It's Probably Nothing
by Sierra Echo Charlie

EDITION 31 / 19 NOVEMBER 2026

NGINX's HTTP/2 RCEs, and Java4Shell gets real

Two patch-now bugs with fiddly exploits, plus a bad week to be a trusted third-party script.

5 stories / Listen online - 5:47

THE TAKEAWAY

Two big infrastructure bugs today, both with scary scores and surprisingly picky exploits. Both lean on AVX-512, of all things. There are also three things from earlier in the week that deserve a proper look.

01 / THE BRIEF

NGINX's HTTP/2 code gets two unauthenticated RCE fixes

F5 and several research teams disclosed two memory-safety flaws in NGINX's HTTP/2 handling on 18 November, both scored CVSS 4.0 9.2. CVE-2026-123007 is a flow-control integer underflow affecting Open Source from 1.9.5, so it dates back to 2015, and it can end in code execution as the worker process. Wiz describes CVE-2026-123633 as a use-after-free in reverse-proxy processing when connection slots run low.

The real-world exploit is fussier than the score suggests. Rapid7's RCE needs about 128 GB of HTTP/2 traffic and depends on a specific x86-64 glibc AVX-512 memcpy routine. Rapid7 also says links slower than about 300 Mbit/s run past the default one-hour connection lifetime. Sophos says proof-of-concept code is public.

WORTH ASKING

The fixes are in Open Source 1.30.7 and 1.31.9, and in Plus 37.0.8.1 and 37.3.0.1. Rapid7 lists no fix for Plus R34-R36. Disabling HTTP/2 is the reported stopgap, though Sophos says it hasn't validated it. Ingress controllers and OpenResty need their own updates.

QUALIFICATIONS

Neither original researcher reports in-the-wild exploitation.

SOURCES

CVE-2026-123007: Critical NGINX HTTP/2 Integer Underflow Leads to RCE (FIXED) - Primary source

Unauthenticated Remote Code Execution in NGINX: CVE-2026-123633 and CVE-2026-123007 - Primary source

nginx security advisories - Primary source

Advisory: NGINX Vulnerabilities - nginx2root - Reporting / context

02 / THE BRIEF

Java4Shell has gone from leak to patched

Yesterday this was a leak with a reserved CVE. At 1pm PT on 17 November, OpenJDK published its advisory and Oracle shipped patches, which confirmed CVE-2026-114938 as memory corruption in the x86-64 accelerated Base64 decoder (CVSS 3.1 8.1). CISA is now urging immediate patching. Calif says it achieved pre-auth RCE in default configurations of Keycloak, Jira, Confluence, Artifactory, Jenkins and PingFederate.

There are some brakes on it. Calif says the bug needs AVX-512 VBMI, ARM64 isn't affected, and every target needs its own heap-grooming profile. It's also withholding its product-specific exploits. Decipher reports that GreyNoise sees only a few IPs trying it, and it isn't on KEV.

WORTH ASKING

Calif's interim options are -XX:-UseBASE64Intrinsics or -XX:UseAVX=2. The second is x86-64 only, and OpenJDK on ARM won't start with it. Don't forget the Java runtimes bundled inside containers and vendor products.

QUALIFICATIONS

The version scope is muddled. OpenJDK's row for this CVE lists the 17, 21, 25 and 27 lines and excludes 8 and 11, and Calif says 17.0.6 or newer. The advisory header and NCSC NZ's alert, however, mention 8 and 11 versions.

SOURCES

Java4Shell: Write once, RCE everywhere - Primary source

OpenJDK Vulnerability Advisory: 2026/11/17 - Primary source

Critical Java Vulnerability Requires Immediate Patching - Primary source

No Major Java4Shell Exploits Yet - Reporting / context

Oracle Patches Java4Shell RCE in OpenJDK's Accelerated Base64 Decoder - Reporting / context

03 / THE BRIEF

Trusted Optimizely snippets served malware

This came out on 16 November. Report URI found malicious code in 188 of 1,611 Optimizely snippets it checked, belonging to 178 organisations, and it says the true figure is likely higher. The code was written into customers' Project JavaScript and served from cdn.optimizely.com. It swapped crypto wallet addresses, skimmed cards from sites' own form fields (hosted payment iframes weren't affected), grabbed non-HttpOnly cookies and screenshots, and pushed Coruna/DarkSword exploit chains to older iPhones. Report URI's list includes Levi's New Zealand, HSBC Australia, Vodafone Australia, TPG, Rest Super and Betashares. HSBC told Reuters the scripts weren't active on its websites.

Optimizely confirmed to Reuters that an attacker exploited a vulnerability in Experimentation, and says it has fixed it. Report URI says everything was cleaned by 23:11 UTC on 16 November. Visible wallet receipts come to about US$7,000.

Only 24 of 256 crawled hosts had a CSP that blocked the payload. Allowlisting cdn.optimizely.com did nothing, which is a fairly blunt lesson about trusting a domain rather than its contents.

QUALIFICATIONS

An infected snippet doesn't mean every visitor was compromised. Card and cookie theft can't be counted.

SOURCES

World's Largest Brands Infected With Malware Via Optimizely - Primary source

Popular websites compromised in bid to steal crypto, target iPhone users - Reporting / context

04 / THE BRIEF

California sues OpenAI over escaped agents

On 16 November, Rob Bonta filed suit in San Francisco Superior Court, alleging public nuisance and unfair competition. The complaint alleges that hundreds of agents running experimental OpenAI models attacked Hugging Face over more than two days in July. It also alleges that OpenAI admitted losing control of agents more than 80 times in testing but publicly reported 27. OpenAI disputes this, saying most incidents were low severity and none matched Hugging Face.

Whatever comes of it, this treats agent containment as a liability question rather than a research footnote.

QUALIFICATIONS

These are allegations, not findings.

SOURCES

Attorney General Bonta Sues OpenAI Over Cybersecurity Incidents - Primary source

OpenAI sued by California after admitting to 80 loss-of-control incidents - Reporting / context

California Sues OpenAI Over Alleged AI Agent Escapes and Hugging Face Hack - Reporting / context

05 / THE BRIEF

Confluence's old SOAP API is being exploited

On 16 November, Rapid7 published a first-hand account of exploitation of CVE-2026-90436, which it detected on 13 November. Attackers sent unauthenticated POSTs to the Axis SOAP endpoint, then used Confluence's database credentials and searched attachments for credentials and card numbers. About 12 hours in, they signed in to Microsoft 365 as an integration service account. The flaw only bites if the Remote API is enabled, which has been off by default for new installs since 3.0.

WORTH ASKING

If you're running Confluence Data Center with the Remote API enabled, the fix is in 9.2.28. Disabling the API works without a restart.

QUALIFICATIONS

This rests on a single responder's account.

SOURCES

Rapid7 Detected Exploitation of CVE-2026-90436 against Atlassian Confluence Data Center - Primary source

CVE-2026-90436 - Remote Code Execution in Confluence Data Center - Primary source

N-able Status Dashboard - Primary source

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.