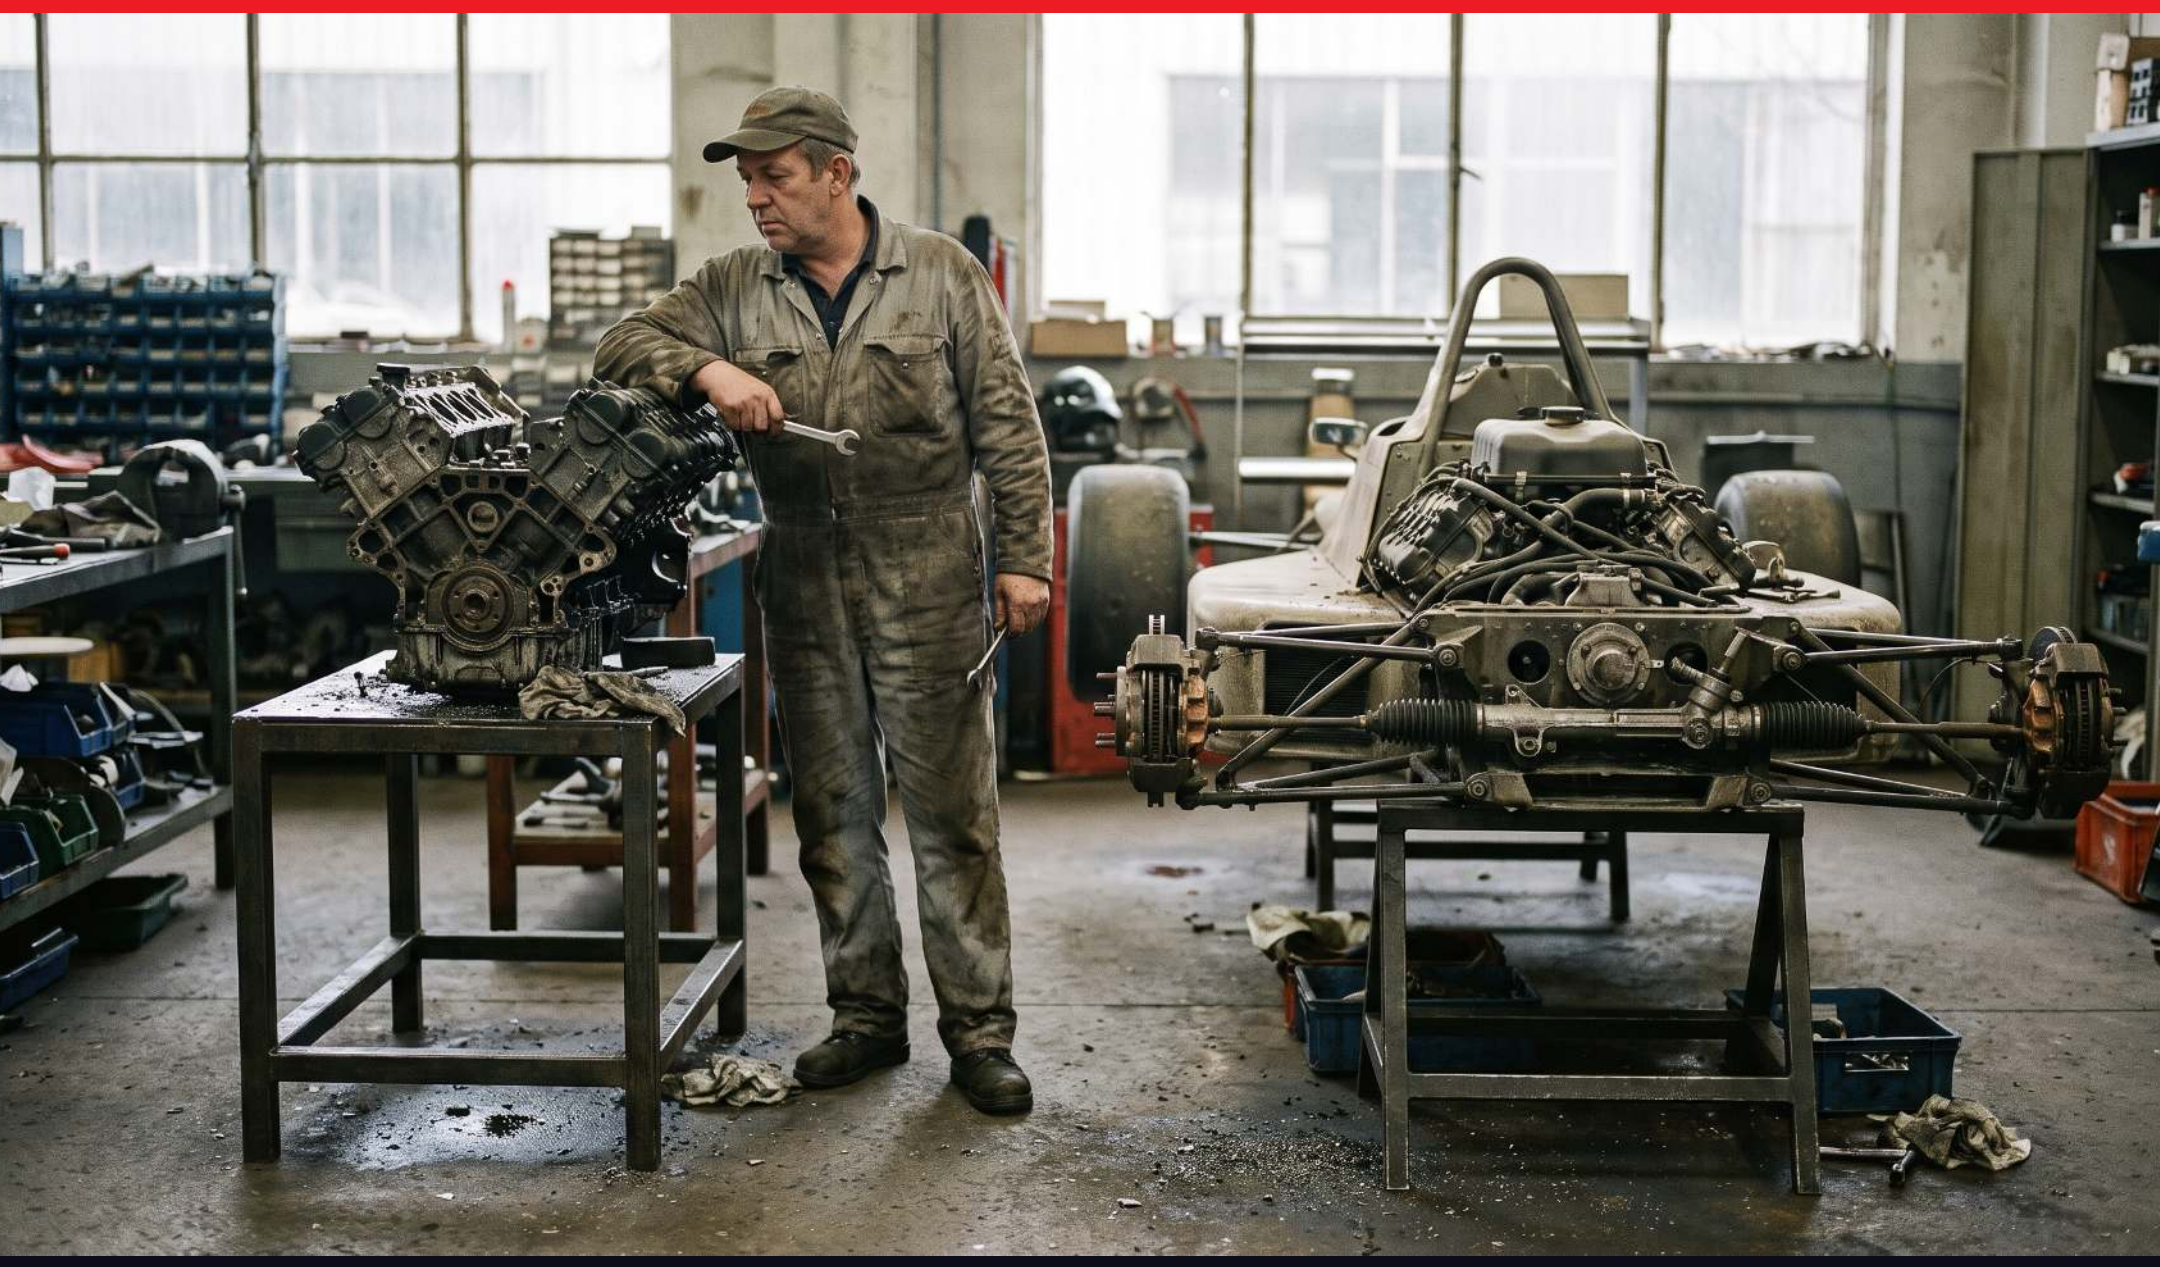

SNOK

WEEKLY DIGEST · W37 · 11 AUG - 10 SEP 2026

The core does not decide What surrounds it does

Fifteen items from four weeks, each with a note on what it changes in your systems

Swipe to read >>>

SNOK

THEME OF THIS ISSUE

Four weeks without an issue produced a pool with one recurring pattern. A model on its own is not a good penetration tester - the structure around it decides. A security note does not protect a system until somebody owns its rollout across four layers at once. A scanner misses a malicious agent skill because it reads code, while the instruction sits in prose.

The same rule governs the invoice. An AI server gets more expensive because of memory, not compute. The largest funding round in European AI history is going into data centres, not into another model. The thing you buy is almost never the thing that decides the outcome.

