

Interviewing for Expertise You Do Not Have

A practical brief for executives interviewing candidates for their first AI ownership role

You are about to interview candidates who know more about AI than anyone on your side of the table.

That is not a hiring mistake. It is why the role exists. You need expertise that your company does not yet have.

But this creates a problem. The candidate who has built production systems and the candidate who has only studied them can use the same vocabulary. Both can cite the same tools. Both can describe successful projects.

The difference between them is the most expensive thing you cannot hear.

AI makes this problem worse than most fields. Titles do not have consistent meanings. “Led AI transformation” can describe an enterprise deployment or a series of workshops. Credentials provide some evidence, but they do not prove production experience.

The interview format also works against you. A candidate speaks with confidence for one hour. People outside the candidate’s domain must then grade the answers.

You cannot solve this problem with one clever technical question.

You can solve it by separating two tasks:

- 1 Verify the candidate’s technical depth.
- 2 Evaluate the candidate’s judgment.

These tasks require different methods. You do not have to perform both yourself.

SEPARATE THE TWO TASKS · SPEND YOUR HOUR ON JUDGMENT

VERIFY TECHNICAL DEPTH

use a qualified evaluator · part 8

EVALUATE JUDGMENT YOURSELF

scar tissue · plain explanation · flip the interview

1

Define the main job first

Before you design the interview, decide which job owns the mandate:

- **Platform:** Create the data and technical foundation that other AI work requires.
- **Adoption:** Change how employees perform recurring work.
- **Governance:** Establish the rules, controls, and accountability for AI use.
- **Build:** Deliver AI systems, automations, or product capabilities.

A strong leader may contribute to all four. One of them must still be the main job.

Then complete this sentence:

Twelve months from now, this person will have succeeded if they have primarily _____.

The answer must describe a business result, not a list of activities.

If you cannot complete the sentence, stop work on the interview. The company has not defined the role well enough to evaluate a candidate for it.

No interview process can correct an undefined mandate.

2

Three substitutes that fail when used alone

The memorized technical question

Someone gives you a “good question to ask AI candidates.” You ask it. The candidate gives a detailed answer.

You still do not know if the answer is correct.

A technical question is useful only when someone qualified can assess the answer and the question reflects the actual work.

The borrowed expert

You ask a technical friend, an investor’s CTO, or a board member’s contact to conduct an interview.

This can work if you define the assessment first.

Without clear instructions, the expert will interview for the job they know. An engineering leader will test engineering depth. A data scientist will test model knowledge.

Give the expert the written mandate. State what the candidate must be able to do. Ask the expert to test those capabilities, not general technical intelligence.

Credential weight

Certifications, titles, and brand-name employers are useful inputs. They are not sufficient evidence.

A successful project at a large company may have depended on a team, data foundation, and budget that your company does not have. A certification can show structured study. It cannot prove that a candidate can move a system from demonstration to daily use.

Use credentials to decide what to verify. Do not use them as a substitute for verification.

3

Separate technical depth from business judgment

You need a reliable way to verify technical depth. The method must match the main job. A qualified evaluator can review architecture, code, data practices, security controls, or production operations as required.

Your part of the interview has a different purpose.

You must determine whether the candidate can make good decisions inside your business.

In a 2024 study, RAND interviewed 65 experienced AI practitioners about why AI projects fail. Poor data, inadequate infrastructure, and technical limits all appeared. But leadership-driven failures were the most frequently cited category.

Companies selected the wrong problem. They optimized the wrong measure. Leaders and technical teams did not share the same objective. Priorities changed before teams could produce results. Systems did not fit the workflows they were supposed to improve. [Read the RAND report.](#)

A technically capable hire can still:

- Build something that employees do not use.
- Automate a task that does not affect a business result.
- Select a complex solution when a simple one would work.
- Ignore an approval, security, or adoption constraint.
- Spend the first year negotiating a mandate that leadership never defined.

These are failures of judgment, mandate, and fit.

You are qualified to evaluate judgment and fit. The mandate is yours to define.

The interview problem is not how to become qualified to grade every AI answer. It is:

How do I move the conversation onto an axis where I am the expert?

4

What AI fluency includes

In July 2026, Coinbase reported that the share of merged code generated by AI had reached roughly 100%. Humans still reviewed the code, but writing code from memory no longer reflected the work.

Coinbase rebuilt its engineering interviews around three dimensions of AI fluency:

DIMENSION	QUESTION
Usage	Can the candidate select and use the right tool to produce a better result?
Application	Does the candidate know when AI is and is not the right solution?
Understanding limits	Does the candidate understand failure, privacy, security, and the need for human judgment?

Coinbase applies this definition to junior and senior engineering candidates, and states that its interview signals ultimately test judgment and taste. [Read the Coinbase account.](#)

Your role may not be an engineering role, and Coinbase is not a model for every company. But its framework makes an important distinction.

Using AI is only one part of AI fluency.

A strong candidate must also decide:

- Which problems are suitable for AI.
- What evidence would justify an investment.
- Where the system can fail.
- Which risks are acceptable.
- Which decisions must remain with people.
- When the company should stop.

These are business decisions with technical inputs. You already evaluate this type of judgment when you hire leaders in other fields.

