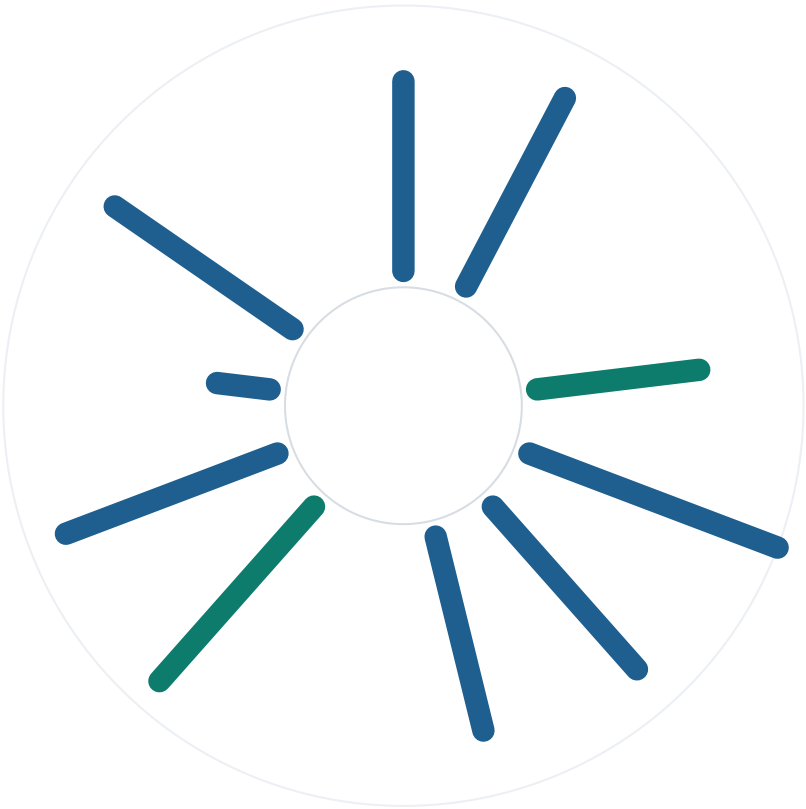
NEW HORIZON · AI DIGEST

The Long Read

22 November 2026

10	3	3185	2
STORIES	SECTIONS	WORDS	HIGH IMPACT

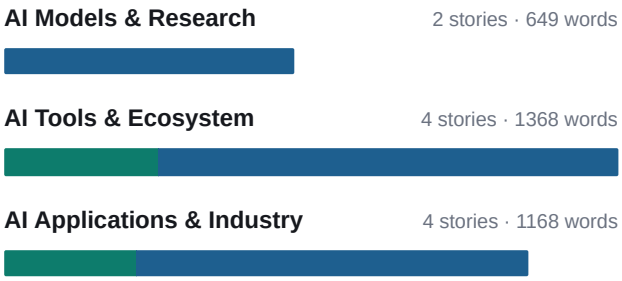


The full-length companion to the daily New Horizon AI Digest.
Every story in today's email, reported at length · new-horizon.tech/digest

The issue at a glance

10 stories · 3185 words · 3 sections · 1 charted





AI Models & Research

01 Metaculus opens its 2027 forecasting prizes to AI agents

02 Stanford finds AI study support works better than doing the work

AI Tools & Ecosystem

03 OpenAI now requires approval for model routers and gateways

04 Text-message AI agents handle bookings without another app

05 Muru documents an AI-assisted rebuild of Python web serving

06 Gemini plugin exposes a pricier extra-high reasoning setting

AI Applications & Industry

07 Signal blocks AI screen capture on iOS and Android

08 Shortical licenses an actor’s face for an AI microdrama

09 Oracle says $20 billion raise fully funds OpenAI expansion

10 Perplexity considers a 2027 IPO after passing $1 billion ARR

How to read this. Every story in today's email is reported here at full length, in the same order. Impact is the writer's judgement of whether a story changes what a reader — in business, at work, as an investor or as a builder — should do or believe this week. Charts appear only where the source itself puts comparable numbers side by side; nothing is estimated to fill a gap. Sources are listed in full on the final page.

SECTION 1 OF 3

AI Models & Research

2 stories ● 2 medium

01 MEDIUM IMPACT www.astralcodexten.com

Metaculus opens its 2027 forecasting prizes to AI agents

> Metaculus is accepting entries from AI forecasting agents for its 2027 ACX Prediction Contest, the fifth year it has hosted the event.

