

COMPASS ADVISORY GROUP

The Compass Guide

A practical guide to deciding where AI belongs in your business

Adam Anderson

15 September 2026 edition

The current reference is maintained at compassadvisory.ai/reference.

Contents

Read in order, or return to the chapter that fits the decision in front of you. The working templates and questions are at the end.

Author's note 3

Introduction: The Question Is Too Big 5

Chapter 1: Stop Asking for an AI Strategy 7

Chapter 2: Find the Costly Friction 9

Chapter 3: The 5x5 Business Map 11

Chapter 4: Appropriate Maturity 13

Chapter 5: Score Opportunities Without Pretending the Score Is Truth 15

Chapter 6: Adoption Is Part of the Product 18

Chapter 7: The Smallest Useful Experiment 20

Chapter 8: Stop, Continue, Invest, or Industrialize 22

Chapter 9: Business Outcomes Are the Portfolio 24

Chapter 10: Governance That Can Be Used 26

Chapter 11: What the CEO Owns 28

Chapter 12: Begin Monday 29

Appendix A: Working Templates 32

Appendix B: Questions for the Room 38

Closing 40

Author's note

You have a business to run, and now you're being told to do something with AI. You may see possibilities without knowing where to begin. That's a reasonable place to start this book. You don't need to arrive with a problem already selected, a project approved, or a tool you want to buy.

We begin by understanding how your company works. Where does it create value? What do people have to work around to get the job done? What could be better for customers or for the team serving them? Looking across those questions usually gives us several opportunities to consider. Then we can compare them and decide which deserves attention now, which needs more investigation, and which can wait.

I can help frame that thinking, ask useful questions, and challenge an assumption that deserves another look. You and the people doing the work bring the knowledge of the business. You develop the answers and own the choices that follow.

AI earns a place when it can improve something that matters, at a cost and risk you can defend, with a change your people can sustain. That might call for a sophisticated system. It might call for a modest tool used by three people. You may also find that an ordinary process improvement is enough, or that this is a poor time to take on anything else.

The benefit could be room to serve more customers, time back for your team, or less spending on work that isn't helping. We will need evidence before treating any of those possibilities as a result. Clarity helps us get there because it makes the next decision more useful.

I haven't written a catalog of tools or a prediction about which model, vendor, or interface will win. Those answers would begin aging while you read them. I've written a way to work through the choices: map the business, discover and compare opportunities, choose a manageable next step, and learn from what happens. Sometimes that means stopping. Sometimes it gives you a reason to continue, invest, or make a capability dependable enough for wider use.

You can use the whole method and the working pages on your own. If you choose to work with me, the value is in the questions, judgment, and conversation we bring to your situation. You remain responsible for the decisions in your business. Advisory work and implementation are separate scopes.

How to use this book

Start with the first executive exercise in Chapter 1. It helps you describe what the business is trying to accomplish and where everyday work deserves a closer look. If you already have a concern in mind, bring it along as one possibility. If you don't, the exercise gives you a way to begin without inventing a problem just to fill in a blank.

Read Chapters 2-4 together as you explore. The friction statement gives you a way to describe what you notice; the 5x5 map helps you look across the company; and the maturity ladder asks how dependably each relevant capability works today. Keep several opportunities open during that exploration. The map shows where work happens. The separate rubric in Chapter 5 helps you compare possible moves before choosing a priority for this quarter or a smaller next step.

Once you've made that comparison, Chapters 6-8 help you understand adoption, prepare a bounded experiment, and decide what its evidence justifies. Preparation, measurement, and a deliberate wait are valid choices too. Chapters 9-12 explain how to keep doing this without asking your organization to absorb more change than it can handle.

Use the appendices during the conversation. Each working page explains what it helps you decide, provides a place to do the work, and points to what follows. By the end, you should have an account of the current situation, several alternatives and the reasons for their ratings, an owner-chosen next move, and a date to revisit it. If you're ready to test, you'll also have an experiment card. The free scaffolding is complete enough to use without hiring me, and the same working sequence is available in the printable kit at compassadvisory.ai/diy-pack.

Introduction: The Question Is Too Big

“How should we use AI?” is an understandable question for a CEO or operational leader. It covers so much ground that answering it directly often produces a tour of possibilities. Someone shows tools, describes competitors, and draws a roadmap. You leave with plenty of ideas, but it can still be difficult to explain which one matters to your company and why it should come first.

I'd start a little closer to the work. How does your business create value today? Where do capable people have to compensate for a process that makes their job harder? What would you like customers to experience differently? You don't have to know the answer before we begin. Those questions help us discover the places worth examining.

As possibilities emerge, we can make each one more specific:

> *Where is the business experiencing a costly friction, and does AI deserve the right to be tested there?*

Now we can ask who experiences it, what happens today, what has already been tried, and what improvement would mean. We can also discover that AI adds little, that another approach is simpler, or that the company needs to prepare before testing anything. Those are useful answers because they change where you put your attention.

AI can touch almost every activity in a company. Your ability to take on change is much more limited. As a leader, you decide which questions deserve investigation, which uses are ready to test, how much learning you can afford, and what evidence would justify more investment. You also decide what to stop and what can remain at a perfectly ordinary level of capability.

The method has four movements:

1. **Map the business.** See where value is created, delayed, lost, protected, or put at risk, and use that view to discover several candidate opportunities.
2. **Choose the strongest opportunity.** Compare the candidates using business impact, likelihood of success, adoption friction, time to evidence, and cost and risk. Explain the reasons before choosing what deserves attention.
3. **Buy the evidence cheaply.** When conditions support a test, use a small piece of real work, a business owner, and a decision rule agreed in advance. When they don't, prepare, measure, or wait deliberately.
4. **Manage what follows.** Decide whether to stop, continue, invest, or industrialize, and keep those choices together as a portfolio of business outcomes.

You can do this without becoming an AI expert. You will need people who know the work, room for facts that challenge the plan, and a willingness to change your mind. The method gives that conversation a structure so you can make one justified decision and learn enough to make the next.

A case we will follow

We'll return throughout the book to Musou Movement, a coaching platform developed from March through September 2026. Sean, the founder and head coach, began with three jobs: managing a community, deepening engagement, and giving each member a path shaped by his teaching. Coaches and clients then helped the team see what deserved attention as they used it.

The case is condensed from project history, approved plans, implementation records, the application structure, and direct coach and client feedback. It shows how use can challenge a plan and make the next business decision clearer. It doesn't establish that every company should build software or follow Musou's sequence. Your situation still needs its own map, comparison, and evidence.

PART ONE: BEGIN WITH THE BUSINESS

Chapter 1: Stop Asking for an AI Strategy

An AI strategy can include sensible things: approved tools, data controls, training, vendors, automation targets, and a sequence of projects. The difficulty comes when that list develops separately from what the company is trying to accomplish.

You still have customers to serve, costs to manage, trust to protect, and a business that needs to be able to act tomorrow. Those responsibilities tell us what a new capability has to contribute. Once we understand them, we have a reason to discuss technology and a way to judge whether it's helping.