5

Interview for scar tissue

Ask:

Tell me about a time when AI made a process, product, or decision worse. What happened, and what did you change afterward?

Do not score the answer because it contains a failure. Failure stories can be prepared as easily as success stories.

Test whether the account holds together:

- What did you believe at the start?
- Which evidence showed that you were wrong?
- Who identified the problem first?
- What did the failure cost?
- What changed afterward?
- Is that change still in place?
- Who else can verify the account?

A strong answer connects evidence to a changed practice. The candidate can explain the original assumption, the failure, and the response. The details remain consistent as you ask follow-up questions.

A weak answer ends with a general lesson such as “people must remain in the loop.” It does not show how the lesson changed a system, a control, or a decision.

This is not a technical evaluation. It is interviewing for scar tissue, which executives have done their whole careers.

6

Require a plain explanation

Ask the candidate to explain an important past decision. If the answer depends on jargon, say:

Make that make sense to me without the technical vocabulary.

A strong explanation covers the options, the constraint, the tradeoff, and the result. It makes the decision clearer without removing necessary detail.

A weak explanation replaces one set of technical terms with another.

This role must explain uncertainty, cost, and risk to employees, executives, and possibly the board. The interview shows whether the candidate can do that work.

7

Flip the interview

The strongest assessment available to you begins with a real workflow.

Select one recurring process from your company. It could be a monthly report, a client brief, or a support escalation.

Describe the process as it actually operates. Include its delays, exceptions, manual steps, and informal workarounds. Do not make the description cleaner than the work.

Then say:

Treat this as your first discussion with us. Conduct the discovery you would need before you recommend any change.

Let the candidate run the next part of the interview.

Do not score the candidate only on the questions asked. Score the sequence of the questions and how the candidate uses each answer.

Listen for questions about:

- The result the process must produce.
- How the company measures that result now.
- Where the required data lives, and who can approve access.
- Which decisions require human approval.
- What the people who perform the work believe.
- What the company has already tried.

Do not provide all the information at once. Let the candidate find what is missing.

After the discovery period, ask for four outputs:

- 1 What do you think the problem might be?
- 2 What important facts are still unknown?
- 3 What would you measure before you change the process?
- 4 What would you do first, and what would you not do yet?

You can now observe how the candidate operates in an uncertain business situation.

Does the candidate separate facts from assumptions? Find the business constraint? Narrow the first step? Protect an important human decision? Select a reversible test? Name a reason to stop?

The candidate who proposes a solution in the first ten minutes is showing you how they will operate in month one: building before mapping, on your payroll.

The candidate who performs discovery but never makes a decision is also showing you how they will operate.

You want disciplined movement: neither immediate construction nor permanent discovery.

8

Verify technical depth with relevant evidence

Technical capability still requires separate verification. Use a combination of:

- A structured interview with an evaluator who has the mandate and knows what capability to test.
- A review of past work, with clear questions about the candidate's contribution and what happened after launch.
- A small, job-relevant work simulation that tests framing, tradeoffs, and communication. Pay for substantial work when appropriate.
- Specific references that verify what the candidate owned, what reached production, who used it, and whether it remains in use.

Do not ask for a general verdict that the candidate is “technical enough.” Define the required evidence before the assessment begins.

Remember that the candidate is also evaluating you

The candidates you most want may have seen AI roles with no real mandate.

They have seen companies hire a person instead of making a decision. They have seen leaders ask for transformation without granting access, authority, budget, or time.

They will look for evidence that your role is different:

- A written mandate.
- A clear main job.
- A named executive sponsor.
- A defined path for data access.
- A realistic first ninety days.
- A clear method for measuring progress.

Strong candidates take prepared jobs. The preparation that makes them easier to evaluate is the same preparation that attracts them.

IN BRIEF

What to listen for

WEIGHT THESE SIGNALS UP

Questions about results, baselines, data ownership, and approval before tools.

Specific failures that led to specific changes.

Plain explanations of technical decisions and business tradeoffs.

Past work described through adoption and results after launch.

A narrow, reversible first step and clear limits on scope.

WEIGHT THESE SIGNALS DOWN

A solution before the candidate understands the process.

Failure stories that remain general after follow-up questions.

Past work described mainly through tools and models.

Enthusiasm for every possible AI use.

No condition that would stop or change the candidate's recommendation.

You do not need to become an AI expert before you conduct these interviews. You need to define the job, separate technical verification from business judgment, and make the candidate operate in a domain you understand. That is enough to make the differences visible.

ABOUT THIS BRIEF

I have built these systems from start to finish: data ingestion, retrieval, user workflows, audit controls, and the governance decisions in between.

I am a Claude Certified Architect. This brief is the interview I would design knowing only that a company is hiring its first AI owner.

Your interviews will have different conditions. Some candidates will surprise you. The method holds either way: define the job, use a qualified evaluator for the technical verification, and keep the judgment questions on ground you know.

If the role is already posted and interviews are approaching, book a time below, then reply to the email that brought you here with the posting. I will come prepared to discuss what your interview needs to distinguish—and what it currently cannot. [Book a time →](#)

Scott Jennings

scott@sufficiently-advanced.ai