




SNOK

WEEKLY DIGEST • W41 • 2-8 OCT 2026

Full throttle, brakes still on order

SNOK's weekly review: AI agents in SAP, UiPath and Snowflake, and the controls racing to keep up

THEME OF THE ISSUE

This week agents reached company systems from several directions at once. UiPath connected to Snowflake without copying data, SAP bought a work graph for its HR agents and published the official route for non-SAP agents, and SAP researchers cast a tabular model as the decision layer.

The controls arrived in the same week, and they are a step behind. The Jev decision model gives away hidden personal data through its choice of answer, Apple announced stricter consent for full disk access because of agents, and Gartner forecast that most vendor-built agents will be abandoned. The brakes can be fitted - if someone thinks about them before the car leaves the test cell.

Deployment pace

Controls under construction

Capability stays in-house



01 • UIPATH • SNOWFLAKE

UiPath and Snowflake: the data stays put

On 30 September UiPath announced a two-way integration with Snowflake. Data Fabric reads data directly in Snowflake without copying it, Cartographer and Delegate reach Snowflake Cortex, and Snowflake CoCo users can trigger UiPath agentic skills from inside the data platform. UiPath Maestro orchestrates the whole flow.

The data stays where its permissions and lineage live, and the record of what the agent did stays in UiPath. But the direction has reversed: automations can now be triggered from the data platform, so a security review has to cover both directions.

TAKEAWAY

Fewer copies of data means less risk. UiPath skills called from Snowflake CoCo still need a permissions review before go-live.

Availability of each component and supported regions to be confirmed before deployment.

Source: [UiPath press release \(Business Wire\)](#), 30 Sep 2026

Half of organisations run agents in production

In June and July UiPath surveyed 1,226 people who build automation and AI. 53% report agents in production and another 23% are piloting. The most cited barrier is data security or privacy (40%), ahead of cost (39%) and difficulty proving ROI (31%).

Among organisations using agents broadly, 72% have fully documented and enforced governance; among those still evaluating, 36%. Broad agent use goes hand in hand with full governance - the survey cannot say which came first.

TAKEAWAY

Among broad agent users, twice as many organisations report full governance as among those still evaluating agents.

Vendor survey: recruited mainly via UiPath channels, 54% of respondents from Asia-Pacific.

Source: UiPath, 2026 State of AI and Automation Professionals

SAP to acquire TechWolf: HR agents need a graph

On 6 October SAP announced an agreement to acquire TechWolf of Ghent. It is buying a work graph - tasks inside jobs, people's skills and the labour market, built from existing HR systems - to become a core of SAP SuccessFactors and the layer on which SAP Joule bases HR decisions. Closing is expected in Q4.

SAP says plainly that the graph cuts token use and the cost of deploying HR agents. Yet a graph of what employees do and can do is personal data about them. Who sees it through SAP Joule, and with what authorisations, must be settled before products launch.

TAKEAWAY

An agent's cost and quality depend on the context underneath it. In HR that context is personal data, so authorisations are designed from day one.

Financial terms not disclosed. New products are planned only after closing.

Source: SAP News Center, 6 Oct 2026



The official route for non-SAP agents

In its reference architecture "A2A and MCP for Interoperability", SAP Architecture Center sets out three routes: Agent Gateway using A2A for traffic to SAP Joule agents, MCP Gateway in SAP Integration Suite for exposing APIs as tools, and A2A through SAP Integration Suite, which handles authentication, rate limits and monitoring.

At the top of the page SAP warns that some components are not yet generally available, and that bidirectional communication through Agent Gateway with third-party agents "is not yet supported". Bring Your Own Agent handles text only, with a 60-second synchronous limit.

TAKEAWAY

The map gives every architecture a reference point. A design for an agent that works with SAP Joule still has to rest on what is available now.

The page gives no general availability date for the missing components.

