

RAISE READINESS 2027: AI DATA ROOM PREP

What investors will ask about your AI costs, and how to be ready before they do.

The same ten questions investors are now putting to AI-native companies, written from your side of the table. For each one, what it is testing, what a defensible answer sounds like, what a plausible answer costs you in a room you cannot re-enter, and what to do before the process starts. It closes with a ninety-day plan and the exhibits that belong in the data room.

WRITTEN FOR

A founder or CEO of an AI-native company preparing a Seed to Series C round, and the CFO, finance lead or CTO who will sit beside them in diligence. You are the person who has to defend the gross margin line when the lead partner rebuilds it from source.

ALSO FOR

Fractional CFOs and finance advisers who prepare data rooms, and existing investors on the board who want the company ready before the next lead arrives. Read it next to the VC portfolio playbook, which is the investor's version of the same questions.

NOT FOR

Companies with no AI in the product, or founders selling to a buyout firm, whose diligence turns on AI cost across a held portfolio and is covered by the private equity playbook. Nothing here is investment, tax or accounting advice.

AI margin became a diligence line, and most companies were not told.

Cost per customer and margin trajectory used to be questions for a Series B. They now come up in seed conversations, because investors have watched enough AI-native companies discover their unit economics after the round rather than before it. The comparison set has moved too, and your deck should show that you know where it moved to.

43%

of companies track AI and other costs by customer at all. 22% track them by transaction. If you can do either, you are already unusual.

CloudZero with Benchmarkit, February 2026

25% vs 60%

gross margin for Bessemer's fastest-scaling AI cohort, the Supernovas, often negative, against about 60% for the slower Shooting Stars.

Bessemer Venture Partners, State of AI, August 2025

50 to 60%

is the gross margin CRV describes as defensible for an AI company at seed, with a demonstrated path to expansion. Below it, bring the explanation.

CRV, July 2026

70% vs 90%

is the gross margin Sierra CEO Bret Taylor says AI agent companies should expect, against the 90% of some pure software companies.

Fortune, April 2026

Read the first card as good news. Fewer than half of companies can attribute cost by customer, so answering even four of the ten questions well sets you apart rather than meeting a baseline. You are not being held to a standard the market already meets. The CloudZero figure is vendor-commissioned research, the Bessemer figure predates 2026, and the Taylor figure is one executive's view in an interview, not survey data. They describe the market, not your company.

The moment a feature ships to a paying customer, its AI cost stops being R&D.

POOL ONE

Experimentation should look messy

Prototypes, evals, model comparisons, engineers using AI coding tools, features that may never ship. Nobody expects this to be predictable, and an investor who pushes on it is asking for less learning.

Know roughly what it costs, who owns it, and whether it grows faster than headcount. That is enough.

POOL TWO

Anything customers pay for should not

The day a feature is in front of a paying account, its AI cost belongs in cost of revenue. It should be attributable to the customer, measured per unit of work, and tied by your finance team to what the providers invoiced.

This is the half you will be asked to defend, and the one the ten questions are about.

Nearly every company reports these as one figure, because the instrumentation it has produces one aggregate number. That is exactly why the split is the fastest way to sound prepared. Walking in with two numbers and a written rule for when spend crosses from one pool to the other tells the partner you understood the question before it was fully asked. A large R&D pool is not a weakness. Being unable to say which experiments became recurring product cost is.

WHAT THE ROOM WILL COMPARE YOU TO

The published benchmarks a partner carries into the meeting. Know them, but do not treat them as your target.

Each row was produced with a different method on a different sample. Expect to hear them quoted, and be ready to explain where your company sits and why.

ICONIQ Growth

Published July 2026

AI product gross margins were 45% in 2025, projected at 53% for 2026 and 59% for 2027, while AI products grow from 32% to a projected 42% of company revenue. Survey of about 300 executives at software companies building AI products, fielded in the second quarter of 2026.

ICONIQ Growth, 2026 State of AI: The Builder's Economy, July 2026

CRV

Published 16 July 2026

Published guidance for AI founders puts a defensible gross margin at seed at 50 to 60%, with a demonstrated path to expansion, against the 70 to 90% traditional SaaS benchmark. The path is the part founders forget. CRV is asking for a trend line, not a single quarter.

