New frontier models, and agents that pay

Anthropic and Google release new top models, OpenAI lets its agents make purchases while United shuts agents out, IBM and Microsoft package agents for enterprises, and HUMAIN sets a December date for its AI stack.

Two of the largest labs shipped new flagship models within a day of each other, and both kept their most capable variants behind vetting or premium tiers. At the same time, the question of what agents may do on a customer's behalf got sharper: OpenAI started letting its agents pay for things with approval, while United became the first US airline to bar outside agents from booking. In Riyadh and Istanbul, the focus was on putting AI to work locally.

MODELS & RELEASES

Anthropic releases Claude Fable 5.2, with Mythos 5.2 limited to vetted organisations

On 9 November Anthropic released Claude Fable 5.2, available on all paid plans, its own platform and the major clouds, at a lower token price than the version it replaces. The more capable Mythos 5.2 is available only to organisations vetted through Anthropic's life-sciences and cyber verification programmes. Anthropic says the public model keeps cybersecurity and biology safeguards in place while refusing fewer ordinary requests.

Why it matters: Top labs now routinely keep their strongest models behind vetting programmes. For regulated firms, access to frontier capability increasingly depends on what they can show about their own use and controls.

Our view: A cheaper model is a chance to re-run your evaluations, not just your budget. We benchmark new releases against each workflow's own test set before switching, so a price cut does not hide a change in behaviour.

Source: CNBC

Google makes Gemini 4 Carbon, its strongest model, generally available

Google made Gemini 4 Carbon generally available on 10 November to paying developers and enterprises through its API and Gemini for business, and to AI Ultra subscribers in its consumer apps. Google positions Carbon above Gemini 4 Argon for long-horizon agentic work and cybersecurity, after first giving it to vetted cyber defenders. An introductory API price runs until the end of November.

Why it matters: The same staged pattern we saw with Argon has repeated within weeks. Release cycles for top models are now short enough that procurement and validation processes need to keep pace.

Our view: Budget for the standard price, not the launch offer, and keep the model choice configurable. Our agent workflows route tasks by difficulty, so a stronger model is used only where it changes the outcome.

Source: Google

AGENTS & TOOLING

OpenAI launches ChatGPT Wallet so its agents can make purchases

OpenAI began rolling out ChatGPT Wallet on 13 November, letting its Dots agents buy things for US Pro users with eligible American Express, Chase, Mastercard or Visa cards. Each purchase needs a biometric approval, and the merchant and agent receive a network token rather than the real card number. OpenAI says it plans to allow purchases within preset spending limits without approving each one.

Why it matters: Agent payments are moving from pilots to mainstream consumer products, built on the card networks' token systems. Banks and card issuers will need to recognise and risk-score agent-initiated transactions.

Our view: Per-transaction approval now, spending limits later: that is the path most agent payment flows will follow. Define in advance which limits, merchants and categories an agent may use without asking.

Source: CNBC

United becomes the first US airline to bar outside AI agents from booking

United announced on 11 November that unauthorised third-party AI agents may no longer book or manage trips through its website, app or indirect channels, and it added AI agents to the prohibited practices in its contract of carriage. The airline cites security and privacy risks from customers sharing account credentials with personal assistants. Bookings made by agents before the change will not be cancelled.

Why it matters: After banks and retailers, airlines are drawing the same line: agents are welcome only through channels the company controls. Any business with customer accounts will face this decision.

Our view: Blocking outside agents is a short-term answer; offering a governed interface is the longer one. Decide which actions you want agents to perform for your customers and expose those, with clear limits, rather than leaving credentials as the only door.

Source: Skift

ENTERPRISE ADOPTION

IBM and Microsoft launch Imagine Enterprise, a platform of prebuilt industry agents

IBM and Microsoft announced Imagine Enterprise on 10 November, a platform for building, running and governing agents tied to industry processes. IBM Consulting contributes more than 6,000 prebuilt agents and assets, Microsoft provides its cloud, data, identity and security stack, and both supply engineers to take deployments into production. Pricing, general availability and production customers have not been disclosed.

Why it matters: Large integrators are packaging agents as a catalogue plus engineers. That can speed up adoption, but it also means a firm's processes run on a vendor's agent library and governance layer.

Our view: A prebuilt agent is a starting point, not a finished control. Whatever platform you choose, keep your own test cases, approval rules and logs so you can verify and, if needed, move what runs on it.

Source: CIO Dive

SOVEREIGN AI & INFRASTRUCTURE

HUMAIN sets December for its Saudi AI stack and a large open-weight model

At an event in Riyadh on 11 November, HUMAIN chief executive Tareq Amin said the company's full AI stack will roll out on 15 December, together with a 3-trillion-parameter open-weight model hosted in Saudi Arabia. HUMAIN+, its broader platform, is now due to launch operationally in the first quarter of 2027. The company has not yet published the model's name, licence terms or language coverage.

Why it matters: In-country models and infrastructure give Saudi banks, insurers and energy firms more options under local data rules. The licence and the operating terms will decide how usable they are.

Our view: Treat a national model like any other candidate: test it on your own Arabic and English tasks and check the licence before planning around it. Sovereign hosting helps only if the model fits the work.

Source: Arab News

REGIONAL WATCH

Türkiye's AI Summit puts agents at the centre

The Turkish Artificial Intelligence Initiative held its ninth Türkiye AI Summit in Istanbul on 11 November under the theme "Agents Among Us", with more than 2,000 participants from 420 organisations, according to the organisers. Two studies were presented for the first time: a report on AI's economic impact on Türkiye and the 2026 TRAI AI survey of more than 350 organisations. The full findings have not yet been widely published.

Why it matters: Agents are now the main topic for Turkish business audiences, not just the tech sector. For regulated firms, that raises the urgency of governance questions alongside pilots.

Our view: Interest is high; the gap is in running agents safely in production. Start with one well-bounded process, with clear approvals and logs, and build internal skills around it before scaling.

Source: Para Dergi

WHAT WE'RE WATCHING

Microsoft Ignite runs on 17–19 November, and Microsoft's Saudi Arabia East cloud region is still due to open this month. The UK ICO's call for evidence on agentic AI closes on 20 November. HUMAIN's AI stack is due on 15 December, and we will watch whether other airlines and retailers follow United on 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 →**