§ FIELD GUIDE · 2026 EDITION — FOR OPERATORS

Is your business ready to put AI to work?

A practical guide to AI readiness for small and mid-sized businesses. What it actually requires, how to score yourself, and where to start. No hype, no jargon.

NUCLEOTTO

THE AI READINESS GUIDE

nucleotto.ai

§ 01 · WHY THIS GUIDE EXISTS

Every business is being told it needs AI. Few are told what AI needs from them.

The pitch is everywhere: agents that answer your customers, draft your reports, run your operations. Some of it is real. We build it for a living.

But here is what the pitch leaves out: **the technology is rarely the constraint.** When an AI deployment fails in a small business, it almost never fails because the model was not clever enough. It fails because the process it was meant to run existed only in someone's head. Because the data lived in six places and agreed in none of them. Because nobody owned the output. Because there was no clear line between what the system could do alone and what needed a human signature.

Readiness is not a technology question. It is an operations question. And it can be assessed honestly in about an hour, with a pen, by the person who runs the business.

> **This guide exists so you can make that assessment yourself, before anyone sends you a proposal.**

Who it is for. Owners and operators of small and mid-sized businesses. You do not need a technical background. If you can describe how your business works, you can use everything in here.

How to use it. Read the model (pages 3 to 7). Score yourself with the checklist (pages 8 and 9). Then read what your score means and where to start (pages 10 and 11). Total time: about an hour.

§ 02 · WHAT READINESS ACTUALLY MEANS

Three things people get wrong.

×01	**"Readiness means buying the right tool."**	Tools are the easy part. A subscription does not make a business ready any more than a treadmill makes a person fit. Most businesses we meet already own software they barely use.
×02	**"Readiness means hiring data scientists."**	You will likely never need one. Modern AI systems are operated, not invented. What you need is what you already have: someone who understands how the work actually flows.
×03	**"We should wait until things calm down."**	Things do not calm down. The businesses that benefit first are not the calmest ones. They are the ones that can describe their own chaos precisely.

THE WORKING DEFINITION

A business is ready for AI when its work is **legible**, its records are **reachable**, and its judgment has an **owner**. Everything in this guide unpacks those three words.

Legible: the work can be written down as steps, inputs, and decisions. Reachable: the systems that hold your records can be read and written by software, with permission. Owned: a named person answers for what the AI produces, and the actions that carry real risk are gated behind a human approval.

§ 03 · THE MODEL

Five dimensions of readiness.

Every assessment we run, in any industry, comes down to the same five questions. Each gets a page in this guide, then a place on your scorecard.

01	**Process legibility**	Can the work be written down? AI runs procedures. If the procedure lives only in someone's head, that head becomes the bottleneck and the single point of failure.
02	**Data condition**	Is there one place where the truth lives? Not clean data. One system of record per kind of fact, so the AI knows which version to believe.
03	**System access**	Can software reach your systems? Exports, APIs, integrations. A system nobody can connect to is a filing cabinet, whatever the subscription costs.
04	**People & ownership**	Does the output have an owner? Someone named, who reviews what the AI does, corrects it, and answers for it. Unowned automation drifts.
05	**Oversight & approval**	Is there a line the AI cannot cross alone? Payments, customer commitments, anything irreversible. The line is drawn before deployment, not after the first incident.

> **Notice what is missing: model choice, vendor choice, budget. Those decisions come after readiness, not before.**

01 / 05 Process legibility *Can the work be written down?*

Pick the task you most want off your plate. Now imagine writing instructions for a competent temp who starts Monday: the steps, the inputs, where the exceptions go. If you can write that page, the task is legible. If the honest answer is "it depends, I just know," you have found your real first project, and it costs nothing but an afternoon.

The good news: this is the cheapest dimension to fix. Document the three processes you would hand to an assistant first. Each is one page. Steps, inputs, outputs, and the exceptions that need a human. That document does double duty: it briefs the AI, and it is the checklist a human uses to verify the AI.