The technology-first trap

You've probably seen some version of this sequence:

1. A leader sees a compelling demonstration.
2. The company asks where the tool could be used.
3. Teams produce a long list of possibilities.
4. A few pilots begin because they're visible or easy to start.
5. Adoption varies, measures remain vague, and ownership moves around.
6. Nobody can say with much confidence whether the effort created value.

The trouble starts with the second step. Once the tool becomes the fixed point, people begin reshaping the business question to fit it. It's easy to stay busy that way.

Give the team a different starting point: understand what needs to improve, find what is getting in the way, and consider which capabilities could help. Keep more than one approach open. That leaves you free to use AI where it earns a place and to choose something else where it doesn't.

A strategy is a pattern of choices

People learn your strategy by watching what you fund, refuse, measure, and stop. Saying AI is important gives them little help with any of those decisions. A useful pattern of choices is more practical:

- We begin with business friction we can describe and investigate.
- We compare alternatives, then work one opportunity at a time until our decision process is reliable.
- We involve an owner and users before selecting a vendor.
- We learn through small experiments before making large commitments.
- We decide whether to act separately from whether to buy, assemble, or build.
- We design for adoption from the beginning.
- We stop when the evidence doesn't support further investment.
- We review the portfolio in terms of business outcomes.

That gives people enough direction to begin. They can explain why an opportunity deserves attention before asking for a technology budget.

The first executive exercise

This exercise helps you begin with the business, even if all you know is that you're supposed to do something with AI. It comes before choosing a project. Use the sentences below to describe how you succeed and surface several areas that may deserve attention. Take them one at a time, listen to the people closest to the work, and leave an unknown visible when nobody has a sound answer yet.

Complete these sentences without using the words *AI*, *automation*, *platform*, or *digital transformation*:

1. Our business wins when we consistently _____.
2. We lose time, money, customers, speed, or trust when _____.
3. The people closest to that problem are _____.
4. We currently know the problem is costly because _____.
5. A meaningful improvement would look like _____.

Repeat the later sentences for the different concerns that emerge. If you can't complete one, you've found something to investigate. Carry those observations into the friction statements and the 5x5 map; keep the possibilities open until you've compared them.

PART ONE: BEGIN WITH THE BUSINESS

Chapter 2: Find the Costly Friction

I use *friction* for a place where useful work becomes slower, more expensive, less reliable, harder to scale, or more dangerous than it needs to be. It gives us something concrete to examine.

A successful company can carry that burden for years. People make the result happen through extra meetings, personal spreadsheets, duplicated checks, memory, apologies, and work after hours. From a distance, the process looks healthy because the job gets done. Spend time with the people doing it and you may discover how much effort they contribute to keep it that way.

Five forms of cost

Looking at five forms of cost helps you describe the consequence of a friction before you decide whether it's worth pursuing.

Money. Look for revenue you aren't capturing, margin given away, avoidable labor, rework, leakage, or excess inventory.

Time. Notice waiting, searching, transferring, reconciling, rewriting, chasing approval, and rebuilding knowledge that already exists somewhere.

Customer experience. Slow responses, inconsistent answers, preventable errors, weak personalization, and handoffs customers have to manage can all make it harder to do business with you.

Speed. Some decisions and actions lose value while they wait. Find where the delay changes the result.

Risk. Examine uncontrolled data, undocumented judgment, unreviewed output, fragile dependencies, compliance exposure, and decisions nobody can explain.

One issue may touch several of these. Describe the connections, but take care with the arithmetic: counting the same consequence twice makes a proposal look stronger without improving it.

Friction before solution

A friction statement lets the team describe the same piece of current work before discussing a remedy. Use it for each candidate you discover through the map. Begin with an event you can recognize, explain what people do in response, and connect that behavior to an observable consequence.

> *When [trigger occurs], [people] must [current behavior], which causes [observable consequence]. We believe this costs [range or proxy], and we will verify that before proposing a solution.*

“We need an AI assistant for sales” skips most of that thinking. Compare it with this:

> *When a qualified inquiry arrives, account executives spend between two and six hours finding prior examples and adapting them into a response. The delay is visible in response time, labor, and inconsistent quality. We need to verify frequency, effort, and effect on conversion.*

The second statement gives us something to observe and leaves the solution open. Once you've written it, test the description with the people doing the work and record what still needs checking.

Put a range on it

You rarely have perfect information at the beginning. Use the records available, direct observation, and estimates from the people doing the job to put a defensible range around the consequence. Label the assumptions so someone else can follow your reasoning.

If the annual cost might be between $40,000 and $120,000, write that range and explain where it came from. The midpoint doesn't become a fact because it's easier to put on a slide.

This helps you judge the size of the opportunity. It can keep a small inconvenience from becoming a strategic program and put a sensible upper boundary on what a solution could cost. If the range is too uncertain to support a decision, measuring the current work may be the next useful step.

Choose one

Keep several frictions in view while you map the business and apply the rubric. Then choose one for your first pass through the rest of the method. Choosing early, before comparison, can leave you improving the first thing somebody mentioned while a more worthwhile opportunity waits elsewhere.

Working through one selected opportunity lets the team learn whether it can define value, explain its ratings, design a useful experiment, observe adoption, and stop when necessary. As that process becomes dependable, you can consider more simultaneous work. The number should follow your ability to learn and support change.

Musou: the broad vision met a specific friction

Musou began with a broad vision of community, engagement, and an individual journey. The early application showed that members could sign in, assess themselves, explore movement content, track progress, participate in a community, and receive coaching plans.

Using it exposed a more specific difficulty. The social experience had grown complicated, while members still lacked an obvious way to begin a routine. Adding more community features wouldn't necessarily help someone decide what to do next.

The team simplified the social network into a coach-led feed. It also changed the movement experience into three steps: choose a category, answer a few questions about the current situation, and receive a filtered playlist. Members had fewer decisions to make, and coaches gained useful control over the experience.

Another friction appeared in delivery. Sean had tracked his one-to-one coaching workflow in Google Sheets. Bringing sessions, movements, nutrition, media, and communication into the platform reduced the administration surrounding each client. More effective delivery and lower overhead created capacity to serve additional paying coaching clients.

Both changes became clearer when the team looked at the behavior that needed to improve. The original vision still mattered, but it didn't require keeping every feature in its original form.

PART TWO: MAP BEFORE YOU MOVE

Chapter 3: The 5x5 Business Map

A department chart tells you who reports to whom. A tool inventory tells you what you own. Neither gives you a complete picture of how work moves through the business, especially when it crosses several teams.

The Compass 5x5 map brings five business systems together with five operating elements. Their intersections create twenty-five cells. That gives people a shared way to describe the company without needing a separate project for every box. Use the view to discover candidate opportunities; you'll compare those with a different instrument, the scoring rubric in Chapter 5.

