A FIELD GUIDE

Ready Enough

What to look at, in what order, before AI is worth the time, money, or attention.

BY DANIEL BERRY

Built by Berry

Contents

Foreword

Most of the operators I talk to are not skeptical of AI. They're skeptical of AI pitches. There's a difference.

The pitches promise transformation. The pitches show a demo that runs cleanly on clean data. The pitches don't mention the part where the demo data was hand-picked, the workflow was simplified for the deck, and the person showing it to you has never had to live with the result.

This book is the opposite of that pitch. It assumes you are competent, that your operation is more complicated than any demo could capture, and that you'd rather know whether AI will work in your business than be told that it will.

The honest answer, for most companies, is: not yet. And not because the technology is missing. Because the operation underneath it is.

The good news is that the work of getting AI-ready is mostly the work of getting operationally ready. The five-layer framework in this guide will help you find the lowest broken layer in your business — the one that's quietly breaking other things, including any AI you'd build on top of it. Fix that, and a lot more than AI starts to work.

Read it through once. Then keep it as a reference. The appendix is designed to be torn out and used at your desk.

— Daniel Berry
Built by Berry

01

Why pilots stall

Walk into ten companies that tried AI in the last two years and you'll see roughly the same story three different ways.

The first story: someone bought a tool. It demoed beautifully. Six months in, three people use it, two stopped after a week, and the rest never logged in. The vendor blames adoption. The team blames the tool. Both are partly right and entirely missing the point.

The second story: someone built a pilot. A small AI workflow, a real use case, a champion who cared about it. The pilot worked. Then the pilot needed to scale. It turned out the workflow it depended on had eighteen variations, four undocumented exceptions, and a person whose tribal knowledge held the whole thing together. The pilot did not scale.

The third story: someone has ChatGPT open on three monitors and treats it as a personal productivity tool. It's useful. It is also invisible to the rest of the business and nothing about how the company runs has changed.

All three are the same failure

The failure isn't the model. The model is fine. In most cases, the model is better than the workflow it was asked to plug into.

The failure is upstream. It's in the data the AI would need and couldn't find. It's in the process that varies too much to describe. It's in the absence of a person whose job it is to make this thing work. It's in the tools that don't talk to each other. It's in the decision the AI was supposed to improve, which turns out to not be a real decision anyone was making.

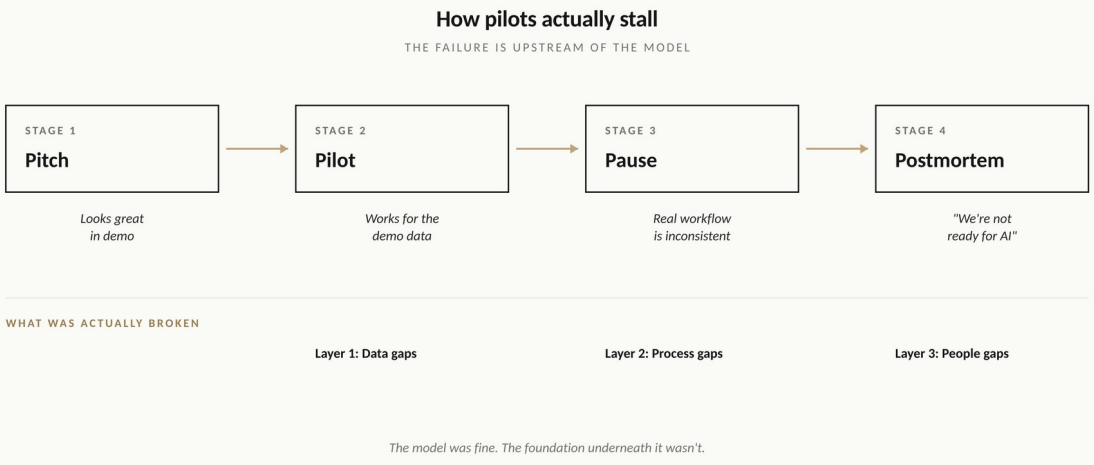

Almost every "AI didn't work for us" story is, on inspection, an operations story wearing AI clothes. The operation wasn't ready. The AI made that visible faster than usual.