20 items in SAP Security Patch Day

10 out of 10 manipulated summaries

61 tools in the MCP server for SAP BTP

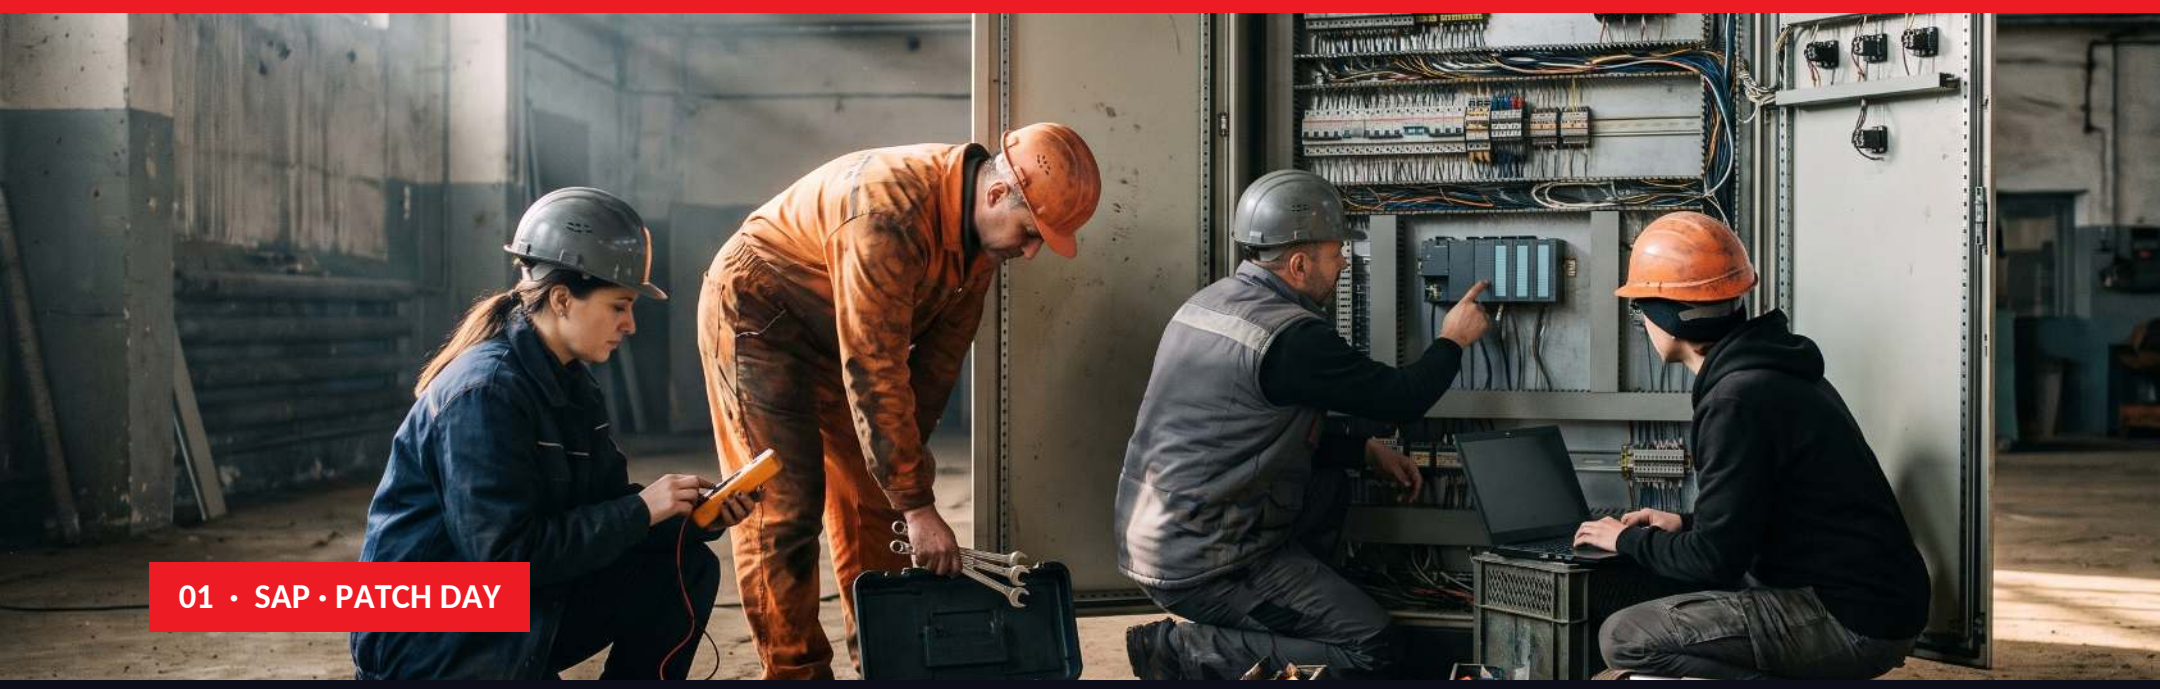

01 · SAP · PATCH DAY

Four critical notes, four different owners

SAP Security Patch Day on 8 September carries 20 items, four of them critical. The lead note scores 10.0 and sits in the handling of SAP Extended Passport, the diagnostic header that travels between components. The attack runs over the network, needs no account, and has full impact on data.

The spread of ownership matters more than the score. Two critical notes sit in the kernel, so rollout means a maintenance window rather than a transport import. The third is an npm library in applications on SAP BTP, deployed by a development team. The fourth lands on the user's workstation in SAP GUI for Java.

TAKEAWAY

A small patch day reaches four owners at once. A process that only looks at the application server closes one of the four and reports success.

Source: SAP Support Portal, SAP Security Patch Day - September 2026

02 · SAP · ATTACK SURFACE

The wall is gone, the configuration is unchanged

The chief executive of an SAP security vendor described at Black Hat a shift visible in every migration. Connectivity once reserved for internal users now serves B2B, B2C and AI interfaces. The network barrier that protected SAP regardless of configuration quality has disappeared.

