

The AI Readiness Framework

A tiered path from unmanaged AI adoption to deliberate AI adoption — for small and mid-sized businesses.

The business problem

What separates AI outcomes is not ambition or spend. MIT research found that roughly 95% of enterprise GenAI pilots deliver no measurable business value. The pattern underneath is not mysterious: companies act on AI before they can see the business clearly enough to know what should be automated, what should be standardized, what should be governed, and what should be left alone.

Every operating company is somewhere on this curve — building deliberately, experimenting, or adopting passively through its employees — and each position fails in its own way.

Deliberate builders fail when they automate eagerly, or outsource business-sensitive decisions to agentic systems that have not yet reached the required level of technical, operational, or adoption maturity.

Experimenters fail when pilots never converge into clean data, stable ownership, measurable baselines, or a defined business process capable of supporting automation.

Passive adopters fail by accumulating invisible risk. Employees are already using AI, but without clear rules, data boundaries, approved tools, or any shared understanding of where AI is touching the business.

In every case the blindness is concrete: they do not know how the work actually moves, which data can be trusted, how performance is measured, where decisions happen, or where AI is already touching the organization.

This framework makes the sequence explicit. It begins with an AI-adjacent business diagnostic: making the company legible to itself first. From there, the path is tiered: standardize the work, automate the predictable, apply AI only where it is justified, and reserve agentic systems for environments mature enough to support them.

Wherever a company enters the curve, the framework answers the same question: what is the right next step from here — including when that step is not AI at all.

How the tiers work

Each tier is two things at once: a state the company can be diagnosed into, and a body of work that moves it up. The boundary between tiers is a shared state: the exit of one tier is the entry of the next. The framework is built from diagnostic work inside small and mid-sized businesses; enterprises face a different version of this problem, outside its scope. Read the entry profiles below and locate yourself honestly. Most companies place themselves one tier too high.

Tier 0 — Operational Legibility

You are here if: processes exist but nobody can fully describe them. Metrics are sparse, scattered, or untrusted. The system of record is informal: a spreadsheet, a shared drive, one person's memory. Nobody knows exactly which AI tools employees already use, or what data goes into them.

What Tier 0 is not: a failing business. Profitable companies live here for years. Tier 0 is also not an AI project — it is the diagnostic that determines whether any AI project can be evaluated at all.

The work: answer four questions.

What are our processes, actually?

How does work really move, and where do decisions happen?

What actually serves as the system of record?

Every operating company has an ERP or something informally functioning as one; the clarifying step is naming what that system actually is, rather than what it's supposed to be.

How do we know how we're doing?

What baseline metrics exist, and are the numbers trusted?

Where is AI already touching our business?

An honest inventory of tools in use, sanctioned or not, and what data flows into them.

The cost of skipping: without a baseline, no future initiative can ever demonstrate value — "it feels faster" and "it makes things easier" are not business cases. And without an inventory, AI exposure keeps accumulating invisibly, on terms nobody agreed to.

Tier 0 is complete when the company can describe how the work actually moves, which data can be trusted, what metrics exist, and where AI already touches the business. The output of Tier 0 is not a fixed business. It is a business that can finally be seen clearly enough to make the next decision.

Tier 1 — Standardization

You are here if: the business is legible but not yet stable. Statistics exist, but not yet good KPIs. The workflow is partly digital, with paper, manual checks, or informal handoffs still in the loop. Conventions exist but depend on individual discipline. Ask three employees to describe the workflow and you get three different answers.

What Tier 1 is not: chaos. It is an operating system that exists, but needs to be stabilized.

The work: this is the human side. Write the SOPs. Assign process ownership explicitly. Tighten definitions and establish shared conventions. Improve handoffs. Turn existing statistics into KPIs the company actually consults. The goal is a single, shared understanding of the process — the answer to "how does this work here" should not depend on who you ask.

The cost of skipping: automating a mess gives you a faster mess. Automation built on an unstandardized process encodes its inconsistencies permanently — and at machine speed.

Tier 1 is complete when any employee gives the same description of the workflow, ownership is assigned, the process is mostly digital (or the remaining paper could be dropped tomorrow without repercussions), and KPIs are visible enough to show where improvement is possible. The process may still be slow. Mistakes may still happen, but they are no longer primarily caused by unclear ownership, inconsistent conventions, or hidden handoffs.

Tier 2 — Deterministic automation

You are here if: the process is standardized and fully visible, everyone describes it the same way, but it runs slower than it should. The remaining problems are technical and operational: redundancies, bottlenecks, manual steps that follow fixed rules.

What Tier 2 is not: a lesser form of AI. Deterministic automation is cheaper, auditable, and predictable — and it should carry most of the load permanently, even in companies that later adopt AI.

The work: this is the technical side. Tier 1 aligned the humans; Tier 2 makes the system hold the line. Automate the rule-based parts of the workflow. Build enforcement into the tools — validation rules, required fields, gated handoffs — so conventions no longer depend on discipline. Remove redundancies, consolidate ownership where possible, and attack bottlenecks. This is also the first tier where architectural decisions genuinely matter: data structure, security, and how automation is introduced to the team.

The cost of skipping: jumping from a standardized-but-manual process straight to AI means asking a probabilistic system to do work a deterministic one would do better, cheaper, and

without hallucination risk. Most of what companies call "AI opportunities" are Tier 2 problems wearing a Tier 3 costume.

Tier 2 is complete when automations carry most of the repetitive labor, each one demonstrably works, and KPIs show improvement against the pre-automation baseline. The clearest signal is cultural: your employees are bored, because the machine took the drudgery.

A position worth stating plainly: this is where most small and mid-sized businesses can — and should — comfortably stop. Well-designed automation producing monitored KPIs is a strong end state, not a waypoint. Moving further is a decision, not a default.

Tier 3 — AI-Assisted Operations

You are here if: Tier 2 is genuinely complete, and there remains work that is judgment-shaped — ambiguous inputs, language-heavy tasks, decisions that rules cannot express — where the business case for AI is concrete rather than aspirational.

What Tier 3 is not: an experiment. The gate into this tier is business necessity: a tangible, measurable advantage that clearly outweighs the cost, not the fact that everyone seems to be adopting agents right now.

The work: apply AI in bounded form. Constrained inputs, narrow loops, human approval gates where stakes are real, fail-closed handling of anything unexpected. Every AI component should have a defined job, a defined boundary, and a defined behavior when it is unsure. Agentic systems — AI that makes decisions and takes actions — are the furthest end of this tier and demand the most: technical maturity, operational maturity, and adoption maturity, tested in stages, including for cost, before final rollout.

The cost of getting here early: an AI layer inherits every weakness of the layers beneath it. Deployed onto unstandardized processes and untrusted data, it doesn't fix the mess — it makes the mess fluent, confident, and hard to audit.

Where to start

The entry point is the same wherever you are on the curve: the diagnostic. Tier 0 is a bounded piece of work with a concrete output — your processes as they actually run, your baseline, and your existing AI exposure — and an honest answer about your next step, including when that step is not AI at all.

If you want to know where your company stands before investing further in AI or automation, get in touch: tanyaomelchukai@gmail.com