CRV, LLM Inference Costs: What AI Founders Should Know, 16 July 2026

CORRECTION ONE

A founder reads 53% as the number to hit before raising. It is a projected average across companies of every size and architecture. A seed company at 40% with cost per customer falling every month, and a clear reason why, can be in a stronger position than a later company at 60% with no per-customer view at all. Trajectory and evidence carry more weight than the level.

CORRECTION TWO

A founder tells the room that falling token prices will fix the margin. Token prices have fallen roughly tenfold a year, from \$60 per million in 2021 to about \$0.06, while consumption has exploded (Forbes citing EpochAI, July 2026). Enterprise agents use 5 to 30 times more tokens than single requests for the same task type (Vista Equity Partners, May 2026). Cheaper tokens and costlier tasks arrive together, and the partner knows it.

WHAT GOES IN THE DATA ROOM

Six exhibits that answer most of the ten questions before they are asked.

The next lead will rebuild your margin from source records rather than accept the deck. These are the documents that let them do it quickly, and each one maps to the questions it closes.

EXHIBIT	WHAT IT CONTAINS	CLOSES
Monthly AI spend, both pools	Last month's total by provider and tool, split into R&D and production, with the written rule for when a workload becomes cost of revenue and the date each one crossed.	Questions 01 and 03
The reconciliation history	Tracked AI cost against each provider invoice, month by month, signed by a named person in finance, with every gap named and the largest one explained.	Question 02
Cost per customer	Production AI cost attributed to each paying account, monthly, as a distribution with median and top decile. Six months at minimum, twelve if you have them, with the reason for each change in direction.	Questions 04 and 09
The pricing to cost map	Each plan, its price, the unit the customer buys, the model cost behind one unit, the break-even usage, and how many accounts are past it today.	Question 05
Controls and provider exposure	Alert thresholds per customer and per feature with the person who receives them, share of spend by provider and model, and the margin effect of a published price move already calculated.	Questions 06, 07 and 08
The owner and the cadence	One name, the monthly review the numbers go to, and the decisions they changed in the last two quarters.	Question 10

The pricing to cost map also covers two questions the investor's playbook asks separately: what unit you sell and what one costs to deliver, and how much of cost of revenue is model spend paid through to providers. Among consumption-priced AI builders, 26% pass all inference cost through, 58% split it and 15% absorb it (ICONIQ, July 2026, a sample of 72). After close, expect the lead to ask for a monthly metric set built from the same records, so build these exhibits to keep running.

THE PLAN

Ninety days is more than enough, if you start before the first meeting is booked.

The order matters. Instrumentation comes first because every later step reads from it, and the invoice tie comes early because it tells you whether the rest can be trusted.

01 Days 1 to 30. Instrument

Tag AI spend by feature and by customer account at the time of the call, not by allocation afterwards. Add up every provider, plus the AI features inside tools you already subscribe to. Decide and write down the pool rule: shipped to a paying customer means cost of revenue from that day.

02 Days 1 to 30. Tie one month to the invoices

Pull last month's provider invoices and compare them line by line to what your tooling reported, by hand if you have to. Name the person in finance who will repeat it every month, and write down the largest gap and its cause.

03 Days 31 to 60. Build the trend and the controls

Plot cost per customer monthly for at least six months. If you were not tagging by account at the time, rebuild the history from usage logs and invoices and say so. Model break-even usage per plan and count the accounts above it. Set one threshold alert per account and per feature with a named recipient, and check company cards and single sign-on logs for AI tools nobody approved.

04 Days 61 to 90. Assemble and rehearse

Calculate what a published provider price move would do to this year's gross margin. Decide whose forecast this is. Put the six exhibits in the data room, then run the ten questions with someone who will push on every plausible answer.

05 In the room. Volunteer the split

Give two numbers before you are asked for one. Name what you do not yet measure and the date you will. Have the reconciliation ready as a page, not a claim.

What they will ask, and what each question is testing.

Nobody expects ten strong answers. What the partner is reading is whether you know which ones you cannot answer, because a founder who knows their own gaps will find the next one before it becomes a problem. Each question comes with a defensible answer, a plausible one, and the preparation that turns the second into the first. The plausible answers are not dishonest. They are what a competent founder says when the instrumentation cannot do better, and they are the answers that become an open item in the investment memo.