> **THE PRINCIPLE**
>
> AI readiness is not a technology question. It's an operations question. The model is the easiest piece. Everything underneath is harder, and the underneath is what determines whether AI works.

Why this is good news

If the bottleneck were the model, you'd be waiting for someone in San Francisco to ship a better one. That's a passive position.

If the bottleneck is your operation, you can do something about it. And the work you do is useful even if you decide AI is not yet the right move. Documented processes, clean data, named owners, real decision points — these are valuable on their own. They make the company run better with or without AI.

That's the orientation of this book. The goal is not to convince you to adopt AI. The goal is to give you a way to look at your operation that tells you, honestly, where you are, and what to fix first. AI works as a side effect of that.

02

The Readiness Stack

The framework is five layers. Each one depends on the ones below it. You can't put a roof on without a foundation, and you can't build the foundation by starting at the roof. This is true of houses, and it is true of how AI plugs into your business.

The five layers split into two tiers. The first three are foundation work — they have to exist before anything else holds weight. The last two are activation — they only matter once the foundation is solid.

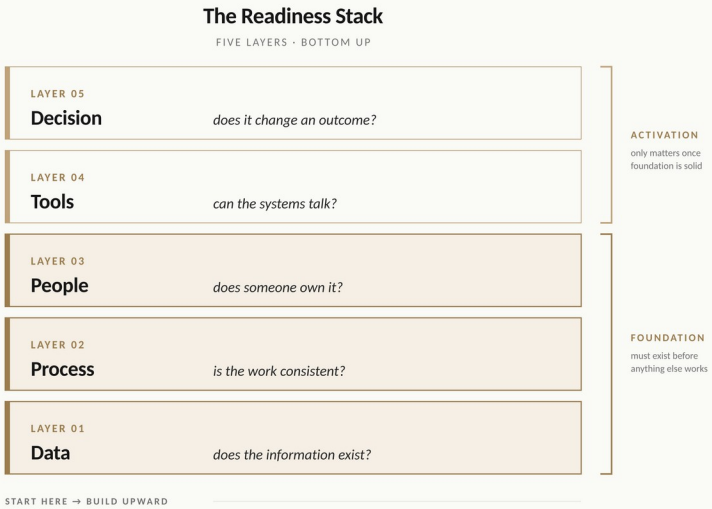

The foundation tier

Three layers. If any one of them is broken, nothing built on top will hold.

- Layer 1 — Data. Does the information AI would need actually exist, somewhere a system can reach? Not in someone's head, not in an email thread, not in a folder on a laptop. Reachable.
- Layer 2 — Process. Is the work consistent enough that you could describe the right answer to a smart outsider? Or is every case its own case?
- Layer 3 — People. Is there one named person with authority to change how this workflow runs? Not a committee, not a vague collective. One person.

The activation tier

Two more layers. They matter, but they only matter after the foundation holds.

- Layer 4 — Tools. Can the systems involved actually exchange information without a human copying and pasting?
- Layer 5 — Decision. Is there a real business decision this AI is supposed to improve? Something that will be made differently because of what it produces?

How one workflow touches all five

Take a workflow that's deceptively boring on purpose: invoice approval at a 50-person services firm. Watch how each of the five layers shows up.

