



COMPASS ADVISORY GROUP

# The Compass Guide

*A practical guide to deciding where AI belongs in your business*

Adam Anderson

---

DOWNLOADABLE EDITION / SEPTEMBER 12, 2026

A snapshot of the Compass body of work.  
The current reference lives at [compassadvisory.ai](https://compassadvisory.ai).

READ IT. USE IT. CONNECT YOUR ASSISTANT.

# The current version lives on the website

This e-book is a dated snapshot of Compass's body of work. The website is the living reference; the MCP server lets your AI assistant use that reference while you work through a business problem.

Keep this copy if you prefer to read, annotate or share a document. For the current method, examples and ways to work with Adam, start at the website.

## THE LIVING REFERENCE

[compassadvisory.ai](https://compassadvisory.ai)

## CONNECT YOUR ASSISTANT

[compassadvisory.ai/connect](https://compassadvisory.ai/connect)

MCP endpoint: <https://compassadvisory.ai/mcp>

## Edition note

Prepared from Adam Anderson's CEO working draft, September 2026. The main text is preserved from that draft. References to "this book" describe this reading copy; the website remains the source for the current Compass method.

# Contents

|                                                                      |    |
|----------------------------------------------------------------------|----|
| Author's note                                                        | 4  |
| 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                                      | 12 |
| Chapter 4: Appropriate Maturity                                      | 14 |
| Chapter 5: Score Opportunities Without Pretending the Score Is Truth | 16 |
| Chapter 6: Adoption Is Part of the Product                           | 19 |
| Chapter 7: The Smallest Useful Experiment                            | 21 |
| Chapter 8: Stop, Continue, Invest, or Industrialize                  | 24 |
| Chapter 9: Business Outcomes Are the Portfolio                       | 26 |
| Chapter 10: Governance That Can Be Used                              | 28 |
| Chapter 11: What the CEO Owns                                        | 30 |
| Chapter 12: Begin Monday                                             | 31 |
| Appendix A: Working Templates                                        | 34 |
| Appendix B: Questions for the Room                                   | 37 |
| Closing                                                              | 39 |

# Author's note

The hard part of AI is not access. It is judgment.

It is not a catalog of AI tools. Any catalog I could give you would begin aging before you finished reading it. It is not a prediction about which model, vendor, or interface will win. And it is not an argument that your company should become “AI-first.”

Your company should be business-first.

AI belongs where it can improve an outcome that matters, under conditions your people can actually support, at a cost and risk the business can defend. Sometimes that will mean a sophisticated system. Sometimes it will mean a modest tool used by three people. Sometimes the correct answer will be: not here, not yet, or not at all.

The goal is not more AI. The goal is better business decisions about where AI belongs - and the discipline to leave it out when it does not.

This book gives you a way to make those decisions. You will map the business before discussing technology. You will choose one costly friction instead of chasing a transformation slogan. You will run the smallest useful experiment that can buy real evidence. Then you will stop, continue, invest, or industrialize based on what actually happened.

That sequence is deliberately restrained. Restraint lets you move faster without becoming reckless.

## How to use this book

Bring one real problem with you. Not an AI ambition. Not a list of tools. Choose work that is already costing the company time, money, customers, speed, or avoidable risk.

Each part narrows the decision. By the end, you should have a named friction, a view of the business capability around it, a scored opportunity, a one-page experiment, and a date on which someone will decide what happens next. The appendices contain the working pages. Use them in the room; they are not homework to admire later.

# Introduction: The Question Is Too Big

“How should we use AI?” sounds responsible. It is usually too large to answer well.

The question invites a tour of possibilities. Someone lists tools, demos automations, describes what competitors are doing, and produces a roadmap. The room feels productive because it contains many ideas. But the ideas are rarely attached to an owner, a number, an adoption condition, or a decision date. The list expands. The decision does not.

