

Ten questions an operating partner should ask every portfolio company about AI, and what a defensible answer sounds like.

Written for the operating partner who has to put AI cost into a diligence memo, a value-creation plan and an exit pack, and defend the number each time. AI now sits in two places in a software or tech-enabled services business. It is a cost of goods sold where it runs inside the product, and an operating expense where the workforce uses it. Both move margin. Neither is evidence until someone can reproduce it from source records.

WRITTEN FOR

The operating partner, value-creation lead or deal partner at a private equity firm with software and tech-enabled services holdings. You own the margin story from the diligence memo to the exit, and a buyer's diligence team will test every line of it.

ALSO FOR

Portfolio company CFOs and CEOs, who will be asked these questions and can prepare the answers before the operating review. The firm's diligence providers, who can use the ten questions and the evidence grades as a working template.

NOT FOR

Venture investors evaluating early AI-native companies, where the questions are about unit economics at scale rather than a hold-period baseline. That is the Crystil VC portfolio playbook. Nothing here is legal, tax or investment advice.

Buyers have started asking about AI in diligence. Almost no sponsor can yet show the evidence.

Two years ago the AI question in a portfolio review was whether a company had a strategy. The question now is what AI costs the company, where that cost sits in the P&L, and whether the number survives a buyer's quality of earnings work. Most portfolio companies can produce a total. Very few can produce a record.

9%

of private equity respondents are confident they could pass an independent AI governance audit within 90 days, the lowest of any sector. Across all industries the figure is 22%. Only 7% of PE respondents have a tested AI incident response plan.

Grant Thornton, 2026 AI Impact Survey, April 2026

44%

of PE operating partners report buyers specifically asking about AI-enabled finance capability during diligence. Only 9% have seen a demonstrable premium for it in a completed transaction.

Accordion and Wakefield Research, August 2026

11%

of organizations forecast their AI costs within plus or minus 10%, down from 15% in 2025. 40% escalated a surprise AI cost to the board.

Mavvrik and Benchmarkit, July 2026 (vendor-commissioned)

66%

of boards now require proof of AI return before approving further AI funding. 78% of finance executives cannot fully tie AI spending to business outcomes.

CloudZero, June 2026 (vendor survey of 260 finance leaders)

Read the third card twice. Nine in ten companies cannot forecast their AI line within 10%, and a buyer will haircut a forward plan built on one.

There are two pools of AI spend, and only one of them belongs in the margin story.

POOL ONE

Research and experimentation

Engineers trying models, prototypes, features that may never ship. Variance here is healthy. It is how the company finds what works, and clamping it down to hit a quarterly number is how a portfolio company stops learning. It sits in R&D expense.

Underwrite its size and its owner. Do not underwrite its volatility.

POOL TWO

Production

Anywhere the company is pricing, forecasting or promising value to a customer. That cost is AI COGS. It should be predictable, attributed per customer, and tied to what the providers invoiced. It drives gross margin, and gross margin drives the multiple.

This is the pool the value-creation plan and the exit pack depend on.

Almost no portfolio company can produce this split today, because its instrumentation reports one aggregate number covering both pools. That is why forward AI cost is so hard to underwrite. A third line sits beside the two pools: internal AI operating expense, the workforce using AI tools to do its job, from coding assistants in engineering to screening tools in HR. It is opex rather than COGS, and it carries most of the governance and employment-law exposure. Every review in this playbook asks for all three lines separately.

STAGE ONE. AT DILIGENCE

Ask the target for the split and the evidence behind it, not the total.

The diligence question has moved from what the company spends on AI to whether it can show the pools separately and the production pool tied to the invoice. Request four things, then ask question 08 on workforce AI.

- **AI cost in three lines.** Production pool, R&D pool and internal AI opex for the last four quarters, with the providers and tools each covers. A single figure is a finding.
- **The production pool per unit.** Cost per customer, per completed task or per priced feature, as a distribution rather than an average.
- **The invoice tie-out.** Whether the finance team reconciles tracked usage to provider invoices, how often, and what the last variance was.
- **The forward AI line.** Which usage drivers and provider prices it assumes, and how last year's forecast compared with actuals.

Grade the evidence, not the confidence, and write the grade into the memo

GRADE	WHAT THE TARGET PRODUCED	HOW TO TREAT IT
A. Reconciled	An attributed record by pool and by customer that the finance team ties to provider invoices on a stated cadence, with variances named	Underwrite it. Ask for the last variance and its cause.
B. Attributed	Cost computed from usage logs and attributed to customers or teams, not yet tied to invoices	Underwrite with a reconciliation condition in the first 100 days.
C. Invoiced	Provider invoice totals, accurate but not split by pool, customer or team	Accurate in total, silent on margin. Treat AI COGS as an estimate.
D. Asserted	A management estimate or a percentage of the cloud bill	Record as unknown. Build the baseline before the value-creation plan relies on it.