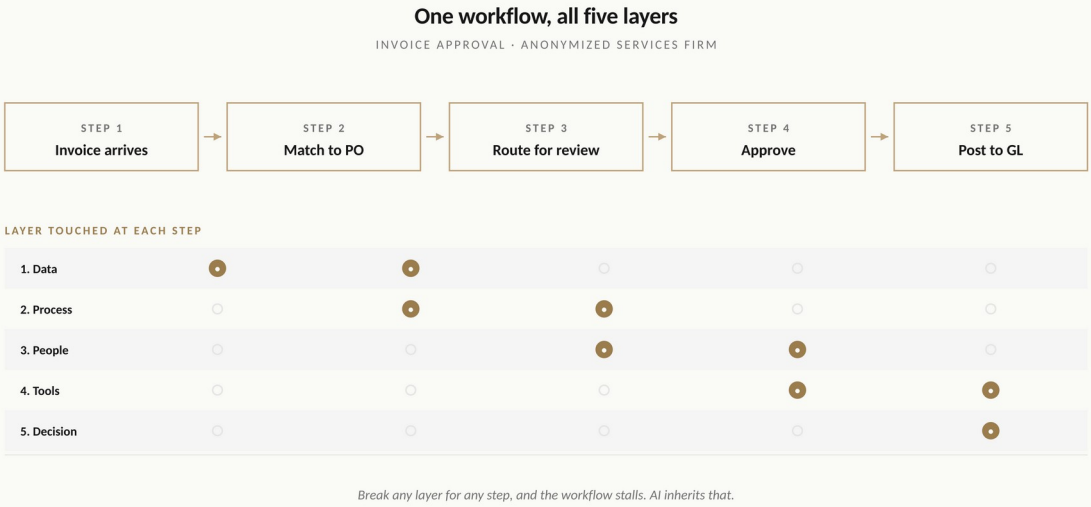

Every step of that workflow depends on at least one layer holding. Most steps depend on two. If your firm can't read the invoice digitally (Layer 1), the AI can't help. If the matching rules vary by client (Layer 2), the AI is guessing. If no one owns the approval step (Layer 3), nothing the AI produces will be acted on. And on up the stack.

> **THE PRINCIPLE**
>
> Find the lowest broken layer first. That's where the work is. Everything above it is wasted effort until that foundation holds.

The rest of this book walks each layer in order. By the end, you'll be able to look at any workflow in your business and ask: which layer breaks first? That's the layer you fix.

03

Data

Does the information exist, in a place a system can reach?

Every AI workflow starts with information. The model has to read something to produce something. That something is your data — and "your data" is a phrase that hides a lot of trouble.

When operators say their company has the data, they usually mean: someone in the company knows the answer. That's a different claim than the one AI needs.

Three ways data quietly fails

In practice, data unreadiness shows up in three patterns. None of them are exotic. All of them are common.

It lives in someone's head.

The most accurate version of how things work is tribal. Sarah knows which vendors give terms. Marcus remembers which clients need approvals routed differently. When Sarah and Marcus are out, the workflow slows down — and AI can't ask them.

It lives in attachments and inboxes.

The information exists, technically. It's in PDFs nested in email threads, in shared drive folders organized by year, in a spreadsheet someone built for one project and everyone now relies on. "Reachable" means a system can pull it out programmatically. A PDF buried six clicks deep in a shared drive isn't reachable.

It lives in tools that don't expose it.

Some software stores your data and gives you no way to get it out. No API, no export, no integration. You can see it on the screen, but nothing else can. The data exists. It is not reachable.

CLAIM	REALITY CHECK
"We have the data — it's all in our system."	Can you produce a CSV, JSON file, or API response containing the relevant fields, for any record, in under five minutes, without asking a person?

If the answer to the reality-check question is no — for any workflow you'd want to apply AI to — Layer 1 is where the work is. Not Layer 5.

What "ready" looks like

- The data lives in one place, or in a small number of places that can be queried together.
- The fields you'd need are consistent — the customer's name isn't "Acme Inc." in one row and "acme" in another.
- A system can read it. There's an API, an export, a database connection — something that doesn't require a person clicking around.
- Someone knows what's in it. Not in the sense of "I built this spreadsheet" — in the sense of being able to define what each field means.

WATCH OUT FOR

"We're a data-driven company" is a culture claim. It does not tell you whether your data is reachable. Audit one specific workflow, end to end, and see what a system would have to do to read what it needs. That's the only honest test.

PAUSE AND CONSIDER

Pick one workflow in your business that you've thought about applying AI to. Before reading on, ask: where does the information that workflow uses actually live? In a system? In an inbox? In a person? You'll know whether Layer 1 holds the moment you try to answer that out loud.