Ask a smaller question:

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

That question begins with the business as it exists. It asks what is slow, expensive, inconsistent, risky, frustrating, or needlessly difficult today. It asks who feels the problem, what it costs, what has already been tried, and what evidence would justify doing more.

It also allows “no” to be a successful answer.

This matters because AI creates a peculiar kind of pressure. The technology is broad enough to touch almost every activity in a company. That does not mean every activity deserves attention. When everything is possible, prioritization becomes more valuable than imagination.

The executive's job is not to keep pace with every capability. It is to allocate attention and capital. That means deciding:

- which business problems deserve investigation;
- which possible uses are ready enough to test;
- which experiments are cheap enough to learn from;
- which results are strong enough to fund;

- which projects should stop;
- and which capabilities should remain deliberately ordinary.

The method has four movements:

1. **Map the business.** See where value is created, delayed, lost, protected, or put at risk.
2. **Choose the strongest opportunity.** Select one costly friction using impact, likelihood, adoption friction, time to evidence, cost, and risk.
3. **Buy the evidence cheaply.** Run a small experiment in real work, with a real owner and a decision rule agreed in advance.
4. **Manage what follows.** Stop, continue, invest, or industrialize. Then manage the set of decisions as a portfolio of business outcomes.

You can use this method without becoming an AI expert. You do need to know your business, invite uncomfortable facts into the room, and refuse to confuse activity with progress.

That is how to do all of this: not all at once, and not everywhere. One justified decision at a time.

## A case we will follow

Throughout the book, we will return to Musou Movement, a coaching platform developed from March through September 2026. It began with three jobs: manage a community, deepen engagement, and give each member a path shaped by Sean's teaching. Its value as a case is not that the plan was perfect. It is that coaches and clients kept showing the team what mattered next.

The case is condensed from project history, approved plans, implementation records, the current application structure, and direct coach and client feedback. It does not stand as proof that every company should build software or follow the same sequence. It shows how business-first decisions become clearer when actual use is allowed to challenge the plan.

**PART ONE: BEGIN WITH THE BUSINESS**

# Chapter 1: Stop Asking for an AI Strategy

An AI strategy built apart from business strategy is usually a list of technology intentions.

It may include sensible elements: approved tools, data controls, training, vendors, automation targets, and a sequence of projects. Those things matter. But they are not the starting point. A company does not exist to deploy tools. It exists to create value, serve customers, manage cost, protect trust, and remain capable of acting tomorrow.

Start with those obligations.

## **The technology-first trap**

Technology-first work often follows a familiar pattern:

1. A leader sees a compelling demonstration.
2. The company asks where the tool could be used.
3. Teams generate a long list of possibilities.
4. A few pilots begin because they are visible or easy.
5. Adoption varies, the measures are vague, and ownership moves around.
6. The organization cannot say whether the work created value.

The mistake occurs in step two. “Where can we use this?” makes the tool the fixed point and the business the variable.

Reverse it.

Ask what the business needs to improve, what is preventing that improvement, and what kind of capability could remove the constraint. AI may be part of the answer. It may be the wrong answer. You should be equally willing to discover either.

## A strategy is a pattern of choices

Strategy is visible in what you fund, what you refuse, what you measure, and what you stop. A useful AI strategy therefore cannot be a statement that AI is important. It must describe a pattern of choices:

- We begin with measurable business friction.
- We work one opportunity at a time until our decision system is reliable.
- We require an owner and users before we require a vendor.
- We run small experiments before making large commitments.
- We separate the decision to act from the decision to buy, assemble, or build.
- We treat adoption as part of the design.
- We stop work when the evidence does not earn further investment.
- We govern the portfolio in business terms.

That is enough strategy to begin. It tells the company what earns attention before any tool earns a budget.

## The first executive exercise

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 \_\_\_\_\_.

If the leadership team cannot complete them, you are not ready to select technology. You are ready to investigate the business.