STAGE TWO. THE FIRST 100 DAYS

Establish an AI cost baseline per portfolio company before anyone sets a target.

75% of portfolio companies arrive at close unprepared or only partially ready for AI deployment, according to a May 2026 survey of 150 PE operating partners. A baseline is not a savings target. It is what every later claim is measured against.

- 01

Name one owner

One person, usually in finance, owns the AI cost record and reports to the CFO and the operating partner. 52% of organizations say there is no clear owner of AI cost (Harness, July 2026, vendor-sponsored).
- 02

Instrument production first, and write the pool rule

Price every model call behind the product and attribute it to the customer it served. Write down when a workload moves from R&D to production and who decides, then apply it the same way every quarter.
- 03

Tie the record to the invoices

Finance reconciles the attributed record against each provider invoice every month and names every gap. This moves a company from grade B to grade A.
- 04

Publish a dated baseline

Three lines, cost per unit the customer buys, the gross margin effect, reconciliation status and a first forecast. Every value-creation claim for the rest of the hold is measured from this document.

AI cost is a margin lever. Treat it like one, measured from the baseline, not asserted in the deck.

66% of senior PE leaders reported AI-related benefits within 12 months in 2026, up from 34% a year earlier, yet only 31% describe their AI implementation as efficient or mostly efficient. 19% of high performers exceeded their AI business case, against 5% of others (FTI Consulting, June 2026).

The headroom is real. AI product gross margins averaged 45% in 2025 and are projected at 53% for 2026 (ICONIQ Growth, July 2026), against roughly 80% for classic SaaS (SaaS Mag, May 2026). Vista Equity Partners has published that disciplined inference management can cut cost 80% or more with accuracy within 1 to 2% (May 2026). That is one sponsor’s framework, not a portfolio average. Test it company by company. And 39% of GPs did not expect AI to have any material financial impact on portfolio companies in 2026 (Bain and StepStone, March 2026). That is a reason to measure, not to skip the line.

LEVER	BASELINE MEASURE	OUTCOME TO PROTECT	EVIDENCE THAT CONFIRMS IT
Model selection and routing	Production cost per unit, by model	Task completion quality, manual review rate	Unit cost down with quality held, from the attributed record
Pricing change	Margin per customer, distribution of heavy users	Retention and expansion	Margin per customer improved without churn in the affected cohort
Workflow redesign	Cost per completed outcome	Customer outcome, not just cost	Cost per outcome falls, outcome rate holds or rises

SAMPLE VALUE-CREATION PLAN LINE

Reduce production AI cost per active customer from the baseline dated [date] by [target] by [quarter], with task completion rate held at or above [level]. Owner: [name], CFO reporting monthly. Evidence: the attributed cost record, reconciled to provider invoices by Finance, and the completion rate from the product analytics system.

Why this wording. A lower bill that raises manual review or hurts retention is a false win. This line names the outcome that must not degrade.

Build the AI evidence pack a buyer will request, from a running record, not a reconstruction.

Buyers are already pricing AI exposure into software assets. Mergermarket reported in July 2026 that software assets which once fetched 10x ARR or more now price at 4.5x to 8x ARR after the AI reset. A margin story nobody can reproduce from source records invites a discount.

ITEM IN THE PACK	WHAT IT CONTAINS, IN QUESTION ORDER
AI cost by pool	Production, R&D and internal opex, by quarter across the hold, with the pool rule
Invoice tie-out	Monthly reconciliation of the attributed record to provider invoices, every variance named
Unit economics	AI COGS as a share of revenue, margin per customer, the heaviest accounts, flat-price exposure
Forecast and provider exposure	Forecast against actuals by quarter, spend share by provider, margin effect of a price change
Workforce AI register	Tools that make or shape decisions about people, jurisdictions, notices given, bias audits
Value-creation results	Each AI initiative with its baseline, owner, measure and result
Ownership and governance	The named owner, the reporting cadence, the incident log, the retention period

Governance and employment-law exposure from workforce AI use

This reaches any portfolio company whose HR team uses AI to screen, rank or decide, and any plan that books AI headcount savings. For a 2026 acquisition, every date below falls inside a typical hold.