04

Process

Is the work consistent enough to describe?

A workflow and a habit are not the same thing. A workflow is something you could hand to a smart outsider with documentation, and they could perform it. A habit is something everyone in the building knows how to do because they've watched it happen for years.

Most operations run on habits more than workflows. That's fine when the people executing them are humans with context. It is fatal when the executor is an AI.

Why undocumented processes break AI worse than they break humans

A human handed an ambiguous workflow does what humans do: they ask someone, they pattern-match on a similar case from last year, they make a judgment call and remember they made it. The work gets done. Sometimes inconsistently — but done.

An AI handed the same workflow can't ask Marcus. It has to produce an answer based on what's written down. If what's written down doesn't cover a real case, the AI guesses confidently. Confidently is the dangerous word. It will be wrong with the same tone as when it's right.

This is why the same firm that says "our process works fine" finds AI "unreliable." The process was leaning on humans to absorb its inconsistency. AI doesn't absorb. It exposes.

HABIT	WORKFLOW
"We approve invoices over $5K with leadership sign-off. Usually. Sometimes the controller can sign. Depends on the client."	Invoices ≥ $5K from clients in tier A route to the CFO. Invoices ≥ $5K from tier B clients route to the controller. Tier is set in the CRM.

Notice what changed. Nothing about the work changed. What changed is that the rule is now describable, the inputs are named, and a system could execute it. That's the difference between a habit and a workflow.

How to tell if a process is ready

- Could you write the rules on one page? Not the executive summary — the actual rules.
- When exceptions happen, do they get documented and added to the rules, or do they get absorbed into someone's memory?
- If the person who normally runs this workflow left tomorrow, would the workflow still produce correct results within a week?

- Could a new hire follow the documentation and get the right answer without asking?

Four yeses means Layer 2 is solid. Two or fewer means you have a habit, not a workflow. AI built on a habit is theater.

WATCH OUT FOR

The instinct to skip ahead and "document as you go" while the AI is being built. This is the same as starting construction before the architect has drawn the plans. You will end up rebuilding both the documentation and the AI, twice.

PAUSE AND CONSIDER

Think of a workflow in your business that has more than one way it gets handled. Before reading on, ask: is the variation deliberate (different paths for different cases) or accidental (different people do it differently)? Deliberate is fine. Accidental is the layer-2 problem.

05

People

Does one named person have authority to change how this works?

This is the layer most operators skip — usually because, on the surface, the workflow does have people involved. People execute it every day. The trap is that "people are involved" is not the same as "someone owns it."

Ownership means: one person can decide that the workflow will change, and that change will stick. Not propose, not suggest, not raise in a meeting. Decide. If that person doesn't exist, your AI rollout has no place to land.

Diffuse vs. owned

Picture two versions of the same workflow.

