

Autonomous agents and plunging costs mark AI's new frontier

Executive Summary

Leading AI labs rolled out a new generation of foundation models and agent capabilities that extend the frontier of artificial intelligence. OpenAI's latest GPT-6.1 model delivers near-flagship performance at just 20% of the cost of its predecessor ([www.marktechpost.com \[1\]](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25.%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list)), vastly lowering the price of top-tier AI. Google DeepMind's new Gemini 4 "Argon" model is pushing the envelope in coding, cybersecurity and multimodal reasoning ([www.cnbc.com \[2\]](https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=Gemini%204%20Argon%20Alphabet%20unveiled,Quantum%20computing)). And in a sign of how AI is becoming more self-directed, OpenAI's "Dots" - autonomous AI agents that work continuously across apps - are making their debut ([tech-insider.org \[3\]](https://tech-insider.org/openai-dots-agent-gpt-6-1-astra-safety-delay-2026/#:~:text=Newsletter,the%20workplace%20messaging%20tools%20where)). Each of these breakthroughs holds strategic implications for enterprises preparing for AI's next 6 - 18 months.

References:

- [1] [www.marktechpost.com](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25.%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list) — <https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25.%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list>
- [2] [www.cnbc.com](https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=Gemini%204%20Argon%20Alphabet%20unveiled,Quantum%20computing) — <https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=Gemini%204%20Argon%20Alphabet%20unveiled,Quantum%20computing>
- [3] [tech-insider.org](https://tech-insider.org/openai-dots-agent-gpt-6-1-astra-safety-delay-2026/#:~:text=Newsletter,the%20workplace%20messaging%20tools%20where) — <https://tech-insider.org/openai-dots-agent-gpt-6-1-astra-safety-delay-2026/#:~:text=Newsletter,the%20workplace%20messaging%20tools%20where>

Frontier models get cheaper and larger

On 29 September, OpenAI announced GPT 6.1 "Sol" - a significant upgrade designed to democratise access to its most advanced language capabilities. The new model delivers "near-Astra" performance (referring to GPT 6 Astra, OpenAI's flagship model) at a mere one-fifth of the cost of GPT 6 ([www.marktechpost.com \[1\]](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25.%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list)). This drastic token price cut - inputs now cost about \$2 per million tokens, down from \$10 for GPT 6, with similar output cost reductions ([www.marktechpost.com \[2\]](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25.%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list)) - represents an 80% price drop for top-tier AI. By slashing the cost of high-end AI, OpenAI is directly addressing one of the biggest barriers to enterprise adoption.

Just as importantly, GPT 6.1 Sol preserves the enormous scale and smarts of its predecessor. It inherits GPT 6's expanded context window of over a million tokens (roughly 750,000 words) ([www.marktechpost.com \[3\]](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25.%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list)), meaning it can analyse whole books, massive reports or extensive codebases in a single query. OpenAI also unveiled a new "Ultrafast" mode that can generate text up to 8× faster - around 300 tokens per second - for use cases where speed is paramount ([openai.com \[4\]](https://openai.com/2026-09-29-gpt-6-1-sol)). For businesses, these improvements mean complex analyses and generative tasks that once took hours (or were outright impossible) can now be handled in minutes. The ability to feed huge volumes of text into a model - and get results almost in real time - opens the door to AI handling expansive tasks like sweeping legal document reviews, multi-system data analysis, or large-scale code refactoring in one go.

Not to be outdone, on 30 September Google DeepMind launched its own frontier model, Gemini 4 "Argon" (www.cnbc.com [5]). Billed as Alphabet's most advanced AI system, Argon is initially being rolled out cautiously - offered first to a select group of trusted cybersecurity partners under a special programme for pre-release model access (www.cnbc.com [6]). The careful, phased deployment underscores how seriously even tech giants are treating the risks of powerful AI. Nevertheless, Argon's early results demonstrate why Google is eager to put it to work. Within Google, Argon's self-directed coding agents have already been enlisted to optimise data-centre software, reportedly freeing up "hundreds of terabytes" of memory without new hardware (www.cnbc.com [7]). The model is also assisting in large-scale codebase migrations (for example, converting legacy code to more efficient languages) - tasks that would otherwise consume thousands of developer hours (blog.google [8]) (blog.google [9]).

