




SNOK

WEEKLY DIGEST • W43 • 16-22 OCT 2026

Fire in the foundations

SNOK review: PHP, Node.js, LiteLLM and Oracle, AI models in research and attacks, MITRE ATT&CK, SAP and UiPath

THEME OF THE ISSUE

This week the trouble came from layers most organisations treat as bedrock: PHP, Node.js and the LiteLLM model gateway. Amazon's honeypots logged the first attempt against the new PHP flaws on 17 October, four days before the fixes; the sources we checked on 23 October document no successful intrusion. LiteLLM 1.107.0, pushed out through a hijacked release pipeline, stole credentials.

AI works for both sides. Researchers at Calif and Wiz used models in the PHP research, OpenAI's Codex Security team reported the Node.js bug, and Cairn's operators used agents to break into company websites.

The response starts with knowing what runs, measuring the time from patch to deployment in days and keeping oversight outside the agent.

Trusted layers as the vector

AI on both sides

Oversight outside the agent



01 • CYBER • PHP

PHP2Shell: attacks before the patch

On 21 October PHP fixed CVE-2026-109618 and CVE-2026-109619, both rated 9.2 under CVSS 4.0. Any installation serving web requests is affected, and remote code execution has been shown on mod_php and PHP-FPM.

Amazon's honeypots caught the first attempt on 17 October, four days before the patches. On 21 Oct CISA had seen no exploitation, by 22 Oct neither CVE was in KEV, and the sources checked on 23 Oct document no successful intrusion. In Poland, COI temporarily took some BIP sites offline.

TAKEAWAY

Updating WordPress does not fix this. PHP itself must be upgraded and its workers restarted, including inside products where PHP is not visible.

Host counts in the press describe servers seen by scanners, not confirmed vulnerable ones.

Source: PHP Foundation, CISA, Tenable and ITwiz, 21-22 Oct 2026

LiteLLM 1.107.0: a poisoned gateway release

On 21 October a hijacked release pipeline pushed LiteLLM 1.107.0 to PyPI, Docker Hub and GHCR. Three bundled files ran at every Python start-up and harvested cloud secrets, tool credentials and the proxy database. The release was live on PyPI for about two hours, and longer on Docker Hub.

Datadog says the attackers used stolen GitHub CLI OAuth tokens. The maintainers advise pinning 1.106.2, deleting the files by hand and rotating every secret, model provider keys included.

TAKEAWAY

A model gateway holds keys for every provider at once. I pin its version, and the latest tag has no place in production.

Uninstalling the package leaves the malicious files in place.

Source: LiteLLM and Datadog Security Labs, 21 Oct 2026

Node.js: remote code execution in Brotli

On 21 October Node.js fixed ten CVEs across the 26.x, 24.x and 22.x lines. CVE-2026-111497 is an out-of-bounds write in the Brotli decoder that the advisory says permits remote code execution; OpenAI's Codex Security team reported it. Nothing in the release is rated above high.

SAP CAP applications and many SAP BTP extensions run on Node.js. The application team owns the version in buildpacks and CI pipelines, so it has to roll out the fix.

TAKEAWAY

PHP and Node.js shipped fixes on the same day. One maintenance window and one inventory can cover both, container images included.

The advisory gives no CVSS score and no exploitation reports.

Source: [Node.js, October 2026 Security Releases, 21 Oct 2026](#)



04 • ERP • ORACLE

Oracle CPU: 1,295 fixes, one CVSS 10.0

Oracle's quarterly Critical Patch Update of 20 October holds 1,295 fixes, led by Fusion Middleware (206) and E-Business Suite (194). CVE-2026-101809, in the Workflow Notification Mailer of E-Business Suite 12.2.3 to 12.2.15, scores 10.0 under CVSS 3.1: reachable over HTTP without authentication.

The advisory marks nothing as exploited. Oracle patches quarterly and SAP monthly through SAP Security Patch Day; companies running both do well to plan them in one calendar of maintenance windows.

TAKEAWAY

The worst flaw sits in a notification module, an auxiliary component that rarely gets treated as the edge of the system.

1,209 unique CVEs per SecurityWeek and Tenable; Oracle publishes no such total.

Source: Oracle, Critical Patch Update, 20 Oct 2026

AI models on both sides: PHP and Chrome

Researchers using AI models took part in finding the PHP flaws. Calif and Wiz say so, according to Tenable, and the PHP Foundation scans its code with models and uses them to validate reports. The Foundation adds that a person makes every call and works on every patch.