SIGNALS YOU ARE READY	SIGNALS YOU ARE NOT
The task repeats weekly or daily in roughly the same shape.	"Every case is different" is the standard answer.
Two different people could do it and produce the same result.	The process changes depending on who does it.
Exceptions are recognisable: you can say what makes a case unusual.	Only one person can do it, and they are hard to interrupt.
Someone has written at least a rough checklist before.	Nobody can say where the task starts and ends.

> **A process you cannot explain to a temp is a process you cannot delegate to software.**

02 / 05 Data condition

Is there one place where the truth lives?

Forget "clean data." The standard is simpler: for each kind of fact, one system holds the authoritative version. Bookings live in the booking system. Contacts live in the CRM. When the spreadsheet and the CRM disagree, everyone knows which one wins. An AI can work with messy records. It cannot work with two systems that both claim to be right.

The test: pick a customer at random. How many places hold a version of their record, and do those versions agree? If the answer is "three places, and no," your first project is not AI. It is electing a winner and retiring the others from authority. They can keep existing. They just stop being the truth.

03 / 05 System access

Can software reach your systems?

Your systems hold the records; the question is whether software is allowed in. Three doors, in descending order of usefulness: an API (best), scheduled exports (workable), and a human copying numbers between screens (the door is closed). Most modern small-business tools have the first or second. The blockers are usually the old ones: the legacy system, the supplier portal, the inbox that functions as a database.

You do not need every door open. You need the doors open on the two or three systems that touch your first project. A closed door elsewhere is a roadmap item, not a verdict.

> **One system of record per fact, and a door software can knock on. That is the whole infrastructure requirement.**

04 / 05 People & ownership

Does the output have an owner?

Every AI system that works in production has the same thing standing next to it: a named human who reads its output, corrects it, and answers for it. Not a committee. Not "the team." A name. In a four-cabin hospitality business, that is the owner reviewing escalations. In a fintech, it is the engineer approving deploys. The shape scales; the principle does not change.

Ownership is also where the improvement comes from. The owner's corrections are the training signal. A system nobody corrects does not stay the same. It quietly gets worse, because the business around it keeps changing.

05 / 05 Oversight & approval

Is there a line the AI cannot cross alone?

Before anything goes live, you write two lists. **List one: actions the system may take alone.** Answer a routine question, sync a calendar, draft a report. **List two: actions that require a human signature.** Move money, make a commitment to a customer, touch production data, send anything irreversible.

This is not a brake on the system. It is what lets you run it fast everywhere else. The businesses that skip this step do not move quicker; they just meet the line for the first time during an incident.

THE HUMAN GATE

Every sensitive action pauses, states what it wants to do, and waits for a named person to approve. We build this into everything we ship. We recommend you accept nothing less, from anyone.

§ 04 · THE SCORECARD

Score yourself honestly.

Twenty statements, four per dimension. Score each: **0** = not true, **1** = partly true, **2** = clearly true. Maximum 40. Do it with a pen, with the person who actually does the work in the room.

STATEMENT	0	1	2
01 · PROCESS LEGIBILITY			
We can name the three tasks we would hand off first.	☐	☐	☐
Each of those tasks repeats in roughly the same shape every time.	☐	☐	☐
At least one is written down as steps a newcomer could follow.	☐	☐	☐
We can say precisely what makes a case an exception.	☐	☐	☐
02 · DATA CONDITION			
For each kind of record, we can name the system that holds the truth.	☐	☐	☐
When two systems disagree, everyone knows which one wins.	☐	☐	☐
A random customer's record agrees across our systems.	☐	☐	☐
No critical record lives only in an inbox or a personal spreadsheet.	☐	☐	☐
03 · SYSTEM ACCESS			
Our main systems offer an API or scheduled exports.	☐	☐	☐
We know who can grant software access, and they would.	☐	☐	☐
Nobody spends hours each week re-typing data between screens.	☐	☐	☐
The systems touching our first project all have an open door.	☐	☐	☐