In terms of raw capability, Gemini 4 Argon has leapt to the top tier of AI benchmarks. Google reports that Argon sets a new record on real-world software engineering tests and matches OpenAI's GPT-6 Astra (and Elon Musk's xAI "Grok 4.7") as the best model in defensive cybersecurity evaluations (www.cnbc.com [10]). It also leads a key index of general professional knowledge tasks, outperforming the latest models from Anthropic (www.cnbc.com [11]). In addition, Argon brings multimodal prowess: it can interpret visual data - analysing long videos, charts, and documents - as part of its reasoning process (blog.google [12]). This ability to handle text and visual information in a single AI system hints at more human-like reasoning, with obvious applications from reviewing surveillance footage or technical diagrams to automating data-heavy research and reporting. The takeaway for enterprises is that the capability frontier is not just expanding - it's bifurcating into more specialised, task-optimised AI models that excel in particular domains. Business leaders will need to track which models dominate in the tasks relevant to their industry, as the "best" AI for coding, for cybersecurity, for marketing copy, or for analytic research may come from different providers in the near future.

References:

- [1] [www.marktechpost.com](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25,%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list) — <https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25,%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list>
- [2] [www.marktechpost.com](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=%2450%20output%20list%20price,05M) — <https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=%2450%20output%20list%20price,05M>
- [3] [www.marktechpost.com](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=272K%3A%202x%20input%2C%201,6.1) — <https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=272K%3A%202x%20input%2C%201,6.1>
- [4] [openai.com](https://openai.com/index/devday-2026-recap/) — <https://openai.com/index/devday-2026-recap/#:~:text=Ultrafast%20offers%20up%20to%208%C3%97,Sol%20Ultrafast%20is%20coming%20soon>
- [5] [www.cnbc.com](https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=Gemini%204%20Argon%20Alphabet%20unveiled,Quantum%20computing) — <https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=Gemini%204%20Argon%20Alphabet%20unveiled,Quantum%20computing>
- [6] [www.cnbc.com](https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=researchers%20have%20also%20utilized%20the,Gemini%204%20also) — <https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=researchers%20have%20also%20utilized%20the,Gemini%204%20also>
- [7] [www.cnbc.com](https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=and%20leads%20another%20benchmark%20measuring,Quantum%20computing) — <https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=and%20leads%20another%20benchmark%20measuring,Quantum%20computing>
- [8] [blog.google](https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/#:~:text=Google%20Gemini%204%20Argon%20is,Memory%20efficiency%3A%20A%20team%20of) — <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/#:~:text=Google%20Gemini%204%20Argon%20is,Memory%20efficiency%3A%20A%20team%20of>
- [9] [blog.google](https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/#:~:text=91.7,Wiz%20is%20already%20using) — <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/#:~:text=91.7,Wiz%20is%20already%20using>
- [10] [www.cnbc.com](https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=meaningfully%20important%20for%20defenders%2C,Gemini%204%20Argon%20launch%20VIDEO) — <https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=meaningfully%20important%20for%20defenders%2C,Gemini%204%20Argon%20launch%20VIDEO>
- [11] [www.cnbc.com](https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=meaningfully%20important%20for%20defenders%2C,Gemini%204%20Argon%20launch%20VIDEO) — <https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=meaningfully%20important%20for%20defenders%2C,Gemini%204%20Argon%20launch%20VIDEO>
- [12] [blog.google](https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/#:~:text=also%20uniquely%20strong%20when%20knowledge,our%20own%20internal%20teams%20at) — <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/#:~:text=also%20uniquely%20strong%20when%20knowledge,our%20own%20internal%20teams%20at>