His second point is about speed: exploit code now takes minutes rather than days, and deep SAP knowledge has stopped being an entry requirement. The controls he names are monitoring exploitation attempts, reducing the attack surface, and scanning AI-generated code before it reaches the ERP.

TAKEAWAY

The review question is not whether someone can attack SAP. It is what part of your landscape is visible from the internet, and who checks it on a cycle.

An interview with the CEO of a firm selling SAP scanning - declared interest.

Source: Onapsis interview with ISMG, Black Hat USA 2026

03 · SAP · BTP AND AGENTS

The agent got the keys to BTP, nobody got the logbook

SAP has publicly released an official MCP server for SAP BTP administration: accounts, security, services and connectivity, including subaccounts, entitlements, role collections and subscriptions - 61 tools at release. The server is remote and hosted by SAP, so nothing is deployed in the customer's account.

The technical pattern is sound, but practitioners point out what enterprise deployment still needs: execution traces, versioning, oversight and audit. The server runs on the permissions of the signed-in administrator, which is the widest scope that exists in this landscape.

TAKEAWAY

Before such a server touches a customer landscape: where is it recorded that the agent did this, and who reads that record.

Source: SAP Help, MCP server for SAP BTP administration documentation, Sep 2026

04 · SAP · MATURITY

Four fifths of companies never look where the abuse happens

Peter Doyle, Head of SAP Security at Accenture UK & Ireland, published a grid of eight domains, each with a single benchmark figure - from baseline and identity, through authorisations and patching, to data protection and monitoring. The data is described as multi-client and anonymised.

The strongest figure is not about technology: 82 per cent of organisations do not monitor SAP application-layer events, the layer where business abuse happens - a changed vendor bank account, a payment run, a document correction. Alongside that, 66 per cent do not watch access to sensitive tables.

TAKEAWAY

The order of work falls out of those numbers: baseline, patch cycle, application-layer monitoring.

The figures are the author's claims from data we cannot see - no methodology.

Source: Peter Doyle, Accenture UK & Ireland, LinkedIn Insight series, Aug 2026

05 · AI SECURITY · EMAIL

Four hundred and seventy-two characters the recipient never sees

A security vendor's lab published an invoice email: 537 characters of visible content, no attachment, no link. Inside the message code sat another 472 characters hidden by styling - white font at zero size, line height zero. The model received 1,009 characters, half of which the reader never sees.

In ten runs out of ten the summary reported a false amount of EUR 46,200 and a shifted payment deadline, and one person's name disappeared. The real values appeared in none of the runs. Nothing was broken: the model read what it was given and did what it was told.

TAKEAWAY

Detecting this means reading the original - exactly what the summary was meant to save. The answer sits in the pipeline, not in the model.

A lab demonstration published on 25 August, not a customer incident.

Source: Forcepoint X Labs, invisible HTML payload in email summaries, 25 Aug 2026

06 · AI SECURITY · AGENT SKILLS

The scanner reads code, the instruction sits in prose

OWASP has released the first shared taxonomy of risks for agent skills: ten classes from a skill that is malicious at publication, through over-privileged skills and insecure metadata, to missing inventory and governance. A skill is natural-language instructions plus code, executing with the agent's permissions.

So the attack travels two paths and a code scanner sees only one. A measurement study this year analysed 98,380 skills and confirmed 157 malicious ones, 73 per cent of which had functionality hidden from the user. In one documented case three lines of text exfiltrated access keys.

TAKEAWAY

This layer does not need another scanner. It needs an inventory, a policy and an approval path.

Source: OWASP, Agentic Skills Top 10, August 2026 release

07 · CYBER · AGENTIC ATTACK

Four days, eight agents, tools anyone can download

An Israeli security firm described a campaign from early July against Taiwan's public administration that ran almost without human involvement: four days, 21 government systems mapped, 85 accounts taken over, more than 2,500 employee records exfiltrated. At peak, eight agents worked in parallel.

