

Strategy First, or Pilot First?

A ten-minute self-assessment for leaders who can't yet tell whether AI activity is producing progress.

For established, regulated, and knowledge-intensive organizations that are past experimentation, but not yet confident enough to scale.

WHAT YOU WILL LEAVE WITH

A defensible next move: build direction, strengthen selection, or learn through bounded pilots, before mistaking AI activity for AI progress.

INSIDE

- 01 The consequence question that can settle the sequence immediately
- 02 A 12-point diagnostic across direction, selection, and execution
- 03 A practical recommendation matched to the layer your organization is missing
- 04 Two executive questions that reveal when activity is being mistaken for progress

JR Key Advisory

Independent AI strategy and value realization advisory

START HERE

The debate hasn't caught up to where you are

You've almost certainly heard both arguments: set the strategy first, because it tells you what's worth doing; or start with a small pilot, because momentum matters.

Both can be reasonable, but neither is universally the right approach.

They share an assumption that your organization is standing at a starting line, choosing a first move. By the time this question reaches an executive team, AI tools are usually already in use, someone has run a pilot, and teams have bought their own subscriptions. Those purchases are now the closest thing the organization has to an AI investment decision. Before choosing strategy-first or pilot-first, it helps to get clarity on where the organization stands right now.

That's also why the research can't provide a definitive answer. No study can tell your organization its next responsible move. That depends on the cost and reversibility of being wrong.

THE VALUE REALIZATION QUESTION UNDERNEATH THE DEBATE IS THIS:

Before you fund the next AI move, can you say what progress would look like, what evidence would prove it, and what would cause you to stop?

Most organizations are dealing with one of three gaps:

Direction gap. AI work is active, but leaders haven't agreed what AI is for, where it should create value, or what the organization is willing to change to get it.

Selection gap. Use cases exist, but initiatives are approved in isolation. And nothing is deliberately stopped, so nothing is truly chosen.

Execution evidence gap. Pilots are underway, but the organization has not defined what each one is meant to prove before it becomes the next request for funding.

THE USEFUL DISTINCTION

"Strategy first" usually means direction. "Pilot first" usually means execution. The layer most often missed is selection: the point where direction becomes a funded, owned portfolio, and initiatives that do not hold up are stopped.

THE THREE LAYERS

Three layers that get called “starting”

1 DIRECTION**PRIMARY OWNER**

Executive team

THE QUESTION What is AI for here?

Direction is agreement on where AI could plausibly improve business outcomes, what the organization is willing to change, and which decisions AI should or shouldn't influence.

When it is missing: Tools accumulate, spending rises, and nobody can say what the activity is meant to produce. The payback window quietly becomes the strategy.

2 SELECTION**PRIMARY OWNER**

Executive sponsor, portfolio or implementation lead

THE QUESTION Which AI initiatives deserve investment before others?

Selection is the discipline of choosing initiatives against direction, assigning owners, naming the evidence required, and defining the criteria that would stop each one.

When it is missing: Work is completed before anyone asks whether it should have been started. Activity creates the appearance of progress as projects accumulate.

3 EXECUTION**PRIMARY OWNER**

Business owner, delivery team, data and technology partners, risk partners where relevant

THE QUESTION What do we need to learn in real conditions before we scale?

Execution is how each chosen initiative is tested, delivered, measured, and integrated into real work.

When it is missing: Pilots succeed in clean conditions and fail against real systems, real data, or real workflows.

This guide identifies which layer is missing, and what should come before the next commitment of attention, budget, or organizational credibility.

STEP 1

The consequence gate

Answer this before anything else.

DECISION GATE

Would this AI system touch a decision that a regulator, auditor, or court could later ask you to justify?