QUESTION 01 ANSWERED BY THE FOUNDER

What did AI cost you last month, in dollars?

An opener, and a filter. The partner is checking whether the number exists at all, not whether it is large. It is harder than it sounds. 68% of AI builders found the true cost of internal AI very or moderately difficult to predict, and only 7% found it easy (ICONIQ, July 2026).

A DEFENSIBLE ANSWER

The figure, without looking it up, with the providers and tools it spans and the split between the two pools.

A PLAUSIBLE ANSWER

"Roughly what our main provider billed." Honest, and it leaves out the AI features inside other tools. If the baseline needs a week, they assume everything below it does too.

GET READY BY

Adding up every provider plus the AI features inside tools you already pay for. The second half is the part most people miss.

QUESTION 02 ANSWERED BY THE CFO OR FINANCE LEAD

Does that number match what your providers actually invoiced you?

The question that separates prepared from confident, and the one the investor's playbook uses to decide whether every other number is a measurement or an estimate. Only 13% of organizations have even basic visibility into actual AI spending (Harness, July 2026). A figure computed from usage logs has no way of knowing it is wrong. The invoice does.

A DEFENSIBLE ANSWER

Yes. A named person in finance compares tracked cost to each provider invoice every month, and you can say when the two last disagreed, by how much, and why.

A PLAUSIBLE ANSWER

"Our dashboard tracks it in real time." Real time is not the same as reconciled, and every figure you give afterwards becomes provisional in their notes.

GET READY BY

Pulling last month's invoices and comparing them line by line to whatever your tooling reported. Do it once before the process starts, then monthly.

QUESTION 03 ANSWERED BY THE CTO

How much of your AI spend sits behind something you sell?

Your gross margin depends on the answer. Software companies building AI products spent 11% of revenue on AI for internal productivity in 2025, projected to reach 16% in 2026 (ICONIQ, July 2026). Mixed into one line with production cost, that spend makes both numbers unreadable, and the margin in your deck becomes an estimate.

A DEFENSIBLE ANSWER

Two monthly figures, an owner for each, and a written rule for when a workload becomes cost of revenue, with the workloads that crossed in the last two quarters and their dates.

A PLAUSIBLE ANSWER

"It is all R&D at this stage."
Possible at pre-seed. Once customers pay for an AI feature, one number means the margin line cannot be verified, which is the fastest route to a discount.

GET READY BY

Tagging spend at the feature level, then writing the rule: shipped to a paying customer means cost of revenue from that day.

QUESTION 04 ANSWERED BY THE FOUNDER

What does your most expensive customer cost you to serve?

They are testing whether you know your tail. Only 43% of companies track costs by customer and 22% by transaction (CloudZero with Benchmarkit, February 2026). Segment economics can diverge sharply inside one product. TechCrunch reported Cursor still losing money on individual developer accounts while holding positive margins on enterprise (April 2026, sourced reporting rather than company disclosure).

A DEFENSIBLE ANSWER

A named account, a figure with its revenue beside it, the share of production cost your top five accounts carry, and a count of any accounts served below cost.

A PLAUSIBLE ANSWER

"Usage is fairly even across the base." An average offered where a distribution was asked for, and it invites the follow-up you least want: is any customer unprofitable, and would you know?

GET READY BY

Ranking every account by AI cost to serve. Knowing your top five by name is most of the work.

QUESTION 05 ANSWERED BY THE FOUNDER

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

The pricing shape that breaks first at scale. Consumption pricing among AI builders rose from 35% to 42% in six months, and outcome-based pricing from 18% to 23% (ICONIQ, July 2026). GitHub moved every Copilot plan to usage-based billing on 1 June 2026, converting flat subscriptions into credits consumed by token usage (GitHub, April 2026). When a category leader changes its pricing shape, the cost shape underneath is usually why.

A DEFENSIBLE ANSWER

The plans with that shape named, the unit each one sells and what it costs to deliver, the break-even usage for each, and the number of accounts already above it.

