

Why Enterprise AI Budgets Keep Growing Despite AI Safety Pledges

ETR Data Suggest Anthropic and OpenAI's Pacing Pledges Widen the Gap with Rivals More Than They Curb Real Risk

Analyst(s): Nick Patience

Publication Date: September 30, 2026

Document #: AIONP202609

What You Need to Know

- **Budgets Are Not Responding to the Alarm:** IT budgets for 2027 open at +5% growth, the strongest initial read since October 2022, with 76% of respondents raising 2026 budgets and 61% of the Global 2000 already ahead of plan (ETR Macro Views, N=1,418, preliminary).
- **Buyers Doubt the Motive but Want the Verification:** Among large-enterprise respondents in ETR's AI Pacing Flash Study (N=50, fielded September 17), 46% call the labs' pacing proposals mostly or entirely competitive positioning, against 20% who see genuine safety measures. Even so, 74% say the pledge itself makes them more confident building on the labs' models.
- **Both Labs Are Growing on New Money:** Among enterprises increasing spend on OpenAI or Anthropic in September's ETR Edge waves, 68% call it net-new or incremental budget, up from 52% in June.
- **Anthropic's First Customers Are Cutting Spend:** Out of 39 respondents, 4 cut Anthropic spend in September, versus 0 of 42 in June, citing price and policy concerns.
- **Nobody Outside the Two Labs Has Joined the Pledge:** Google, Meta, and xAI have made none of the same commitments, which keeps the pacing story confined to the two vendors already furthest ahead.

The Futurum View

Anthropic and OpenAI's pacing pledges look less like safety measures and more like a way to raise the price of competing with them.

This is a familiar pattern outside AI, too: a company with money to spare backs a costly rule or standard, public-spirited framing included, that a smaller rival can't absorb as easily. Doug McMillon, Walmart's CEO until his retirement in January 2026, repeatedly called for a higher federal minimum wage while in the role, most recently in 2021; Axios has reported that a wage

floor Walmart absorbs easily is often enough to put smaller, local competitors out of business. Tesla has spent years pushing regulators to keep strict vehicle-emission rules in place, including in comments to the Trump administration in September 2025, warning that any rollback would cost it billions of dollars in the regulatory credits other automakers pay it to help them comply.

But both of those were government mandates applied to the whole industry, and that's a real difference: nothing forces Google, Meta, or xAI to match Anthropic and OpenAI's pledge today, which makes it closer to a costly certification than a rule imposed on the market.

That could change, as Amodei's essay calls for an antitrust waiver that would let labs coordinate on pacing together, which would tighten the parallel considerably if it happens. Both labs are now offering something a smaller rival cannot easily match: standing access for outside evaluators, a public promise to slow down, and the infrastructure needed to prove it. The barrier isn't the size of the balance sheet – Google and Meta both have deeper pockets than either lab. It's how concentrated the bet is. Anthropic and OpenAI have staked the whole company on being at the frontier, so the cost of the pledge is trivial next to that bet: Anthropic alone has locked in multi-gigawatt compute commitments this year with Amazon, Broadcom, and SpaceX worth tens of billions of dollars, and funding a handful of embedded evaluators barely registers against outlays of that size. For Google and Meta, the same commitment has to compete against search, cloud, ads, and every other business line for budget and attention, which makes an open-ended compliance program a harder sell internally, even when the money is there. The practical effect, whatever the intent behind it, is to widen the gap between the two labs already at the front and everyone else trying to reach them – Google and Meta included, China's labs even more so.

The clearest test of how costly this is actually meant to be is what each lab has done with its own listing plans. OpenAI has publicly deferred its 2026 IPO to 2027, with Altman calling it an ill-advised moment, given the safety questions in play. That's costly, as it obviously postpones access to public capital. Anthropic has made no equivalent move; by its own reporting, it is still pressing ahead toward a listing while making the identical pacing commitment. If the pledge were principally about buying time for alignment work to catch up, both labs carrying the same risk would be expected to make the same trade-off, but only one has.

Futurum's ETR September data supports this reading in two ways. First, enterprise buyers already doubt the labs' motives, without any prompting: 46% call the pacing proposals mostly or entirely competitive positioning, against just 20% who see them as genuine safety measures. Second, that doubt hasn't cost either lab anything. Enterprise AI budgets and vendor funding kept climbing throughout the debate.