Illinois HB 3773 In force 1 January 2026	Amends the Illinois Human Rights Act. Prohibits employer use of AI that has the effect of discriminating on a protected class, and requires notice to employees before AI is used in recruitment, hiring, promotion, discharge, discipline and the other listed purposes. Crowell and Moring, September 2024. National Law Review, May 2025
Colorado SB 26-189 Employer obligations from 1 January 2027	Replaces the 2024 Colorado AI Act. Deployers of automated decision-making technology in consequential decisions, including employment, owe notice at the point of interaction, a plain-language explanation within 30 days of an adverse outcome, human review on request, and compliance records kept for at least three years. Colorado General Assembly, May 2026. Seyfarth, May 2026

Connecticut PA 26-15

WARN and no-defense provisions from 1 October 2026

WARN notices filed with the Connecticut Department of Labor must disclose whether the layoffs are related to the employer’s use of artificial intelligence. The anti-discrimination statute now states that using an automated employment decision tool is not a defense to a discriminatory practice complaint. Disclosure duties for those tools follow from 1 October 2027.

Paul Hastings, 2026. Mintz, June 2026

EU AI Act, Annex III

Applies from 2 December 2027

High-risk obligations for AI used in employment and HR apply to deployers, including informing workers' representatives and affected workers before a system goes live and keeping logs for six months. Article 2 reaches a US parent whose tool output is used at an EU subsidiary.

European Commission, July 2026. Regulation (EU) 2024/1689, Articles 2 and 26

THE PORTFOLIO STANDARD

Ask the same ten questions at every holding, and roll the answers into one table.

For a fund, the value is not any single answer. It is the comparison. Standardize the questions and definitions, keep the pools separate, and let each company keep its own unit.

COMPANY	QUESTIONS ANSWERED DEFENSIBLY	AI COGS AS % OF REVENUE	PRODUCTION POOL SHARE OF AI SPEND	OWNER	EVIDENCE GRADE
Portfolio company A	__ of 10	__%	__%	Name, title	A, B, C or D
Portfolio company B	__ of 10	__%	__%	Name, title	A, B, C or D
Portfolio company C	__ of 10	__%	__%	Name, title	A, B, C or D

Definitions. Defensible means the answer arrived with a document, a name or a date. AI COGS is the production pool only, over trailing quarter revenue. Production pool share is production AI cost over total AI cost across all three lines. Mark a cell unknown rather than estimate it. **Cadence.** Ask the ten at close, at day 100 and quarterly after that, and put questions 02 and 03 into every new diligence template. Read direction across quarters, not rank across companies.

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

Asked by the operating partner of a portfolio company. The plausible answers are not dishonest. They are what a competent person says when the instrumentation cannot do better, and they become open items in a buyer's data room.

QUESTION 01 ASK THE CFO

What did AI cost us last quarter, split into production, R&D and internal use?

The baseline and the split in one question. A single figure hides the moment a prototype became a feature customers depend on. 78% of finance executives cannot fully tie AI spending to business outcomes (CloudZero, June 2026).

A DEFENSIBLE ANSWER

Three figures, produced without notice, with the providers and tools each covers and the date of the baseline report they come from.

A PLAUSIBLE ANSWER

One total, or "it is inside the cloud bill," with an offer to break it out later.

THEN ASK

Does that include the AI features inside software we already pay for, and which line are they in?

QUESTION 02 ASK THE CFO

Does tracked AI cost tie to the provider invoices, and who reconciles it?

The question that separates grade A from grade B. A figure computed from logs is a calculation. A figure finance ties to the invoice every month is evidence, and a buyer's quality of earnings team will ask for it.

A DEFENSIBLE ANSWER

Yes, monthly, by a named person in finance, with the last variance and its cause on file.

A PLAUSIBLE ANSWER

"Our dashboard tracks it in real time." A dashboard has no way to know it is wrong.

THEN ASK

What was the largest variance in the last six months, and what caused it?

QUESTION 03 ASK THE CTO

Who decides when a workload moves from R&D into production?

Without a written rule the pools drift, and part of gross margin becomes an estimate.

A DEFENSIBLE ANSWER

A written rule, a named owner, and the list of workloads that moved last quarter.

A PLAUSIBLE ANSWER

"Engineering knows." Knowing is not a rule finance can apply.

THEN ASK

Which feature moved last, and when did its cost reach COGS?

QUESTION 04 ASK THE CFO AND CTO

What is AI COGS as a share of revenue, and what does our most expensive customer cost to serve?

AI cost concentrates in a small number of heavy users, so the average hides the exposure. AI product gross margins averaged 45% in 2025 and are projected at 53% for 2026 (ICONIQ Growth, July 2026), well below classic SaaS. Margin per customer is where that gap shows up first.