# Chapter 2: Find the Costly Friction

A friction is a point where useful work becomes slower, more expensive, less reliable, harder to scale, or more dangerous than it should be.

Friction is not limited to broken processes. It can hide inside a successful company for years. People compensate for it with extra meetings, personal spreadsheets, heroic employees, duplicated checks, memory, apology, and work performed after hours. The result arrives, so the system appears healthy. The cost is absorbed by people.

## Five forms of cost

Look for friction in five forms:

**Money.** Revenue not captured, margin surrendered, avoidable labor, rework, leakage, or excess inventory.

**Time.** Waiting, searching, transferring, reconciling, rewriting, chasing approval, or rebuilding knowledge.

**Customer experience.** Slow response, inconsistent answers, preventable error, weak personalization, or a handoff customers must manage themselves.

**Speed.** Decisions or actions that arrive after their value has declined.

**Risk.** Uncontrolled data, undocumented judgment, unreviewed output, fragile dependencies, compliance exposure, or decisions no one can explain.

A problem can touch several forms at once. Do not count them twice. The point is to see the full consequence, not inflate the case.

## Friction before solution

Write the friction in current-state language:

*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.*

A weak statement says: “We need an AI assistant for sales.”

A stronger statement says: “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 stronger statement does not assume the answer.

## **Put a range on it**

False precision is not required. A defensible range is.

Use available records, direct observation, and the estimates of the people doing the work. Separate what you know from what you assume. If the annual cost might be between \$40,000 and \$120,000, write the range and its assumptions. Do not turn the midpoint into a fact.

The range has two purposes. It prevents a small inconvenience from becoming a strategic program, and it places an upper boundary on what a solution could rationally cost.

## **Choose one**

The first pass through the method should select one friction. This is not because the company has only one problem. It is because a decision system becomes trustworthy through use.

One opportunity lets the team learn whether it can define value, score honestly, design a real experiment, observe adoption, and stop when needed. After that, the organization can increase the number of simultaneous bets without multiplying confusion.

## **Musou: the broad vision met a specific friction**

Musou's first product vision was intentionally broad: community, engagement, and an individual journey. The early application proved that members could sign in, assess themselves, explore movement content, track progress, participate in a community, and receive coaching plans.

Use then exposed a more precise problem. The social experience had become too complicated, and members did not have an obvious way to begin a routine. The important friction was not “we need more community features.” It was decision overload at the point where a member needed a clear next action.

The response was subtraction. The broad social network became a simpler coach-led feed. The movement experience became a three-step flow: choose a category, answer a few questions about the current situation, and receive a filtered playlist. Musou reduced decisions for members while increasing useful control for coaches.

A second friction sat inside delivery. Sean's one-to-one coaching workflow had been tracked in Google Sheets. Moving sessions, movements, nutrition, media, and communication into the platform reduced the administrative work surrounding each client. The important result was not tidier administration by itself. More effective delivery and lower overhead created capacity for Sean to take on more paying coaching clients.

That correction is the method in miniature. Start with the behavior that is failing, not with attachment to the feature list.

## PART TWO: MAP BEFORE YOU MOVE

# Chapter 3: The 5x5 Business Map

Most AI inventories are organized by department or by tool. Both views are incomplete.

Departments hide work that crosses boundaries. Tool inventories tell you what the company owns, not what the business can do. The 5x5 map creates a shared view of business systems and the capabilities supporting them.

The Compass map puts five business systems on one axis and five operating elements on the other. Their intersections create twenty-five cells: small enough to discuss, broad enough to expose patterns.

### Five business systems

The five systems are:

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

Across those systems, examine:

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.

The labels keep the conversation grounded. The map is a view of connected business capabilities, not a shopping list.

## How to run the mapping conversation

Put the map where everyone can see it. For each cell, ask:

- 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?

Record plain observations. Do not diagnose while collecting. If the room immediately starts proposing software, return to the current state.