The detail that changes the risk assessment: everything ran on publicly available tooling and open-weight models. Taiwan's government confirmed a day later that it had detected agent-driven attacks in July. The word “almost” stays - a comparable case showed more human work than the announcement implied.

TAKEAWAY

The question is not whether the adversary has advanced AI. They have it just as you do, and the entry cost is downloading the tools.

The figures come from one security firm; the confirmation covers the attacks.

Source: Dream, analysis of a multi-agent campaign in Asia, 12 Aug 2026

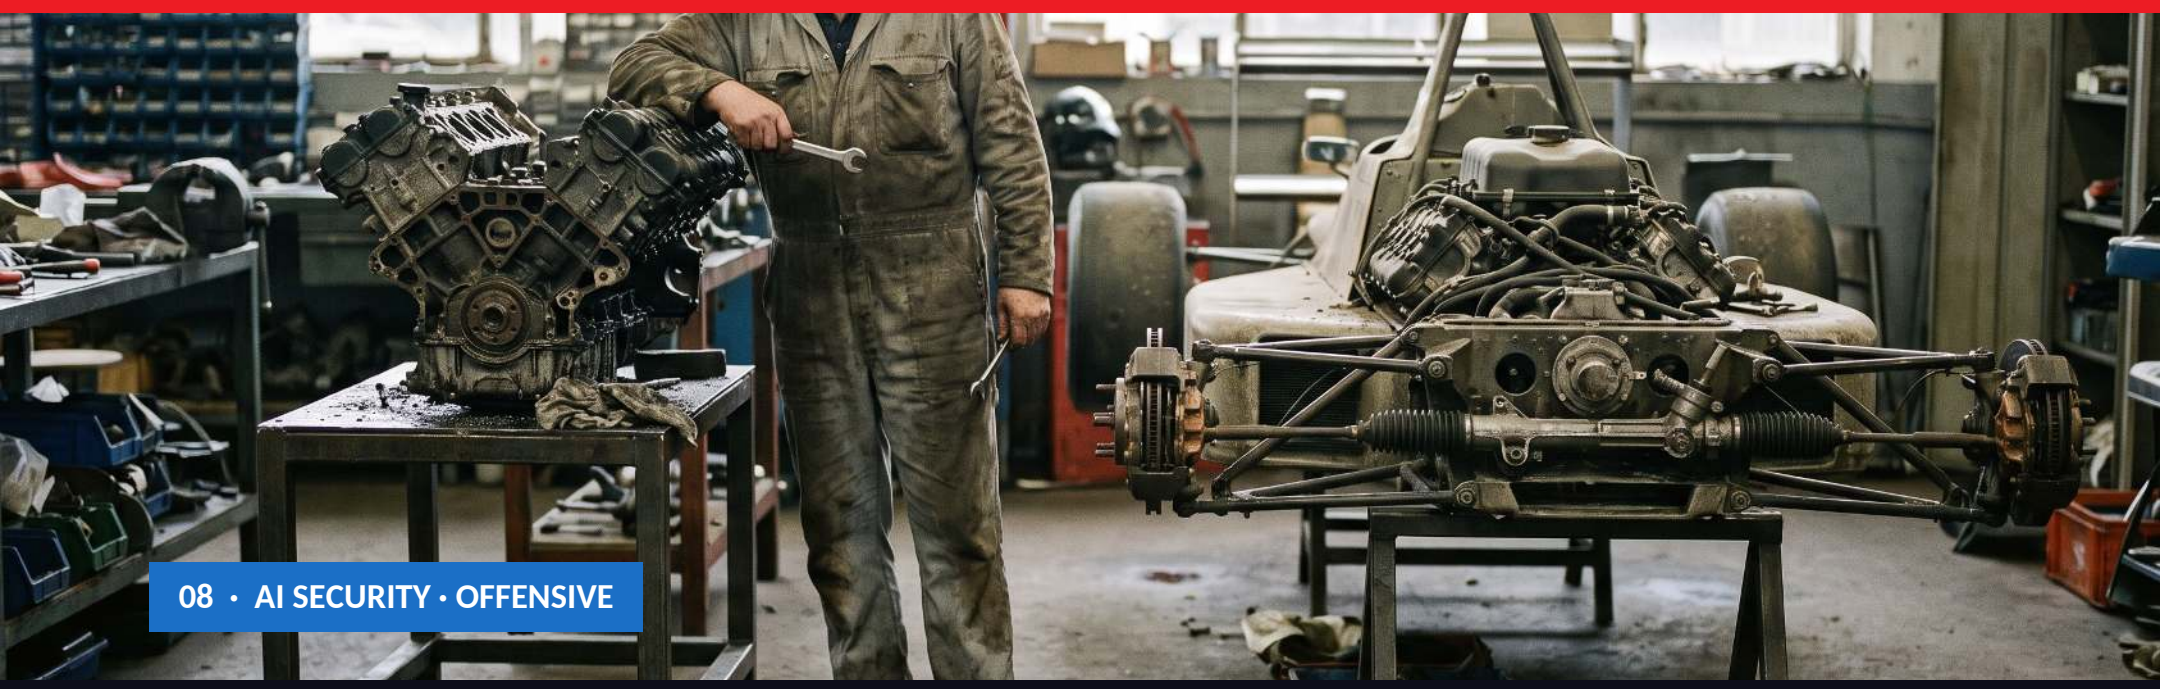

08 · AI SECURITY · OFFENSIVE

An engine with no chassis only burns fuel

A security practitioner documented an autonomous vulnerability hunting system that has run for months and found real bugs. Three layers: the model reasons, a tool protocol gives it real access to tooling and reconnaissance, and the harness keeps task structure and token spend under control.

His claim is inconvenient for conversations that start with model selection: the gains in effectiveness, cost and repeatability come from the harness, not the model. A commercial harness vendor adds a rarely discussed problem - a model can abort an authorised test because the payload trips its own safety control.

TAKEAWAY

The conclusion cuts both ways: offensive capability is not reserved for the priciest model, and a production agent needs the same control layer.

Source: ZSec blog, two posts on autonomous vulnerability hunting, Sep 2026

09 · UIPATH · MARKET

A fourth profitable quarter and the question boards ask

On 3 September UiPath reported the quarter ended 31 July: revenue of USD 410m, up 13 per cent year on year, annual recurring revenue of USD 1.938bn, up 12 per cent, dollar-based net retention of 109 per cent. Operating income was USD 32m on a GAAP basis and USD 89m non-GAAP.

