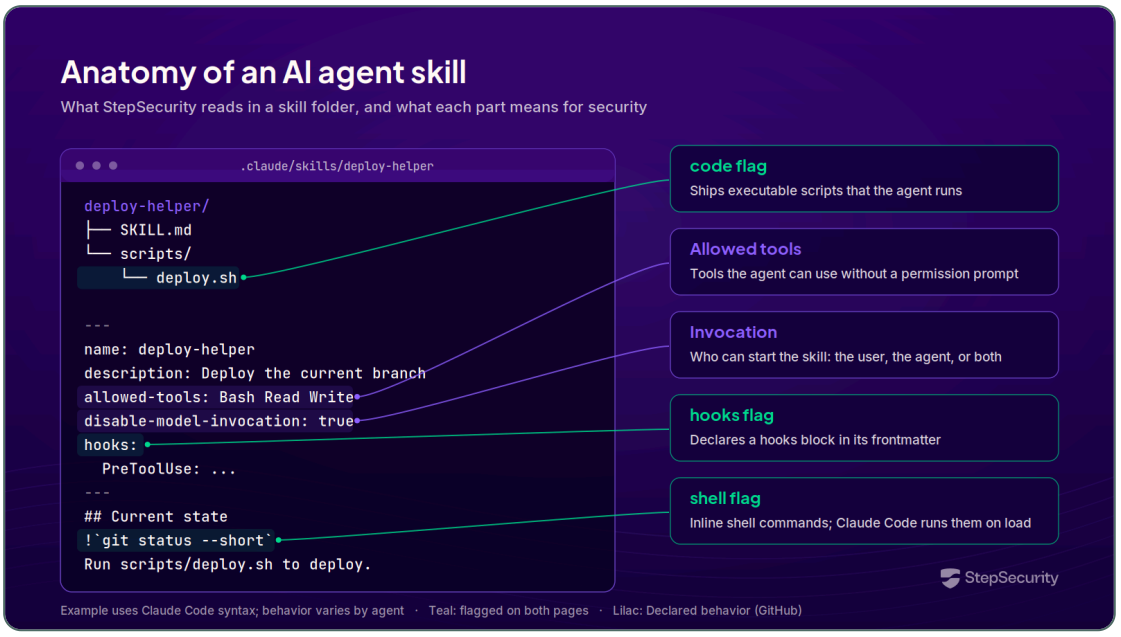

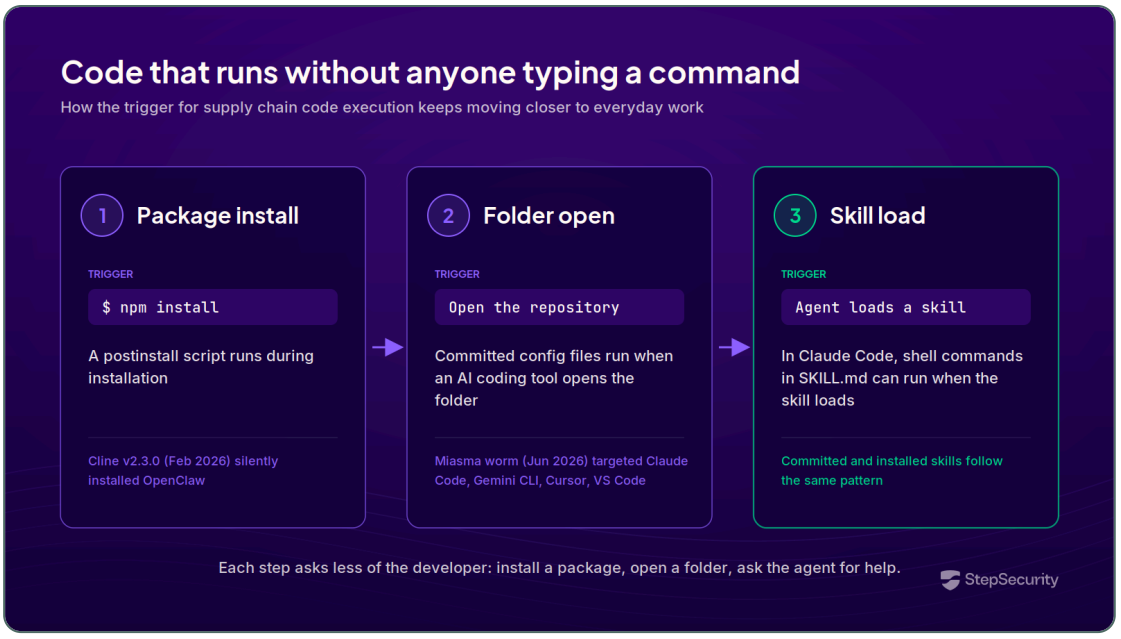

Immediately surface every repository skill that pre-approves shell execution — the highest-risk class, equivalent in impact to a CI/CD workflow change.

> 📍 In the StepSecurity dashboard, go to GitHub › Agent Skills, then apply the `Shell execution` filter. Open each flagged skill's details panel and review `Declared behavior` › `Allowed tools` for entries such as `Bash`.

