2027 AI budget: a guide for an informed conversation between CEO, CFO & CTO.

We deep dive into the challenges for organizations from 10 to 400 individuals.



Most AI cost research measures companies above $1B in revenue. This guide is for companies under 400 people. Those companies in which 2027 will likely be the year they'll need to start using Frontier Labs enterprise deals, or will run out of credits. The numbers come from public research and calls with founders and engineering leaders at companies with 6 to 150 AI seats.[6]

2026: Unjustified spending.



"Cheaper than people, yes, but we are talking about millions of euros a year."[6]

2.5x more expensive.

Menlo Ventures surveyed nearly 400 US enterprises. The average company aimed for 1.9x AI spend in 2026, and ended at 4.7x. The median was about 22% over plan.

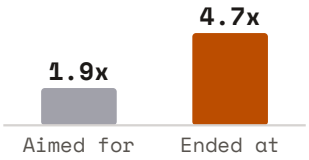

What should you learn

- Every organization has a small group of power users.
- People leave, Agents stay. Most experiments and individual setups are not cleaned up.
- New budgets funded 55% of that incremental spend, money is being moved from other lines.[1]

Flipping a ROI coin.



PwC asked 4,454 CEOs. 56% saw neither revenue nor cost benefit.
McKinsey's 1,719 respondents: only 37% report any positive EBIT contribution.
BCG's 1,330 CxOs: 48.5% generate value.

Even if not all AI projects bring direct contributions, when it works, the benefits are asymmetrical.[2,3,4]

The bill keeps rising.



In 2026 every major vendor moved away from flat seats.
Claude Team stops at 150 seats, and above that every token bills at API rates.
ChatGPT Business stops at 200.
GitHub Copilot moved to token-metered credits with no free fallback.
Cursor bills third-party models at API list price, plus a token fee on Teams.
Microsoft 365 Copilot bills advanced AI by usage, on by default for new Copilot Business.[5]

We've seen teams leaping from 30k/mo to 180k/mo with the switch to Enterprise. It's not surprising to hear teams talking about a 20x bill increase.
Others are implementing whitelists, or limiting access to each provider.

Where you are

- **Small team growing, 5 to 20.** You stay under every seat cap. Your risk is overhiring before you are fully leveraging what's available for you.
- **Established unicorn scaling, 20 to 50, holding headcount.** You may start having pockets of AI-pilled individuals. The risk is ROI diffuses.
- **Series B run, about 50, doubling in 2027.** At about 100 people you'll likely start getting close to hitting limits. 2027 is the year to setup your AI native infrastructure.
- **Well-established, 5+ years or 200+ people.** Change management, learning and development, and intra-organization dynamics are your bottleneck.

2027 budget shape: avoid the 2026 failure modes.



> "not everyone consumes it, but the heavy users consume and would love to overconsume."[6]

Your budget aggregates individuals as seats, and it averages out. It's an easier way to keep track of spending, but the reality of organizations is different. Top performers keep consuming most of the AI budget, whereas mid/low performers are being outpaced. If you keep that methodology, you are masking the potential exponential increase on budget you'll see.

The challenge is worse than you expect. Deloitte asked 200 CFOs at $1B+ companies: 53% review AI costs retrospectively, only 20% use advanced AI-specific controls, and 74% need days to several weeks to pause usage. That's not considering the long lasting business disruption it may bring. The risk on lacking control means your budget will be a suggestion, unless the organization takes it seriously.[8]

Per-seat pricing assumes every seat costs the same. Let me give you an example, Ramp's study across 70,000+ US businesses puts median AI spend at $12.50 employee/month. That's ridiculously low. In fact, the top 10% spend $675.60, the top 1% spent $9,568 in September. Our telemetry over the last 30 days has the same shape: our top 10 known installs account for 73% of all model cost, whereas the median install spent $18.

Prosus found it with agents ROI follows that same curve, too. About 2% of its 60,000 internal agents drive most of the business impact.[11]

Make sure your budget is accompanied by a clear plan to enforce controls that allow to overspend in the areas that prove ROI, while also allowing you to move across providers to actually realize those gains fully.

Note: Be careful, two challenges we see often: Time gains are not needed if your organization lacks growth and some wins are difficult to account for. That's the challenge.

How we see teams doing it.

Per team. Pools per team, with manager discretion. You also need to track individuals, but this is now a management challenge. They'll need to justify the ROI of having superstars producing like teams.

Per top individual.



Heavy users need to be tracked separately. They need their own guardrails and budgets. These innovators need to carry one extra responsibility: Upskill the team.
Teams we talk to land near $3,000/month for a tech or product role, and ~$200 for everyone else. At a 20% tech share that is $9,120 employee/year. Is that doubling your organization or did you just triple headcount?

