

Ten questions about AI cost, and how to tell a real answer from a plausible one.

AI spend used to be an engineering line item nobody questioned. It is now large enough to move gross margin and volatile enough to reach the board without warning. This playbook gives directors ten questions a competent management team should be able to answer, what a defensible answer sounds like, what a plausible answer sounds like, and the follow-up that separates them. Most management teams cannot yet answer more than a few, and that is the normal condition of the market rather than evidence of a weak team.

WRITTEN FOR

Directors and audit committee chairs about to ask about AI spend for the first time, and the board chair who has to decide whether the answers are good enough to fund the next round of AI investment.

ALSO FOR

Investors who sit on several boards and want one standard across a portfolio, and the CEO who wants to know what the board is about to ask before the meeting does.

NOT FOR

The finance team that builds the record. That is the CFO playbook. Counsel reviewing AI use against employment and data law has its own compliance playbook. This document describes what published sources say. It is not legal, tax or investment advice.

The board is now the one being asked, and the one being surprised.

Two years ago nobody asked a director about AI cost. It became a governance question the moment it got big enough to move a margin, and the questions now arrive from above as often as from below. Boards are making funding conditional on a return, and AI cost surprises are reaching the boardroom before management has a way to explain them.

66%

of boards now require proof of AI return before approving further funding. 43% of finance leaders are being asked for return metrics they cannot provide.

CloudZero, June 2026

40%

of organizations had an unexpected AI cost that required board-level escalation in the past year. Only 11% forecast AI spending within plus or minus 10%.

Mavvrik with Benchmarkit, July 2026

72%

of organizations had an unexpected AI cost spike in the past year, 33% more than once. Only 20% could find the cause within hours if spend doubled.

Harness, July 2026

43%

of companies track AI and cloud cost by customer at all. 22% track it by transaction.

CloudZero with Benchmarkit, February 2026

Read the last card twice. It sets your expectations for the meeting. If management cannot answer several of the questions in this playbook, that is where most of the market sits. What matters is whether they know they cannot, and whether someone has a date to fix it. Gartner warned in 2024 that a naive calculation of AI cost can be off by 500 to 1,000%, which is why a confident number with nothing behind it is the answer to probe hardest.

There are two kinds of AI spend, and only one of them should be predictable.

POOL ONE

Spend that should run loose

Research, experimentation, prototypes, engineers trying things. The variance here is healthy. It is how a company finds what works, and clamping down on it is how a company stops learning.

A board should not police this number. It should know roughly how big it is and who owns the envelope.

POOL TWO

Spend that should be locked down

Anywhere the company prices, forecasts or promises value to a customer. That cost belongs in gross margin. It should be predictable, attributed to the customer and team that caused it, and tied by finance to the provider invoices every month.

This is the number the board is accountable for.

Almost every company reports both pools as one figure, and that is why so many boards are surprised. It hides the moment a research project became a production feature the company now pays for at scale. The useful question is never what is our AI spend. It is which pool that dollar came from, and whether management can prove it.

WHAT THE BOARD SHOULD RECEIVE

Ask for this evidence in the next board pack and see which pages arrive.

WHAT THE BOARD IS ASKING	THE EVIDENCE THAT ANSWERS IT	QUESTIONS
What does AI cost us?	Last month's spend in dollars, split by pool, covering providers, subscriptions and expensed tools.	One and three
Are these numbers right?	Finance's monthly tie-out of the attributed record to each provider invoice, with every gap named.	Two
Is AI eroding our margin?	Cost against revenue per customer, the costliest accounts, and flat-priced products with usage-based cost.	Four, five and nine
Will we be surprised, and by what?	Named alerts, measured unapproved use, and the margin effect of a price rise from the largest provider.	Six, seven and eight
Who is accountable?	One named owner and a standing report on a fixed date.	Ten

Ask in this order. The first three decide whether any of the others can be answered.

Each question comes with what a defensible answer sounds like, what a plausible answer sounds like, and the follow-up that separates them. The plausible answers are not dishonest. They are what a competent person says when the instrumentation cannot do better.

