Architecture at the Table



A board briefing

7 Questions That Expose a Weak AI Plan

Seven questions any board member can ask in one meeting, the answers to listen for, and a simple score out of 14 to decide: fund, fund with conditions, or send back.

The plan that looks perfect

Sooner or later, an AI plan will land in front of you. Confident numbers. A slick demo. A vendor with a famous logo. And it will ask for a lot of money.

Here is the problem. Most weak AI plans do not look weak. They look exactly like strong ones, until about six months in. So how do you tell the difference, in one meeting, without being a technologist? You ask the right questions.

This briefing gives you seven of them. Each comes with the answer you want to hear, and the red-flag answer that should make you pause. Whether you sit on the board, sign the budget, or are the one presenting the plan, these questions will save you a great deal of money.

> **How to use it.** Ask the questions as a partner, not a prosecutor. The goal is not to catch the team out. It is to help them build a plan that survives contact with reality. Teams treated that way start answering the questions before you ask them.

The seven questions

1 What business outcome will this change, and how will we measure it?

If the plan cannot name a number that moves, it is not a plan. It is a hope.

The answer you want
Specific. Reduce the time to resolve a customer complaint from five days to two. Measured monthly. Owned by the head of service.

The red flag
Vague. We will become AI first. We will improve productivity. We will transform the customer experience.

2 What will this cost per outcome, at full scale?

Not in the pilot. At full scale. AI costs grow with usage, so a pilot that costs very little can become a significant line in the budget once ten thousand people use it.

The answer you want
A cost per outcome, compared with doing it today.

The red flag
A licence fee and a shrug.

3 How will we know it is good enough?

Demos are chosen to succeed. Real customers are not.

The answer you want
We have a test set of real cases, with agreed correct answers, and the system must pass a defined mark before it goes live, and again after every change.

The red flag
The demo went really well, or our users love it.

4 What data does this need, and do we have the right to use it?

Three parts to listen for. Is the data good enough? Can the system reach it in time? And are we legally and contractually allowed to use it this way? Personal data, customer contracts and data from partners often come with restrictions nobody checked.

The answer you want
Clear answers on quality, access and rights, owned inside the organisation.

The red flag
The vendor handles all of that. They do not. You are accountable for your data, whoever processes it.

5 What happens when it is wrong?

Because it will be. Every AI system makes mistakes. Accuracy is an average. Your reputation is decided by the worst case.

The answer you want
It describes the failure, and the safety net. When the system is unsure, it hands over to a person. High-risk decisions always get human review. There are limits on what it can do on its own. And it is monitored every week.

The red flag
The model is very accurate.

6 Who owns this on day two?

After the project team celebrates and moves on, who watches the quality, manages the cost, handles incidents, and decides when to update or retire it?

The answer you want
A named person, with a budget and a team.

The red flag
IT will support it. That usually means nobody will.

7 What would make us stop?

The one almost nobody asks. A strong plan has kill criteria, agreed up front, and an exit plan. If this vendor raises prices, or falls behind, can we move? Are our prompts, test sets and integrations ours?

The answer you want
If cost per outcome has not fallen by a certain date, we stop. If quality drops below the mark, we pause. Our data is portable and our assets are ours.

The red flag
Silence. A plan with no exit is not a plan. It is a commitment you have not admitted to yet.

The bonus question: does this need AI at all?

It catches the most expensive mistake of all. Some AI plans would be solved faster, cheaper and more reliably by a simple rule, a better form, or fixing a broken process. If customers cannot find the returns policy, the answer might be a clearer web page, not a chatbot.

AI is extraordinary at handling ambiguity, language and judgement at scale. If the problem has none of those, you are paying a premium for a capability you do not need. The best AI teams ask this question first, and they are not embarrassed when the answer is no.

Scoring the answers

For each question, give **0** if there is no answer, **1** if the answer is partial, and **2** if it is clear, specific and owned by someone. That gives a score out of 14.

11 to 14
Fund it. The team knows what it is doing.

7 to 10
Fund a smaller first phase, with the gaps written down as conditions and a date to come back.

Below 7
Send it back. Not as a rejection, but with a clear list of what a fundable version looks like.

Worked example: two plans, one meeting

Plan A promises to transform customer service with a leading AI vendor. Strong demo, vague answers: the outcome is "better experiences", the cost is "a licence", quality is proven by the demo, ownership sits with IT, and nobody has thought about stopping. Plan B is smaller: cut the time to answer supplier invoice queries from three days to one, with a cost per query compared with today, a test set of 500 real queries with a pass mark, a named owner in finance, and a rule to stop if the saving is not visible in three months. Plan B gets funded, because it is the one that can prove it worked.

Scoring sheet

Plan: _____ Date: _____ Presented by: _____

Question	None	Partial	Clear	Notes and conditions
1. What business outcome will this change, and how will we measure it?	0	1	2	
2. What will this cost per outcome, at full scale?	0	1	2	
3. How will we know it is good enough?	0	1	2	
4. What data does this need, and do we have the right to use it?	0	1	2	
5. What happens when it is wrong?	0	1	2	
6. Who owns this on day two?	0	1	2	
7. What would make us stop?	0	1	2	

Total: _____ / 14 **Decision:** ☐ Fund ☐ Fund a first phase with conditions ☐ Send back

If you are presenting the plan

Answer these seven questions before anyone asks them. Put them on one page at the front of your proposal. Nothing builds a board's confidence faster than a team that has already found its own weak spots.

Your Monday morning move. Take the largest AI initiative currently funded in your organisation. Ask the seven questions. Count how many get a clear answer. Fewer than five? That is your conversation for this week.

Architecture at the Table

A weekly show on enterprise AI, architecture and technology strategy, for architects, CTOs and the business leaders who fund them. New episodes every Tuesday and Thursday.

Watch the episode 7 Questions That Expose a Weak AI Plan · youtu.be/hPkDXC4pvCw

Subscribe youtube.com/@ArchitectureAtTheTable

More guides and case studies architecture.vinava.co.uk

Free to share inside your organisation. Please keep this page attached.