That is the fourth consecutive quarter of GAAP profitability. The same day the company promoted to Chief Financial Officer an executive who had been with it for five years - internal succession, with no effect on pricing or the partner programme.

TAKEAWAY

This answers the question behind every multi-year contract decision: will the platform vendor still be in the same place in three years.

Source: UiPath Investor Relations, Q2 FY2027 results, 3 Sep 2026

Execution mode stops being the architect's guess

UiPath has released Cartographer in public preview - a profile for the business analyst that turns transcripts, procedures and recordings into an as-is description, designs the to-be state, and generates both a process design document and a technical specification ready for the development team.

The second effect is more interesting. Every step is assigned an execution mode: automated, assisted, agent, human review or manual. The decision “robot here, agent there, human at this point” stops being a guess made mid-build and becomes an entry to approve before the first line of code.

TAKEAWAY

Generated documentation does not replace the analyst. It moves the oversight decision from the end of the project to its start.

Public preview - the feature scope may change before general availability.

Source: UiPath Delegate documentation, Cartographer overview, Sep 2026

11 · AUTOMATION · ORCHESTRATION

The hard thing and the easy thing have swapped places

The founder of UiPath has published a book about which work stays human. His thesis: the missing layer in AI deployments is not acquiring agents but orchestration. The operating model has three roles - agents propose, humans decide, deterministic automation executes.