Five business systems

These five systems help you follow value through the company. Use them as prompts to notice work you might miss if you only discussed departments:

1. **Generate Revenue** - how the company creates demand, wins business, prices, sells, and collects.
2. **Deliver Value** - how the product or service is produced, fulfilled, supported, and improved.
3. **Support Operations** - the recurring work that keeps delivery reliable and the organization functioning.
4. **Enable Everything** - the capabilities that help people coordinate, learn, communicate, and act.
5. **Govern & Protect** - how the company manages money, risk, compliance, access, accountability, and continuity.

Five operating elements

Within each system, look through these five elements. Together they help you see what supports the work and where it depends on something fragile:

1. **Process** - the way work moves, including policy, handoffs, exceptions, and delay.
2. **People** - roles, judgment, skills, behavior, and accountability.
3. **Technology** - the systems through which people perform and coordinate work.
4. **Data** - records, documents, definitions, history, access, and quality.
5. **Vendors** - outside capabilities, contracts, dependencies, and incentives.

Keep the labels consistent so people can connect their observations across the map. A cell describes part of a business capability; it doesn't create a requirement to buy something.

How to run the mapping conversation

The mapping conversation gives you a shared account of today's work before you assess maturity or compare opportunities. Put the grid where everyone can see it and record plain observations in the relevant cells. Begin broadly enough to notice different possibilities, then spend more time where the consequences appear to matter.

For each cell, work through the questions that help you understand it, one at a time:

- What happens here today?
- Who owns the result?
- Who actually performs the work?
- What information enters and leaves?
- Where does work wait, repeat, or fail?
- What workaround have people created?
- What number would reveal improvement?

Let the answer guide the next question. If nobody knows, record what needs observing and who could help. You don't need to manufacture an answer to finish the grid. When someone proposes software, capture the idea and return to describing the current situation. Afterward, look across the cells for candidate improvements and the evidence needed to assess them.

Patterns matter more than boxes

A friction can appear in several cells because work crosses boundaries. A slow proposal might first look like a sales-process issue. The map may show that approved language sits in documents, pricing rules live with one executive, prior examples are hard to find, and nobody measures response time.

That combination matters. A tool that speeds up writing may leave the wait for pricing untouched. Looking at people, process, knowledge, workflow, and governance together gives you a better chance of understanding the whole capability.

Leave the mapping conversation with several worthwhile possibilities to examine and a clearer account of their connections. Twenty-five cells should help you see the business; they don't need to become twenty-five projects.

PART TWO: MAP BEFORE YOU MOVE

Chapter 4: Appropriate Maturity

Maturity gives us a way to discuss how dependably a capability works and how much the business needs from it. It loses that usefulness when people treat the highest level as the goal for every activity.

A low-frequency, low-risk task may work perfectly well as an individual practice. A high-volume, regulated, customer-facing decision may need an integrated system with monitoring. Both can be appropriate. Pushing every part of the company toward maximum maturity can add expense and risk without improving the outcome.

A six-level ladder

Use this operational ladder to describe what people can depend on today:

0. **None** - the capability does not exist in a dependable form.
1. **Individual** - a person can perform it, often using personal knowledge or tools.
2. **Approved** - the organization has accepted a method, tool, or policy.
3. **Integrated** - the capability is built into normal workflow.
4. **Connected** - it uses relevant data and coordinates across systems or teams.
5. **AI-native** - learning, generation, prediction, or intelligent action is fundamental to how the capability works.

These levels describe dependable operational capability, with or without AI. A documented invoicing checklist with named responsibility, routinely used in the normal workflow, may be level 3 even though it contains no AI. The evidence has to support the assessment. A pilot or vendor promise alone doesn't establish level 5.

Keep an unknown assessment marked *unassessed*. Zero means you have determined that no dependable capability exists; it doesn't mean you haven't looked yet. This ladder isn't a percentage of AI use, and its numbers are separate from the opportunity scores in the next chapter.

Current, appropriate, and gap

This assessment helps you decide whether a capability needs attention after you've described the work. For each relevant cell, record its current maturity, the level appropriate for the business now, and the evidence supporting both judgments. Use observations from the map and check the description with the people who rely on that capability.

- Current maturity: _____, or unassessed.
- Appropriate maturity now: _____, or not yet determined.
- Evidence and assumptions for each judgment: _____.

A current level of one and an appropriate level of three gives you a question worth investigating. It doesn't automatically authorize a project. If both are two, the capability may already meet the need. Levels 2 or 3 are often enough.

Carry any meaningful gap into your list of opportunities, with its reasons and unknowns attached. Then compare it with the other things the company could do.

Maturity is not sourcing

Once you know the outcome and the capability required, you can consider how to obtain it. You may **buy** a product, **assemble** existing services with internal workflow, or **build** something distinctive.

Keep that decision separate from maturity and priority. A strategically important capability doesn't automatically require custom software. An ordinary one still needs sensible governance if you obtain it from a vendor. Clarify what the business needs, use evidence to support the decision to proceed, and then choose the source that fits.

Musou: appropriate maturity depends on the job

Musou began as a mobile web application rather than a native app. That reduced cost and complexity while giving real users a way to test the member and coaching workflows. It was enough for what the team needed to learn at that stage.

The content system faced different demands as it grew. A small movement library could begin with hand-entered content. Hundreds of rows, duplicates, inconsistent names, and exposed video links made that approach harder to maintain. The team added import tools, update-and-insert behavior, name normalization, multi-category support, private storage, and temporary viewing links.

The required maturity changed because the operational requirement changed. Scale, privacy, and repeatability justified more capability in that part of the business while other parts could remain simpler.

PART THREE: CHOOSE THE BET

Chapter 5: Score Opportunities Without Pretending the Score Is Truth

Once the map has surfaced several opportunities, the scorecard helps you put them beside each other and explain your judgment. It makes assumptions easier to discuss. A total cannot establish that those assumptions are true or choose on the owner's behalf.

Describe each candidate and the first step you're comparing, then give it a whole-number score from zero to five on each of these dimensions:

1. **Business impact** - if it works, how much does the outcome matter?
2. **Likelihood of success** - given the process, data, technology, and constraints, how plausible is improvement?
3. **Adoption friction** - how likely are the people involved to use the changed capability? Five means low friction and easy adoption.
4. **Time to evidence** - how quickly can the company learn something decision-worthy? Five means evidence arrives quickly.
5. **Cost and risk** - how favorable is the experiment's cost and risk? Five means low cost and contained risk.

Beside every rating, write the reason, the evidence you have, and what remains unknown. An estimate can be useful if everyone understands what it is. Missing information should stay visible; it isn't a reason to invent a score or silently enter zero. You can prepare or measure before completing the comparison.

Keep the possible outcome separate from the scope of the first learning step. A large potential improvement may be worth investigating through a small record review, but the low cost of that review tells you little about the cost of a full rollout. If you later consider implementation, scope and compare that commitment afresh.

After the ratings are explained, apply the gates and bands below. Use them to focus the conversation on what you should learn or do next, with the owner making the decision.