But that skepticism begs the question of whether or not buyers are right to feel this relaxed. The same survey found 74% say a lab's pacing commitment makes them more confident building on its models, and only 1 of the 36 respondents who want more AI regulation gave pacing as a reason to want less of it. If frontier capability really is advancing as fast as Amodei's essay claims, a market that treats an unaudited pledge as enough reassurance isn't showing good judgment about the risk. It may simply show that a pledge, audited or not, is enough to

keep the money coming. Both explanations fit the same numbers. What the data actually proves is narrower than either one: whichever reading is closer to true, it hasn't moved a single budget line so far.

Budgets Aren't Listening to the Debate

The political argument over pacing has been loud – the White House now calls concern about AI risk a hoax at worst or a partisan position at best. But the number that would show real damage to demand moved the other way. The 2027 IT budgets open at +5% growth – the strongest first read ETR has recorded since October 2022. The 2026 budget expectations rose to +4.0% from +3.8% in July. Among respondents, 76% are raising budgets, a five-survey high. Global 2000 expectations improved 60 basis points to +3.6%.

Spending is also shifting from hardware to software. Software spend is up 3.9%, led by Security (+7.8%) and Data/BI (+6.5%). Cloud growth eased to +7.1% from 7.5%. Hardware growth fell to +2.7% from +4.0%, as memory prices rose (+7.3%) and enterprises pushed out refresh cycles.

Table 1: 2027 IT Budget Growth Against Recent History

Metric	Value
2027 IT Budget Growth (Initial Read)	+5.0%
2026 IT Budget Growth (Current)	+4.0% (up from +3.8% in July)
Global 2000, Sequential Change	+3.6% (+60 bps)
Software Spend Growth	+3.9%
— Security	+7.8%
— Data/BI	+6.5%
Cloud Spend Growth	+7.1% (down from 7.5%)
Hardware Spend Growth	+2.7% (down from 4.0%)

Source: Futurum Research (ETR), Macro Views survey, N=1,418 (preliminary), September 2026


What Buyers Say They'll Pay For

Buyers don't trust the labs' reasons, but they still want the verification. Among flash-study respondents, 70% say they'd be more willing to deploy AI in production if third-party evaluation were in place. A total of 60% would pay a premium for a certified model. Many of

these same buyers doubt why the labs offered evaluation in the first place. In other words, they want the check regardless of the motive behind it.

That check is harder to build than it sounds. A good evaluator needs to be independent, knowledgeable, and well-funded, and it's hard to be all three at once. METR can't cover two labs' full output forever, and most people qualified to evaluate a frontier model have already worked at one of the labs they'd be checking. Whether a second and third evaluator shows up, funded in a way that doesn't recreate that same conflict, will decide whether 'certified' ends up meaning something real or just becomes a label with nothing behind it.

Figure 1: AI Pacing Flash Study: Buyer Views on Motive and Mechanism




Source: Futurum Research (ETR), AI Pacing Flash Study, N=50 large enterprises (54% Global 2000), fielded September 17, 2026

Just 3% of the 36 respondents who want more AI regulation cited pacing as a reason to want less of it (1 of 36).

Both Labs Are Growing on New Money

Rather than relying on reallocated capital, expansion across both labs is driven by fresh funding sources. September's Edge waves (Figure 3), 68% of respondents increasing spend on OpenAI or Anthropic call it net-new or incremental money, up from 52% in June. OpenAI's increases are 82% net-new, up from 65%. Anthropic's are 60%. When Anthropic's budget does come from somewhere else, it's mostly coming from contractors and consultants (59%) and deferred hardware purchases (57%), not from OpenAI. In March, a third of Anthropic's increases came out of another AI vendor's budget, almost always OpenAI's. That share has dropped to 23%.

Figure 2: Funding Source of Combined OpenAI + Anthropic Spend Increases



Source: Futurum Research (ETR), Edge Commentary Waves, June/September 2026

More customers are still switching from OpenAI to Anthropic than the other way, though the numbers are small (Figure 4): 12 of the 14 respondents cutting OpenAI spend named Anthropic as the reason, mostly citing Claude Code, model quality, and cost. OpenAI has moved upmarket since March – customers spending \$100K to \$1M are now 47% of its base, up from 24% – and it's the second-best sequential improver on ETR's preliminary October TSIS Z-score screen, at +1.64, against Anthropic's seventh-worst reading of –1.21.

Anthropic also saw its first decrease in cohort this cycle: 4 respondents out of 39, versus 0 of 42 in June. They cited a pricing change, an IT policy restricting Claude over data-retention risk, unpredictable behavior in new releases, and doubts about ROI. Cost and pricing now show up in 31% of Anthropic's Edge responses, up from 19%. But the broader Macro data cuts against the ROI complaint: 52% of respondents report positive AI ROI with external AI solutions, and 17% now see sustained ROI at scale, up from 13%.