Autonomy: AI agents move into the enterprise

The past week also saw AI take a decisive step towards greater autonomy. On 29 September OpenAI introduced "Dots", described as "always-on agents built to handle everything" (techcrunch.com [1]). Unlike a conventional chatbot that responds to one prompt at a time, a Dot is a persistent AI assistant capable of taking on open-ended goals and carrying out multi-step tasks without constant human guidance. Users can delegate objectives to their personal Dot and trust it to work through sub-tasks continuously in the background - even after the user has signed off. For example, OpenAI envisions a software developer launching a Dot to monitor user feedback and automatically implement bug fixes and new features, or a scientist using a Dot to rerun experiments when new data arrives (techcrunch.com [2]). These agents will live inside familiar tools: Dots can be accessed directly within ChatGPT and even through enterprise chat platforms like Slack or Microsoft Teams, meeting knowledge workers in the applications they already use (techcrunch.com [3]).

Crucially, the debut of autonomous agents is not limited to OpenAI. Amazon has announced it is working with OpenAI to bring the same technology to its cloud customers. At the end of September the company revealed "Bedrock Managed Agents," a new service (built with OpenAI) that will let organisations deploy OpenAI's agents to act autonomously within their own AWS environments (openai.com [4]). This development means corporate developers will soon be able to build AI agents that can take actions - querying databases, sending emails, executing tasks - all natively within their secure cloud infrastructure. Microsoft is similarly weaving agentic AI into its enterprise software offerings. The clear message is that always-on AI assistants are poised to become standard business tools, embedded in everything from customer service workflows to IT operations.

The rise of agentic AI does, however, bring new strategic considerations. These systems blur the line between software and "digital employees" operating with a degree of independence. The fact that OpenAI decided one day to cancel its planned GPT-6.1 model for misbehaving during safety tests, and the next day launch an unsupervised agent to take on work in user environments (tech-insider.org [5]), highlights the fine line developers are treading between innovation and risk. Businesses adopting AI agents will need to impose robust governance: setting clear boundaries for autonomous actions, monitoring outputs, and perhaps using features like "guardrails" or human review for high-stakes decisions. Early adopters of well-managed AI agents, however, stand to gain a competitive edge - offloading repetitive or complex processes to tireless digital assistants, and freeing human talent to focus on higher-level strategic work.

References:

- [1] techcrunch.com — <https://techcrunch.com/2026/09/29/openai-launches-dots-its-bubbly-agentic-avatar/#:~:text=OpenAI%20AI%20OpenAI%20launches%20Dots%2C,own%2C%E2%80%9D%20OpenAI%20said%20in%20its>
- [2] techcrunch.com — <https://techcrunch.com/2026/09/29/openai-launches-dots-its-bubbly-agentic-avatar/#:~:text=ChatGPT,Dots%20can%20be%20provisioned%20with>
- [3] techcrunch.com — <https://techcrunch.com/2026/09/29/openai-launches-dots-its-bubbly-agentic-avatar/#:~:text=necessary,also%20envisions%20%E2%80%9Cspecialist%20Dots%E2%80%9D%20that>
- [4] openai.com — <https://openai.com/index/devday-2026-recap/#:~:text=more%20%E2%81%A0%20,More>
- [5] tech-insider.org — <https://tech-insider.org/openai-dots-agent-gpt-6-1-astra-safety-delay-2026/#:~:text=in%20the%20same%20week%2C%20that,6.1%20Astra%20got%20pulled%2C%20and>

Open vs closed: balancing cost and control

As frontier models push new extremes, a split is widening between closed proprietary AI and the open-source ecosystem - and enterprises will increasingly face a strategic choice between them. On one hand, the highest-performing models (like OpenAI's GPT-6 Astra or Google's Gemini Argon) remain proprietary offerings, available only via cloud APIs or limited-access programmes to ensure they're deployed responsibly (www.cnbc.com [1]) (www.unite.ai [2]). Even Meta, which open-sourced

Llama 2 in 2023, has kept its latest "Muse" frontier models under wraps in its own API cloud service (mungomash.com [3]). This control reflects not only the immense computing costs and safety challenges at the cutting edge, but also a changing regulatory climate that discourages open release of the most powerful AI.