Keep pushing to use the right model for the right activities, and creating software as much as possible. That's where you can see savings. Median engineers run $252 a month[12], heavy users report $600 to $700[6], and the 90th percentile reaches $1,666[12].

Per use case.

Track Agentic efforts like you track product efforts. Give them a budget, a deadline, and invest in 'cleaning-up' failed efforts. Prosus found most agents save under 20 hours a month while a tiny fraction deliver thousands of hours. Your budget should account for two or three efforts that carried that type of return in 2026. Those are your successful bets.[11]

Where you are

- **Small team growing, 5 to 20.** Don't waste time. Have a conversation, ask for a couple of lines on general direction, move on.
- **Established unicorn scaling, 20 to 50, holding headcount.** Start thinking on teams, and combine growing headcount together with AI spending when unlocking revenue.
- **Series B run, about 50, doubling in 2027.** Each team gets its own budget pool from day one. It's survival of the fittest with monthly oversight.
- **Well-established, 5+ years or 200+ people.** The per-use-case line is the only one that survives politics. Start enforcing the same discipline for AI efforts as you do with product or business.

Use more AI. Pay less.

Most 2027 plans ask for a bigger budget, but AI is being commoditized. It's the tension between adoption and maturity. You need KPIs to track price per task, while task types rise. This is going to reflect on the following:

Lever one: the price per token.

Routine work should be using cheaper models and deterministic software.
In the past month, one team paid $0.11 per million tokens and another paid $0.63. A 5.6x gap. The former ran 74% of its tokens on a mid-price coding model. The latter ran 61% on frontier models. A single premium model accounts for 4% of tokens & 23% of the bill.[10]

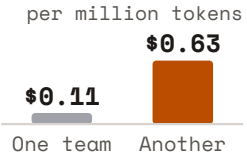

Lever two: the tokens per task.

Recurring work needs to be automated, and improved.
Our own nightly review cut 52% of its token spend to scripts, worth $17k and 450 agent hours in a month. Spending didn't drop, it just moved to allow for more innovation.[20]



If you halve tokens per task and halve the price per token, the bill falls 75%.

Our numbers.

From 15 teams measured six weeks apart, tokens went down 8.0%, completed tasks flat, cost down 45.1%. That makes sense, the teams didn't grow, so human generated tasks stay flat. Their automated workflows are also stable, at least the high volume ones, so making your system efficient will get you there quite quickly. For the orgs that actually increased task volume, you can deduce a similar impact.

For us, tokens up 9.6%, cost down 43.6%, cost per task down 78.2%. But we grew significantly, so it went up overall.[10,19]

Adoption conundrum

Price per token fell 99% in two years, and total spend still climbed.
Menlo's average company spent 4.7x. Jellyfish puts developer tool spend up 10x this year. In Pragmatic Engineer's 27 companies study, one went from $500K in January to $8M in September.[1,18,29,13]

It's all about spending per unit of work.

The number for the plan.

Software companies surveyed by ICONIQ put internal AI spend at 16% of revenue in 2026, heading to 19% in 2027. They include salaries of AI roles and change management, together with providers. Other surveys account for 2% to 3% of revenue: BCG's AI Radar 2.1% for tech companies[15], Aleph x Benchmarkit about 2.7% for median B2B SaaS[17]. All numbers are general averages, not truly useful for smaller teams.

Headcount.

If revenue doesn't grow, headcount won't either, and the saving will appear as hiring you didn't do. ICONIQ's growth bands show it: companies growing above 100% added 133% to 146% more headcount, while companies under 25% growth held flat. Job cuts concentrate in large firms.[21,16]

Where you are

- **Small team growing, 5 to 20.** Lever one pays best. Unless you have a system to automatically improve your workflows, changing models is the fast switch.
- **Established unicorn scaling, 20 to 50, holding headcount.** Time to think on a system that will help you scale these levers as you grow. You need to set the basis.
- **Series B run, about 50, doubling in 2027.** Same as above, but you need to rush. Otherwise you'll be dealing with added politics soon.
- **Well-established, 5+ years or 200+ people.** You need to find a transparent way to do this at scale. Your whole organization needs to follow these principles, and training is too slow. This is a major CTO challenge.

Grilling the budget, and ten questions for your CEO and CTO

This is what went wrong during the past year.

Top issues during 2026.



Implementation runs late. Gartner found 63% of leaders said it was slower than expected.[22]
Nobody prices the outcome. KPMG's Q3 2026 survey: only 12% consistently assess ROI.[23]
Operating cost constrains use. ~20% of McKinsey's respondents say operating costs, tokens included, limited what they could do.[3]
The instrumentation lies. Licences activated and weekly active users are easy to pull and tell you nothing.

Caps could kneecap you.

Blanket caps can block the best user, and incentivize others to overspend. Cognition calls them "akin to a hiring freeze in the AI era," and describes an engineer blocked for four days whose merged PRs fell to a third. Vendors now sell pooled team budgets for the same reason.