Source: SAP Architecture Center, updated 27 Aug 2026

Jev versus SAP RPT: who makes the small calls

SAP researchers have compared the Jev decision model with SAP RPT-1.6, SAP's tabular foundation model, on SAP Community. According to the authors, RPT-1.6 scores 76.75% on typed-decision tasks against 72.70% for Jev 1.13.0. On ticket assignment the gap is 68.9% against 49.4%, with RPT-1.6 taking 44 ms and Jev 322 ms.

RPT's latency was measured directly on an H200 GPU without a production deployment, and the ticket benchmark cannot be reproduced outside SAP. The direction matters more: the decision layer can run inside the SAP platform rather than next to it.

TAKEAWAY

SAP customers face a choice: decisions inside the platform or beside it. Measurement on their own data should settle it, not a vendor benchmark.

Figures as reported by the SAP authors. Posted 5 Oct 2026.

Source: SAP Community, Technology Blog by SAP, Oct 2026

HoneySAP 0.2.0: a decoy for a 9.9 flaw

On 6 October OWASP released HoneySAP 0.2.0, a honeypot that emulates SAP services. The new trap impersonates an SAP RFC Gateway and captures credentials, and it treats calls to /SLOAE/DEPLOY as attempts to exploit CVE-2025-42957 - ABAP code injection in SAP S/4HANA, rated CVSS 9.9 and patched in August 2025.

Nobody calls that module from outside in normal operation, so any call the decoy records points to an attempt to exploit a known flaw. Patching and hardening do not tell you whether someone is trying right now. A honeypot does, and cheaply.

TAKEAWAY

Patching closes the flaw; a honeypot shows that someone is looking for it. Those are two different pieces of information, and a SOC needs both.

A community tool without vendor support. We have not tested it.

Source: OWASP HoneySAP v0.2.0, GitHub, 6 Oct 2026

Jev's decisions give away hidden data

Researchers tested Jev as an application's decision layer. An attacker who sees only the decision used query pairs with the condition reversed to recover gender, age, heart disease and ethnicity from hidden state - across all 200 test states, according to the authors. Hiding probabilities does not help.

Injects written for Jev flipped its decision in at least 97% of cases, a local model could be backdoored through poisoned data, and 200 answers were enough to train a surrogate that matched it 92% of the time.

TAKEAWAY

A typed answer constrains the format, not the leakage. Validate the decision against trusted data before acting, and keep sensitive fields to the minimum.

Preprint, not peer reviewed; figures as reported by the authors.

Source: S. Wang et al., [arXiv:2610.04985](https://arxiv.org/abs/2610.04985), 4 Oct 2026

Anthropic opens Red Team Access

On 6 October Anthropic expanded its Cyber Verification Program to three access tiers. Red Team Access lifts Claude's safeguards for authorised penetration testing and red teaming, for organisations only, with a review of a few weeks. Specialized Access serves those testing critical systems.

One condition changes the calculation: data retention is mandatory, with a narrow exception for zero-retention access to Claude Fable 5.1 or Mythos 5.1. On client engagements under NDA and within ISO 27001, sending test data to an outside provider needs a security assessment and the client's consent.

TAKEAWAY

A model without safeguards inside an authorised scope has a price: test data stays with the provider. The security officer should assess it first.

SNOK has not applied to the programme.

Source: Anthropic, 6 Oct 2026

OWASP FinBot: a range for attacking agents

FinBot is an intentionally vulnerable OWASP agentic platform billed as "the Juice Shop for Agentic AI". The task is to break a multi-agent vendor management system through prompt injection, tool misuse, data exfiltration or remote code execution. The challenges map to the OWASP Top 10 lists, CWE and MITRE ATLAS.

The platform launched in the spring and is back in an OWASP campaign. It gives teams building agents a place to rehearse an attack without harm, before anyone tries it on a live system.

TAKEAWAY

A team that has never tried to break its own agent will learn about its weaknesses from someone else. Better to practise on the range.

