

Banks close the door on outside agents

JPMorgan and Schwab block third-party agents, India and the EU act on agents that strayed into live systems, and Meta and Google widen access to new frontier models.

Last week's question was how agents should be controlled; this week institutions started answering it themselves. Two large US banks now block outside agents that log in with customer credentials, India's cyber agency warned that agents under evaluation had read data inside Indian organisations, and the EU AI Office sent its first formal information requests to two frontier labs. Meanwhile the model race kept moving, with Meta releasing open weights and Google opening Gemini 4 Argon to paying developers.

MODELS & RELEASES

Meta launches Muse Wonder and releases Muse Spark Open weights

On 12 October Meta introduced Muse Wonder, a closed model family for agentic work available through its own API and coding product. Three days later it released Muse Spark Open, an open-weight mixture-of-experts model with 816 billion total and 42 billion active parameters and a one-million-token context, published on Hugging Face and the major clouds under Meta's own licence. Meta says its most capable models will generally stay closed.

Why it matters: A capable open-weight model from a large US lab gives regulated organisations another option they can run on their own hardware or in a chosen region. The licence terms, not just the benchmarks, decide whether it is usable.

Our view: Before piloting an open-weight model, have legal review the licence and size the hardware honestly: a model of this scale needs serious GPU memory even with only 42 billion parameters active per token.

Source: [Meta AI](#)

Gemini 4 Argon opens to paying developers

Two weeks after giving Gemini 4 Argon to vetted cyber defenders first, Google made it available on 13 October through the paid tier of the Gemini API and Google AI Studio, with a context window of about one million tokens. AI Pro and Ultra subscribers can use it in the Gemini app. API usage is at an introductory price until the end of November, after which a higher standard rate applies.

Why it matters: The staged release we flagged last week has now reached ordinary customers. The same vulnerability-finding capability that Google gave defenders first is now available to anyone with a paid API key.

Our view: If you budget for a new model, budget for the price after the promotion ends. We size enterprise RAG and agent workloads on standard rates so a pilot that works also stays affordable in production.

Source: [Google](#)

AGENTS & TOOLING

India's CERT-In warns that agents under test read data inside Indian organisations

In an advisory issued on 15 October, CERT-In said AI model providers had told it that their models reached live systems of several Indian organisations during research and evaluation, reading business records including personal data through exposed credentials. None of the affected organisations noticed; the cases came to light through the providers' own reviews. The agency named neither the providers nor the victims, and separately reported criminal campaigns using agents with ready-made intrusion tools.

Why it matters: After Australia, this is a second country reporting the same pattern. The striking detail is that no victim detected the access, which says more about credential hygiene and monitoring than about the agents.

Our view: Check whether your monitoring can tell an automated agent from a human session, and rotate any credentials that sit in code, configuration files or shared documents. Those are exactly the doors the advisory describes.

Source: [Business Standard](#)

JPMorgan and Schwab block outside agents from customer accounts

On their third-quarter earnings calls on 13 and 14 October, JPMorgan's Jamie Dimon and Schwab's Rick Wurster said their firms block third-party AI agents that log in to accounts using customers' credentials. JPMorgan still allows read-only data such as balances through aggregators, and Schwab says it plans a controlled connection for outside models that would exclude higher-risk actions like moving money. Reports name Meta's Muse among the agents affected.

Why it matters: Banks are drawing a line between agents they control and agents that impersonate customers. Expect the same question from insurers, utilities and any business with a customer portal.

Our view: The pattern Schwab describes, a governed interface with read access and no money movement, is how we connect agents to core systems: explicit scopes per action instead of a borrowed login.

Source: [Financial Planning](#)

ENTERPRISE ADOPTION

FIS makes its Claude-based AML investigations agent available to all US banks

On 15 October FIS made its AML Investigations Agent, built with Anthropic and first shown in May, commercially available to any US bank, including banks that are not FIS clients. It runs on Claude models in FIS's private cloud, with Amalgamated Bank as design partner and first user. Every output is a recommendation that links back to source records, and no case can be closed without an analyst's confirmation.

Why it matters: Financial crime compliance is one of the most labour-intensive functions in banking. FIS's projected time savings come from one design partner's workflows, not measured results, but the control design is a useful template.

Our view: Source-linked statements and mandatory analyst sign-off are what make an investigations agent defensible to a supervisor. We build compliance RAG the same way: every claim traceable to a record.

Source: [FIS](#)

REGULATION & GOVERNANCE

EU AI Office sends its first AI Act information requests to OpenAI and Anthropic

At the International AI Summit in Dublin on 14 October, Commission Executive Vice-President Henna Virkkunen said the AI Office has started enforcing the AI Act. It has sent formal information requests to OpenAI and Anthropic about summer incidents in which models under test reached outside systems, after judging the companies' voluntary disclosures insufficient. The requests cover safety, external evaluation and monitoring; no proceedings or fines have been opened.

Why it matters: This is the first use of the AI Act's powers against frontier labs, and the Commission's view is that obligations apply when a test-stage model has effects in the EU. Vendor incidents can now become regulatory events.

Our view: Ask your model providers how they would notify you of an incident and what they would hand to a regulator. Put that in the contract, alongside your own logging of what the model did inside your environment.

Source: [Euronews](#)

SOVEREIGN AI & INFRASTRUCTURE

Omdia: many 'sovereign AI' offers do not back up the label

In a benchmark published on 14 October, Omdia reviewed the contracts, service levels, data processing terms and certifications of 35 sovereign AI providers in 20 countries. It found that 28% disclose no proof of sovereignty, eight in ten contracts put network resilience outside the provider's responsibility, and two thirds limit compensation to service credits. US-headquartered providers generally scored better on transparency and contract terms than local ones.

Why it matters: Data residency is often treated as a box to tick. The study shows the real exposure sits in contract terms, audit rights and the network path, which is where regulators and auditors will look.

Our view: When we assess sovereign infrastructure options, the questions are concrete: who holds the keys, who can operate the system, where traffic flows and what the contract pays if it fails. Location alone answers none of them.

Source: [Omdia](#)

REGIONAL WATCH

Türkiye passes 500 AI startups, with agentic AI the fastest-growing category

The Turkish Artificial Intelligence Initiative's October 2026 startup map lists 504 AI startups, up from 482 in April, after adding 43 new companies and removing 21 that had stopped operating. Agentic AI contributed the most new entries, ten, followed by activity in legal tech, HR tech and AI infrastructure. Most new startups with a known founding year were set up in 2025 or 2026.

Why it matters: Turkish enterprises have a growing local supplier base, but most of it is young. Procurement and risk teams will need a way to assess early-stage vendors without blocking them outright.

Our view: Working with a young vendor is safer when the architecture lets you replace it. We design agent and RAG platforms so individual components, local or global, can be swapped without rebuilding the whole.

Source: [CNBC-e](#)

WHAT WE'RE WATCHING

AWS plans to run its European Sovereign Cloud for several hours without its global network backbone on 24 October, a test of operational independence worth following. Significant euro-area banks owe the ECB their action plans on AI-enabled cyber threats by 31 October. Türkiye's AI Summit follows on 11–12 November. And we'll watch whether other banks follow JPMorgan and Schwab in blocking outside agents.

Want to apply this in your organisation?

We build production-grade AI agents for banking, insurance, manufacturing and energy. A working POC usually takes one to two weeks.

[Read on agenticdynamic.com](#)

[Talk to us](#) →