A PLAUSIBLE ANSWER

"Heavy users are rare, and margins are healthy." An aggregate answer to a question about a cost that varies by customer and by plan. It reads as not having modelled the thing that breaks first.

GET READY BY

Modelling break-even usage per plan, then counting the accounts already above it. Bring the count, not the reassurance.

QUESTION 06 ANSWERED BY THE CTO

If a customer doubled usage tomorrow, would you find out before the invoice?

A surprise you caught is a story about controls. A surprise your provider told you about is a story about the absence of them. 72% of organizations had an unexpected AI cost spike in the past year, a third were caught off guard more than once, and only 20% could identify the cause of a spend doubling within hours (Harness, July 2026, a vendor-sponsored survey of 700 leaders).

A DEFENSIBLE ANSWER

A threshold per customer and per feature, a named person who receives the alert, and an example of a spike you caught early.

A PLAUSIBLE ANSWER

"We review spend every month." Monthly review finds the spike after it has been paid for, and invites a question about the largest one you have had.

GET READY BY

Putting one threshold alert per account and per feature in place. Even a crude one changes the answer.

QUESTION 07 ANSWERED BY THE CTO

How much AI are your people using that you did not approve?

This one usually arrives through the security questionnaire rather than the margin conversation, and it touches both. In a survey of 1,250 office professionals at large companies, 66% said they had used an AI tool at work they believed violated company policy, and 88% had shared work information with public AI tools (PagerDuty, June 2026). Your team is smaller, but the habit is the same, and unapproved tools are spend and data exposure nobody is tracking.

A DEFENSIBLE ANSWER

You have measured it, even roughly, and can say how. A small honest number beats a confident zero.

A PLAUSIBLE ANSWER

"We have a policy." It answers a governance question with an intention, and in a data room it becomes an open item rather than a closed one.

GET READY BY

Checking which AI tools appear on company cards and in single sign-on logs. Most teams find something.

QUESTION 08 ANSWERED BY THE CTO

What happens to you if one provider changes its prices?

Concentration risk, and it moves in both directions. AI builders now run about 3.3 model providers on average, and 81% cite Anthropic (ICONIQ, July 2026). On 30 July 2026 OpenAI cut one model's price by 80%, and Anthropic's introductory Sonnet 5 pricing ended on 31 August 2026, reverting up 50% (CloudZero and Digital Applied trackers, August 2026). A rate limit does not show on the invoice. It shows up as a customer outage.

A DEFENSIBLE ANSWER

Share of production spend by provider and model, the gross margin effect of a stated price move already calculated, and a tested fallback when a limit is hit.

A PLAUSIBLE ANSWER

"We are model-agnostic and could switch in a week." Probably true, and silent on how long it takes, what it costs and what breaks. They will price the uncertainty.

GET READY BY

Running one number: what a move the size of the 50% Sonnet 5 reversion, applied to your largest provider, does to this year's gross margin.

Is your cost per customer going up or down?

The question the valuation turns on. Total spend rising is expected as you grow. Cost per customer rising means the model gets worse at scale, and that is a different company. 66% of AI builders improved cost per query or task by at least 10% over six months, 19% by more than 30% (ICONIQ, July 2026). The same report describes one workflow projected at \$0.10 per run that reached \$1.50 or more as agents retried. Only a per-customer, per-outcome view would have shown it.

A DEFENSIBLE ANSWER

A per-customer figure with at least six months of trend, the explanation for its direction, and the three changes that moved it most.

A PLAUSIBLE ANSWER

"Margins will expand as model prices fall." A forecast resting on the provider's roadmap. Showing totals only leaves them to assume the trend is bad, because founders who like their trend always show it.

GET READY BY

Plotting cost per customer monthly for six months. If it is rising, bring the reason and the plan rather than leaving the chart out.

Who owns this number?

The close. AI cost sits between the engineers who create it and the finance function that reports it. 52% of organizations say there is no clear owner of AI cost (Harness, July 2026), and 78% of FinOps teams report to a CTO or CIO rather than to finance (FinOps Foundation, February 2026). In a startup, without one owner, cost per customer is nobody's forecast.

A DEFENSIBLE ANSWER

One name, the monthly review the numbers go to, and a decision they changed last quarter. If the company is small enough that the name is yours, say so plainly.