Hard gates

If business impact scores 0 or 1, stop the candidate as scored. Technical feasibility doesn't justify a project with too little business value.

If adoption friction scores 0 or 1, stop it as scored as well. The people involved cannot absorb the change dependably enough. A system can work technically and still fail to help because it never becomes part of the job. If likelihood of success is 0, stop: the proposed approach has no credible path to the outcome under the conditions being scored.

Any of these conditions overrides the total. Record what would have to change before reconsidering, or compare another candidate. You may choose preparation or a deliberate wait, but don't use a high score elsewhere to push the same blocked experiment ahead. A genuinely changed scope needs a fresh assessment.

The scoring conversation

This conversation helps you understand why people see the same opportunity differently. Bring together the process owner, representative users, an economic owner, and someone responsible for risk. Let users speak before the most senior leader so the group can hear how the proposed change would meet actual work.

Discuss one dimension at a time and ask for the reason behind the rating. If finance sees impact at five and operations sees it at two, find out what each group expects to improve. If leadership sees adoption at five and users see it at one, spend time there. That difference may explain more about the opportunity than the final total.

Record the disagreement and its uncertainty instead of averaging it away. When two candidates are close, consider whether a one-point change in a disputed rating would change their order. The owner can then explain the choice, including any preparation, narrower scope, or deferral. Keep that explanation with the comparison.

Four portfolio positions

The Compass model adds the five whole-number ratings for a total from 0 to 25. It applies the hard gates first, then assigns one of four positions. The calculation is:

```
total = business_impact + likelihood + adoption_friction + time_to_evidence + cost_and_risk
```

Condition, checked in this order	Provisional band
Business impact or adoption friction = 0 or 1, or likelihood of success = 0, whatever the total	NOT YET
Total 21-25, with all gates passed	RUN - THIS WEEK
Total 17-20, with all gates passed	RUN - 2 WEEKS
Total 13-16 and business impact at least 4, with all gates passed	INVEST AFTER EVIDENCE
All remaining combinations	NOT YET

The working meaning of those positions is:

- **Run - this week.** The ratings suggest an important, plausible, adoptable opportunity with a contained way to learn quickly. Check that evidence, permission, and capacity support the proposed step before starting it.
- **Run - two weeks.** The candidate looks promising, with preparation to consider before a valid experiment can begin.
- **Invest after evidence.** The potential outcome matters, but the next task is to design evidence that could justify a later commitment.
- **Not yet.** The present conditions don't support proceeding as scored. Record what must become true before reconsidering.

Every band is provisional. It reflects the supplied ratings, verifies nothing, and gives no investment approval. A high total can still lead you to gather a baseline, prepare, or wait. The dates in the labels guide attention; they don't override missing access, permission, evidence, or room in the team's week.

Record both the calculated band and the owner's actual next decision. Include the alternatives left for later and their re-entry conditions. That lets the numbers help you think without concealing the judgment that matters.

Musou: select the business move, not the most interesting feature

When Musou introduced live workshops, it could broadcast video while members watched and used text chat. Using the feature showed that coaches needed considerably more around that moment: event scheduling, attendee management, communication, replay access, a public link, and ticket and revenue information.

Workshop operations became the stronger opportunity. Payment also had to agree with tickets, subscriptions, discounts, complimentary access, entitlements, and payment status so the application could determine who was entitled to enter.

Sean began offering one-to-many online workshops at $35 per workshop. That opened a meaningful revenue and cash-flow opportunity and became his first major “aha” moment about the platform. What had begun as community infrastructure could support a repeatable paid offer without requiring every customer interaction to happen one-to-one.

Coach feedback also exposed a sales problem in signup. A visitor interested in one workshop had to understand too much of the membership platform before buying. A public workshop page and an email-first guest purchase flow let the buyer establish fuller account access later.

Through the scorecard's lens, these moves combined visible business impact, fast evidence, and an adoption path grounded in what coaches and buyers were trying to do. The case doesn't supply numerical ratings here, so we shouldn't pretend to know an exact total.

It does help explain the different kinds of impact. Workshops opened a direct revenue path. Bringing one-to-one coaching into the platform improved delivery and allowed more paying clients to be served. Replacing Google Sheets reduced overhead, and the released capacity gave Sean room to grow. The business reason for the work extended beyond tidier administration.

PART THREE: CHOOSE THE BET

Chapter 6: Adoption Is Part of the Product

By the time a system is built, many of the decisions that affect adoption have already been made. That's why I want the people doing the work involved while we're still choosing and shaping the experiment.

They need a change that fits the job, solves a problem they experience, and preserves judgment they still have to exercise. If the new process asks them to carry more hidden effort than it gives back, enthusiasm for the technology won't make that burden disappear.

Study the workaround

A personal spreadsheet, saved prompt, private checklist, or unofficial database tells you that someone has already spent effort closing a gap. Looking at that workaround can help you understand the requirement and why an approved system isn't meeting it. Use this conversation during discovery and adoption planning, before deciding what should replace it.

Ask the person these questions one at a time, starting with what their approach makes possible:

- What problem does this solve for you?
- When do you choose it over the official process?
- What information does it contain that the formal system lacks?
- What would make you stop using it?

Begin with curiosity so the person can explain the job they're trying to do. Carry what you learn into the proposed behavior change and the experiment's adoption conditions.

Name the behavior change

This step turns a proposed capability into something you can observe people using. Do it before writing the experiment card so the team can see what has to change for value to appear.

“Deploy a knowledge assistant” describes implementation. It leaves the working behavior unstated. A more useful description is: “Customer-support agents use the approved assistant before escalating a policy question, and they correct or flag answers that lack a source.”

Now you can observe whether that happens. You can also discuss the access, trust, response quality, source visibility, correction, and escalation it requires. Write the equivalent behavior for your candidate, check it with the users, and carry those conditions into the scope and measures.

Protect human signatures

Before testing, identify outputs and actions that need a person's name attached to the review. This makes accountability clear while you still have time to design the workflow around it.

Common candidates include prices, legal commitments, employment decisions, regulated communications, safety decisions, sensitive customer responses, and statements that could materially change trust. Decide which apply to your work and what the reviewer must examine.

Keep the list specific enough to enforce and short enough for people to remember. “Use AI responsibly” leaves too much for someone to figure out at the moment of use. Take the agreed review requirements into the experiment's guardrails and the human-signature page in Chapter 10.

PART FOUR: BUY EVIDENCE

Chapter 7: The Smallest Useful Experiment

A proof of concept can show that something is possible to build. Before making a larger business commitment, you also need to learn how it performs in the work and whether people can use it under the conditions you actually have.

Design the experiment around the uncertainty blocking that decision. You may need a narrow trial of a changed process. You may first need to observe today's workflow and establish a baseline. The useful size is the smallest honest test that can tell you enough to decide what follows.

The experiment card

The experiment card brings the team's agreement into one place before work starts. It follows the comparison and adoption conversation, helping everyone see what they're testing, what it could mean, and when they'll decide. Aim for one page, using Appendix A.3; keep supporting evidence alongside it when needed.