QUESTION 01 ASK THE CFO

What did AI cost us last month, in dollars?

The baseline. You are not testing whether the number is big. You are testing whether it exists, and whether anyone can produce it without a week of notice. 81% of organizations say they cannot fully account for their AI costs, and only 13% report even basic visibility into actual AI spending.

A DEFENSIBLE ANSWER

A specific figure, produced without hesitation, with the providers, subscriptions and tools it covers named.

A PLAUSIBLE ANSWER

A range, or "it sits inside the cloud bill," or a promise to come back with it.

THEN ASK

Does that include the AI features inside the software we already subscribe to, and the tools people expense?

QUESTION 02 ASK THE CFO AND CONTROLLER

Does that number tie to what the providers actually invoiced us?

The most important question in the playbook. Most tools compute a figure from usage logs. The invoice is what was paid. A computed figure and a figure that has been tied to the bill are different things, and only the second belongs in a board pack. The gap between them is where untracked usage and price changes hide.

A DEFENSIBLE ANSWER

Yes. Finance ties the attributed record to each provider invoice on a stated cadence, and can tell you the last time the two did not match and why.

A PLAUSIBLE ANSWER

"Our dashboard tracks it in real time." A dashboard is a calculation. It is not a tie-out, and it has no way to know it is wrong.

THEN ASK

Who performs that tie-out, and what was the largest gap in the last six months?

QUESTION 03 ASK THE CFO AND CTO

How much of our AI spend sits behind something we sell?

The two-pool split. If the company cannot separate experimentation from production, part of gross margin is currently an estimate and nobody knows which part. It also explains forecast error. Only 11% of organizations forecast AI spending within plus or minus 10%, down from 15% the year before, and a blended number cannot be diagnosed.

A DEFENSIBLE ANSWER

Two numbers, and a written rule for when a workload moves from one pool to the other, usually the day it is priced or promised to a customer.

A PLAUSIBLE ANSWER

One number, with an offer to break it out later. Later is doing a lot of work in that sentence.

THEN ASK

If we cannot split it, how is the customer-facing half currently appearing in our gross margin?

QUESTION 04 ASK THE CFO

What does our most expensive customer cost us to serve?

AI cost concentrates, so an average across the base tells you little about your exposure. Only 43% of companies track these costs by customer at all, and 22% by transaction. Organizations estimate 26% of their AI spend is wasted, and waste cannot be found in a total. It is found in the customer, team or feature that produced it.

A DEFENSIBLE ANSWER

A named account and a figure, plus the spread between the cheapest and most expensive customers to serve.

A PLAUSIBLE ANSWER

An average cost per customer. Averages hide exactly the tail you are asking about.

THEN ASK

Are any customers unprofitable to serve today, and would we know next month if one became so?

QUESTION 05 ASK THE CEO

Are we charging a flat price for anything whose cost underneath is usage-based?

The classic trap. A flat price set before anyone knew what heavy use looked like turns the heaviest users into the least profitable ones. 49% of organizations were forced to reprice AI-powered products after unforeseen cost increases. GitHub moved every Copilot plan to usage-based billing on 1 June 2026. Among AI builders that price by consumption, only 26% pass the full cost of inference through to customers.

A DEFENSIBLE ANSWER

They name the products with that shape and the usage level at which each one breaks even.

A PLAUSIBLE ANSWER

"Our margins are healthy." That is an aggregate statement about a per-customer problem.

THEN ASK

How many customers are already past break-even, and is that number growing?

QUESTION 06 ASK THE CTO

If usage doubled tomorrow, would anyone find out before the invoice?

Surprise is the failure mode. 72% of organizations had an unexpected AI cost spike in the past year, a third more than once, and only 20% could identify the cause within hours if spend doubled overnight. Separately, 40% of organizations had an unexpected AI cost that required board-level escalation. The question is whether there is a tripwire or only a monthly review.

A DEFENSIBLE ANSWER

A named alert, a named threshold, and a named person who receives it the same day.

A PLAUSIBLE ANSWER

"We review spend monthly." A monthly review finds the spike after it has been paid for.