EXAMPLES INCLUDE

- Credit, lending, pricing, eligibility, or underwriting decisions
- Clinical, diagnostic, or care-related support
- Anything producing a record you must be able to defend
- Compliance-sensitive recommendations
- Hiring, promotion, benefits, or performance decisions
- Safety-critical operations
- Customer-impacting decisions the organization may need to explain

IF YES

Before you test or deploy a high-consequence AI system, document why it was selected, who owns the decisions it supports, how it'll be validated, where its boundaries sit, and what evidence would justify continued investment. In regulated settings, some of this may be required rather than recommended, and testing may need to happen through an approved governance process or regulatory sandbox.

Legal, compliance, and risk teams can assess the system against those decisions. They can't supply the business purpose in place of leadership. If nobody has made those decisions, the defaults are set by whoever built or bought the system.

IF NO, OR ONLY PARTLY

A bounded pilot may be appropriate, especially when the organization needs evidence from its own data, systems, workflows, and exception rates. The pilot still needs a learning question, a stopping date, and a standard for what the evidence will be used to decide. Then continue to the diagnostic.

STEP 2

The diagnostic

Tick every statement that's broadly true today. Don't agonize over borderline answers. The pattern matters more than any single response. The blocks are intentionally unlabeled so the pattern is easier to see after you answer.

You may recognize your organization in more than one block. Score for weight, not perfection. The heaviest block should govern the next move. If two blocks are tied, use the recommendation map on the next page.

- BLOCK A**
- ☐ Two of our executives would give different answers to "what is AI for at this company."
 - ☐ We can name the tools we're using, but not the business outcome they're meant to move.
 - ☐ Nobody has said out loud what we'd be willing to change to get value from AI.
 - ☐ What we do next is shaped more by what's easy to buy than by where we're trying to go.

- BLOCK B**
- ☐ Initiatives get approved one at a time by whoever asks.
 - ☐ We couldn't explain why the projects in flight were chosen over the ones that weren't.
 - ☐ The main test an idea has to pass is whether it pays back within a single budget cycle.
 - ☐ No AI initiative here has ever been deliberately stopped.

- BLOCK C**
- ☐ Almost nobody on the leadership team has used these tools on real work.
 - ☐ We would be guessing about what's feasible if we tried to plan this today.
 - ☐ We have no internal reference point for what "working well" looks like.
 - ☐ We don't yet know enough to judge which of our ideas are the most promising.

COUNT YOUR TICKS

Carry the heaviest block, or tied blocks, to Step 3.

A _____ **B** _____ **C** _____

STEP 3

Read the pattern

Use your heaviest block. If A and C are both heavy, read that section first.

IF

**Block A
heaviest**

Build direction

The missing layer is agreement. Before the next AI investment decision, define what AI is for, where it could plausibly improve business outcomes, and what leadership is willing to change to get value from it. If executives disagree, find out why before booking the offsite; people who haven't agreed on the problem will produce a document, not a decision.

If Block B is also heavy, resist starting with portfolio discipline until leadership decides what matters. Otherwise each choice may look sensible on its own while the overall portfolio lacks a clear purpose.

IF

**Block B
heaviest**

Strengthen selection

Some direction exists. The missing layer is portfolio discipline. Before approving more initiatives one at a time, create a prioritized set of AI opportunities with an owner, evidence standard, review date, and stopping criteria for each. The test is whether the organization can explain why these initiatives deserve attention before the others.

IF

**Block C
heaviest**

Buy execution evidence with bounded pilots

Direction is broadly agreed. What you can't yet judge is which ideas will hold up in your environment. That's a question about your data, systems, and exception rates. Reading won't answer it. Run two or three bounded pilots, each with a stopping date and a question it's designed to answer. The goal is evidence you can select against, not a final portfolio.

If Block B is also heavy, treat the pilots as the first act of selection: name an owner, review date, and stopping criteria for each before it starts. Then take this assessment again.

IF

**Blocks A and C
both heavy**

Ground the room, then set direction

The organization needs direction, and the instinct may be to get more experience first. Executives already hold the business context a pilot can't supply: margins, constraints, bottlenecks, risks, and what the organization can't afford to get wrong. What may be missing is a current, vendor-neutral understanding of what these systems can and can't do. Close that knowledge gap first. Then set direction. Pilots answer which initiative works here, and that's a selection question that needs direction before it can be answered well.

