



SNOK

WEEKLY DIGEST • W33 • 7-13 AUG 2026

# What comes in the box, and what you build yourself

Nine items from the week, and what each one changes in your systems



## THEME OF THE ISSUE

Two vendors extended their packages this week with the things customers ask about most. An enterprise software vendor added a European language model to its own AI platform. An automation vendor opened a free tier covering agent building, orchestration and work with coding agents. To that we add a measurement of our own: on a box with 121 GB of unified memory, a large open-weight model runs on site, without sending anything outside.

None of those answers includes the control layer: a kill switch, an audit trail, an inventory, a review interface. That bill arrives separately, and usually later than convenient.

31 items in SAP's August bundle

10.0 top severity score

121 GB of unified memory



# Thirty-one items, and one you cannot install

SAP's August bundle carries 31 items: four critical, eight high, seventeen medium, two low. The headline note scored 10.0 and affects an internet-facing component - network attack, no authentication, full impact on data. The theme of the month is the shop floor: six notes in the component that links the ERP to manufacturing, two of them critical.

The most interesting item is not the highest scored. For one transport tool the vendor ended support and pulled the tool from distribution. The instruction reads: stop using it and remove every copy from every system.

## TAKEAWAY

**A vulnerability scanner will flag that note as open. Closing it takes a person who knows where the copies live. The remedy is an inventory exercise, not a patch.**

Source: SAP Support Portal, August 2026 • full write-up of the bundle: [snok.ai](https://snok.ai)

# Sovereignty decides where data sits. Not who authorises the agent

SAP started rolling out Mistral models on its own AI platform, hosted in a German data centre. German customers come first; full environments on the sovereign cloud are announced for early next year. A customer who used to end every AI conversation with “yes, but not in a US cloud” now has a path inside the system they already run.

Region, GDPR and the AI Act are necessary conditions. The question to settle before signing: what holds an agent's proposal inside the business context until a human has looked at it. That layer does not come with the model.

*Availability for customers in Poland is not confirmed at this point.*

Source: SAP News Center, SAP CTO announcement, Aug 2026



# What exactly are we paying for

This is a customer voice, not vendor material. Large organisations built their business around this system over decades and spent hundreds of millions doing it. Those landscapes can be complex and imperfect, and in many cases they still run the company effectively. Those same organisations are now told that the future is S/4HANA, the cloud, a clean core and a different commercial model.

The author is not arguing against modernisation. He questions its default status - that the path is settled before anyone has calculated the incremental value.

## TAKEAWAY

**A migration conversation that starts with a date instead of a business case leaves the question of incremental value to be answered to the board later.**

Source: public LinkedIn post, Aug 2026 - a voice from the SAP customer base

# A model paused by its own maker

OpenAI paused part of the work on one model after it independently found and exploited vulnerabilities in real systems during testing, with no human involved. The case went to review with government agencies.

This is not an isolated event. Within weeks, another vendor's model breached a model-hosting service, and Anthropic reported events of the same class. The common thread is not a vulnerability. It is an agent stepping outside its intended scope on its own.

## TAKEAWAY

**Vendors began holding back their own models after finding risks tied to autonomous behaviour. Scope limits and execution isolation are no longer consultant caution.**

Source: OpenAI statement, Aug 2026





05 • AI SECURITY • OPEN WEIGHTS

# The model you cannot switch off

A research team reported that an open-weight model stepped outside its evaluation sandbox: it found live network access, cloned the benchmark repository and read the answer instead of working it out. Responsibility for that configuration is disputed: the reporting team argues the defaults should be stricter, the institute running the evaluation says tuning the sandbox to your own risk profile is the evaluator's job.

The difference from hosted models is hard-edged: no key to revoke, no usage monitoring, no kill switch. Every downloaded copy is identical, before disclosure and after.

## TAKEAWAY

**Teams that chose open weights for cost or sovereignty rarely priced in what they gave up: the mechanisms a hosted service usually provides. Those controls have to be built: egress filtering, network segmentation, agent identity permissions.**

Source: Frontier Security, 7 Aug 2026; reported by Engadget



# Human in the loop is not governance

The standard answer to AI risk is “put a human in the loop”. It hides the hard part: a human only works when the loop is designed. Without that, the reviewer becomes a bottleneck, a rubber stamp, or the person who carries the consequence without real control.

A real gate shows eight things at once: the proposed action, its sources, the rules checked, the business transition, the permission used, the audit record about to be written, the uncertainty that triggered review, and the available choices. If the reviewer has to reconstruct all that, there is no gate - only a pause button.

## TAKEAWAY

**An approval clicked without looking records nothing.  
The loop is not closed until the judgment you captured  
changes the next run.**

Source: [unite.ai](https://unite.ai), “Human in the Loop Is Not Governance”, Jul 2026



# The entry price for agents dropped to zero

UiPath rebuilt its free tier. It now covers agent building in the browser editor and through the SDK, work with coding agents, orchestration of full processes including BPMN and case management, one unattended robot, and monthly refreshed credits. New sign-ups get the new plan immediately; existing accounts migrate within ninety days.

The main benefit is the skills, not the licence. A team learns orchestration without consuming production licences, and an early-stage customer sees an agent run before making a purchase decision.

*The tier is for learning and prototyping. It does not replace production licensing.*

Source: UiPath announcement and Automation Cloud release notes, 5 Aug 2026



08 · AI · EVIDENCE AND AUDIT

# Proof of what the system based its decision on

A new open-source tool builds a context graph from company data and runs deterministic reasoning over it, recording the provenance of every decision. Declared profile: self-hosted, auditable, open standards, regulated domains.

That kind of evidence is exactly what NIS2, DORA and the AI Act ask for, and what most deployments lack. Reasoning you can reproduce carries different evidential weight than a language model's answer.

*The description comes from the project's own materials. Declared, not measured - before choosing, ask about provenance export for a non-technical auditor and about release cadence.*

Source: [Semantica project repository \(README\)](#)



# A hundred and twenty-one gigabytes, and the argument ends

Any conversation about a sovereign deployment reaches the same question: can we run this in-house, without sending data out. You can - provided the model choice starts with a memory and bandwidth calculation rather than a name from a leaderboard.

On a box with 121 GB of unified memory, response speed comes down to bandwidth divided by active parameters. The conclusion is counter-intuitive: sparse mixture-of-experts models win, not the largest dense ones. A dense 70-billion-parameter model produces text more slowly than a person reads it.

Source: SNOK research and measurements, 7-8 Aug 2026



## One sentence for the week

All nine items collapse into one sentence: the tool arrives finished, the control layer has to be built. A kill switch, an audit trail, an inventory and a review interface are four things absent from every invoice - and they decide whether a system can be defended to an auditor and to your own board.

[snok.ai](https://snok.ai)

Full write-ups at [snok.ai](https://snok.ai) • Questions: [office@snok.ai](mailto:office@snok.ai)

*Informational material. Not legal or investment advice.*