## Patterns matter more than boxes

The value of a grid is not that every activity fits neatly into one square. It is that a visible friction often crosses several squares.

A slow proposal may look like a sales-process problem. The map may reveal that approved language lives in documents, pricing rules live with one executive, prior examples are difficult to search, and no measure exists for response time. The issue is not one missing tool. It is a capability that crosses people, process, knowledge, workflow, and governance.

That is the difference between automating a task and improving a business. Do not leave the map with twenty-five projects. Leave with a clearer view of the few places worth examining.

# Chapter 4: Appropriate Maturity

Maturity is useful only if it helps you make a decision. It becomes dangerous when it turns into a race toward the highest score.

Higher maturity is not automatically better.

A capability should be as mature as the business needs it to be - no more and no less. A low-frequency, low-risk activity may be perfectly healthy as an individual practice. A high-volume, regulated, customer-facing decision may require an integrated, monitored system. Treating both as candidates for maximum maturity wastes money and may create new risk.

## A six-level ladder

Use this operational ladder:

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.

The labels describe current operational reality, not ambition. A pilot does not move a process to level five. A vendor promise does not count. Score what people can depend on today.

## Current, appropriate, and gap

For each relevant cell, record three things:

- current maturity;
- appropriate maturity for the business now;
- evidence supporting both judgments.

The gap is not automatically a project. It is a question. A current level of one and an appropriate level of three may deserve attention. A current level of two and an appropriate level of two is not a failure. It is alignment.

## **Maturity is not sourcing**

Do not confuse what a capability must do with how you obtain it.

After the target is clear, the company may **buy** a product, **assemble** existing services and internal workflow, or **build** something distinctive. Sourcing is a separate decision. A strategically important capability is not automatically custom software. A commodity capability is not automatically safe to outsource without governance.

First define the outcome and required maturity. Then select the source.

## **Musou: appropriate maturity depends on the job**

Musou did not need every capability to mature at the same rate. It began as a mobile web application rather than a native app, reducing cost and complexity while allowing real-user testing. That was an appropriate early choice for the job: prove the member and coaching workflows.

Other capabilities demanded more maturity as their consequences grew. A small movement library could begin with hand-entered content. Once hundreds of rows, duplicates, inconsistent names, and exposed video links appeared, “content display” was no longer the job. The business needed a maintainable content system. Import tools, update-and-insert behavior, name normalization, multi-category support, private storage, and temporary viewing links followed.

The target moved because the operational requirement moved. The company did not pursue maximum maturity everywhere. It increased maturity where scale, privacy, and repeatability made the current level inadequate.

**PART THREE: CHOOSE THE BET**

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

A score makes assumptions visible. It does not make them true.

Score each candidate opportunity from zero to five on five 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.

Write one sentence of evidence beside every score. A number without a reason is decoration.

## Two hard gates

If business impact is zero, stop. A feasible project with no meaningful outcome is still a bad investment.

If adoption friction scores zero, stop. A technically successful system that people will not or cannot use is not a successful system.

These gates protect the organization from two common temptations: building something impressive that does not matter, and building something valuable in theory that cannot survive contact with work.



## The scoring conversation

Have the process owner, representative users, an economic owner, and someone responsible for risk participate. Let users speak before the most senior leader. Ask for disagreement rather than averaging it away.

If finance scores impact at five and operations scores it at two, the gap is information. If leadership scores adoption at five and users score it at one, the gap may be the entire project.

## Four portfolio positions

The current Compass scoring model produces four practical positions:

- **Run - this week.** Important, plausible, adoptable, fast to learn from, and contained.
- **Run - two weeks.** Promising, with preparation required before a valid experiment can begin.
- **Invest after evidence.** Meaningful but too uncertain or consequential for an immediate commitment; design evidence first.
- **Not yet.** Conditions are not sufficient. Record what must become true before reconsidering.

Do not let the labels create false precision. Their job is to force a decision about attention: now, after preparation, after more evidence, or not yet.