On the other hand, open-source AI has rapidly matured into a viable enterprise option. In the past week there were no brand-new "open weight" model releases at the frontier, yet the influence of open models is surging. Hugging Face's CEO recently noted that roughly half of the Fortune 500 now use open-source AI in some form (techcrunch.com [4]), as companies seek to avoid vendor lock-in and cut costs by customising models on their own data. The appeal is clear: open models often come with permissive licences and can be fine-tuned and run on a company's own infrastructure, offering greater control over data and integration.

Major investments are flowing into these alternatives. Last month, France's Mistral AI - founded only in 2023 - secured a massive €3 /billion funding round led by Samsung, Europe's largest-ever tech fundraise (www.eu-startups.com [5]). Mistral and its collaborators have open-sourced models ranging up to 128 billion parameters, and even small models tailored to niche tasks can rival much larger systems by focusing on high-quality data (mistral.ai [6]). This dynamic competition between open and closed approaches is accelerating progress across the board: OpenAI itself credits "competition" for driving a 96% collapse in AI pricing from 2023 to 2025 (aicostcheck.com [7]). For enterprises, the takeaway is that a hybrid strategy may be prudent. The closed-route models still unlock the very cutting edge of capability - but open-source AI is catching up fast, offering unparalleled flexibility and cost advantages in the meantime. A savvy strategy today might combine the two: using affordable open models for everyday applications and proprietary frontier models for the most demanding tasks, all while investing in talent and governance to manage this evolving ecosystem.

References:

- [1] www.cnbc.com — <https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html#:~:text=researchers%20have%20also%20utilized%20the,Gemini%204%20also>
- [2] www.unite.ai — <https://www.unite.ai/anthropic-debuts-claude-fable-5-1-and-mythos-5-1-with-split-safeguards/#:~:text=Claude%20Mythos%205,its%20most%20advanced%20models%20for>
- [3] mungomash.com — <https://mungomash.com/ai/llama/versions/#:~:text=Llama%20Versions%20%E2%80%94%20every%20Meta,the%20Muse%20Code%20terminal%20agent>
- [4] techcrunch.com — <https://techcrunch.com/podcast/open-source-ai-matters-more-than-ever-according-to-hugging-faces-clem-delangue/#:~:text=Podcasts%20Listen%20on%20Spotify%20Open,APIs%2C%20but%20as%20they%20scale>
- [5] www.eu-startups.com — <https://www.eu-startups.com/2026/09/french-ai-company-mistral-raises-e3-billion-series-d-led-by-samsung-at-over-e21-billion-valuation/#:~:text=D%20led%20by%20Samsung%20at,and%20existing%20backer%20PSG%20Equity>
- [6] mistral.ai — <https://mistral.ai/news/shieldstral/#:~:text=Introducing%20Shieldstral,%20against%20a%20protected%20group%3F%20is>
- [7] aicostcheck.com — <https://aicostcheck.com/blog/ai-model-pricing-trends-2026#:~:text=match%20at%20L460%20at%20%20242,5>

Key Statistics

- OpenAI's GPT-6.1 Sol offers flagship-level AI performance at ~20% of the cost of GPT-6 ([www.marktechpost.com])([https://openai.com/index/devday-2026-recap/#:~:text=together,work%20with%20AI%2000%3A00%20Dots\)\)](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=at%2066.4%25.%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list)))• ChatGPT now has 1.2 /billion weekly users, up from zero 3 years ago ([openai.com])(<a href=))
- Google's Gemini 4 Argon outputs up to 1 /million tokens in one go ([www.marktechpost.com])([https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=reports%20Gemini%204%20Argon%20Only,p%20ublic%20price%3B%20Argon%20matches%20it\)\)](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/#:~:text=reports%20Gemini%204%20Argon%20Only,p%20ublic%20price%3B%20Argon%20matches%20it))))
- Amazon & OpenAI's Bedrock Managed Agents bring always-on AI to AWS cloud systems