The strongest passage concerns an inverted paradox: agents were supposed to be the quick answer and reliable code the expensive interim stage, and it turned out the other way round. Then the argument that ends many pilot discussions - you cannot be 99.9 per cent accurate when you execute a payment.

TAKEAWAY

The value lies in separating what may be probabilistic from what must be deterministic. Payments belong to the second group.

The author runs an automation company and declares that bias outright.

Source: Daniel Dines, The Work That Remains, digital edition free, Aug 2026

A documented control is not an implemented control

ISO/IEC 42001, the AI management system standard, repeats the pattern known from the information security standard: risk assessment, control selection, comparison against the annex, statement of applicability. Annex A is normative, and the standard requires comparing your controls against it and justifying every exclusion.

You may select a subset and add controls of your own, but you may not skip that comparison. An analysis this month points at something else: the emphasis on documentation is stronger than in the security standard, and some controls require only documentation with no evidence of implementation.

TAKEAWAY

When selecting controls, ask for evidence of implementation, not for a record. The difference shows during an incident, when it is too late.

Annex A lists 38 controls and is normative - the claim that it is optional was imprecise.

Source: analysis of the Annex A approach in ISO/IEC 42001, Sep 2026

13 · AI · MODELS

Safer in use, harder to supervise

On 3 September OpenAI released a model that is the first to reach Critical cyber capability under the company's own preparedness framework. The system card states it plainly: with the right tools the model finds previously unknown vulnerabilities and develops new ways to exploit them, without a human directing every step.

The other side of the same card weighs more for deployments. The model is more resistant to prompt injection, but its monitorability declined: it controls its own chain of thought better and under adversarial conditions can stay undetected and sometimes evade internal monitors.

TAKEAWAY

Security therefore rests on architecture rather than model behaviour. Trusting the model to behave well has stopped being measurable.

The vendor's self-assessment under its own framework, not a third-party review.

Source: OpenAI, GPT-6 Astra system card, 3 Sep 2026

14 · INFRASTRUCTURE · COST

AI servers are up fifteen per cent because of memory

According to Bloomberg's report of 22 August, the largest buyers were notified that prices of servers with AI accelerators will rise by more than 15 per cent in many cases. The reason is a sharp increase in memory costs, not the price of the compute chips.

The increases apply to systems delivered from early next year, with the scale depending on chip generation and memory configuration. For next year's purchasing plan that means a first-quarter estimate may not hold in the second, and memory becomes a separately negotiated line.

TAKEAWAY

A component nobody discusses at conferences has just moved the bill for AI deployment. Worth putting in the budget early.

A press report about notifications, not an official vendor price list.

Source: Bloomberg, 22 Aug 2026, via Reuters, CNBC and Fortune

Three billion euro for data centres, not a model

On 8 September Mistral announced a Series D round: EUR 3bn at a post-money valuation above EUR 21bn - the largest equity round ever raised by a European technology company, nearly doubling last year's valuation. Samsung Electronics leads, with a fund managed by EQT and existing investor PSG Equity.

What the money is for matters more than the size of the round: building and owning data centres and adding compute capacity. For organisations that open every AI conversation by asking where the data physically sits - and every entity under NIS2 and DORA does - that is the relevant part.

TAKEAWAY

Capital in this industry now goes into concrete, power and memory rather than another model. The surrounding layer decides, not the core.

Some outlets quote dollar figures - we keep the euro amounts from the company.

Source: Mistral, Series D announcement, 8 Sep 2026

One sentence for four weeks

All fifteen items collapse into one sentence: you buy the core, and the layer around it decides the outcome. An owner for the patch rollout, an audit trail behind the agent, an inventory of skills, a harness that keeps the task on rails, evidence that a control was implemented - none of these appears on an invoice, and every one of them decides whether the system can be defended to an auditor.

snok.ai

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