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

When Musou introduced live workshops, it could technically broadcast video while members watched and used text chat. Real use showed that “go live” was not the business job. Coaches needed to schedule an event, manage attendees, communicate, control replay access, share a public link, and see ticket and revenue information.

The stronger opportunity was therefore workshop operations, not a more elaborate broadcast experience. Monetization also revealed that payment was only one piece of the capability. Tickets, subscriptions, discounts, complimentary access, entitlements, and payment status all had to agree about who could enter.

Sean began offering one-to-many online workshops at \$35 per workshop. This created a meaningful new revenue and cash-flow opportunity and became his first major “aha” moment about the platform. A capability that had begun as community infrastructure could now support a repeatable paid offer without requiring Sean to deliver every customer interaction one-to-one.

The same discipline simplified acquisition. A visitor who wanted one workshop was initially asked to understand too much of the membership platform. Coach feedback identified signup friction as a sales problem. Musou introduced a public workshop page and an email-first guest purchase flow, allowing the buyer to establish fuller account access later.

Viewed through the scorecard, these moves combined visible business impact with fast evidence and a clear adoption path. They addressed what coaches and buyers were already trying to do.

The sequence also clarifies “impact.” Workshop operations opened a direct revenue path. Bringing the platform into one-to-one coaching improved delivery and allowed more paying clients to be served. Replacing Google Sheets reduced administrative overhead, but the value was the capacity released for growth. Cost reduction was not the destination.

# Chapter 6: Adoption Is Part of the Product

Adoption is often postponed until after the system is built. By then, the decisions that made adoption easy or impossible have already been made.

People do not adopt a capability because it is innovative. They adopt it because it fits the work, solves a felt problem, preserves necessary judgment, and does not impose a hidden tax larger than the benefit.

## Study the workaround

Workarounds are evidence. A personal spreadsheet, saved prompt, private checklist, or unofficial database tells you that someone has already found a gap worth closing. It may also reveal why the approved system failed.

Do not begin by punishing the workaround. Ask:

- 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?

The answer may be a better product requirement than anything produced in a conference room.

## Name the behavior change

Every experiment should state the behavior required for value to appear.

“Deploy a knowledge assistant” is an implementation statement.

“Customer-support agents use the approved assistant before escalating a policy question, and they correct or flag answers that lack a source” is a behavior statement.

The second can be observed. It also exposes the work required: access, trust, response quality, source visibility, correction, and escalation.

### **Protect human signatures**

Some outputs should never leave the company without a person's name attached. Define the list before testing.

Common candidates include prices, legal commitments, employment decisions, regulated communications, safety decisions, sensitive customer responses, and statements that can materially alter trust.

The list should be short enough to remember and specific enough to enforce. “Use AI responsibly” is not a control.

**PART FOUR: BUY EVIDENCE**

# Chapter 7: The Smallest Useful Experiment

A proof of concept shows that something can be built. A useful experiment tests whether it improves real work under real conditions.

The experiment is not a miniature version of the eventual system. It is the cheapest honest way to reduce the uncertainty blocking the next decision.

## **The experiment card**

Every experiment should fit on one page:

**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.

If the team cannot complete the card, it is not ready to experiment.

## **Real work, small boundary**

Use real inputs, representative users, and a live workflow whenever risk allows. A demonstration with invented data answers whether a demo can be performed. It rarely answers whether the organization can create value.

Keep the boundary small. One product line, one queue, one document type, one customer segment, one internal team, or one recurring decision can be enough.

The right size is not the smallest technical build. It is the smallest test capable of producing credible business evidence.

## **Agree before seeing the result**

Set the evidence standard before the experiment begins. Otherwise enthusiasm and sunk cost will move the goalposts.

The standard can be a range. It should account for quality, adoption, and risk, not only speed. A process that becomes faster while producing more expensive mistakes did not improve.

Write what happens under each result:

- 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.