Without enforced caps, you'll see unexpected spikes. On a survey of 300 companies, 77% reported a negative consequence from ungoverned AI spending.[28] Test that your caps are enforced in realtime.

Don't forget to code.



For us, a human-started task costs ~5 times more than a workflow one. That means workflows are 33.6% of our workload and represent 10.6% of budget. When looking at these numbers always bake in failure rate, and retries.[25]

Change management



"I am doing this role of trying to be the champion and try to externalize things to other teams. And it has been really hard. We're still struggling."

"there is a lot of fear. Some people are already picturing the Terminator walking into their house."[6]

Even when you hire or promote an explicit AI role, the challenge stays there. A few conversations and adoption KPIs won't do, unless people are seeing meaningful change that benefits them. Your CEO & CTO need to find those win-win scenarios for each role, and do most of the changes transparently.

Ten questions.

CFO to CEO

1. Which two or three use cases drove the AI ROI in 2026, and who owned each one?[11]
2. Does our headcount plan follow revenue growth? If revenue is flat, where does the saving show? What about our AI budget?[21]
3. Can we leave our main AI vendor within a quarter, and at what cost?[6]

CFO to CTO

4. Who are our top 1% and top 10% AI users, and how much do they cost per month?[9]
5. For what efforts will we have to stop using frontier models?[10]
6. How do we track AI token efficiency per task? Is this improving (fewer tokens per task)?[10,19]
7. Are we using our seats fully or are we leaving money on the table?[5]

All three

8. What's AI Opex or Capex?[26,28]
9. How do we know we are using AI for new challenges and not repetitive work?[20]
10. Which team owns each AI budget and how do we reward them if they are responsible with it?[6]

Where you are

- **Small team growing, 5 to 20.** Focus on the last 9 and 3. That will unlock the rest.
- **Established unicorn scaling, 20 to 50, holding headcount.** 2 and 4 matter most. One is your growth story, the other is your multiplying lever.
- **Series B run, about 50, doubling in 2027.** 6 & 10 are going to be key. Both set the right incentives for your growing team.
- **Well-established, 5+ years or 200+ people.** 2, 8 and 10. These high level discussions can be resolved in that room.

Agent Swarm

agent-swarm.dev

Sources

1. https://menlovc.com/perspective/2026-the-state-of-enterprise-ai/
2. https://www.pwc.com/gx/en/ceo-survey/2026/pwc-ceo-survey-2026.pdf
3. McKinsey, *The state of AI in 2026: On the road to ROI*, 2026-08-25. 1,719 respondents, 97 countries.
4. BCG, *Applied AI Index 2026*, 2026-09-30. 1,330 CxOs and senior leaders, 20+ sectors.
5. Vendor pricing pages.
6. Multiple calls.
8. Deloitte, *3Q 2026 CFO Signals*, 2026-10-21. 200 North American CFOs at $1B+ companies.
9. https://ramp.com/data/ai-index-oct-2026
11. Prosus, *The Coming Age of AI Colleagues*, May 2026. Prosus group, 40,000 employees, more than 60,000 internal agents.
12. Jellyfish, via R&D World, 2026-10-27. Developers using AI coding tools.
13. Pragmatic Engineer, *What's the ROI of software engineers*, 2026-10-20. 27 companies.
14. ICONIQ, *2026 State of AI: The Builder's Economy*, July 2026. N=303 software companies building AI, $1M-5M to $1B+ revenue, 87% North America.
15. BCG, *AI Radar*, 2026-01-15. 2,360 executives, 16 markets.
16. *The CFO Survey* (Duke, Richmond Fed, Atlanta Fed), fielded near year-end 2025, paper 2026-05-27. 734 US executives weighted to Census by sector and size.
17. Aleph x Benchmarkit 2026. 342 B2B SaaS and AI-native companies.
18. BOND (Meeker), *Trends in Artificial Intelligence*, 2025-05-30. Inference cost per token.
19. https://www.agent-swarm.dev/blog/our-agents-dream-every-night
20. https://www.agent-swarm.dev/blog/automation-savings
21. ICONIQ, *2026 State of Scaling*, September 2026. Median headcount change by revenue growth band; N from 57 to 390 by year.
22. Gartner, 2026-05-28.
23. KPMG, *AI Pulse*, Q3 2026.
25. https://www.pleasedontdeploy.com/p/the-case-against-ai-workstations
28. Redwood, survey of 300.
29. https://techcrunch.com/2026/10/28/anthropic-just-launched-tools-to-show-companies-what-their-ai-spending-is-buying/

Agent Swarm (https://agent-swarm.dev) is an open-source (MIT) AI operating system for companies. We are here to help you upskill your team, while decoupling your company from all AI providers.