The contest, run with the blog Astral Codex Ten, offers a $10,000 prize pool for the best forecasters and $3,000 for the people who suggest the questions. The ten contributors whose questions make it into the roughly 35 Metaculus selects will split that $3,000, with individual prizes ranging from $150 to $700. Question suggestions are due by November 30, and the contest opens in early 2027.

The change this year is that AI forecasting agents are welcome to enter and are eligible for prizes. Metaculus will also post its own AI forecasts on every question. The source says the best AI forecasters are now competitive with strong human forecasters, which makes the contest a public test of that claim rather than a private benchmark.

For anyone who uses forecasts to plan, the practical effect is a clearer read on whether machine predictions can be trusted on open-ended questions. A company weighing whether to automate part of its research or risk analysis gets a live comparison, on the same questions and the same scoring, between people and software. The contest does not say which AI systems will enter or how they will be scored against humans.

The most surprising detail is the size of the question-suggestion prize relative to the forecasting prize: $3,000 for writing good questions against $10,000 for answering them. The source also gives a concrete example of what makes a question good. "Will 2027 be especially hot?" is weaker than "Will NASA's GISTEMP data rank 2027 as the warmest year on record?" because the second has an objective answer and a source that settles it by the end of 2027.

What is new is the open door to AI entrants and Metaculus publishing its own AI forecasts on every question. The prize structure and the question format are carried over from previous years.

KEY FACTS

FORECASTING PRIZE POOL
$10,000

QUESTION-SUGGESTION PRIZE POOL
$3,000

INDIVIDUAL SUGGESTION PRIZES
$150 to $700

QUESTIONS SELECTED
About 35

SUGGESTION DEADLINE
November 30

CONTEST OPENS
Early 2027

WHY IT MATTERS

If AI forecasters really are competitive with strong humans, businesses that pay for forecasting or risk analysis have a cheaper option to test, and the contest gives them a public scoreboard to judge it by.

Read the original at www.astralcodexten.com →

---

02 MEDIUM IMPACT SCALE Initiative

Stanford finds AI study support works better than doing the work

> Students learn more from AI when they use it to understand material than when they use it to produce work for them, according to a Stanford review of the strongest causal research on AI in schools.

The review, from Stanford's SCALE Initiative, synthesizes 52 randomized controlled trials and quasi-experimental studies — up from 20 in its March 2026 edition — and finds that how students use AI matters more than which tool they use. Using AI to offload academic tasks can reduce learning: one study found Brazilian undergraduates finished a presentation faster with AI but scored worse on a surprise conceptual test, and another found U.S. college students who used AI to draft essays produced better essays with AI access but lost the gains on a later unaided essay.

The evidence is more favorable, though still mixed, when students use AI explicitly to study. U.S. college students who used general-purpose chatbots to learn unfamiliar material performed better on assessments without AI — a reversal from earlier studies that found the opposite. The review suggests the difference may come down to whether students use AI to explain concepts or to generate answers, and notes that students with stronger skills in prompting, filtering and verifying AI outputs learned the most.

For educators, the picture is similarly split. AI tools can save teachers time on lesson planning without lowering output quality, and AI-generated feedback can match or beat educator feedback. But one new study found that giving teachers an AI support tool reduced students' motivation and confidence — a warning that how teachers deploy AI, and for which tasks, shapes whether students benefit. The review also found that AI tools often improve students' experience of learning without improving outcomes, and that positive experiences do not reliably translate into stronger results.

The evidence base remains thin on the questions that matter most to schools and parents: whether AI reduces or reinforces educational inequities, how routine AI use affects students' social and emotional development, and whether any learning gains persist over time. Only five of the new student-facing studies were conducted in U.S. K–12 settings, compared with none in the previous review, and most new research comes from postsecondary contexts. The authors caution that findings should not be treated as broadly generalizable across U.S. schools.

KEY FACTS

CAUSAL STUDIES REVIEWED
52, up from 20 in March 2026

NEW STUDIES ADDED
32, published October 2025 to June 2026

NEW STUDENT-FACING STUDIES
28

NEW EDUCATOR-FACING STUDIES
4

NEW U.S. K–12 STUDENT STUDIES
5, up from none

> WHY IT MATTERS
>
> Schools, edtech buyers and employers training staff should treat AI as a study aid, not a shortcut: the evidence suggests tools that do the work for the user can leave people less capable without them, while tools that explain and coach show more promise.