THEN ASK

What was our largest surprise so far, and how did we find out about it?

QUESTION 07 ASK THE CIO

How much AI are our people using that we did not approve?

Governance rather than margin, but it reaches the board through the security questionnaire and the audit committee. 66% of office professionals have used AI tools at work despite believing it breached company policy, and 88% have shared work information with public AI tools, including 34% who shared customer data.

A DEFENSIBLE ANSWER

They have measured it, even roughly, can describe how, and can say what changed after they did.

A PLAUSIBLE ANSWER

"We have a policy." A policy is a statement of intent. It is not a measurement.

THEN ASK

How would we know if someone were using a corporate AI account for work we never sanctioned?

QUESTION 08 ASK THE CFO

What happens to us if one provider changes its prices?

Concentration risk, and it runs in both directions. Within 32 days in mid 2026 one provider cut a model's price by 80% and another provider's rates went back up by 50% when introductory pricing ended. 61% of IT leaders cut projects because of unplanned cost increases. A cut is only a saving if someone notices it and the forecast reflects it.

A DEFENSIBLE ANSWER

They know the share of production spend by provider and have run the gross margin effect of a material increase from the largest one.

A PLAUSIBLE ANSWER

"We could switch providers." Possibly true, and it says nothing about how long that would take or what it would cost to re-test.

THEN ASK

What would a material increase from our largest provider do to gross margin this year?

QUESTION 09 ASK THE CEO AND CFO

Is our cost per customer going up or down?

Direction matters more than level. AI costs do not scale the way software costs scale, and the gap shows in margin. ICONIQ's survey of companies building AI products puts AI product gross margin at 45% in 2025, projected at 53% for 2026. Bessemer found its fastest-scaling AI cohort running gross margins around 25% against about 60% for slower-growing peers. Improvement is possible. 66% of AI builders cut cost per query or task by at least 10% in six months.

A DEFENSIBLE ANSWER

A per-customer figure with six months of trend behind it, and an explanation for the direction.

A PLAUSIBLE ANSWER

Absolute totals only. Total spend rising is expected. Cost per customer rising is the warning.

THEN ASK

At our planned growth rate, where does this number sit in twelve months?

QUESTION 10 ASK THE CEO

Who owns this number?

Close on accountability. AI cost sits between engineering, who create it, and finance, who report it. 52% of organizations say there is no clear owner of AI cost. 78% of FinOps teams report to a CTO or CIO, which means the number is often owned by the people spending it.

A DEFENSIBLE ANSWER

A single name, and that person is in the room or sends the board a standing report on a fixed date.

A PLAUSIBLE ANSWER

"It is shared between finance and engineering." Shared ownership of a forecast is unowned.

THEN ASK

Whose forecast is it, and what happens if it turns out to be wrong?

Count the defensible ones, not the confident ones.

Confidence is not evidence. Score an answer as defensible only if it came with a document, a name, or a date attached.

EIGHT TO TEN

Unusually well run.

Ahead of most peers. The remaining gap is usually question two, because tying the record to every provider invoice is the piece most finance teams have not yet made routine.

Ask for the evidence in the board pack quarterly and move on.

FOUR TO SEVEN

Normal, and worth a decision.

Do not try to close every gap. Pick the two that matter for where the company is. Pricing and cost per customer if it sells AI. Governance and unapproved use if it mostly runs AI internally.

Put a name and a date against those two in the minutes.

THREE OR FEWER

An instrumentation problem, not a competence one.

Fewer than half of companies track AI cost by customer, so most of the market sits here. The risk is that nobody finds the number until an outside party asks for it.

Set a date by which the company can answer questions one, two and three.

One caution. A management team that says "we do not know yet, and here is when we will" is in a better position than one that produces a confident number nobody has tied to an invoice. The second answer is the one that becomes a surprise later.

IF YOU SIT ON SEVERAL BOARDS

The same ten questions work across a portfolio.

For a fund, the value is the comparison. Ask all ten at every company, not only the one that had the incident. One company's margin surprise is information about the others, and the pattern shows where the risk sits, which is usually not where the noise is.