Friction. What current problem are we addressing?

Hypothesis. What do we believe will improve, for whom, and why?

Owner. Who is accountable for the business outcome?

Users. Who will use the changed process in real work?

Scope. What narrow slice will be tested, and what is excluded?

Measure. What one primary measure will decide whether the result matters?

Evidence standard. What result would justify continuing or investing?

Guardrails. What data, actions, or outputs are prohibited or require review?

Decision date. When will we stop testing and decide?

Decision choices. Stop, continue, invest, or industrialize.

Make the learning plan complete by also recording the outcome's link to the P&L or balance sheet, the bounded change, the baseline or the condition for obtaining one, a target or decision range, the time or sample boundary, cost and capacity limits, and the stop rule. Leave space for the observed result and next decision. If the financial link is still uncertain, say so; a plausible connection isn't a measured return.

If the card exposes information or agreement you don't have, choose the preparation needed to obtain it. A test intended to show improvement needs an appropriate baseline or measurement; supplying a narrative about the work doesn't establish one. A baseline-only step can record what will be observed and what would make a later test possible. It doesn't need an invented improvement target. Once the conditions are in place, use the card to keep the experiment within its agreed boundary.

Real work, small boundary

Use real inputs, representative users, and a live workflow when the risk allows it. Invented data can help you demonstrate a feature, but it usually leaves important questions about the organization unanswered.

Keep the scope manageable: one product line, queue, document type, customer segment, internal team, or recurring decision may be enough. Choose the boundary for the business evidence it can produce. A tiny build that tells you nothing useful is still more work than you needed to do.

Agree before seeing the result

The evidence standard gives everyone a fair way to interpret the experiment. Agree on it before starting, while you're still able to discuss disappointing results without defending the effort already spent.

Use a range if that's more honest than a single threshold, and include quality, adoption, and risk alongside speed. A faster process that creates more expensive mistakes may leave the business worse off.

Write down what each result would mean:

- Below the lower boundary: stop or redesign.
- Inside the uncertain band: continue only to answer a named uncertainty.
- Above the decision boundary: prepare the investment case.
- Strong result plus repeatable adoption: consider industrialization.

Carry those rules onto the card and into the decision meeting. If the first step is baseline gathering, agree on what you need to observe and the decision it will inform before claiming an improvement standard is ready.

Evidence has a price

Experimentation takes time and effort. The question is how much uncertainty it's worth reducing before you commit more of either.

A two-week test can be valuable if it prevents a year-long mistake, even when the idea being tested fails. Learning early gives you choices while the commitment is still manageable. Write down the cost and capacity you're willing to spend, so “we're still learning” doesn't become an unlimited extension.

Musou: real use purchased the evidence

Musou's most useful product decisions followed concrete attempts by coaches and clients. They used real devices during coaching sessions and workshops, with actual media and purchases. Their reports identified particular moments of confusion or lost trust: a routine was hard to begin, signup blocked a buyer, a replay didn't appear, a long recording saved only a few seconds, a microphone failed after being unplugged and reconnected, or a client couldn't scrub through a video.

Those reports gave the team different questions to investigate. The work moved into storage inspection, upload limits, media metadata, device-track lifecycle, permission boundaries, and recording-state logs because each could help explain an observed failure.

The value of that evidence was its connection to the next decision. It showed where the capability met real work and what needed to change for people to depend on it.

PART FOUR: BUY EVIDENCE

Chapter 8: Stop, Continue, Invest, or Industrialize

Put the decision meeting on the calendar before anyone builds. That meeting completes the experiment because it connects the learning to what the business will do next.

Begin with the card and the agreed evidence standard. Then examine what happened, including adoption and guardrail exceptions. A demonstration may help explain a finding, but the decision needs to rest on the evidence you've collected.

Stop

Stop when the outcome is too small, adoption doesn't occur, guardrails can't be maintained, cost exceeds plausible value, or the experiment fails to reduce the uncertainty that mattered.

Explain the expected range, the observed result, and why further work isn't justified. Thank the owner and users for helping the company learn while the commitment was still bounded. Record a re-entry condition if one exists, so the decision is available when circumstances change.

Continue

Continuing makes sense when another bounded test can answer a specific unresolved question. “We need more time” doesn't tell you what the business will learn in return.

Name the remaining uncertainty, explain why a further experiment is the least costly useful way to address it, and agree on a new limit and decision date. That gives the next round a purpose everyone can assess.

Invest

Invest when the evidence supports a larger commitment and the capability still needs implementation, integration, training, process change, or controls. Carry the original outcome measure and guardrails into the investment case, then account for the scope and cost of the larger decision.

Keep measuring the business result as the work progresses. Configuration, purchased licenses, and completed training tell you that activities happened. You still need to learn whether those activities improved what the company set out to change.

Industrialize

Industrialization is justified when the capability works, people adopt it repeatedly, the economics hold, controls are dependable, and broader use makes sense. It may require stronger infrastructure, integration, monitoring, support, documentation, and governance.

A successful pilot gives you something to examine; it doesn't make expansion automatic. If a small workflow creates its best value at that size, you can keep it small. Decide how much dependability and scale the business actually needs.

Musou: continue, invest, and industrialize were different decisions

Musou's recording workflow began by showing that a coach could record from a laptop and put video into a client's session. Feedback then showed that coaches wanted the polished workflow used for workshops. Continuing meant reusing a shared recording studio across workshops, private sessions, and permanent member routines.

Workshop failures exposed consequences that justified a larger investment. Browser recording was vulnerable to long files, lost connections, closed tabs, and device changes. The team responded with resumable uploads, smaller files, save-before-finish behavior, continuous local backup, manual replay attachment, server-side recording as the primary copy, and capture logs. Later hardware testing led to stable internal audio and video tracks so replacing a physical device didn't necessarily terminate the recording.

A lost hour-long workshop could erase irreplaceable work and damage customer trust. That made reliability, recovery, and visibility into failure part of the required capability. The broader investment followed the consequence people had encountered and the dependability the job now needed.

PART FIVE: MANAGE A PORTFOLIO, NOT A PARADE

Chapter 9: Business Outcomes Are the Portfolio

When you review a list of AI projects, tools, and vendors, it's easy to spend the meeting talking about what has started. Organizing the same work around business outcomes helps you ask what you've learned, which decisions are due, and what the results justify.

The outcome card in Appendix A.4 gives each item a consistent record. Use it after choosing an opportunity, then update it as the decision changes. Include:

- the friction and outcome sought;
- the business owner;
- current maturity and appropriate maturity, with unknowns visible;
- the current decision position;
- committed cost and attention;
- the primary measure and evidence range;
- the next decision date;
- the sourcing choice, if made;
- the current status: stop, test, invest, industrialize, defer, or complete.

Those records let you review the set of commitments together. You can then decide whether each still deserves the people, money, and attention allocated to it.