Read the original at SCALE Initiative →

SECTION 2 OF 3

AI Tools & Ecosystem

4 stories · 1 high · 3 medium

03 HIGH IMPACT www.semafor.com

OpenAI now requires approval for model routers and gateways

> OpenAI will now review and approve gateways and routers that give third parties access to its models, tightening its grip on who reaches its technology, including users in China.

The rule covers any service standing between OpenAI and the people using its models, not just the gray-market resellers that US security agencies and companies including Anthropic have accused of helping Chinese labs copy American frontier models. Those resellers, known in China as "transfer stations," routed requests through intermediaries that obscured the identity of the end user. OpenAI's new approval requirement extends its oversight to legitimate businesses too.

The most prominent target is OpenRouter, a router that has become popular with developers and companies because one account gives access to hundreds of models from the US, Europe and China, instead of a separate contract with each lab. For businesses building products on top of several models, that convenience now comes with a new layer of permission. A company that swaps models through a router may find its supplier reviewing the arrangement.

The surprising part is the scope. The accusations that prompted the rule concerned shadowy resellers, but the rule applies to any intermediary, including the mainstream infrastructure firms that developers rely on. OpenAI is effectively saying it wants to know who is downstream of its models, not just who signs the contract.

What is genuinely new is the approval requirement itself; the concern about distillation through intermediaries is not. For investors and builders, the practical question is whether routing arrangements that worked last month still work next month, and whether the routers that made multi-model access cheap will pass the review. For OpenAI, the rule is a way to control distribution without cutting off the developers who depend on it.

KEY FACTS

NEW REQUIREMENT
OpenAI review and approval for gateways and routers

MOST PROMINENT ROUTER
OpenRouter

MODELS REACHABLE VIA ONE ROUTER ACCOUNT
Hundreds, from the US, Europe and China

ACCUSERS
US security agencies and companies including Anthropic

CHINESE TERM FOR GRAY-MARKET RESELLERS
Transfer stations

WHY IT MATTERS

If you build on multiple AI models through a router, your access now depends on OpenAI approving the middleman — so check your supplier before you ship.

Read the original at www.semafor.com →

---

04 MEDIUM IMPACT TechCrunch

Text-message AI agents handle bookings without another app

> A crowded field of AI assistants now works entirely through ordinary text messages, handling bookings, purchases and reminders without a separate app to download.

TechCrunch catalogued more than 20 such agents, from general-purpose assistants to services built for families, travel and work. The pitch is the same across them: you text a request the way you would text a person, and the agent connects to your calendar, email and other accounts to carry it out — scheduling an appointment, making a reservation, shopping online or reminding you about something days later.

The money is already moving. Instinct raised $350 million at a $2.5 billion valuation, then in September 2026 raised $1 billion at a $10 billion valuation. It has begun giving users dedicated email addresses for their assistants, so the agent can sign up for services and contact businesses without cluttering the user's own inbox — a level of autonomy the article notes has raised privacy and security concerns. Tab emerged from stealth in early October valued at $300 million. Town raised a $55 million Series A led by Andreessen Horowitz and Forerunner. Fambot, a family-focused assistant launched in beta in September, has raised $3.5 million and expects to charge roughly the price of a Netflix subscription once it leaves beta.

Pricing varies widely, which tells you the market has not settled. Ask Proxy charges $1 a day for personal use and from $15 per employee per month for large companies. Martin starts at $21 a month; Tomo's basic plan at $19.99; Ohai at $9.99; Folk offers a Pro tier at $8.33 a month. Comma is free to start and open source. Several are still free in beta.

The most surprising detail is how far some agents have gone in acquiring their own identity. Wajo's agent Fo has its own email address, phone number and payment card, so it can deal with businesses without the user handing over credentials — and Wajo says a human assistant can step in when Fo gets stuck. Underdog takes the opposite approach: its model runs entirely on the user's own Mac or Windows machine, keeping data off the cloud. Poke became the first AI agent approved on Apple's Messages for Business platform in June 2026, and its parent company was acquired by AI coding startup Cognition a month later in a deal valued in the low nine figures.