A PLAUSIBLE ANSWER

"Engineering and finance share it." Shared ownership of a forecast reads as unowned, and it undermines every figure you gave in the previous nine answers.

GET READY BY

Deciding whose forecast it is before the meeting, not during it.

WHERE YOU STAND TODAY

Count what you could answer right now, without preparation.

Score an answer as defensible only if it comes with a document, a name or a date attached. Confidence is not evidence. "We do not measure that yet, and here is when we will" is a stronger answer than a confident figure that falls apart under one follow-up, because the second turns a question about instrumentation into a question about judgement.

EIGHT TO TEN

Lead with it.

This is a genuine differentiator and almost nobody has it. Put the two-pool split and the reconciliation history in the deck rather than waiting to be asked. Volunteering it is worth more than surviving the question.

Offer the monthly metric set to the lead before they ask for it.

FOUR TO SEVEN

Fix two, disclose the rest.

Pick questions two and three. They make the other answers credible, and they are the two most likely to be probed. For the rest, say what you do not measure yet and when you will.

Put a name and a date against every gap in the data room.

THREE OR FEWER

Normal, and fixable in weeks.

Fewer than half of companies attribute AI cost by customer at all, so this is where most of the market sits. It is an instrumentation gap, not a judgement on the business. The only risk is discovering it in the room rather than before it.

Start the ninety-day plan with the invoice tie and the pool rule.

One thing worth saying plainly. Investors are not expecting perfect numbers from an early company. What ends conversations is a confident figure that falls apart under one follow-up. A founder who names the gap and the date it closes has given the partner something to write in the memo besides an open item.

BEFORE THE MEETING

AI cost. Ten questions they will ask you.

Mark each one honestly, then work on the blanks. Send it to your CTO and finance lead so all three of you rehearse from the same sheet.

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

Count an answer as ready only if you could give it in the room, today, with a document, a name or a date attached.

WHERE THIS CAME FROM

We measure AI cost at the unit your investors underwrite.

Crystil is an independent third party, SOC 2 compliant. Its 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 or team beneath it.

Workflows group many model calls into one logical task, so a completed outcome has one cost.

Customers sets cost per customer beside revenue per customer to give margin per customer.

Pulse routes a team's AI coding tools, such as Claude Code, Codex and Gemini CLI, through Crystil with no SDK, which makes the R&D pool visible.

We do not reconcile provider invoices, set target margins or value companies. We provide the attributed record a finance team reconciles against its invoices. Questions three, four and nine are the ones our record is built to answer.

[Talk to an AI measurement expert](#)

crystil.com

COMPANION GUIDES

For investors: the VC portfolio playbook, the questions and monthly metric set your lead will use. **For your CTO:** the AI product cost playbook.

Sources

1. CloudZero with Benchmarkit, FinOps In The AI Era, 12 February 2026. Vendor-commissioned.
2. Bessemer Venture Partners, The State of AI 2025, 13 August 2025, published before 2026. CRV, LLM Inference Costs: What AI Founders Should Know, 16 July 2026.
3. Fortune, interview with Sierra CEO Bret Taylor, 21 April 2026. A named executive's view, not survey data.
4. ICONIQ Growth, 2026 State of AI: The Builder's Economy, July 2026. Survey of about 300 executives fielded in the second quarter of 2026. The pass-through figure rests on 72 respondents.
5. Forbes citing EpochAI, 28 July 2026. Vista Equity Partners, AI Inference Economics: A Framework for Investors, 27 May 2026.
6. Harness, 2026 State of AI in FinOps, 29 July 2026. Survey of 700 engineering and FinOps leaders by Sapio Research, vendor-sponsored.
7. TechCrunch, reporting on Cursor, 17 April 2026, sourced reporting, not company disclosure. GitHub, GitHub Copilot is moving to usage-based billing, 27 April 2026.
8. PagerDuty, survey of 1,250 office professionals in non-IT roles at companies with \$500M or more in revenue, 11 June 2026.
9. CloudZero and Digital Applied pricing trackers, 5 August 2026. FinOps Foundation, State of FinOps 2026, 19 February 2026.

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.