## **Evidence has a price**

You are not trying to make experimentation free. You are deciding how much uncertainty is worth buying down before committing more.

A two-week test that prevents a year-long mistake can be valuable even when the tested idea fails. Failure is expensive when it arrives late and political. It is useful when it arrives early and is attached to a decision.

## **Musou: real use purchased the evidence**

Musou's strongest product decisions came from concrete attempts rather than speculative requirements. Coaches and clients used the application on real devices, during real coaching sessions and workshop workflows, with real media and purchases. Their reports described the exact moment trust broke: a routine was hard to start, a signup flow blocked a buyer, a replay did not appear, a long recording saved only a few seconds, a microphone failed after being unplugged and reconnected, or a client could not scrub through a video.

Each report reduced a different uncertainty. The work then moved beneath the surface symptom: storage inspection, upload limits, media metadata, device-track lifecycle, permission boundaries, and recording-state logs.

That is what buying evidence looks like. “Users liked the demo” is not the finding. The finding is what happened when a bounded capability met real work closely enough to expose the next decision.

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

The decision meeting is part of the experiment. Put it on the calendar before anyone builds.

Do not let the meeting become a demonstration. Start with the card and the evidence.

## **Stop**

Stop when the outcome is too small, adoption does not occur, guardrails cannot be maintained, cost exceeds the plausible value, or the experiment fails to reduce the important uncertainty.

Stopping is not an accusation. State the expected range, the observed result, and the reason. Thank the owner and users for producing evidence. Record the re-entry condition if one exists.

## **Continue**

Continue only when the next test answers a specific unresolved question. “We need more time” is not enough. Name what remains unknown and why another bounded experiment is the cheapest way to learn it.

## **Invest**

Invest when the evidence supports a larger commitment but the capability still requires implementation, integration, training, process change, or controls. The investment case should inherit the experiment's outcome measure and guardrails.

Do not replace the business measure with deployment milestones. Configuration completed, licenses purchased, and users trained describe activity. They do not prove value.



## Industrialize

Industrialize when the capability works, adoption is repeatable, economics hold, controls are dependable, and broader use is justified.

Industrialization may involve stronger infrastructure, integration, monitoring, support, documentation, and governance. It should not be automatic merely because a pilot succeeded. A small workflow can remain small if that is where it creates the most value.

### **Musou: continue, invest, and industrialize were different decisions**

Musou's recording workflow makes the distinction visible. The first implementation showed that a coach could record from a laptop and place video into a client's session. Feedback then established that coaches wanted the same polished workflow used for workshops. Continuing meant reusing a shared recording studio across workshops, private sessions, and permanent member routines.

Real workshop failures justified a larger investment. Browser recording was vulnerable to long files, lost connections, closed tabs, and device changes. The response was layered: 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 a physical device could be replaced without necessarily terminating the recording.

That was industrialization of a high-consequence capability. A lost hour-long workshop could erase irreplaceable work and damage customer trust. Reliability, recovery, and visibility into failure became part of the product, not cleanup deferred until later.

**PART FIVE: MANAGE A PORTFOLIO, NOT A PARADE**

# Chapter 9: Business Outcomes Are the Portfolio

Companies often build an “AI portfolio” made of projects, tools, and vendors. That view rewards starts. A business-outcome portfolio rewards evidence, decisions, and results.

Each portfolio item should describe:

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

## **Attention is the limiting resource**

Capital is visible. Executive attention, user patience, implementation capacity, and appetite for change are easier to overspend.

Set a work-in-progress limit. A company capable of funding twenty experiments may be capable of learning honestly from only three. Start from the number of owners and user groups who can sustain real participation.

## **Keep a real “not now” list**

Ideas that do not enter the portfolio should not disappear into a backlog with no rules. Record the re-entry condition:

- 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.

Review the condition, not the age of the idea. Time passing does not make a weak opportunity stronger.