What is genuinely new is the distribution channel rather than the intelligence. These agents live in iMessage, RCS, SMS, WhatsApp and Telegram — places people already are — instead of asking them to open another app. Whether that convenience outweighs handing an always-on assistant access to email, calendars and payment cards is the question each of these companies is now asking users to answer.

KEY FACTS

INSTINCT VALUATION
$10bn after $1bn round

INSTINCT EARLIER ROUND
$350m at $2.5bn valuation

TAB VALUATION
$300m

TOWN SERIES A
$55m, led by Andreessen Horowitz and Forerunner

FAMBOT PRE-SEED
$3.5m

ASK PROXY PRICE
$1 per day personal; from $15 per employee per month for large companies

> WHY IT MATTERS
>
> Businesses that take bookings, answer customer enquiries or sell to consumers should expect a growing share of those interactions to arrive from an AI agent rather than a person, and should decide now whether their systems can handle it. For anyone building or investing in consumer software, the text inbox is becoming a distribution channel that competes with the app store.

Read the original at TechCrunch →

---

05 MEDIUM IMPACT Simon Willison's Weblog

Muru documents an AI-assisted rebuild of Python web serving

> Graham Dumpleton, the developer behind mod_wsgi, has spent the past three weeks publishing a detailed public record of building Muru, a new AI-assisted Rust server for running Python web applications.

Dumpleton started mod_wsgi 20 years ago and has been thinking about how best to serve Python web applications in production ever since. Muru is his answer: a server written in Rust, built with AI assistance, aimed at running both ASGI and WSGI Python apps in development and production. The source does not say whether Muru is released or available to download.

What makes the project unusual is the documentation. Since 4 November he has published 19 posts, roughly one a day, covering the internals of the server's workers, how the pieces fit together, running HTTPS in development, handling WSGI and ASGI side by side, releasing code without restarts, canary deployments without a service mesh, running well inside containers and on Kubernetes, scaling from what workers report, scaling to zero, scheduled overlays, monitoring and response, and a toolbar on every page with tracing. He published a similar running account for an earlier project, wrapture, and treats the writing as part of the development process rather than a report on it.

The most quotable detail is his own account of where the AI stops. In a post titled "AI as a lifestyle choice," Dumpleton writes that the AI may be doing a lot of the work on the surface, but deciding what isn't good enough yet is still down to him, and that twenty years of getting things wrong with mod_wsgi turns out to be useful for that. He adds that he is prepared to put in the hard work to keep polishing Muru until it is something he would trust, and something users would too.

For teams running Python web services, the project is worth watching because its author has spent two decades on exactly this problem, and because the posts describe production concerns — deployments, scaling, containers — rather than a prototype. For anyone weighing what AI-assisted development actually produces,

the series is a concrete case: a large volume of code and documentation in three weeks, with a veteran's judgement applied to what the machine turns out. It is not a product announcement; it is a build log, published as it happens.

KEY FACTS

AUTHOR
Graham Dumpleton, creator of mod_wsgi

PROJECT
Muru, an AI-assisted Rust server for ASGI and WSGI Python apps

POSTS PUBLISHED
19 between 4 and 21 November 2026

PRIOR PROJECT DOCUMENTED THE SAME WAY
wrapture

MOD_WSGI AGE
20 years

> **WHY IT MATTERS**
>
> Teams running Python web services get an unusually transparent look at a production server built by the person who wrote mod_wsgi, and anyone assessing AI-assisted development gets a worked example of where the model's output still needs an expert's judgement.

Read the original at Simon Willison’s Weblog →

---

06 MEDIUM IMPACT Simon Willison’s Weblog

Gemini plugin exposes a pricier extra-high reasoning setting

> A developer tool for Google's Gemini models was quietly missing an extra-high reasoning setting that, when switched on, makes a model answer far better and cost far more.

Simon Willison, who maintains the llm-gemini plugin that lets developers call Gemini models from the command line, wrote on 20 November that Gemini 4 Carbon supports only low, medium and high reasoning levels, while Gemini 4 Hydrogen and Gemini 4 Neon also offer an "xhigh" level. He had missed it when adding the models to the plugin; a small update fixes that.