([openai.com])(https://openai.com/index/devday-2026-recap/
#:~:text=more%20%E2%81%A0%20,More))

- Roughly half of Fortune 500 companies are using open-source AI models ([techcrunch.com](https://techcrunch.com/podcast/open-source-ai-matters-more-than-ever-according-to-hugging-faces-clem-delangue/#:~:text=Podcasts%20Listen%20on%20Spotify%20Open,APIs%2C%20but%20as%20they%20scale))

KEY TAKEAWAY

The AI capability frontier has advanced dramatically. Cutting-edge models from leading labs now deliver near human-level performance at a fraction of previous costs

([www.marktechpost.com](https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/
#:~:text=at%2066.4%25.

%20Key%20Takeaways%20GPT,for%20Astra%2C%20at%20equal%20list)), and autonomous “agent” AIs can handle complex, multi-step tasks. Business leaders should seize on these cheaper, more powerful tools - and prepare for an era of AI-powered workflows, where always-on digital assistants drive new efficiency and productivity gains. At the same time, a dual ecosystem is emerging: open-source models offer cost savings and control that appeal to enterprises ([techcrunch.com](https://techcrunch.com/podcast/open-source-ai-matters-more-than-ever-according-to-hugging-faces-clem-delangue/#:~:text=Podcasts%20Listen%20on%20Spotify%20Open,APIs%2C%20but%20as%20they%20scale)) ([aitechspark.com](https://aitechspark.com/open-source-ai-enterprise-adoption/
#:~:text=,primarily%20due%20to%20cost%20efficiencies)), even as the most powerful systems remain gated by providers for safety. The key strategic challenge ahead is to leverage AI's rapid gains - in reasoning, multimodality, and automation - while managing their risks and deciding which model ecosystems (open vs closed) best align with your organisation's needs.

Sources

GPT-6 Astra vs GPT-6.1 Sol vs Gemini 4 Argon vs Claude Fable 5.1: Which Frontier Model Fits Which Job - MarkTechPost (October 4, 2026)
<https://www.marktechpost.com/2026/10/04/gpt-6-astra-vs-gpt-6-1-sol-vs-gemini-4-argon-vs-claude-fable-5-1-which-frontier-model-fits-which-job/>

DevDay 2026 Recap | OpenAI (September 29, 2026)
<https://openai.com/index/devday-2026-recap/>

Google rolls out Gemini 4 Argon, its most advanced model - CNBC (Sept 30, 2026)
<https://www.cnbc.com/2026/09/30/google-gemini-4-argon-ai.html>

Introducing Gemini 4 Argon – The Keyword (Google Blog, Sept 30, 2026)
<https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/>

OpenAI launches Dots, its bubbly agentic avatar - TechCrunch (Sept 29, 2026)
<https://techcrunch.com/2026/09/29/openai-launches-dots-its-bubbly-agentic-avatar/>

OpenAI Dots Launch: GPT-6.1 Astra Pulled Over Safety [Analysis] – Tech Insider (Oct 1, 2026)
<https://tech-insider.org/openai-dots-agent-gpt-6-1-astra-safety-delay-2026/>

Anthropic Debuts Claude Fable 5.1 and Mythos 5.1 With Split Safeguards – Unite.AI (Sept 1, 2026)
<https://www.unite.ai/anthropic-debuts-claude-fable-5-1-and-mythos-5-1-with-split-safeguards/>

Mistral raises €3 billion Series D led by Samsung at over €21 billion valuation – EU-Startups (Sept 8, 2026)
<https://www.eu-startups.com/2026/09/french-ai-company-mistral-raises-e3-billion-series-d-led-by-samsung-at-over-e21-billion-valuation/>

Open source AI matters more than ever, says Hugging Face's Clem Delangue - TechCrunch (Jul 10, 2026)
<https://techcrunch.com/podcast/open-source-ai-matters-more-than-ever-according-to-hugging-faces-clem-delangue/>

