

AI Engagement Scoping Guide

How to scope an AI engagement so it actually ships, process map first, tool selection second, vendor pitch third or never.

FOR

CEOs, COOs, GMs, PE deal teams

FROM

Gordon Mickel • mickel.tech

PRACTICE

AI systems & transformation • independent
practice

USE

Decide whether an AI initiative is real or
premature

The pattern is not technical. It is procedural.

AI initiatives that fail rarely fail because the model was wrong. They fail because the engagement was scoped against the wrong question, usually a tool or a vendor, instead of against a workflow, a constraint and a measurable outcome.

Across operations, support, finance, sales and engineering, the four failure patterns below show up again and again. If your current AI initiative looks like one of them, the highest-ROI intervention is not a faster model, it is a re-scope.

The vendor pitch

A vendor demo is gorgeous. The platform is real. You buy. Six months later there is a tab in the application no one opens. Pattern: tool was selected before the workflow was mapped.

The strategy paper

A consulting deck arrives with a 28-page AI strategy and a roadmap. It is well-written. Nothing ships from it. Pattern: strategy detached from process; no system actually instrumented.

The bottoms-up sprawl

Engineering picks up Cursor, marketing picks up a copywriting tool, sales picks up an SDR agent. Twelve subscriptions, no governance, no shared data. Pattern: no central capability function.

The board pressure pilot

"We need to be doing something with AI by Q2." Pilot stood up to satisfy the board. KPIs unclear. Outcome: a thing exists; nobody can say if it worked. Pattern: governance ahead of clarity on what is being governed.

The right first question is never "which AI tool should we buy." It is "which workflow is most expensive, and which part of it is the bottleneck."

Map the workflow before the tool.

The cheapest, fastest way to surface real AI opportunities is to map the end-to-end workflow on a wall before any vendor enters the conversation. The map answers four questions in twenty minutes that no demo can answer in a week.

The four questions

Where does the data live? Which systems hold the inputs, where are they stale, who owns the writes. Answers determine whether an AI use case is even feasible inside the next two quarters.

Where are the bottlenecks? The workflow has a slowest step, and that step is rarely where attention is currently going. AI work is highest-value when it dissolves the bottleneck, not when it accelerates a step that was already fast.

What breaks at scale? If volume doubled tomorrow, which step is the first to crack? The brittleness map drives prioritisation.

Where is the human glue? Every workflow has humans manually stitching together tools that should be connected. That glue is usually the strongest signal of where automation pays back fastest.

Output of a process map

A workflow diagram (not a strategy deck) annotated with the three highest-value automation candidates, the data dependencies of each, the rough effort estimate, and the measurable KPI each would move. This is the input to scoping the actual engagement.

Workflows worth scoring first

The domains where the map most often surfaces a payback candidate: finance operations (AP/AR triage, reconciliation, close readiness), revenue and GTM operations (CRM hygiene, forecast assembly, proposal drafting), procurement and vendor management, customer service and intake, HR and admin, compliance and reporting. Yours may not be on this list; that is what the mapping is for.

A go / no-go test before any AI engagement.

Use these four selection gates before committing budget. An engagement that fails any one of them is not ready -- irrespective of how exciting the opportunity looks.

GATE	NAME	WHAT PASS LOOKS LIKE
D1	KPI clarity	Named, baselined, measurable KPIs the engagement is supposed to move. Examples: deflection rate, FCR, MTTR, gross profit per minute, NRR, ARPU, cycle time. If you cannot name the KPI in one sentence, defer the engagement.
D2	Reversible path	A defined rollout sequence: shadow → partial → full, with rollback plans at each step. AI work that cannot be rolled back without engineering heroics is high-risk; design the reversibility before the build.
D3	Data and stack access	APIs, keys, service accounts available within two weeks, or credible stubs that mirror real shape. Engagements that wait six weeks for a data pull lose momentum and shed sponsors.
D4	Sponsor and budget	A named executive sponsor with budget authority, not a project champion three layers down. KPI ownership at the same level. Without this, the engagement does not survive its first reprioritisation.

Three caps that override good intentions

If the workflow is not stable (procedure changes monthly), defer AI work until it is. Automating a moving target burns budget without compounding.

If governance is below 2/5 (no audit trail, no residency commitments, no banned-intent ledger), fix that first. AI in production without governance is reputational debt accruing in the background.

If the cycle time of the underlying organisation is quarterly, AI cannot fix that, it will only make the quarterly cycle produce more output. Address the cadence first or pair the AI work with a delivery-cadence change.