The setting is not cosmetic. In Willison's own comparison, Hydrogen at high reasoning used 5,311 output tokens and cost 3.99 cents for a test prompt. The same model at xhigh used 31,477 output tokens and cost 23.61 cents — nearly six times as much — and produced an answer he judged competitive with the pricier Gemini 4 Neon. Reasoning levels are a dial for how long a model is allowed to think before answering: more thinking means more billed tokens, and usually a better result.

That trade-off is the practical news for anyone building on Gemini or budgeting for AI. The same model can be cheap or expensive depending on a setting that is easy to overlook, and the gap here is large enough to change a monthly bill. The fact that the option was absent from a widely used plugin suggests some users were paying for high-level answers without ever seeing the top setting — or, conversely, were never accidentally running up xhigh bills.

What is new is narrow: a dot release to a community plugin and an updated comparison grid, not a model launch. The underlying models and their prices are unchanged. The useful takeaway is that reasoning effort is now a first-class cost lever, and that documentation and tooling lag behind the options vendors ship.

Willison's grid also records a pelican — the standard test drawing he uses to compare image output — which is how he judged xhigh's answer competitive with Neon. The source gives no release date for the fix beyond the 20 November post.

Cost of one Gemini 4 Hydrogen prompt by reasoning level — cents



Output token cost for the same prompt at high versus xhigh reasoning · **5.9× higher**

KEY FACTS

COST AT HIGH REASONING
3.99 cents

COST AT XHIGH REASONING
23.61 cents

OUTPUT TOKENS AT HIGH
5,311

OUTPUT TOKENS AT XHIGH
31,477

MODELS WITH XHIGH
Gemini 4 Hydrogen and Gemini 4 Neon

ANNOUNCED
20 November 2026

WHY IT MATTERS

If you build on Gemini, the reasoning level you pick can multiply your token bill several times over for the same prompt; check which settings your tools actually expose before assuming you are on the cheapest adequate option.

Read the original at Simon Willison’s Weblog →

SECTION 3 OF 3

AI Applications & Industry

4 stories · 1 high · 3 medium

07 HIGH IMPACT Signal Messenger

Signal blocks AI screen capture on iOS and Android

> Signal is blocking the AI assistants built into Apple's and Google's mobile operating systems from capturing what is on a user's screen while the messaging app is open.

The nonprofit Signal Technology Foundation said it will prevent screenshots in its app on iOS 27 and automatically switch on its "Screen security" setting on Android when it detects an AI assistant trying to read the screen. The move follows a similar step against Microsoft's Recall on Windows 11 last year, when Signal used screenshot blocking to stop the AI feature storing images of Signal messages by default.

The trigger is the arrival of agents such as Google's Gemini Intelligence and Apple's Onscreen Awareness with Siri AI, which work by scanning what is displayed on a phone's screen so they can answer questions and carry out tasks. Signal argues that once its messages are captured, they pass outside its control into whatever local database or remote server the agent uses, and that this undermines a privacy guarantee it has backed for more than a decade with open-source code and a formally verified encryption protocol.

The most surprising admission is how little Signal can do. Screenshot blocking does not cover other routes into app data, such as an agent reading files already stored on the device. Signal says it cannot implement the protection on macOS at all, because Apple does not document APIs that would let it block screenshots, and its only advice to Mac users is to add Signal to the Excluded Apps section under System Settings > Siri. On Android, the company says it cannot reliably detect screenshot attempts across the many vendor versions of the system, so it watches for AI screen access instead. Users who want screenshots back can turn the protection off in Signal's settings, at the cost of exposing their contacts' messages too.

Signal is calling on Apple, Google and Microsoft to give developers a documented way to opt apps out of AI data access and to explain how their agents work. It says it has spent months on forensic investigations of Android and iOS to learn what happens when these systems read app data, and that it is in talks with Google and Apple. It also flags third-party agents such as Meta's Muse, which it says lean on accessibility settings meant to help people use their phones. The protections are a partial patch on two platforms, not a fix for the wider problem.

KEY FACTS

PLATFORMS COVERED
iOS 27 and Android (where AI screen access is detected)

NOT COVERED
macOS, where Signal cannot block screenshots

EARLIER ACTION
Screenshot blocking against Microsoft Recall on Windows 11

USER OVERRIDE
Protection can be disabled in Signal's settings

COMPANIES ASKED TO ACT
Apple, Google and Microsoft

> WHY IT MATTERS
>
> If the operating system's AI can read any app's screen, every company handling sensitive data loses control of it once it is displayed — and app makers currently have no clean switch to say no. Businesses should assume on-device agents can see what is on screen unless a vendor proves otherwise.