— Names scan scope, folders, and flags; only a UI path given, not a runnable command.

`launch-20261008-9cfac2cd`

WAS THIS USEFUL? 👍 👎

REVERSE ENGINEERING, MALWARE & FORENSICS

◆ Digital Forensics & Incident Response

Belkasoft

announced · Oct 9, 2026

BLOG ↗

Accelerate your digital forensics and incident response investigations with Belkasoft

Belkasoft X adds MTK dump acquisition via BootROM mode: physical imaging, FBE userdata decryption, and offline brute-force passcode recovery for MediaTek Android devices.

- Adds MTK dump acquisition method using BootROM (BROM) mode to create a physical image and extract cryptographic material from supported MediaTek-based Android devices.
- Adds decryption of the `userdata` partition (file-based encryption / FBE) using cryptographic material obtained during acquisition.
- Adds offline brute-force passcode recovery via the Brute-force options dialog: built-in Belkasoft dictionary (4- and 6-digit numeric), sequential numeric brute-force, and a custom User dictionary (.txt, one entry per line) supporting alphanumeric passcodes — running at ~60 candidates/second without device lock-screen limits.
- Supports importing MTKClient-acquired dumps as a Mobile image data source after key-format conversion.

COVERAGE

NEW supported platform MediaTek chipset physical acquisition support

All of Belkasoft's coverage changes →

WAS THIS USEFUL? 👍 👎

ElcomSoft

announced · Oct 9, 2026

BLOG ↗

Try professional password recovery, data decryption, mobile and cloud forensic tools from a manufacturer with 30+ years of expertise, providing tools and training to law enforcement, financial and intelligence agencies.

iOS Forensic Toolkit 10.12 gains bootloader-level full file system and keychain extraction for Apple HomePod mini (macOS and Linux editions).

- Adds bootloader-level extraction for Apple HomePod mini using the `usbliter8` SecureROM exploit, returning a full file system image and decrypted keychain regardless of installed audioOS version.
- HomePod mini extraction is available in the macOS and Linux editions of iOS Forensic Toolkit 10.12.

COVERAGE

NEW supported platform Bootloader extraction for Apple HomePod mini

All of ElcomSoft's coverage changes →

WAS THIS USEFUL? 👍 👎

DETECTION & RESPONSE

◆ Threat Intelligence

Recorded Future

announced · Oct 9, 2026 BLOG ↗ RANK

Get started with the Recorded Future API. Covers authentication, base URLs, API key setup, and how to integrate threat intelligence into your security workflows.

> Recorded Future added two curated AI Infrastructure Indicator Lists, giving customers a free way to discover and enforce policy against shadow AI network traffic via the List API.

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

01 **AI Infrastructure Indicator Lists** NEW 85

Two new curated indicator lists — AI Infrastructure and Chinese AI Infrastructure — contain IP addresses, CIDR ranges, and domains with Risk Scores and Comment fields (provider, verification date, association method: Live JSON, Official documentation, BGP/WHOIS, or Community/OSINT) for identifying AI-related network traffic. They cover general-purpose AI assistants, AI coding assistants, vibe-coding tools, AI search tools, voice/audio AI tools, AI video tools, Chinese-domiciled AI tools, and autonomous AI agents, and are exposed via the List API with daily retrieval recommended and weekly list refresh cadence. The lists are available at no additional cost to customers with a Cyber Operations or Threat Intelligence Module license.

Before enforcing blocks, baseline 30-90 days of DNS and proxy logs against the full list to discover shadow AI traffic — then segment by service provider in the Comment field and Risk Score to prioritise action.

> 📍 In the Recorded Future platform, navigate to the AI Infrastructure Indicator List, then run it against your historical DNS and proxy logs. Filter the results by the provider in the Comment field and sort by Risk Score to identify vibe-coding and Chinese AI tools as the highest-priority undiscovered categories.

— Names fields, categories, API, cadence but UI path only launch-20261009-70b423d2

WAS THIS USEFUL? 👍 👎

GOVERN

◆ AI/LLM Security

Knostic ⓘ

Auto-patch for OpenAnt: Generating trust in autonomous patching (2026-10-11)

BLOG ↗ RANK

Knostic enforces need-to-know access controls for enterprise AI assistants such as Microsoft Copilot.

> Knostic shipped Auto Patcher for OpenAnt, an open-source harness that generates AI patches for vulnerabilities and attaches a Trust Report documenting the evidence behind them.

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

01 **Auto Patcher for OpenAnt with Trust Report** NEW 58

Auto Patcher for OpenAnt is an open-source vulnerability discovery and remediation harness that investigates a vulnerability and its repository, generates a candidate fix, and challenges its own conclusions. It produces a Trust Report bundling the patch, collected evidence, and a recommendation, surfacing what was verified and what remains unknown after autonomous patch generation.

— Explains mechanism and output but no commands, flags, or install steps given.

Auto-patch for OpenAnt: Generating trus…

WAS THIS USEFUL? 👍 👎

// END OF TRANSMISSION

The Cyber Toolchain · curated daily · October 11, 2026
Releases & features from the open-source and commercial cyber tools that practitioners actually run.

▷ Subscribe