Four shapes engagements take.

Most engagements fit one of four shapes. Picking the right shape early prevents the most common scoping failure: paying for a build before the audit, or for an audit when a day would have answered it.

01	Strategy Session 1 day, on-site or remote • your side: 4-8 hours Half- to full-day working session on one specific question, vendor selection, build-vs-buy, architecture review, "is this the right AI use case." Written notes and a clear next-step recommendation by end of day. Fee credited against any follow-on engagement within 90 days.	CHF 3,000-5,000
02	Process Map + AI Opportunity Audit 2-3 weeks fixed • your side: under 20 hours total Map workflows end-to-end. Score automation candidates by impact and feasibility. Propose 3-5 prioritised initiatives with ROI estimates and a 90-day plan. Board-ready report. The lowest-risk way to find out where AI actually pays for itself in your specific operations.	CHF 15,000-25,000
03	AI System Build 4-10 weeks fixed • your side: sponsor + 2-4 h/week Production deployment of one AI system class: operations agents, knowledge platform with verifiable retrieval, voice, sovereign or private LLM. Requirements through go-live with operations handover, then a 30/60/90-day hardening window. Built for the day after the demo, not the demo.	CHF 40,000-150,000
04	Fractional AI Lead 1-2 days/week, monthly retainer • your side: 1 steering h/week Programme governance, vendor selection, hands-on architecture, internal capability building. For organisations that need senior AI leadership without making the full hire. Designed to make the team self-sufficient, not dependent; engagements end when the team owns the discipline.	Quarterly minimum

Which shape, when.

A short decision tree. Read top-down. Stop at the first branch that fits.

Q1. Do you have one specific question (vendor selection, build-vs-buy, "is this the right use case") that you can articulate in one sentence?

→ **YES:** book a Strategy Session. Get a defensible answer in a day.

→ **NO:** continue.

Q2. Do you know *which* workflow you want AI to address, with a baseline KPI you would expect it to move?

→ **NO:** book a Process Map + Audit. Two to three weeks. Identifies which 3-5 workflows actually move P&L. Output drives every subsequent decision.

→ **YES:** continue.

Q3. Have all four selection gates (D1-D4 on page 3) been passed for the workflow you want to address?

→ **NO:** do not commit budget yet. The Process Map + Audit is still the right starting point, fix the gates as part of it.

→ **YES:** scope an AI System Build. Production deployment of one system class. Fixed-fee.

Q4. Do you need ongoing senior AI leadership for a portfolio of initiatives without making the full executive hire?

→ **YES:** Fractional AI Lead. Quarterly minimum. Designed to make the team self-sufficient, that is the test of whether it worked.

Eight questions for the first call.

A discovery call goes better when both sides have done twenty minutes of preparation. These are the questions I ask on every first call, answers in advance compress the call from discovery to scoping.

What is the workflow you want AI to address? One paragraph. No jargon.

What KPI would you expect AI to move? Name it. State the current baseline. State a target.

Where does the underlying data live? Systems, owners, access posture, latency.

What is the current process at the bottleneck? Who, how often, with what tools, in what time.

What have you tried so far? Tools, vendors, internal pilots. What worked, what did not, why.

Who is the executive sponsor? Title, budget authority, time available, what success looks like to them.

What is the rollback plan if it does not work? How reversible is the deployment.

What does "done" look like in 90 days? Concrete, observable, measurable.

When to walk, both sides.

Some patterns mean the engagement should not proceed yet, or at all. From either side. Honest red flags save both parties three months of expensive disappointment.

From the buyer side

The advisor cannot say no. If the answer to every scoping question is "yes, I can do that," scope is being padded. A good advisor declines work that is not a fit.

The pitch leads with a tool. Tools are implementation, not strategy. If the first conversation is about a model or a platform, the workflow has not been mapped.

No anti-retainer position. Engagements designed to make you self-sufficient are built differently from engagements designed to extend. Ask the question.

No reversibility design. If the rollout has no rollback plan, the work is being scoped optimistically. Pause.

From the advisor side

No named sponsor. If the deal is brokered by someone three layers down with no executive air cover, the engagement does not survive its first reprioritisation.

KPI cannot be named. If neither side can name the KPI in one sentence, the work cannot be measured. Stop.

Workflow is in flux. Procedure changing monthly means automation will be obsolete on arrival.

Adversarial scoping conversation. Negotiation is normal; adversarial framing in scoping is a preview of the engagement. If it is uncomfortable now it will not get easier.