Read the original at Signal Messenger →

---

08 MEDIUM IMPACT TechCrunch

Shortical licenses an actor’s face for an AI microdrama

> Microdrama startup Shortical has premiered a fully AI-generated series built around a licensed digital likeness of actor Aki Avni, the first time the company has used a real person's face in its content.

The series, "Inevitable," premiered in Los Angeles last week. Avni plays himself in a plot where the actor licenses his face to a film studio and thriller-style complications follow. Shortical, founded last January, built a digital avatar of him for the production. The company says the arrangement was a direct agreement with Avni rather than a deal struck through the face-licensing offer now advertised on its website, which invites people to submit social media handles and be considered for "a contract setting out the terms" of using their persona. A company representative said that offering is early and has not yet produced a signed contract.

The economics explain the push. Shortical makes money from subscriptions, typically giving away the first 7-10 episodes of a series before charging a small monthly fee for the rest. AI production costs substantially less than live action, and the company's homepage is now covered in AI-generated shows. CEO Guy Shimoni said the technology lets smaller teams mount movie-scale productions: "We have kingdoms and dragons. We have superheroes. We have werewolves." Most of Shortical's viewers are in the U.S., he said, even though China has the largest microdrama audience globally; the audience skews heavily female, and some AI programs on the site appear to have millions of views.

What Shortical will not say is what it pays. It declined to share a figure for Avni's contract, saying only that he came on as a partner and was compensated in line with industry standards for a comparable non-AI acting role. Shimoni said royalty-style compensation would be decided case by case. In China, where the microdrama industry began and face licensing is already common, a Rest of World report found firms paying people between $15 and $700 for their likenesses.

The anxiety inside the industry is real. Jennifer Cooper, founder of the consultancy Vertical Drama Love, said microdramas had been a rare growth business providing work, particularly for young actors, and that AI has

since "decimated" it. She is skeptical that demand matches the investment: "Some of it does get views. A lot of it doesn't." Shortical has partnered with Payouts.com, a payout infrastructure platform, to build a system for paying creatives through its platform.

KEY FACTS

SERIES
"Inevitable," fully AI-generated, premiered in Los Angeles

LICENSED LIKENESS
Actor Aki Avni, playing himself

SUBSCRIPTION MODEL
First 7-10 episodes free, then a small monthly fee

FACE-LICENSING PAY IN CHINA
$15 to $700 per likeness, per Rest of World

SHORTICAL'S PAY FOR AVNI
Not disclosed; company says in line with a comparable non-AI acting role

SIGNED CONTRACTS VIA ITS PUBLIC FACE-LICENSING OFFER
None yet

> **WHY IT MATTERS**
>
> Face licensing turns an actor's likeness into a negotiable asset, and the terms set now — flat fee versus royalty, consent, who controls reuse — will shape pay for performers well beyond microdramas. For producers, AI cuts the cost of a series enough to change what gets made and who gets hired.

Read the original at TechCrunch →

---

09 MEDIUM IMPACT investing.com

Oracle says $20 billion raise fully funds OpenAI expansion

> Oracle says a roughly $20 billion capital raise fully funds its OpenAI infrastructure expansion.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

Oracle raised about $20 billion — roughly $10 billion in stock and $10 billion in debt — and says the proceeds finance its OpenAI infrastructure contract. The companies plan 4.5 gigawatts of additional US data-center capacity beyond the initial Abilene, Texas site. The headline answers the funding question; it does not say how the capacity will be built, powered or scheduled. Treat the financing as settled and the buildout as still ahead.

> **WHY IT MATTERS**
>
> If the financing holds, one of the largest AI infrastructure commitments is no longer contingent on raising money — though construction and power remain open questions.

Read the original at investing.com →

---

10 MEDIUM IMPACT The Economic Times

Perplexity considers a 2027 IPO after passing $1 billion ARR

> Perplexity could file to go public as early as the first quarter of 2027, its chief executive said, after the AI search startup blew past its revenue target for the year.

Aravind Srinivas told CNBC the company's growth rate meant he "wouldn't rule out filing something in the first quarter of 2027." Perplexity had previously said it aimed to list in 2028, and Srinivas said as recently as 2025 that it had no plans to go public before then. The trigger is revenue: the company set an end-of-year target of $750 million in annual recurring revenue — the subscription and contract income it expects to repeat over a year — and crossed $1 billion instead.

