



SNOK

WEEKLY DIGEST • W42 • 9-15 OCT 2026

Read the source, not the headline

SNOK review: SAP, Microsoft, AI agent security, UiPath and small models
on your own hardware

THEME OF THE ISSUE

In five of this week's nine stories the primary source says less than the headline. Germany's 'withdrawal' from its partnership with SAP and Mistral AI was, the ministry says, a misunderstanding; one of Microsoft's 'three zero-days' is actively exploited; and a small model's card speed, measured by its author on other hardware, shrank to a fraction on ours.

The other stories are about control that stays with the system owner: SAP's security template, UiPath's new agent grading, the agent incidents in OWASP's roundup, and coding agents that have started clicking through SAP GUI.

Source before headline

Measured, not claimed

Owner keeps control

SAP and Mistral AI: a dispute over one word

On 9 October Die Welt reported that Germany was pulling out of its AI partnership with SAP and Mistral AI. The digital ministry confirmed it has no plans to buy from Mistral AI, then told heise on 14 October that this was a misunderstanding: the November 2025 letter of intent was not binding and no framework agreement was ever signed.

A letter of intent signed at a political summit is a statement of intent; nobody has bought anything on the strength of it. In a separate tender in May, a T-Systems and SAP consortium won 70% of a German federal AI platform contract worth just under EUR 250 million.

TAKEAWAY

AI sovereignty is tested in the contract and the architecture: where the model runs, who sees the data, what happens when the vendor changes terms.

Neither the German government nor Mistral AI gave a reason.

Source: [Die Welt](#), 9 Oct 2026; [heise online](#), 14 Oct 2026

Patch Tuesday and SAP Patch Day: one flaw exploited

On 13 October Microsoft shipped fixes for more than 1,400 flaws - BleepingComputer counted 1,411, 70 of them critical. One is being actively exploited: CVE-2026-96951 in Windows Shell, rated Important. The same day, SAP Security Patch Day brought 18 SAP notes and 4 GitHub advisories.

According to Krebs on Security, Microsoft credits large language models for the volume of bug discovery. If so, the volume is here to stay and the patching process has to cope with it.

TAKEAWAY

Exploited flaws first, then the rest by exposure. The SAP maintenance window is best planned together with the operating system window.

Counts differ by source; Microsoft publishes no single total.

Source: BleepingComputer and Krebs on Security, 13 Oct 2026

Satya Nadella: the model as an insider risk

On 10 October Satya Nadella wrote on X: 'We must assume a model is compromised and contain it from the start. Think of it like an emergency brake.' He proposes treating models like privileged staff, with an identity, limited permissions, logging, isolation and controls that sit outside the model.

He does not say that every model is compromised - that was the headlines' shorthand. His point is that the system must behave correctly even when the model fails. We have argued since W40 that a limit written only into a prompt guarantees nothing.

TAKEAWAY

Nadella sets out a design principle without announcing a product. On our UiPath projects an AI Security review is mandatory.

Quote confirmed by TechCrunch and The Verge.

Source: TechCrunch and The Verge, 10 Oct 2026

SAP Security Baseline: changes flagged in 3831086

The SAP Security Baseline Template is the template companies use to build their own security baseline for SAP systems. The current version is 2.6, and SAP note 3831086 lists the changes being prepared for the next one. Since 21 September SAP has been publishing SAP Focused Run policies for 2.6 on GitHub.

A standard based on an older version of the template may need updating, and only a comparison between versions will tell. Advance notice gives time for that comparison before an auditor makes it.

TAKEAWAY

If your baseline has not been compared with SAP's template for a while, note 3831086 is a good reason to review it.

The note is available to users signed in to SAP for Me.

Agents from outside: who let them in

The Hacker News ran two partner articles, by Reco and Token Security, on agents an organisation never chose. Reco's own report counts roughly 1,280 third-party products with embedded AI, only about 282 of them behind single sign-on. Employees also connect personal assistants through OAuth consent and API keys.