Attention is the limiting resource

Capital is relatively easy to see in a plan. Executive attention, user patience, implementation capacity, and appetite for change are easier to spend without noticing the balance.

Set a work-in-progress limit based on the owners and user groups who can sustain real participation. A company may be able to fund twenty experiments while having the capacity to learn honestly from only three. Give people enough room to observe, discuss, and adopt the change before adding another commitment.

Keep a real “not now” list

A deferred idea needs a reason and a condition for returning to it. This record helps you keep worthwhile possibilities available without allowing an unexplained backlog to claim attention every quarter.

For each deferred opportunity, write what must become true. For example:

- a data source becomes reliable;
- a process owner is assigned;
- a contract expires;

- volume reaches a stated threshold;
- a risk control exists;
- another experiment resolves a dependency.

Name who will notice the change or when the condition will be checked. At review time, examine that condition. An idea doesn't become stronger simply because it's been waiting a long time. If it now deserves reconsideration, update the comparison using the new facts.

Retire completed work

Move completed and stopped work out of the active view so the portfolio reflects commitments people still have to manage. Preserve the final evidence and decision where the team can find them.

That history helps the company avoid repeating an experiment without realizing it or reopening an idea without understanding why it stopped. The active list stays useful while the learning remains available.

PART FIVE: MANAGE A PORTFOLIO, NOT A PARADE

Chapter 10: Governance That Can Be Used

Governance helps when people can use it during the work. If the official rules live in a binder while employees privately invent practical ones, the organization has missed the conversation it needed to have.

Begin with three one-page controls: approved use, human signatures, and a check on inputs behind the metrics. They make expectations visible and give managers a way to handle uncertainty before it turns into a larger problem.

Approved-use page

This page tells people what they can use and where they need help or approval. Prepare it before relevant tools enter the workflow, using language a manager can explain during an ordinary working day.

State:

- which tools may be used on company accounts;
- what kinds of data may and may not enter them;
- which activities require additional approval;
- who receives training and when;
- how an employee reports a mistake or concern.

Check the page against a real task so you can see whether an employee could follow it. Then make it available alongside the workflow and review it as the tools or uses change.

Human-signature page

Use this page to turn the review requirements identified in Chapter 6 into a working control. List the outputs and actions requiring a named person, then define what that person needs to examine before approving them.

A reviewer needs access to the relevant source and enough understanding to make the judgment. A checkbox alone won't give them that. Test the arrangement on the bounded workflow and carry any gaps into the next decision about readiness.

Gamed-inputs check

Metrics influence behavior, sometimes in ways you didn't intend. A regular check helps you keep an improving number connected to the result it was meant to represent.

Each quarter, choose one result that improved unusually quickly and trace it back to its inputs. Confirm that the measure still means what it meant when you set the decision boundary. Treat this as a routine part of understanding the system: people tend to optimize what the organization rewards.

Bring any change in meaning to the quarterly review before relying on the number to justify further investment.

The quarterly review

The quarterly review connects individual experiments to the company's capacity and priorities. Use the same agenda so people know which evidence to bring and which decisions the meeting needs to produce:

1. Compare actual results with promised ranges.
2. Retire completed work.
3. Review re-entry conditions for deferred opportunities.
4. Reconsider sourcing and defensibility where the market changed.
5. Review approved use, human signatures, and metric integrity.
6. Ask: **What should stop?**

Record the decisions and assign the next actions. The final question needs a real answer when the evidence calls for one; otherwise the review can become a way to renew every commitment without examining it.

Musou: reliability and security grew with the relationships

As Musou added private coaching, payments, workshops, media, and messaging, a simple role label could no longer describe the access people needed. The team had to ask more specific questions: coach of which client, host of which workshop, buyer of which replay?

Database permissions were tightened around those relationships. Coaches were limited to assigned clients. Workshop hosts could see purchase information for their own customers rather than everyone. Private workshop media and coaching records were protected from unrelated users, and subscriber content was enforced below the interface. Roles were separated from profiles, movement videos used temporary viewing links, and payment status was reconciled through signed server callbacks.

Recording reliability needed similar care. Capture logs made the start, duration, upload, backup, and failure states observable. Local backup was removed only after a confirmed upload. If a recording didn't exist, the team had to say so and improve the safeguards. Implying that a missing asset could be recovered would only make the next conversation harder.

The controls became useful as they followed the actual relationships and consequences of the work.

PART SIX: LEAD THE DECISION SYSTEM

Chapter 11: What the CEO Owns

You don't need to own every AI project. As CEO, you do own the conditions under which work receives attention, money, legitimacy, and the right to continue. Operational leaders help make those conditions workable in the parts of the business they lead.

The five responsibilities below give you a way to support the people making decisions without becoming the answer to every question. An adviser can help you think them through and challenge assumptions. The business still owns its priorities and the choices about what to do next.

The business language

Keep the conversation connected to outcomes, friction, evidence, and choices. When a proposal arrives with urgency about AI, ask what it would change in the company and how you would know it helped. The same business discipline should apply as it would to any other use of attention and capital.

The right room

Bring together the economic owner, process owner, representative users, relevant technology capability, and risk judgment. Each sees constraints the others may miss. Give people close to the work room to explain what happens before the most senior person offers an answer.

Permission to tell the truth

People need to be able to say that the data isn't ready, the process is poorly understood, the tool isn't being used, or the expected value hasn't appeared. That includes telling you when your favorite idea should stop.

How you respond determines what you'll hear next time. If candor is punished, people can continue producing polished reports while the evidence you need disappears from the conversation.

Clear decision rights

Before starting, name who can authorize an experiment, approve data use, accept risk, fund investment, and stop work. People can then act within a clear boundary and know where to take an issue that exceeds it. Leaving those responsibilities ambiguous makes it harder to resolve a disappointing project when everyone has already invested in it.

The kill question

Ask what should stop as part of the ordinary review, before the portfolio gets crowded. Waiting for a budget crisis makes the conversation more difficult and gives weak work longer to consume capacity. The purpose is to return attention to opportunities that have earned it, including the ones you haven't had room to investigate yet.

PART SIX: LEAD THE DECISION SYSTEM

Chapter 12: Begin Monday

You can begin with the people who know the business and time set aside to learn together. You don't need a transformation office or a problem already selected. The schedule below shows how discovery can lead to comparison, preparation, and a bounded test.

Treat the dates as a working sequence, not a commitment to outrun the evidence. If access, a baseline, agreement, or capacity is missing, the useful next move may be to obtain it or wait. The calendar should make the decision visible, not make it for you.

Day 1: Name the friction

Begin with the first executive exercise and a broad look at the 5x5 map. This helps you notice candidate frictions before deciding which matters most. Describe what the business does, where people encounter difficulty, and what an improvement could mean.

Use the friction statement for the different concerns that emerge. Record the people involved, any plausible economic consequence, and what is known or assumed. Carry several possibilities into the closer mapping work over the next two days.