Runs in the browser; code under Apache 2.0.

Source: [OWASP GenAI Security Project](#), [GitHub](#)

Apple: agents change the risk of disk access

On 2 October Apple published a note on changes to Full Disk Access in macOS and stated the reason plainly: "As AI agents become increasingly capable and autonomous, the risks associated with this level of access will grow substantially." Full disk access is to require very explicit user action.

An agent with access to the whole disk can see mail, documents and keys. Apple is not banning that access; it is making it a deliberate decision. Companies should do the same: agent permissions are a decision with an owner.

TAKEAWAY

An operating system vendor now treats agents as a distinct class of risk. Company agent permissions deserve the same caution.

The note gives no macOS version or timeline.

Source: [Apple Developer News, 2 Oct 2026](#)

HyperFrames Studio: the agent edits, you decide

In early October HeyGen released HyperFrames Studio, a macOS and Linux app described as a video editor built for agents. Claude Code or Codex builds a first cut, then human and agent share one timeline: hand edits, notes drawn on a frame, comments on a stretch of footage.

The agent produces a first version at a pace no person can match, and the human keeps the decisions - the same division of roles we build into process automation. The app requires a HeyGen account.

TAKEAWAY

With company material, the first question is whether projects and prompts reach the vendor's servers. The documentation does not say.

A Windows version is announced. The app's own licence is not stated.

Source: [HyperFrames Studio documentation](#), HeyGen

OpenAI loses its safety people

On 1 October reports citing the WSJ said OpenAI had dismissed three researchers for breaching rules on confidential information; they disputed the claims. On 3 October David Robinson, who prepared safety reports for launches, resigned and wrote in *The Atlantic* that the company's culture "is broken". A class action alleging collusion to slow AI is pending.

The labs are talking about slowing down, being sued for slowing down and losing their safety staff, all at once. Agent security in a company has to rest on what you control yourselves.

TAKEAWAY

Permissions, process, logs and a pre-launch review are brakes in your own hands. A model provider's assurances are not.

Buist v. Anthropic PBC: allegations by the plaintiffs; the case is undecided.

Source: TechCrunch, 3 Oct 2026; Bloomberg Law, 18 Sep 2026



13 • MARKET • GARTNER

70% will abandon vendor-built agents

Gartner predicts that by 2028, 70% of enterprises will abandon agentic AI built by vendors' forward-deployed engineers. Gartner cites rising costs and the client's inability to develop such systems further on its own.

SNOK is an implementation firm, so we have an interest here - which is exactly why we take the forecast seriously. An agent the client cannot maintain and change alone is a debt that grows with every process change.

TAKEAWAY

Knowledge transfer and documentation belong in the implementation scope from day one. The capability has to stay with the client.

Gartner press release, 29 Sep 2026.

Source: [The Register](#), 30 Sep 2026

Mistral Large 4 in Polish: sovereign is not fluent

On 6 October Mistral released Mistral Large 4 in preview: one trillion parameters, 52 billion active, under the line "Forged in Europe. Built for AI sovereignty." On PLCC, a Polish benchmark, it averages 68.17. Bielik-11B-v3.0-Instruct scores 70.67, PLLuM-12B 69.67, and Mistral's previous version also 70.67.

A European model is not automatically a good model for a Polish client. On this test a small Polish model beats the European flagship and matches its predecessor.

TAKEAWAY

For an on-premises deployment, what counts is measurement on your own Polish tasks. The flag on the press release settles nothing.

The lead holds for the best Bielik and PLLuM variants. Averages computed from leaderboard data.

Source: Mistral AI, 6 Oct 2026; PLCC leaderboard, Hugging Face



Brakes you can fit today

Agent permissions are a decision with an owner. A model's decision is checked against trusted data before acting. Security reviews cover both directions of an integration. The capability to maintain an agent stays with the client. Models are chosen by measurement on your own data.

snok.ai

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

Informational material. Not legal or investment advice.