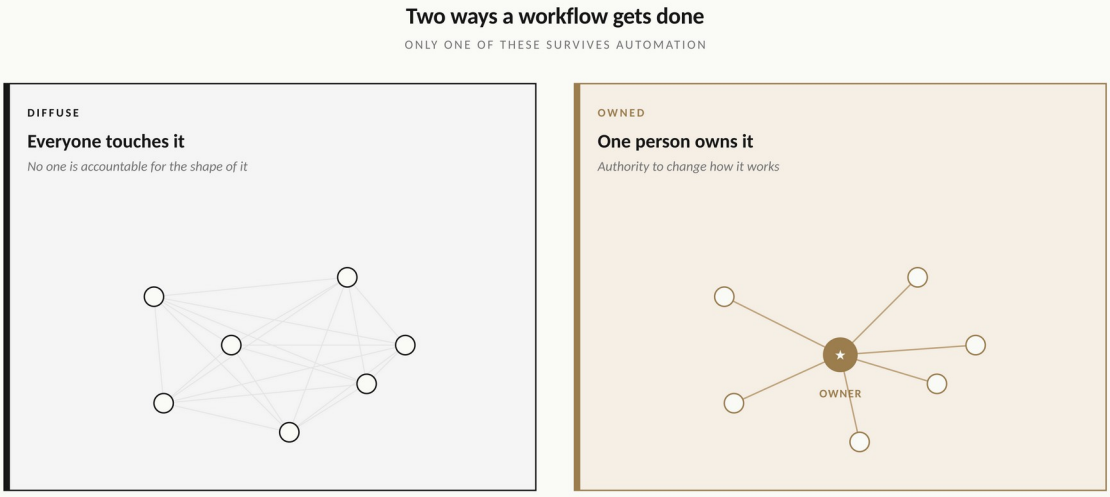

On the left: everyone touches the workflow. No single person is accountable for the shape of it. When something changes — a tool, an exception, a new client type — the change happens in fragments, in different people's heads, at different times. The workflow drifts.

On the right: one person owns it. They may not execute every step. But they decide how the workflow runs, they're the one who updates it when reality changes, and when an AI tool is introduced, they're the one who makes it stick or fail.

AI rollouts go to the right. Diffuse workflows reject them — not loudly, but silently. The tool gets installed, the training happens, the rollout calls itself successful, and three months later no one is using it. There's no one whose job was to make sure they did.

The signs you don't have an owner

- When you ask "who owns this process?" the answer is a department, not a person.
- Changes to the workflow happen by accident — someone tries a new approach, others copy it eventually, and after a year the workflow has shifted without anyone noticing.
- Exceptions get handled by whoever's nearest, and the answer depends on who that was.
- If you wanted to change how the workflow runs starting Monday, you couldn't say who to tell.

> **WATCH OUT FOR**
>
> The committee. "Operations and Finance jointly own this." That sentence means no one owns it. Two-person ownership is fine if one is the lead. Genuine co-ownership of an operational workflow is a recipe for stalemate.

What "ready" looks like

- One named person owns the workflow.
- They have authority — not just responsibility — to change it.
- Other people executing the workflow know who the owner is.
- When AI gets introduced, the owner is the one who decides whether it stays.

> **PAUSE AND CONSIDER**
>
> *Pick a workflow you'd want AI to support. Before reading on, name the person — by name, not by title — who would decide that the AI version stays or goes after 90 days. If you can't name them, Layer 3 is the layer.*

06

Tools

Can the systems involved exchange information?

The activation tier starts here. If the foundation holds — data is reachable, processes are described, owners exist — the tools layer is where the work becomes concrete.

Tools readiness is a simple question with a complicated answer: can the systems involved in this workflow exchange data without a human in the middle copying and pasting? The simple part is the question. The complicated part is that most companies discover, when they look, that the answer is no — and they've been quietly absorbing that cost for years in the form of human copy-paste.

What "the systems can talk" means in practice

- The CRM has an API, or at minimum a clean export.
- The accounting system can receive data programmatically, or has a webhook to send it.
- The document storage can be queried — files are findable by content, not just by where someone filed them.
- Custom internal tools have endpoints, not just user interfaces.

None of these are exotic. They're the baseline for any business that wants to automate anything. AI included.

The copy-paste tax

Watch any operations team for a day. Count the times someone copies data from one tool, pastes it into another, and edits it slightly to fit. That number is the copy-paste tax. It is paid in salary, hours, and accumulated errors. Every minute of it represents a layer-4 failure that's been absorbed by humans instead of fixed.

AI doesn't reduce the copy-paste tax by being smart. It reduces it by replacing the copy-paste with a connection. If the connection doesn't exist, the AI is just another tool people copy-paste into.

> **WATCH OUT FOR**
>
> Tools that claim to be "AI-ready" but expose no API. The label is marketing. If your systems can't programmatically exchange data, no amount of AI bolted onto them will change what your team has to do by hand.

What "ready" looks like

- Every system involved in the workflow has a programmatic interface.
- The data exchange is documented and predictable — same shape, same fields, same place.
- The places where humans are still in the loop are deliberate, not accidental.
- Adding or replacing a tool is a project, not a crisis.