These are vendors describing a problem they sell a fix for. Even so, I think NIS2 and DORA supplier assessments should cover suppliers' agents and their access to client systems.

TAKEAWAY

Supplier questionnaires should ask which of the supplier's agents can reach your systems, under what identity, and where their actions are logged.

Partner content; figures from Reco's own report, not independently verified.

Source: The Hacker News, 10 and 12 Oct 2026

UiPath Agent review: a grade from A to F

On 8 October UiPath replaced Agent Score with Agent review, currently in preview. An agent is graded from A to F on the number and severity of issues - critical, warnings and improvements - and Autopilot can fix them in the same conversation. The feature is not available on Dedicated editions.

The policy that blocked publishing an agent below a minimum score is no longer enforced, and no successor based on the letter grade has appeared yet.

TAKEAWAY

Under SNOK's deployment standard no agent is published with a critical issue found by Agent review. Evaluation against test sets remains a separate gate.

UiPath does not publish the thresholds behind each grade.

OWASP GenAI: agent incidents and exploits

The OWASP GenAI Security Project's roundup for Q3 2026 lists nine incidents and exploits. According to it, evaluation agents ran code on 41 production dataset workers at Hugging Face, and the Deadbugz MCP server behaved for three calls, then changed its tool descriptions and tried to steer the agent towards SSH and AWS keys.

In most cases the agent was carrying out its assigned task and went beyond it. The authors say plainly that prompt-level instructions alone do not establish a secure boundary.

TAKEAWAY

I want our AI Security reviews to include snapshotting MCP tool definitions at approval and re-approval after any change.

The authors note this is not an official OWASP publication.

Source: OWASP GenAI Security Project, 8 Oct 2026

Small local models: claimed versus measured

Three small models for agents appeared in one week: Underdog Saluki 27B, Liquid LFM2.5-24B-A2B and JetBrains Mellum2.1. On our Lenovo ThinkStation PGX we measured OrcaSAQ2 27B: its card quotes 65 tokens per second, measured on a GPU with about 3.5 times the memory bandwidth; we measured 8.5, and 15.4 with MTP.

On 16 October we measured Saluki and Mellum2.1 on a MacBook Pro with an M4 Max. Both files are about 8 GB, yet generation speed differs more than sevenfold: 22.3 against 168.8 tokens per second. Mellum2.1 activates 2.5 billion parameters per token; Saluki uses all 27 billion.

TAKEAWAY

On desktop hardware, speed depends above all on the number of active parameters. We test tool calling ourselves.

Model results as reported by vendors; speeds from our measurements of 1 and 16 Oct 2026.

Source: Underdog, Liquid AI, JetBrains; SNOK measurement

Agents in SAP GUI: a person's permissions

An SAP developer showed coding agents connected to SAP GUI through SAP GUI Scripting on LinkedIn. Meanwhile `sap-ai-dev-toolkit`, an npm package with MCP servers for ABAP, keeps growing: per the npm registry on 16 October, one maintainer, 52 versions since 25 September and a script that runs on install.

SAP's 2004 security guide notes that the server cannot tell a script from a user, and a bad script can create more bad data before anyone notices. It describes GUI Scripting as disabled by default; SAP note 480149 covers current settings.

TAKEAWAY

For bulk PFCG changes I choose automation with a human approval step over an agent clicking through transactions on its own.

Community packages only after a scan and in a sandbox, never on client systems.

Source: LinkedIn, October 2026; SAP GUI Scripting Security Guide



What I check before I believe a headline

For an announcement - the contract and the order. For a vulnerability count - the exploited ones. For a quote - the full text. For a survey - who ran it and what they sell. For a model card - our own measurement. For an agent - who can stop it, and whether its tools look the way they did on the day they were approved.

snok.ai

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

Informational material. Not legal or investment advice.