Days 2-3: Map the capability

Now use the relevant cells of the 5x5 map to understand the capabilities around those possibilities. Observe the work and speak with users about where information, judgment, process, workflow, and governance support or compensate for one another.

Keep the current situation separate from proposed solutions. Record unknowns and the observations needed to resolve them. That gives the maturity discussion a basis and may reveal additional alternatives worth comparing.

Day 4: Set current and appropriate maturity

This assessment asks whether the capabilities you've mapped are dependable enough for their business purpose. Describe how each relevant one operates today, use the ladder where evidence supports a level, and leave the rest unassessed.

Decide what maturity the business needs now and record why. Keep sourcing for a later decision. Bring the meaningful gaps, adequate capabilities, and remaining uncertainties into the opportunity comparison.

Day 5: Score the opportunity

Put the candidate opportunities beside each other so the company can choose where to focus. Use business impact, likelihood of success, adoption friction, time to evidence, and cost and risk, with a reason, evidence, and unknowns for every rating. Apply the hard gates and exact provisional bands.

Compare the explanations as well as the totals. The owner then chooses a test, preparation or measurement first, or a deliberate wait. Record why, which alternatives are deferred, and what would bring them back. If a test is appropriate, carry the selected scope into the experiment card.

Week 2: Design the experiment

Use the card to agree on the learning step before starting it. Name the owner and users, scope and exclusions, primary measure, evidence standard, guardrails, and decision date. Include the baseline or a plan to obtain it, target or decision range, sample or time limit, cost and capacity boundary, and stop rule.

Choose the smallest real slice capable of producing useful evidence. If the conditions for a test aren't present, record the preparation or deferral and its re-entry condition. Begin only when the selected step has the access, permission, and participation it needs.

Weeks 3-4: Run and observe

The purpose of this period is to learn from the agreed scope of real work. Track the primary measure, adoption behavior, errors, and guardrail exceptions. If you're gathering a baseline first, record the current situation without presenting it as an improvement.

Keep the test within its boundary even if early results are exciting, and raise disappointing findings while there's still time to respond. Bring the observations and any limitations to the scheduled decision meeting.

Decision day

Begin by reading the hypothesis and evidence standard. Present the observed result, explain any limits, and let the owner decide whether to stop, continue, invest, or industrialize. Use the decision record so the reason and next responsibility remain clear.

If the step was preparation or measurement, decide whether it has made a valid experiment possible or whether more preparation or a wait is warranted. Then return to the alternatives with what you've learned and choose the next friction that deserves attention.

Musou: the end-to-end loop

Looking across Musou's development history, you can see a repeated decision loop:

1. Begin with existing work: coaching members, prescribing movement, running workshops, communicating, and maintaining a practice.
2. Build the smallest version people can actually use.
3. Put it in front of coaches and clients quickly.
4. Capture the precise point of confusion, failure, or lost trust.
5. Diagnose the underlying cause instead of addressing only the visible symptom.
6. Generalize a successful solution when the same need appears elsewhere.
7. Test again under more realistic conditions.

The broad vision remained while the roadmap changed. Community became simpler and coaching more personal. Workshops developed into an operating and revenue system. Video grew to include creation, delivery, recovery, and feedback. Security moved from broad roles to access tied to particular relationships.

Sean's first major aha came with one-to-many online workshops at $35 each, creating a meaningful revenue and cash-flow opportunity. The platform then entered his one-to-one coaching workflow. Delivery became more effective, Google Sheets no longer carried the process, administrative overhead fell, and capacity for additional paying clients expanded.

That progression helps explain why efficiency mattered in this case. Removing friction released capacity, which gave Sean room to serve and sell in ways that had been difficult. The next business question is whether growth follows, and by how much. The case supplies no quantified growth result here, so the opportunity and the measured outcome must remain distinct.

The method in one page

Use this page when you want to place a current question back into the whole process. Begin where your evidence supports you, and move back a step if the conversation reveals something you haven't understood yet.

1. Start by understanding the business and its outcomes.
2. Map how work runs and discover several costly frictions or opportunities.
3. Describe the capabilities around them and the unknowns to investigate.
4. Name current and appropriate maturity where evidence supports an assessment.
5. Compare candidates on business impact, likelihood of success, adoption friction, time to evidence, and cost and risk, with reasons for the ratings.
6. Stop a candidate as scored if impact or adoption viability is 0 or 1, or likelihood of success is 0; record what would have to change.
7. Let the owner choose a priority and the smallest useful experiment, preparation step, or deliberate wait.
8. Agree on evidence, scope, capacity, guardrails, and the decision date before starting.
9. Use real work and representative users where conditions permit, and record what happens.
10. Decide whether to stop, continue, invest, or industrialize; distinguish preparation from demonstrated improvement.
11. Manage those decisions as a portfolio of outcomes, including deferred alternatives and their re-entry conditions.
12. Ask what should stop so there is room for the work that deserves attention next.

The tools will keep changing. Repeating this process gives the company a way to exercise judgment as they do, while keeping the people who own the business in charge of its choices.

Appendix A: Working Templates

A.1 Friction statement

This page describes a current situation before you select a solution. Use it as the map reveals candidate opportunities, filling out a separate statement for each. Start with an observable event and what people do in response; label estimates and leave unknowns visible so the next conversation can test the account.

Trigger: When _____ occurs,

People and current work: _____ must _____,

Observable consequence: which causes _____.

Estimated cost or proxy: We believe this costs _____ because _____.

Unknowns to verify: _____.

Check the statement with the people doing the work, connect it to the relevant map cells, and carry it into the comparison. You can keep several statements open until you're ready to choose.

A.2 Opportunity scorecard

The scorecard helps you compare the opportunities you've discovered. Complete one for each candidate, then place them together. Give each rating a reason, an evidence basis, and its remaining uncertainty. Keep the possible business outcome distinct from the cost and time of the first learning step. If a rating can't yet be supported, mark it unknown and gather what you need before calculating a complete total.

Candidate and relevant map cells: _____

Possible business outcome: _____

First step being compared, including scope and exclusions: _____

Dimension	Score 0-5	Reason and evidence for the score	Unknowns or disagreement
Business impact			
Likelihood of success			
Adoption friction, where 5 is favorable			
Time to evidence, where 5 is fast			
Cost and risk, where 5 is favorable			

Hard gates: If impact or adoption friction is 0 or 1, or likelihood of success is 0, stop the candidate as scored, whatever the total.

Total: _____ / 25, calculated only when all five whole-number ratings are supplied.

Provisional portfolio position: Run - this week / Run - two weeks / Invest after evidence / Not yet

Apply the gates first. Otherwise, 21-25 gives RUN - THIS WEEK; 17-20 gives RUN - 2 WEEKS; 13-16 with business impact at least 4 gives INVEST AFTER EVIDENCE; all other combinations give NOT YET. These are arithmetic bands on the supplied ratings. They verify nothing and approve no investment.