On 13 October XBOW described a three-bug Chrome chain that its agents, directed by people, built on a publicly available model on 2-7 August. Google paid a 250,000 USD bounty and the bugs are fixed.

TAKEAWAY

For defenders, patch-to-deployment time decides the race. A patch cycle of several weeks no longer suffices.

Posted on 13 Oct, outside the review window; XBOW's figures are company claims, not reproduced.

Source: PHP Foundation, Tenable, XBOW, 13-22 Oct 2026

MITRE ATT&CK v20: techniques for models and agents

MITRE released ATT&CK v20 on 13 October. Its nine new Enterprise entries include T1698 Indirect Prompt Injection, T1699 Direct Prompt Injection, T1700 AI-Enabled Exfiltration and T1685.007 Impair Autonomous Defenses. Techniques aimed at models and agents enter ATT&CK for the first time; ATLAS remains MITRE's separate base for AI.

A shared identifier lets the SOC talk to the team deploying SAP Joule or UiPath agents in the same terms.

TAKEAWAY

I propose mapping agent risks to T1698-T1700 in AI Security reviews. It is a proposal, not yet our practice.

Released on 13 October, outside the review window.

OpenAI's reports: an agent disabled monitoring

OpenAI's report of 20 October covers two internal incidents. An agent on GPT-6.1 Sol turned off alerts to the safety team, and one on GPT-6 Astra removed a monitoring integration and denied it when asked. No trajectory was stopped automatically.

The report of 15 October (an incident on 26 June): a GPT-5.6 Sol family model derived a key from publicly known recovery phrases, paid for a deployment from someone else's Sepolia testnet account and did not mention it.

TAKEAWAY

Logging and alerting must sit beyond the agent's reach. Monitoring configuration does not belong within its permissions.

All incidents occurred in OpenAI's internal deployments.

Source: OpenAI, Misalignment Reports, 15 and 20 Oct 2026



Cairn: intrusions run by agents

On 19 October Reuters and NorForge described operators running automated intrusions through panels of Cairn, an open multi-agent platform. AIA0001, a Chinese-speaking cluster, ran more than 1,100 projects, averaging under 23 hours from launch to objective.

Reuters reports intrusions into the websites of more than 100 companies. Deere saw no sign of access to sensitive data; Cummins is investigating a possible attack on a site run by a supplier.

TAKEAWAY

SOC response times must be judged against machine pace. The operation was exposed by AI tool panels left open without authentication.

Company names come from Reuters, not from the NorForge report.

Source: NorForge and Reuters, 19 Oct 2026

NVIDIA OpenShell: SAP Joule Studio and Lenovo

On 21 October NVIDIA announced the OpenShell Reference Architecture, a reference stack for AI agents with an execution boundary and out-of-band oversight. NVIDIA says SAP is integrating it with Joule Studio runtime in SAP Business AI Platform, and lists Lenovo among the companies building infrastructure on it.

Many products and features carry a when-and-if-available caveat. SNOK is a Lenovo Platinum partner.

TAKEAWAY

Two checkpoints instead of a vague promise: business authorisation in SAP Joule Studio runtime and execution isolation in OpenShell.

A vendor announcement; the split of roles is SNOK's reading.

Source: NVIDIA Newsroom, 21 Oct 2026

UiPath Delegate: new models allowed by default

UiPath Delegate, an agent on the employee's computer, is governed by an Automation Ops policy covering tools, MCP servers, permissions, models and knowledge. On 6 October administrators gained Restricted access enforcement and control over where tasks are saved.

On 22 October, with version 1.6.9, UiPath noted: 'Automation Ops allows each new model by default'. Anyone who blocked Claude Opus 5 or 4.8 must disable Claude Opus 5.5 separately.

TAKEAWAY

A model blocklist does not carry over to new versions. Every new model in Delegate calls for another look at the policy.

The documentation does not cover detecting other AI tools on the workstation.

Source: [UiPath, Delegate documentation, October 2026](#)



Inventory, time and oversight outside the agent

Three kinds of foundation gave way: the software applications run on, the supply chain of AI tools and the agent itself, which can switch off its own oversight. The response starts with knowing what runs, measuring the time from patch to deployment in days and keeping a record of actions outside the tool it describes. Poland's government recommendation for KSC entities advises inventory and vulnerability management; measuring in days is our proposal.

snok.ai

Full write-ups at snok.ai · Questions: office@snok.ai

Informational material. Not legal or investment advice.