07

Decision

Does the AI actually change an outcome?

This is the layer almost nobody starts with, and the one that decides whether the entire stack was worth building. AI that doesn't change a decision is theater. It can be very expensive theater.

The test for Layer 5 isn't whether the AI produces output. It's whether the output causes a real business decision to be made — faster, better, or at all — that wouldn't have been made otherwise.

Output vs. decision change

Most failed AI pilots produce output. That's not the problem. The problem is that the output doesn't change anyone's behavior. It gets generated, looked at, occasionally pasted into a slide, and ignored. The decision the AI was supposed to inform either gets made the same way it was always made, or it turns out no one was making that decision in the first place.

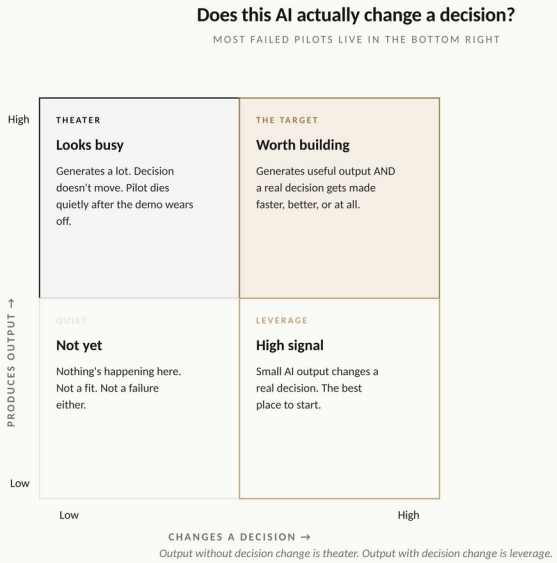

Output without decision change is theater. Output with decision change is leverage.

Look at the bottom-right quadrant. That's where the leverage is — small AI output, big decision impact. A weekly report that nudges one resourcing call. A flag that changes which deals get reviewed. A summary that makes a meeting unnecessary. These are tiny outputs that move real decisions, and they are the highest-return AI investments most businesses can make.

The top-left is theater. A wall of generated content that doesn't change what anyone does. It's the easiest place to end up. Almost every AI pitch you'll hear is selling something in this quadrant.

> **THE PRINCIPLE**
>
> Output without decision change is theater. Output with decision change is leverage. Build for leverage.

What "ready" looks like

- You can name the specific decision the AI is supposed to improve.
- You can name the person who makes that decision today.
- You can describe how that decision would be made differently with the AI's output.
- If the AI vanished, someone would notice — because a real decision would be harder to make.

If you can't answer any one of those, you don't yet have a use case. You have a feature looking for a problem.

> **WATCH OUT FOR**
>
> The pilot that proves "AI can do X." Of course it can. The question that matters is whether anyone will make a different decision because it did. "Can" is not "should." "Works" is not "useful."

08

The self-assessment

Reading the framework is the easy part. Using it on your own business is where the work is. This chapter walks you through that exercise — with a pen, not a download. The point is to make you stop and look honestly at three workflows.

How to use the scorecard

Pick three workflows in your business. Don't pick the obvious AI candidates. Pick the ones that hurt — the ones where you've thought "this should be easier." The scorecard works for any candidate; pick what's actually expensive.

For each workflow, score each of the five layers from 0 to 3.

- 0 — Not in place. We don't have this.
- 1 — In place for some cases, ad hoc.
- 2 — Mostly in place, with known gaps.
- 3 — Solid. Documented, owned, working.

Be honest. The exercise is useless if you score what you want the answer to be.

WORKFLOW	DATA	PROCESS	PEOPLE	TOOLS	DECISION
Workflow 1					
Workflow 2					
Workflow 3					

What the scores tell you

The pattern matters more than the total. Look for the lowest score across the foundation tier — Data, Process, People. That's the lowest broken layer. That's where the work is.