## **Retire completed work**

A portfolio that only grows is hiding something. Remove completed and stopped work from the active view, but preserve the final evidence and decision. The organizational memory matters more than the badge that says “AI initiative.”

# Chapter 10: Governance That Can Be Used

Governance should make good action easier and dangerous action harder. It should not be a binder written for an imaginary audit while employees invent the working rules in private.

Begin with three one-page controls.

## **Approved-use page**

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.

Write in language a manager can use during work.

## **Human-signature page**

List outputs and actions that require a named human review. Define what review means. A checkbox clicked by someone who cannot see the source is not meaningful oversight.

## **Gamed-inputs check**

Metrics change behavior. Each quarter, choose one result that improved unusually quickly and trace it back to the inputs. Confirm that the measure still means what it meant when the decision boundary was set.

This is not an accusation of cheating. People optimize what the organization rewards. A routine check protects the meaning of the number.

## The quarterly review

Use the same agenda each quarter:

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?**

The final question is the proof that governance is real.

## Musou: reliability and security grew with the relationships

As Musou added private coaching, payments, workshops, media, and messaging, a simple role label was no longer a sufficient control. The relevant questions became relational: 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. 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.

Reliability required the same specificity. Capture logs made recording start, duration, upload, backup, and failure states observable. Local backup was removed only after confirmed upload. When a recording did not exist, the correct response was to say so and improve the safeguards - not imply that a missing asset could be recovered.

Governance became useful because it followed the real consequences of the workflow.

**PART SIX: LEAD THE DECISION SYSTEM**

# Chapter 11: What the CEO Owns

The CEO should not own every AI project. The CEO owns the conditions under which work receives attention, money, legitimacy, and the right to continue.

That means owning five things.

## **The business language**

Keep the conversation anchored in outcomes, friction, evidence, and choices. Do not allow urgency about AI to exempt a project from ordinary business discipline.

## **The right room**

Include the economic owner, process owner, representative users, relevant technology capability, and risk judgment. Do not let the most senior person answer every question first.

## **Permission to tell the truth**

People must be able to say that the data is not ready, the process is poorly understood, the tool is not being used, the expected value is absent, or the leader's favorite idea should stop.

If those statements are punished, the portfolio will still produce reports. It will stop producing evidence.

## **Clear decision rights**

Name who can start an experiment, who approves data use, who accepts risk, who funds investment, and who can stop work. Ambiguity makes every disappointing project politically permanent.

## **The kill question**

Ask what should stop before the portfolio becomes crowded. Ask it routinely, not only during a budget crisis. The purpose is not austerity. It is to return attention to opportunities that have earned it.

# Chapter 12: Begin Monday

You do not need a transformation office to start. You need one costly friction, the right people, and a calendar.

## **Day 1: Name the friction**

Choose a current problem with a visible owner and plausible economic consequence. Write the friction statement. List what is known and assumed.

## **Days 2-3: Map the capability**

Use the relevant cells of the 5x5 map. Observe the work. Speak with users. Identify where information, judgment, process, workflow, and governance break or compensate for one another.

## **Day 4: Set current and appropriate maturity**

Score the capability as it operates today. Decide what maturity the business actually needs. Keep sourcing out of the conversation.

## **Day 5: Score the opportunity**

Use impact, likelihood, adoption friction, time to evidence, and cost/risk. Apply the two hard gates. Decide whether the opportunity runs now, after preparation, after more evidence, or not yet.

## **Week 2: Design the experiment**

Complete the one-page experiment card. Name the owner, users, measure, evidence standard, guardrails, scope, and decision date. Select the smallest real boundary capable of producing credible evidence.

## **Weeks 3-4: Run and observe**

Use real work. Track the primary measure, adoption behavior, error, and guardrail exceptions. Do not expand the test because early results are exciting. Do not hide a disappointing result until the final meeting.

## Decision day

Read the hypothesis and evidence standard first. Present the observed result without adjectives. Decide: stop, continue, invest, or industrialize. Record why.