STATEMENT	0	1	2
04 · PEOPLE & OWNERSHIP			
A named person would own each automated output.	☐	☐	☐
That person has time to review and correct, not just forward.	☐	☐	☐
The team doing the work today knows this is on the table.	☐	☐	☐
Someone internally is curious about AI rather than cornered by it.	☐	☐	☐
05 · OVERSIGHT & APPROVAL			
We can list the actions that must never happen without a signature.	☐	☐	☐
We know which mistakes are reversible and which are not.	☐	☐	☐
We would insist on a test environment before anything touches production.	☐	☐	☐
We would measure results against a baseline, not against a feeling.	☐	☐	☐

What your score means.

0 – 14

FOUNDATIONS FIRST

Do not buy anything yet. Document three processes, elect your systems of record, and re-score in a month. This is unglamorous work that costs almost nothing and moves every number on this page.

15 – 27

READY FOR A FIRST CONTAINED PROJECT

Pick one frequent, rule-bound, low-risk task and automate that alone. Prove the loop: run, review, correct, measure. Expand only after the first one holds.

28 – 40

READY TO EMBED AN OPERATOR

Your operations can support AI doing real work across several functions, with human gates on the sensitive actions. The constraint now is sequencing, not readiness.

§ 05 · WHERE TO START

Pick boring.
Pick frequent.
Pick reversible.

The best first project is never the most impressive one. It is the one that runs every day, follows rules, and cannot do real damage if it gets something wrong.

Good first candidates from businesses we work with: answering routine customer questions, syncing records between systems, producing the weekly report, scanning a market and ranking what changed overnight. Bad first candidates: anything touching money, anything that makes promises to customers, anything you cannot undo.

WEEKS 1–2	**Map.** Write the one-page procedure. Elect the systems of record. Draw the two lists: acts alone / needs a signature.
WEEKS 3–6	**Build.** The automation is built against your real systems, in a test environment, against the procedure you wrote.
WEEKS 7–10	**Run supervised.** It works; a human reviews everything it produces. Corrections feed back in. The gate gets exercised for real.
WEEKS 11–13	**Measure.** Hours returned, response times, error rates, against the baseline you kept from before. Verified numbers or none.

> **One automation in production beats five in slides. Prove the loop once, then compound.**

§ 06 · FAILURE MODES

Six ways this goes wrong.

We have seen each of these more than once. None of them is a technology failure.

×01	**Tool-first thinking**	Buying the platform, then hunting for a problem worthy of it. The subscription outlives the enthusiasm by years.
×02	**Automating fog**	Pointing AI at a process nobody can describe. The system faithfully reproduces the confusion, faster and at scale.
×03	**No owner**	"The AI handles it now." Output that nobody reads degrades quietly until a customer reads it for you.
×04	**No gate**	Letting the system act alone on sensitive actions because approvals felt slow. The line gets drawn after the incident instead, at far higher cost.
×05	**Big-bang launch**	Automating six functions at once. When something breaks, nobody can say which change caused it. Sequenced deployments fail small and teach fast.
×06	**Measuring nothing**	No baseline kept, so "is it working?" gets answered by mood. Six months in, nobody can defend the spend, including the people who were right.

> Every failure on this page is avoidable with the scorecard you just filled in.

We embed AI operators inside running businesses.

Nucleotto is an AI enablement agency. We build and run Otto: an AI operations layer that works inside small and mid-sized businesses, from fintech to hospitality to recruitment. Every engagement runs on the same principles this guide is built on.

01 Staging first. Nothing meets production untested.

02 The human gate. Sensitive actions wait for a named approver.

03 Verified numbers or none. Results are measured against your baseline.

NUCLEOTTO

READ THE CASE FILES

nucleotto.ai