Put questions two and three in the diligence template

The tie-out and the two-pool split are the foundation the other eight rest on. Without them, cost per customer, pricing and forecast are estimates.

Expect the answers before a process, not during it

Only 9% of private equity respondents are very confident they could pass an AI governance audit within 90 days, the lowest of any sector surveyed.

TAKE THIS INTO THE ROOM

AI cost. Ten questions for the board.

Print this, or send it to the other directors and to management before the meeting so everyone is working from the same sheet.

01	What did AI cost us last month, in dollars?	CFO
02	Does that number tie to what the providers actually invoiced us?	CFO, CONTROLLER
03	How much of our AI spend sits behind something we sell?	CFO, CTO
04	What does our most expensive customer cost us to serve?	CFO
05	Are we charging a flat price for anything whose cost underneath is usage-based?	CEO
06	If usage doubled tomorrow, would anyone find out before the invoice?	CTO
07	How much AI are our people using that we did not approve?	CIO
08	What happens to us if one provider changes its prices?	CFO
09	Is our cost per customer going up or down?	CEO, CFO
10	Who owns this number?	CEO

Score an answer as defensible only if it arrived with a document, a name, or a date attached.

WHERE THIS CAME FROM

We are the people who measure this.

Crystil is an independent third party and SOC 2 compliant. Our Cost Tracking SDK is one line of code. It prices every model call by input, cached and output tokens and the model used, and attributes the cost to a customer, organization or team, with an optional department beneath it. Workflows gives a multi-call task one cost. Customers sets cost per customer against revenue to give margin per customer. Pulse routes a team's AI coding tools, Claude Code, Codex and Gemini CLI, through Crystil with no SDK.

What we do not do matters as much. Crystil provides the attributed record a finance team reconciles against its provider invoices. We do not reconcile them for you, and we certify nothing. Our record helps management answer questions one, three, four, six and nine.

[Talk to an AI measurement expert](#)

crystil.com

COMPANION GUIDES

The **CFO playbook** is for the finance team that produces these answers. The **AI compliance playbook** is for counsel, and covers the legal side of question seven.

Sources

1. CloudZero, CloudZero Survey Says: 78% of Finance Execs Can't Fully Tie AI Spending to Business Outcomes, press release, 24 June 2026. 260 senior finance leaders. Vendor research.
2. Mavrik with Benchmarkit, 2026 State of AI Cost Governance Report, press release, 29 July 2026. 396 enterprises. Vendor research.
3. Harness, 2026 State of AI in FinOps, press release, 29 July 2026. 700 engineering and FinOps leaders. Vendor research.
4. CloudZero with Benchmarkit, FinOps in the AI Era: A Critical Recalibration, press release, 12 February 2026. Vendor research.
5. Gartner, remarks by Mary Mesaglio, Gartner IT Symposium, Gold Coast, September 2024, as reported by The Register, 9 September 2024.
6. GitHub, GitHub Copilot is moving to usage-based billing, 27 April 2026.
7. ICONIQ Growth, 2026 State of AI: The Builder's Economy, July 2026. About 300 executives. The 26% figure is from 72 consumption-pricing respondents.
8. PagerDuty, Two-Thirds of Office Professionals Have Used Unauthorized AI Tools at Work, 11 June 2026. 1,250 office professionals. Vendor research.
9. Digital Applied, AI API Pricing, August 2026: Cuts, Promos, and Traps, 5 August 2026, with the CloudZero pricing tracker.
10. Zylo, 2026 SaaS Management Index, 29 January 2026. Vendor research.
11. Bessemer Venture Partners, The State of AI 2025, 13 August 2025.
12. FinOps Foundation, State of FinOps 2026, 19 February 2026. 1,192 respondents.
13. Grant Thornton, Private Equity insights: 2026 AI Impact Survey Report, 21 April 2026. 100 private equity respondents.

Crystil. v1.0, 7 October 2026. Every figure is published and dated. Nothing here is modelled, estimated or invented. This document describes what published sources say. It is not legal, tax or investment advice.