A defensible number for the board.

A useful AI engagement should pay itself back inside four quarters in the worst case. If you cannot rough it on the back of an envelope, the work has not been scoped to a measurable outcome.

The arithmetic, in three lines

Baseline cost of the workflow today: hours × loaded cost × frequency, plus the loss when the bottleneck blocks downstream revenue or capacity.

Realistic recovery after the AI work ships: commonly 30-60% of the baseline cost in the first 90 days, climbing as the system matures and adoption deepens.

Engagement cost divided by recovered run-rate equals payback period. Anything below four quarters in the worst case clears the bar.

Two numbers worth tracking separately

Cost-out: hours and dollars no longer consumed by the old process.

Capacity-up: revenue or output the previously-bottlenecked function can now produce. This is where the biggest second-order effect lives, features that were not on the roadmap because they were not viable at baseline cost.

The engagements that compound save cost and then open product surface that was previously uneconomic. Plan against that second-order effect, not only the first.

Six things, not sixty.

Most AI governance frameworks are written for boards that want to feel responsible. Six controls cover the operationally important ground; the rest is documentation theatre.

Residency policy. Where data is allowed to travel. What stays on-prem or in your VPC. Sovereign inference for regulated workflows.

Audit log. Every AI invocation logged with inputs, outputs, model, prompt version, user, time. Stored for the retention period your industry requires.

Banned-intent ledger. What the system is not allowed to do. Reviewed quarterly.

Eval suites and drift checks. Quality is measured continuously, not assessed at procurement.

Rollback plan with SLO. 15 minutes to partial rollback, 60 minutes to full. Drilled, not assumed.

Vendor and model review cadence. Quarterly. Models drift. Pricing changes. Better options emerge. Revisit on a calendar, not a crisis.

What lands at the end of each engagement.

ENGAGEMENT	WHAT YOU GET
Strategy Session	Written notes from the working session, a defensible answer to your specific question, a one-page next-step recommendation, optional follow-up call within two weeks.
Process Map + Audit	Workflow map (annotated), 3-5 prioritised AI initiatives scored on impact and feasibility, ROI estimate per initiative, 90-day rollout plan, board-ready executive readout.
AI System Build	Production AI system live in your environment. Architecture documentation. Eval suite. Operations handover. Rollback drilled. Internal team trained on the system. A 30/60/90-day hardening window after go-live: eval regressions, decision-log reviews, capability transfer until the team runs it alone.
Fractional AI Lead	Programme governance, vendor / model selection, hands-on architecture, capability building. Monthly steering. Engagement designed to end when the team can run it without you.

What the first six months look like.

MONTH	FOCUS	EXPECTED OUTPUT
0	Discovery + intake	30-min call, fit decision, NDA, scope agreed.
1	Process map + audit	Workflow mapped, opportunities scored, plan agreed.
2	Build kicks off	Architecture set, integration plumbing in place, team onboarded.
3	Build continues	System in shadow mode against real traffic, eval suite running.
4	Rollout	Partial rollout (10-30% of volume), KPIs measured, rollback drilled.
5	Full rollout + handover	Full traffic, operations team owns the system, governance live.
6	Steering or close	Monthly steering for the next quarter, or engagement closes if the team can run it alone.

The shape varies, a Process Map + Audit alone closes at Month 1; a Fractional Lead engagement runs continuously with quarterly review points. The trajectory above describes the typical end-to-end shape when the engagement starts at intake and runs through to a production system in operations.

Three ways to use this guide.

Run the gates and the decision tree internally. If the answer is clear, the guide has done its job.

Bring me in for a Strategy Session on one specific question, vendor selection, build-vs-buy, architecture review. CHF 3-5,000, full-day, written output. Credited against any follow-on engagement.

Engage a Process Map + Audit if there is no single specific question yet, if what you need first is to see the workflow honestly and surface where AI actually pays back. CHF 15-25,000, 2-3 weeks, board-ready output.

ENGAGE

A 30-minute discovery call is the cheapest way in.

On the call we map your highest-friction workflow, talk through what an AI system would actually do for you, and agree whether it is worth a process audit. No vendor pitch. No 18-month programme, no rip-and-replace: systems run on the stack you already have, and your team owns them at handover.

BOOK

cal.com/gmickel/discovery30

EMAIL

gordon@mickel.tech

READ

mickel.tech/ai-transformation

PRACTICE

Binningen, CH · DE / EN