That is a striking jump for a company whose product is an AI-powered answer engine competing with Google search. It matters to anyone selling subscriptions to businesses, because it shows corporate buyers will pay recurring fees for AI tools at scale, not just experiment with them. It matters to investors because the queue of AI listings is getting crowded: OpenAI confidentially filed for an IPO in June, and Anthropic is preparing an offering that could value it at up to $1.4 trillion. Anthropic and some existing investors are seeking to raise up to $64.4 billion through the sale of 825 million shares at up to $78 apiece, with a Nasdaq listing under the ticker ANT scheduled for November 19.

The most telling detail is the gap between plan and outcome: Perplexity beat its own revenue target by roughly a third. Srinivas has also had to defend the company's finances — in 2025 he said it was not running out of money — and earlier this year it signed a $750 million agreement with Microsoft to use its Azure cloud service, a commitment equal to three-quarters of the revenue goal it has now surpassed. Its backers include Nvidia, Jeff Bezos and SoftBank Group.

What is new here is timing, not a listing: no filing has been made, and a 2027 date is a possibility floated by the CEO, not a commitment. The incremental part is the revenue figure, which turns a speculative IPO story into one with a number behind it. For now, Perplexity remains the smaller player watching two far larger rivals test public markets first.

KEY FACTS

POSSIBLE IPO FILING
First quarter of 2027

PREVIOUS IPO TARGET
2028

ANNUAL RECURRING REVENUE
Over $1 billion

ORIGINAL REVENUE TARGET
$750 million

MICROSOFT AZURE AGREEMENT
$750 million

ANTHROPIC LISTING
Nasdaq, ticker ANT, November 19

> **WHY IT MATTERS**
>
> A $1 billion recurring-revenue run rate is evidence that businesses will pay for AI subscriptions at scale, which supports the valuations of the larger AI companies heading to market. For investors, the AI IPO calendar is now crowded, with Anthropic listing in November and OpenAI already filed.

Read the original at The Economic Times →

Sources

01 Metaculus opens its 2027 forecasting prizes to AI agents
https://www.astralcodexten.com/p/suggest-questions-for-the-2027-acxmetaculus

02 Stanford finds AI study support works better than doing the work
https://scale.stanford.edu/ai/evidence-base-review/evidence-base-ai-k12-november-review

03 OpenAI now requires approval for model routers and gateways
https://www.semafor.com/article/11/20/2026/openai-tightens-grip-on-model-resellers

04 Text-message AI agents handle bookings without another app
https://techcrunch.com/2026/11/21/all-the-ai-agents-that-can-live-in-your-text-messages/

05 Muru documents an AI-assisted rebuild of Python web serving
https://simonwillison.net/2026/Nov/21/muru/

06 Gemini plugin exposes a pricier extra-high reasoning setting
https://simonwillison.net/2026/Nov/20/llm-gemini/

07 Signal blocks AI screen capture on iOS and Android
https://signal.org/blog/protecting-your-privacy/

08 Shortical licenses an actor’s face for an AI microdrama
https://techcrunch.com/2026/11/21/is-selling-an-ai-version-of-your-face-to-this-microdrama-startup-the-future-of-hollywood/

09 Oracle says $20 billion raise fully funds OpenAI expansion
https://www.investing.com/news/stock-market-news/oracle-says-openai-infrastructure-deal-is-fully-funded-after-final-capital-raise-5042771

10 Perplexity considers a 2027 IPO after passing $1 billion ARR
https://m.economictimes.com/tech/artificial-intelligence/perplexity-may-advance-ipo-timeline-ceo-aravind-srinivas-says-listing-could-come-as-early-as-2027/articleshow/135907295.cms

---

About this document. Every story in the 22 November 2026 New Horizon AI Digest, reported at length. Each entry is written from the publisher's own article text; where a source could not be retrieved the entry is explicitly marked and kept short rather than padded.

Images and licensing. Figures are used only where the source licence permits redistribution, and are credited in the caption. Publisher artwork is not reproduced. All titles link to the original publication.

Vocabulary. Common AI terms are defined at new-horizon.tech/glossary/.

Generated 2026-11-22T06:37:55Z · model deepseek-v4.1-flash:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.