It's Probably Nothing
by Sierra Echo Charlie

EDITION 28 / 16 NOVEMBER 2026

Old agent traces, fresh supply-chain trouble

Cursor explains its extension compromise; N-central's hotfix has a configuration-dependent gap.

4 stories / Listen online - 3:59

THE TAKEAWAY

Cursor's post-mortem is a particularly awkward example of retained agent data becoming something much more consequential than a privacy problem.

01 / THE BRIEF

The credentials were still in the traces

Cursor's 13 November post-mortem traces its extension compromise back to a key exposed at Fireworks. Cursor says attackers used it to copy historical coding-agent traces, mostly from 2025 and outside Privacy Mode. Unredacted credentials then opened access to selected SaaS tools and more internal information.

On 29 October, a stolen credential let attackers place malicious Claude Code copies directly into Cursor's marketplace cache. At most 5,000 users potentially received them over approximately 11 hours; that isn't a count of infections. Cursor says neither Anthropic nor Open VSX distributed the malicious copies.

Cursor reports trace exposure for fewer than 0.05% of users and API-credential exposure for less than 0.01% of customers: different denominators. It says Privacy Mode prompts and code were unaffected and believes the incident is contained. Retained debugging material became a route into software distribution. That's quite a second life.

WORTH ASKING

If you received Anthropic.claude-code 2.1.312 or 2.1.313 through Cursor, its warning is worth reading: uninstalling alone may be insufficient.

SOURCES

An update on our recent security incident - Primary source

Cursor confirms security breach led to malicious Claude Code extension - Reporting / context

Leaked Fireworks AI Key Enabled Backdoored Claude Code Extensions in Cursor - Reporting / context

02 / THE BRIEF

N-central: HF2 isn’t the whole answer

The change since our earlier coverage is reported successful compromise despite HF2. SecurityWeek reports that Lupovis and Huntback observed compromises and XMRig deployment; Huntback reportedly identified 19 victims running the latest patches.

The prerequisite matters: Huntress says the unauthenticated route to server root, including HF2 build 2026.4.1.26, requires the non-default Legacy API to be enabled and reachable. Its endpoint command-execution paths also require Legacy API. Huntress says it hadn’t observed exploitation in its partners’ environments.

N-able published compromise indicators on 13 November. HF3 was still in final validation on 15 November, with availability unestablished at cutoff.

WORTH ASKING

For N-central operators, the vendor’s guidance is to disable Legacy API, retain HF2 remediation and upgrade Agents and Probes.

SOURCES

N-able Status Dashboard - November 12 Update: Critical N-central Security Update - Primary source

N-able Disables Cloud Services Following Disclosure of "Able2Shell" Exploit Chains - Primary source

Hackers Exploit New N-central Zero-Day Hours After PoC Publication - Reporting / context

N-able Takes Hosted N-central Offline After Critical Vulnerability Disclosure - Reporting / context

03 / THE BRIEF

Check Point’s root stage now comes with forensic triage

CISA’s 13 November addition of CVE-2026-123039 requires forensic triage, with a 16 November deadline for covered US federal agencies. This isn’t standalone remote entry: DynamicCLI turns existing local execution into root through a world-writable, unauthenticated socket.

Check Point’s investigation describes PHP2Shell entry through unpatched Mobile Access portals, followed by root escalation and persistent gateway backdoors. In the investigated intrusion, attackers then reached an ESXi host over SSH and ran the destroy_v6 wiper. The broader second wave had no evidenced destructive activity. This is more than a patch-version conversation.

SOURCES

CISA Catalog of Known Exploited Vulnerabilities - Primary source

CVE-2026-123039 - Local privilege escalation to root in the DynamicCLI process via local world-writable control socket - Primary source

PHP2Shell Exploitation on Check Point Gateway - Primary source

CISA Flags Exploited UnrealIRCd Command Injection and Check Point Root Escalation Flaws - Reporting / context

04 / THE BRIEF

The outage notices acquired a phishing edition

The University of Auckland documented Microsoft 365 disruption on 13 November, then warned about fake Microsoft and IT-support outage messages.

iTnews reports Microsoft observed fake service-health emails citing genuine incident EX1550543, threatening deletion of held messages within 24 hours unless recipients followed recovery links. Credential collection wasn't confirmed.

Microsoft told ABC there was no indication the outage itself was a cyber attack. The useful distinction: the service failure wasn't established as malicious; the opportunistic recovery messages were a separate problem. A real incident number is apparently available to everyone.

SOURCES

Global Microsoft outage - University of Auckland - Primary source

Global outage hits Microsoft's Outlook, 365 services - Reporting / context

Access to Microsoft Outlook and other services being restored as cyber security chief investigates - Reporting / context

Phishers quick to exploit extended M365 outage - Reporting / context

Phishing Campaign Exploits Microsoft 365 Outage With Fake Email Recovery Notices - Reporting / context

THAT'S IT FOR TODAY.

Listen to this edition or browse the archive. Narrated using an AI voice.

Researched by our AI agents using Mallory and original sources. Reviewed and published by Sierra Echo Charlie.

Want intelligence your team can act on? Reach out for a chat.