A DEFENSIBLE ANSWER

A percentage with its definition, a named account and its cost to serve, and the spread from cheapest to most expensive.

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 if they became so?

QUESTION 05 ASK THE CEO AND CFO

Where do we charge a flat price for something whose cost underneath is usage-based?

The pricing trap, and the lever the value-creation plan most often needs. 49% of organizations were affected by margin compression that forced them to reprice AI-powered products after unforeseen cost increases (Mavvrik and Benchmarkit, July 2026, vendor-commissioned). Repricing by plan is a decision. Repricing because margin broke is a finding.

A DEFENSIBLE ANSWER

The products with that shape named, the break-even usage level modelled, and the count of customers past it.

A PLAUSIBLE ANSWER

"Our margins are healthy." An aggregate statement about a per-customer problem.

THEN ASK

Is the number of customers past break-even growing, and what does the next price review assume?

QUESTION 06 ASK THE CFO

How close was last quarter's AI forecast to actual, and what is the forward line built on?

Forward AI cost is the hardest line in the plan to underwrite. Only 11% of organizations forecast AI cost within plus or minus 10% (Mavvrik and Benchmarkit, July 2026). 42% of decision-makers admit approving AI spending without reliable cost projections (CloudZero, July 2026).

A DEFENSIBLE ANSWER

Forecast against actual for each of the last four quarters, the variance explained, and the usage and price drivers behind next year.

A PLAUSIBLE ANSWER

"We grew last year's number by our growth rate." Usage does not scale with revenue that neatly.

THEN ASK

What would make next year's AI line miss by more than 25%, and would we see it coming?

QUESTION 07 ASK THE CTO

If our largest provider changed its prices, what would happen to gross margin?

Concentration risk, and it moves in both directions. Within 32 days in mid-2026 one major provider cut a model's price by 80% while another's introductory pricing ended with a 50% increase (CloudZero and Digital Applied pricing trackers, August 2026). A company that cannot model either move cannot claim one or plan for the other.

A DEFENSIBLE ANSWER

Spend share by provider and model, and the gross margin effect of a material price change, already run.

A PLAUSIBLE ANSWER

"We could switch providers."
Possibly true, and silent on how long it would take or what it would cost in quality.

THEN ASK

When did we last change a model for cost reasons, and what happened to quality and unit cost?

QUESTION 08 ASK THE CHRO AND GENERAL COUNSEL

Which of our AI uses make or shape decisions about employees or candidates, and where?

The question the employment instruments turn on. Illinois has applied since 1 January 2026, Connecticut's WARN and no-defense provisions since 1 October 2026, Colorado's employer obligations start 1 January 2027, and the EU's Annex III duties on 2 December 2027. Each depends on knowing which tools touch which people in which jurisdiction.

A DEFENSIBLE ANSWER

A named list of tools that screen, rank, score or recommend on people, tagged by state and country, with the notice given and its date.

A PLAUSIBLE ANSWER

"HR uses an applicant tracking system with some AI features."
Which features, on whom, where, is the whole question.

THEN ASK

If we ran a reduction in force in Connecticut next quarter, what would the WARN notice say about AI?

QUESTION 09 ASK THE CFO

Which AI value-creation actions have a baseline, an owner and a measured result?

The test of whether the plan is measured or asserted. Only 31% of senior PE leaders describe their AI implementation as efficient or mostly efficient, and high performers are nearly four times as likely to exceed their AI business case (FTI Consulting, June 2026). The ones that do can show the starting point.

A DEFENSIBLE ANSWER

Each initiative with its baseline date, target, outcome protected, owner and result to date, from source records.

A PLAUSIBLE ANSWER

"We cut our AI bill by a third."
Measured against which baseline, and what happened to quality?

THEN ASK

Which initiative did not work, and how did we find out?

If a buyer asked for the AI evidence pack tomorrow, what would we send, and who owns it?

The production test, and accountability in the same breath. Only 9% of PE respondents are confident they could pass an AI governance audit within 90 days (Grant Thornton, April 2026). 52% of organizations say there is no clear owner of AI cost (Harness, July 2026).

A DEFENSIBLE ANSWER

The eight items in the stage four table, produced from a running record within a week, and one named owner with a monthly report.

A PLAUSIBLE ANSWER

"We would pull something together." That sentence is the finding.

THEN ASK

If that owner left next month, would the record keep running?

READING THE ANSWERS

Count the defensible ones, not the confident ones.

Score an answer as defensible only if it came with a document, a name or a date attached. "We do not have that record yet, and here is the date we will" beats a confident number nobody can reproduce.

EIGHT TO TEN

Exit-ready on AI.