If Block B is also heavy, the same sequence applies; add the selection discipline from the Block B section once direction is set.

STEP 4

Close the loop

Wherever you start, the work is incomplete until you define a goal you can revisit.

Direction without a measure is an opinion. Write down what should be different if the direction is working.

A portfolio decision without a review date is a preference. Set a date to review the evidence and reconsider the decision.

A pilot without a learning question is activity. State what uncertainty the pilot exists to reduce.

THE SIX-MONTH TEST

Before the next AI investment decision, write down what you expect to be true in six months, then set a date on the calendar to compare it with what actually happened.

EXPECTED CHANGE

EVIDENCE OR SIGNAL

DECISION OWNER

REVIEW DATE

WHAT WOULD CAUSE US TO STOP OR CHANGE COURSE?

The review is what allows leaders to answer the question: **What did we deliberately stop?**

EXECUTIVE SHORTCUT

If you do nothing else, use these two questions and listen for answers that are specific, consistent, and supported by a real example.

01 Ask three leaders separately: “What is AI for here?”

If the answers don't match, the teams below them can't make a coherent set of project choices. That's a direction problem, and leadership owns it.

02 Name an AI initiative you deliberately stopped, and the criteria that stopped it.

If nothing comes to mind, selection isn't happening. Initiatives are being launched, but the organization isn't making real choices among them.

For CEOs, these questions test decision quality. For sponsors and transformation leads, they test whether the recommendation is ready to defend.

SCOPE

What this guide doesn't answer

This guide answers one question: where to start before the next significant AI investment.

ORGANIZATIONAL READINESS

It does not assess organizational readiness. Whether your organization is equipped to deliver on a direction once it has set one is a separate question.

VALUE REALIZATION

It does not measure value realization. How you will measure realized value against what you said you wanted is another distinct question, and many organizations discover the gap only after money has been spent.

Those questions should not be skipped. They sit beyond what a ten-minute self-assessment can responsibly answer.

The purpose of the boundary is to keep a starting-point decision from being mistaken for the whole journey.

OPTIONAL RESOURCE

Free Executive AI Alignment Diagnostic

You just answered twelve statements alone. The free Executive AI Alignment Diagnostic is a way to put a related set of questions to multiple leaders independently, without requiring another meeting first. It takes about five minutes per participant and surfaces where leadership alignment appears clear, where it diverges, and which conditions may need attention before the next AI investment decision.

WHAT IT IS

A directional self-assessment with 24 questions across six organizational conditions.

WHAT IT ISN'T

A verdict, readiness certification, or disguised recommendation to buy a service.

Start the free diagnostic at jrkeyadvisory.com

EVIDENCE NOTE

This guide draws on management research rather than vendor surveys: three meta-analyses of the planning–performance literature covering over 250 study samples, a direct comparison of planned and iterative development across 72 product-development projects, and a randomized field experiment across 515 high-growth firms. The meta-analytic evidence is correlational. The field experiment is causal but is a working paper with a startup sample. Full sources and limitations are available on request.

AI activity is not AI progress.

JR Key Advisory is an independent AI strategy and value realization advisory for established, regulated, and knowledge-intensive mid-market organizations.

I work upstream of implementation vendors, before platform selection, staffing models, and delivery momentum turn early AI activity into a portfolio the organization must defend.

That perspective comes from 13+ years in UX research and design, 10+ years in software development, and work as a certified facilitator and instructional designer. I've worked across federal government, telecommunications, insurance, pharmacy, and finance, and there's a pattern that has repeated in each: technology creates options faster than organizations create the shared judgment to use them well. AI accelerates analysis, production, and experimentation, which widens that gap. Closing it takes shared judgment, which is why, even in the age of AI, human collaboration is still the ultimate competitive advantage.

JR Key Advisory helps leadership teams assess where they are with AI currently, align on what AI is for, prioritize what deserves investment, implement with client-side discipline, and calibrate based on evidence.