- If any foundation layer scores 0 or 1, start there. Don't even think about AI yet — fix the operational foundation first, because it's holding everything else back.
- If the foundation tier is mostly 2s and 3s, but the activation tier has 0s or 1s, you have a tooling or decision-design problem. These are solvable with effort, often without buying new software.

- If everything scores 2 or 3, you have a real candidate. Scope a project around the lowest-scoring layer first; that's where your readiness gap closes.

What to do next

You now have a ranked list of where your operation is and isn't ready. Three honest paths from here:

1. Fix it yourself. The foundation work is mostly operational discipline, not technical capability. If you have the time and the team, this is real work that pays for itself with or without AI.
2. Bring in help. If the readiness gap is one you'd rather not close alone, this is the kind of work Built by Berry was built for. We work with operators on exactly the readiness gaps this guide names, then build the operational AI on top once the foundation holds.
3. Wait. There's no shame in deciding the readiness work isn't this year's priority. The framework is still useful — it tells you what you'd need to fix when it is.

> **THE PRINCIPLE**
>
> You don't have to be fully ready. You have to know where you are. The companies that get AI right are the ones who do the readiness work before they do the AI work. That's it. The order matters more than the speed.

If you want to talk through what the scorecard surfaced — or scope work to close the gap — Built by Berry takes operational AI engagements where the readiness work and the build are the same project. We are not a permanent line item. The work is useful even if you don't engage us afterward.

Scope a project at builtbyberry.com.

APPENDIX A

The Readiness Stack, one page

A tear-out reference. The framework, the questions to ask at each layer, and the smell of a layer that isn't ready.

#	LAYER	THE QUESTION	SMELL OF UNREADY
05	**Decision**	Does this AI change a decision someone owns?	Output is generated but no one's behavior changes.
04	**Tools**	Can the systems involved exchange data without a human in the loop?	Copying between tools is required for the workflow to function.
03	**People**	Does one named person have authority to change how this workflow runs?	Multiple people touch it; no one can decide.
02	**Process**	Could a smart outsider perform this workflow from documentation alone?	Every case is a special case.
01	**Data**	Is the information AI would need stored somewhere a system can reach?	The answer is "someone knows it" or "it's in an email."

HOW TO USE THIS PAGE

Walk a specific workflow up the stack from Layer 1 to Layer 5. At each layer, answer the question honestly. The lowest layer where the answer is "no" is where your work is. Fix that layer before anything else.

APPENDIX B

Glossary

Terms operators encounter when AI vendors pitch them. Honest definitions, short.

Agent

Software that uses an AI model to take actions across multiple steps, often calling other tools. Useful when the steps are well-defined. Failure-prone when they aren't. "Agent" is also a marketing word; ask what specific actions it takes.

Fine-tuning

Training an AI model on your own data so it learns patterns specific to your business. Expensive and often unnecessary — most useful results come from giving the model better context at the moment of use, not from training it further.

RAG (Retrieval-Augmented Generation)

The pattern where an AI model is given relevant documents from your business at the moment of answering, instead of being trained on them. Almost always the right starting point. The phrase is jargon; the idea is "give the AI your stuff when it's working."

Human-in-the-loop

A workflow where AI produces output and a human reviews or approves it before the output causes a real-world action. Useful for high-stakes decisions. A way of saying "the AI suggests, the human decides."

Model

The actual AI — the thing that takes text in and produces text out. The model is the easiest piece of the stack to swap. Most readiness work isn't about which model.

Pilot

A small, scoped trial of an AI use case. The right kind of pilot has clear success criteria and an explicit path to either kill it or scale it. The wrong kind of pilot has neither and dies quietly.

Prompt

The instructions you give the AI. Good prompts are specific about the input, the desired output, and the rules. Most "AI didn't work" stories are partly prompt stories.

Workflow

A repeatable sequence of steps a business performs to produce a result. AI plugs into workflows. Workflows that aren't documented don't have anywhere for AI to plug in.

READY ENOUGH

A Field Guide to AI Readiness

Daniel Berry · Built by Berry

builtbyberry.com