The remaining gap is usually question two or question six, because the invoice tie-out and forecast history take quarters to build. Close them and the pack stands up to a buyer.

Put the AI evidence pack in the quarterly operating review and keep it current.

FOUR TO SEVEN

Normal, and worth a decision.

Do not try to close every gap. Questions one to three are the foundation. Question eight carries the nearest legal dates. Put a name and a date against those four.

Write the four into the value-creation plan with owners.

THREE OR FEWER

An instrumentation problem, not a competence one.

Most of the market sits here. The risk is not that management is weak. It is that the buyer's diligence team discovers the gap first, on its timetable, with a price attached.

Start the 100-day baseline now, beginning with the production pool.

TAKE THIS INTO THE ROOM

AI cost and evidence. Ten questions for every portfolio company.

- | | | |
|----|---|----------|
| 01 | What did AI cost us last quarter, split into production, R&D and internal use? | CFO |
| 02 | Does tracked AI cost tie to the provider invoices, and who reconciles it? | CFO |
| 03 | Who decides when a workload moves from R&D into production? | CTO |
| 04 | What is AI COGS as a share of revenue, and what does our most expensive customer cost to serve? | CFO, CTO |
| 05 | Where do we charge a flat price for something whose cost underneath is usage-based? | CEO, CFO |
| 06 | How close was last quarter's AI forecast to actual, and what is the forward line built on? | CFO |
| 07 | If our largest provider changed its prices, what would happen to gross margin? | CTO |
| 08 | Which of our AI uses make or shape decisions about employees or candidates, and where? | CHRO, GC |
| 09 | Which AI value-creation actions have a baseline, an owner and a measured result? | CFO |
| 10 | If a buyer asked for the AI evidence pack tomorrow, what would we send, and who owns it? | CEO |

Score an answer as defensible only if it arrived with a document, a name, or a date attached. Roll the count into the portfolio table with AI COGS as a share of revenue, production pool share and the owner.

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 model, and attributes it to a customer, organization or team, with an optional department. Workflows gives a whole multi-call task one cost. Customers compares cost with revenue per customer to give margin per customer. Pulse routes a team's AI coding tools, such as Claude Code, through Crystil with no SDK. Sentinel is opt-in and off by default. Switched on, it applies rules the company sets, and can stop prompts carrying material the company has ruled out, such as employee records, and notify HR and Legal.

What we do not do. We do not reconcile invoices for you. We provide the attributed record your finance team reconciles against them. We certify nothing and give no legal, tax or investment advice. Our record is built to help answer questions one, two, four, five and nine.

[Talk to an AI measurement expert](#)

crystil.com

COMPANION GUIDES

Directors: the board playbook. Venture investors: the VC playbook. Question eight: our webinar, "Before the Notice Is Due: AI, HR Law and Board Liability Heading Into 2027", Thursday 5 November 2026, 11:00 AM ET.

Sources

1. Grant Thornton, 2026 AI Impact Survey, private equity and manufacturing reports (cross-industry figure), 21 April 2026.
2. Accordion and Wakefield Research, The PE AI Adoption Benchmark, 21 August 2026. 150 PE operating partners, fielded May 2026.
3. Mavvrik and Benchmarkit, 2026 State of AI Cost Governance Report, 29 July 2026, via Channel Insider. Mavvrik-commissioned.
4. CloudZero, press release, 24 June 2026, survey of 260 finance professionals. CloudZero, OpenAI pricing guide, 10 July 2026. CloudZero and Digital Applied, AI API pricing tracker, 5 August 2026.
5. Harness, State of AI in FinOps 2026, 29 July 2026. Vendor-sponsored.
6. FTI Consulting, 2026 Value Creation Index, press release, 4 June 2026. 555 senior PE leaders.
7. ICONIQ Growth, State of AI 2026, July 2026. SaaS Mag, AI COGS and SaaS gross margin compression, 15 May 2026.
8. Vista Equity Partners, Inference economics in enterprise AI, 27 May 2026. Bain and Company and StepStone Group, GP Outlook 2026, 2 March 2026.
9. Mergermarket (ION Analytics), AI reshapes software dealmaking, 8 July 2026.
10. Crowell and Moring, 23 September 2024, and National Law Review, 1 May 2025, on Illinois HB 3773.
11. Colorado General Assembly, SB 26-189, signed 14 May 2026. Seyfarth, 22 May 2026.
12. Paul Hastings, client alert, 2026, on Connecticut Public Act 26-15, signed 27 May 2026. Mintz, 23 June 2026.
13. European Commission, AI Omnibus enters into force, 27 July 2026. Regulation (EU) 2024/1689, Articles 2 and 26.

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