Table 2: September Vendor Breakdown

Vendor	Net-New Share (Sept)	Reallocated from Another AI Vendor (Sept)	For Comparison: March
OpenAI	82% (18 of 22)	0%	0% (never has been)
ANTHROPIC	60% (21 of 35)	23%	~33%

Source: Futurum Research (ETR), Edge Commentary Waves, March/June/September 2026

What to Watch

- **Whether More Evaluators Show Up:** It's already happening, but not the way the pledge implied. Anthropic named its first embedded evaluator on September 18 in the form of Faculty, Accenture's AI consulting arm, and is funding Faculty's work directly, with no published standards yet for what it can see or how it reports. However, that's a paid contractor, not an independent nonprofit such as METR. Whether a second nonprofit evaluator shows up, funded in a way that doesn't recreate this same conflict, will decide whether 'certified' means independent or just audited by someone on the payroll.
- **Whether Confidence Holds After an Incident:** The 74% confidence boost rests on a voluntary pledge, not an audited one. A containment failure during a sanctioned evaluation, in the mold of the Hugging Face incident, probably won't test it, as that's the kind of thing the evaluator pledge is designed to catch. An incident where a model coordinates or acts on its own initiative, unprompted, is the one that would actually show whether the number is real or just untested.
- **Whether Other Labs Join In:** Google, Meta, and xAI haven't committed to their own evaluators yet. Until they do, the competitive-distance effect stays limited to the two labs already ahead.
- **Whether Anthropic's Decreasers Grow:** Out of 39, 4 in September, up from zero in June – worth tracking whether this becomes a real trend or stays a rounding error against the 35 increasers.
- **Whether Buyers Switch to Open Weights:** Of flash-study respondents, 52% say they'd get more interested in open-source models if frontier progress slowed, against just 26% who'd use more compute. That's a way around the moat that doesn't depend on the frontier labs at all.

Other Insights from Futurum

[Anthropic's \\$65B Raise: Can Claude's Enterprise Surge Justify a \\$965B Valuation?](#)

[OpenAI's GPT-6 Astra: Benchmarks, Cyber Risks, and Market Impact](#)

[The US Just Switched Off Anthropic's Frontier Model: What Happens Next?](#)

About Us

About the Authors

Nick Patience, VP & Practice Lead, AI Platforms at Futurum, is a thought leader in the development, deployment, and adoption of AI – an area he has researched for 25 years. He is a sought-after speaker and advisor, known for his expertise in the drivers of AI adoption, industry use cases, and the infrastructure behind its development and deployment. Nick is based in London and can be reached at npatience@futurumgroup.com.

About The Futurum Group

Every day, The Futurum Group's analysts, researchers, and advisors help business leaders worldwide anticipate tectonic shifts in their industries and leverage disruptive innovation. Unlike traditional analysts, The Futurum Group works not only in analysis and research but also takes that insight and knowledge even further, engaging all the way through the go-to-market process.

Futurum Research provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing. It also creates, distributes, and amplifies rich media content that all stakeholders read, watch, and listen to.

See more details on The Futurum Group at futurumgroup.com.

Copyright & Use License

Copyright Notice

Copyright ©2026 by The Futurum Group, LLC. All rights, including that of translation into other languages, are specifically reserved. No part of this publication may be reproduced in any form, stored in a retrieval system, or transmitted by any method or means, electrical, mechanical, photographic, or otherwise, without the express written permission of The Futurum Group futurumgroup.com. United States copyright laws and international treaties protect this publication. Unauthorized distribution or reproduction of this publication, or any portion of it, may result in severe civil and criminal penalties and will be prosecuted to the maximum extent necessary to protect the publisher's rights.

License Notice

This document may be distributed within the licensed organization only.

The following acts are prohibited:

Transmittal to others outside your immediate organization, including partners, resellers, external consultants, etc., in any media format

Posting on a website that is accessible to others outside your immediate organization

The possession or use within an unlicensed organization

Limitation of Liability Notice

The information contained herein has been obtained from sources believed to be reliable. The Futurum Group shall have no liability for errors, omissions, or inadequacies in the information contained herein or for interpretations thereof. The reader assumes sole responsibility for the selection of these materials to achieve their intended results. The opinions expressed herein are subject to change without notice.