Then choose the next friction.

## Musou: the end-to-end loop

Musou's development history can be read as a repeated decision loop:

1. Begin with an existing workflow: 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 rather than patching 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 survived, but the roadmap changed. Community became simpler. Coaching became more personal. Workshops became an operating and revenue system. Video became a creation, delivery, recovery, and feedback capability. Security moved from broad roles to relationship-specific access.

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

The lesson is not that Musou accumulated features. It is not even that the business became more efficient. The useful chain is more precise:



*Remove a costly friction. Release capacity. Use that capacity to serve and sell in ways that were previously difficult. Measure whether growth follows.*

Efficiency was the lever. Revenue growth was the business reason to pull it.

### **The method in one page**

1. Start with the business, not the tool.
2. Find one costly friction.
3. Map the capability around it.
4. Name current and appropriate maturity.
5. Score impact, plausibility, adoption, evidence speed, cost, and risk.
6. Stop immediately at zero impact or zero adoption viability.
7. Design the smallest useful experiment.
8. Agree on evidence and guardrails before starting.
9. Use real work and representative users.
10. Stop, continue, invest, or industrialize.
11. Manage decisions as a portfolio of outcomes.
12. Ask what should stop.

This is not a one-time AI exercise. It is an operating discipline for making business choices while the technology keeps changing.

The tools will change. The need for judgment will not.

# Appendix A: Working Templates

## A.1 Friction statement

**Trigger:** When \_\_\_\_\_ occurs,

**People and current work:** \_\_\_\_\_ must \_\_\_\_\_,

**Observable consequence:** which causes \_\_\_\_\_.

**Estimated cost or proxy:** We believe this costs \_\_\_\_\_ because \_\_\_\_\_.

**Unknowns to verify:** \_\_\_\_\_.

## A.2 Opportunity scorecard

| Dimension                               | Score 0-5 | Evidence for the score |
|-----------------------------------------|-----------|------------------------|
| 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 zero, stop.

**Portfolio position:** Run - this week / Run - two weeks / Invest after evidence / Not yet

**Reason:** \_\_\_\_\_

## A.3 Experiment card

- Friction:
- Hypothesis:
- Business owner:
- Representative users:
- Scope:
- Exclusions:
- Primary measure:
- Evidence standard:
- Quality check:
- Guardrails:
- Start date:
- Decision date:
- Decision choices:

## A.4 Outcome card

- Business outcome:
- Friction addressed:
- Owner:
- Current maturity:
- Appropriate maturity:
- Experiment or investment status:
- Committed cost:
- Attention required:
- Primary measure and range:
- Adoption measure:
- Next decision date:
- Sourcing choice:
- Current decision:
- Re-entry condition, if deferred:

## A.5 Decision record

**Expected:** \_\_\_\_\_

**Observed:** \_\_\_\_\_

**Decision:** Stop / Continue / Invest / Industrialize

**Why:** \_\_\_\_\_

**Next action and owner:** \_\_\_\_\_

**Next decision date:** \_\_\_\_\_

# Appendix B: Questions for the Room

## To find friction

- 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?

## To test importance

- 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?

## To test adoption

- 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?

## **To design evidence**

- 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?

## **To govern the portfolio**

- 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?

# Closing

The pressure to “do something with AI” will keep changing shape. New tools will make old demonstrations ordinary. Capabilities that look expensive today will become cheap. Risks that seem theoretical will become familiar. Some early moves will work. Others will not.

You cannot eliminate that uncertainty with a longer roadmap.

You can build a company that makes better decisions inside that uncertainty.

Map the business. Choose one costly friction. Buy evidence cheaply. Manage what follows. Keep the maturity appropriate to the outcome. Treat adoption as part of the product. Separate priority from sourcing. Fund what earns the right to expand. Stop what does not.

Then do it again.

That is how to do all of this.