Reason for the comparison and any close or disputed ratings: _____

Owner's next choice: Bounded test / Prepare or measure first / Defer

Why this choice, given evidence, permission, and capacity: _____

Alternatives deferred and each re-entry condition: _____

Decision owner and review date: _____

Use the choice to complete A.3 if you're ready for a learning step. A high total can still lead to a baseline-only decision or a wait. Record a deferral and its condition in A.4 rather than treating the band as an instruction to begin.

A.3 Experiment card

This card turns the selected opportunity into an agreed learning step. Complete it after comparing candidates and discussing adoption, before anyone starts work. Aim for one page with supporting records alongside it. When the baseline is missing, describe the observation needed and the gate for a later test instead of inventing an improvement target.

- Friction:
- Hypothesis:
- Business outcome and link to the P&L or balance sheet, or why that link is still uncertain:
- Business owner:
- Representative users:
- Selected step: test / preparation or baseline measurement:
- Scope and bounded change:
- Exclusions:
- Primary measure:
- Current baseline, or plan and gate for obtaining one:
- Target or decision range, or what the baseline step must establish:
- Evidence standard:
- Quality check:
- Adoption behavior and measure:
- Guardrails, access, permissions, and required human review:
- Time or sample boundary:
- Cost limit:

- Capacity limit and available participants:
- Stop rule:
- Start date:
- Decision date:
- Decision choices: stop / continue / invest / industrialize; after preparation, decide whether a test is ready or should wait:
- Observed result and limitations, completed after the step:
- Next decision and reason:

Check that the selected step has the participation and permission it needs. If it doesn't, record the preparation or wait and the condition for returning. Review the completed evidence with A.5 on the agreed date.

A.4 Outcome card

The outcome card keeps the business purpose visible as an opportunity moves through testing, investment, or deferral. Start it when the owner chooses a direction and update it after each decision. It gives the portfolio review a consistent account of what the company has committed and what should happen next.

- Business outcome:
- Friction addressed:
- Owner:
- Relevant map cells:
- Current maturity, or unassessed:
- Appropriate maturity, or not yet determined:
- Evidence and assumptions for maturity judgments:
- Provisional score position, if scored:

- Experiment or investment status:
- Committed cost:
- Attention required:
- Primary measure and range:
- Adoption measure:
- Next decision date:
- Sourcing choice, if made:
- Current decision:
- Alternatives deferred and reasons:
- Re-entry condition, if deferred:
- Person or date responsible for checking the re-entry condition:
- Final evidence and decision-record location:

Bring the current card to the quarterly review. Move completed or stopped work out of the active list while preserving its evidence and decision.

A.5 Decision record

Use this page at the scheduled decision meeting to connect what you expected with what actually happened. Read the experiment's agreed standard first, record the observations and their limits, and let the owner explain the choice. If the step only gathered a baseline or prepared the work, say so before deciding whether an experiment can begin.

Expected: _____

Observed: _____

Evidence source, limitations, and remaining unknowns: _____

Decision: Stop / Continue / Invest / Industrialize

For preparation or deferral, next state: Test ready / More preparation or measurement / Wait

Why: _____

Next action and owner: _____

Re-entry condition if waiting: _____

Next decision date: _____

Update the outcome card with this decision, retain the supporting evidence, and return to the candidate comparison when it's time to choose the next priority.

Appendix B: Questions for the Room

This appendix gives you questions to use throughout the method, from discovering how the business works to reviewing the portfolio. Choose the group that fits the current activity, briefly explain what you're trying to understand, and ask one question on its own. Listen to the whole answer before deciding what follows; you don't need to work through every prompt to have a useful conversation.

When someone says they don't know, record what needs looking at instead of filling in an answer for them. The questions should help people develop their own account of the work and decide what deserves attention next.

To find friction

Use these questions when you need to understand the business and discover possibilities, including when nobody has a project in mind. Choose the question that fits what you've just heard, ask it, and listen before moving on. Record observations and unknowns in the map as the conversation develops.

- Where do capable people spend time compensating for the system?
- What regularly waits for one person?
- What is copied, reconciled, or rewritten?
- What do customers have to chase?
- Where do errors become visible only after they are expensive?
- Which decisions depend on knowledge that is difficult to find?
- Where has the team created an unofficial workaround?

Turn the different concerns into friction statements, then examine their importance before selecting one.

To test importance

These questions help you understand whether an observed friction deserves attention alongside the other candidates. Use them after describing the current work, keeping estimates distinct from facts. Take one at a time so people can explain the consequence behind their judgment.

- What changes economically if this improves?
- What happens if nothing changes for twelve months?
- Who owns the outcome?
- Which available number is closest to the value?
- What would make this problem too small to pursue?

Carry the reasons and uncertainty into the business-impact rating. If the basis is missing, identify the observation needed before relying on a total.

To test adoption

This conversation helps you see whether the people involved can absorb the proposed change. Use it while comparing opportunities and before designing the experiment. Ask representative users about their own work, one question at a time, and allow their answers to change the scope.

- Who must behave differently?
- What do they do today that already works for them?
- What new effort are we asking them to carry?
- What judgment must remain human?
- What would make the changed process easier to abandon than use?

Record the required behavior, support, and limits on the experiment card. If adoption viability is 0 or 1, stop the candidate as scored and record what would need to change.

To design evidence

Use these questions once the owner has selected an opportunity or a preparation step. They help you keep the learning small enough to manage and specific enough to inform a decision. Work through the relevant questions in conversation and use the answers to complete the experiment card.

- What uncertainty blocks the next decision?
- What is the smallest real test that can reduce it?
- What one measure matters most?
- What result would cause us to stop?
- What result would justify investment?
- When will we decide?

If you don't yet have a baseline, agree on how to establish it before claiming a target is supported. Finish by checking scope, participation, cost, capacity, guardrails, and the stop rule.

To govern the portfolio

These questions help the leadership team review its commitments as a whole. Use them with the outcome cards and decision records during the quarterly review. Take the relevant issues in turn and record an owner for any follow-up so the discussion changes what happens next.

- Which items have no active owner?
- Which items have no upcoming decision date?
- Which results are activities rather than outcomes?
- Which deferred item now meets its re-entry condition?
- Which completed item should leave the active list?
- What should stop?

Update the active portfolio, retain the learning from completed work, and choose the next priorities within the capacity people actually have.

Closing

The pressure to do something with AI will keep changing as tools and capabilities develop. Demonstrations that look remarkable today will become ordinary. Some expensive capabilities will become cheaper, and risks that once seemed theoretical will become familiar. Some choices you make will work; others will need to change.

You can give your company a dependable way to make decisions while that uncertainty remains. Start with how the business works, discover several opportunities, and compare them with reasons people can examine. Choose what deserves attention now and be clear about what will wait.

Then take a step small enough to learn from. Keep maturity appropriate to the job, include adoption in the design, and decide sourcing separately. Use the evidence to determine whether to stop, continue, invest, or make the capability dependable for broader use. Keep the pace within what your people can sustain.

You and your team own those answers. The method gives you a way to work through them together and carry the learning into the next choice. That's how we begin, and how we keep going.