toodlebox

5 Things to Prepare Your Business for AI

Five ideas worth working through before your next AI decision — grounded in what we've actually seen go right, and wrong, inside businesses like yours. Work through them with your leadership team before you sign anything, build anything, or ban anything.

toodlebox

The AI decision has already been made without you.

Here's the pattern we see almost every time we sit down with a business owner or team lead: they think they're at the start of an AI decision. In reality, three or four people on their team made that decision months ago — quietly, tool by tool, with no one coordinating it.

It usually looks something like this. Someone in ops has been pasting client emails into ChatGPT to draft responses faster. Someone in finance turned on the AI summarization feature inside your accounting software without asking whether that vendor now processes your financial data through a third-party model. Someone in sales installed a browser extension that reads every page they visit, including the CRM, to auto-draft follow-ups. None of these people did anything malicious. Each decision made sense in isolation. Nobody was ever in the room to see all three at once.

That's not a criticism. It's just how adoption works when a technology gets this useful this fast. The problem isn't that AI is in use — that ship has sailed at almost every company we've walked into, including yours, whether you've confirmed it yet or not. The problem is that most businesses can't currently answer basic questions about how: which tools, touching what data, producing what kind of trail.

That gap used to be tolerable. It's getting less tolerable by the quarter. The standard vendor security questionnaires clients send before signing or renewing a contract have started adding entire sections on AI governance — which tools you use, what data they touch, whether a human reviews the output. Cyber insurers are moving the same direction: some now ask about AI controls at renewal, and a few carriers have introduced coverage add-ons that require proof of AI governance before they'll extend certain protections. None of this is hypothetical — it's showing up on real forms in real renewal cycles right now. The businesses that get asked and can't answer usually aren't losing the deal or the policy because they're doing something wrong. They're losing it because they can't prove they're doing anything at all.

The five ideas below aren't a product pitch — they're the questions worth answering regardless of who helps you build the answer. We'll share where we've seen each one go wrong, and a small exercise you can run this week for each.

01 Map Where AI Already Lives in Your Business

In discovery conversations, we routinely find two to three times more AI tools in active use than leadership expected.

Personal ChatGPT logins, a Copilot trial someone enabled, a browser extension that summarizes email — none of it shows up on an IT asset list, because nobody installed it through IT. It just accumulated. The goal here isn't to catch anyone doing something wrong. It's to replace assumption with an actual inventory.

The shadow footprint tends to show up in three places at once. Frontline and admin staff adopt consumer tools for speed — email drafting, meeting notes, quick research — usually on personal accounts with free-tier terms that let the vendor train on whatever gets typed in. Managers adopt AI features already sitting inside tools you've already approved: the summarize button in your helpdesk, the smart-reply in your CRM, the meeting-notes bot your video platform turned on by default in an update. And a technical minority — usually one or two people — connect AI tools directly to your systems through APIs or automation platforms, which is where the exposure gets serious fastest, because those integrations often carry standing access to whatever they touch.

- Ask a handful of people across different teams what AI tool they used most recently, and for what.
- Note anything connected to email, shared drives, CRM records, or customer data — that's your highest-exposure category, regardless of how small the tool seems.
- Check for AI features already switched on inside tools you already pay for — many vendors turn these on by default with an update, not a new purchase.
- Ask specifically about browser extensions and desktop plug-ins; they're the easiest to install and the easiest to forget about.

- Treat the list as a living document, not a one-time audit — it will be out of date in a month.

A real inventory captures four things for each tool: who's using it, what data it can see, whether it's a personal or business account, and what the vendor's data-retention terms actually say. If you can't answer all four for a tool, you don't have visibility into it yet — you have a rumor about it.

TRY THIS

Pick three employees in different roles and ask them what AI tool they used last week, and what they fed into it. If any answer surprises you, that's a sign your real footprint is bigger than your assumed one — and it's worth repeating with three more people next month.

02 Decide What AI Should Never Touch — Before You Decide What It Can

A one-page boundary is worth more than a forty-page policy nobody reads.

Most AI policies fail because they try to describe everything AI is allowed to do. That list is always incomplete and always out of date. It's far more durable to define the short list of things AI should never touch without a person in the loop — certain financial thresholds, certain categories of customer data, certain irreversible actions — and build outward from there.

It rarely goes wrong in an obviously reckless way. It goes wrong like this: an AI tool drafts a reply to what looks like a routine vendor email, someone approves it on autopilot because the draft reads well, and the underlying request — a change to payment details, a discount promised to a customer, a scope change on a contract — turns out to be exactly the kind of thing that needed a second set of human eyes before anything left the building. The tool didn't do anything wrong. It did precisely what it was asked. Nobody had decided in advance that this category of decision required a stop sign.

- Name three to five decisions in your business that should always require a human sign-off — payment or account changes, anything touching a legal commitment, anything irreversible.
- Set the threshold in advance, not in the moment — for example, any financial change above your normal approval limit, any HR action, any customer-facing promise.
- Assign one owner for exceptions, so “who approved this?” always has an answer.
- Write the list down somewhere everyone touching AI tools can actually find it — a policy nobody's seen isn't a policy.
- Revisit the list quarterly — what's off-limits today may need to change as tools mature and your comfort with them grows.

TRY THIS

Write down the three decisions in your business you'd never want made by AI alone, and who has to sign off instead. That short list is a stronger starting policy than most companies write in a month — and it's the first thing we ask for in an assessment.

03 Make AI-Assisted Decisions Provable, Not Just Explainable

If you can't produce the record in under two minutes, you don't have one — you have a memory of one.

There's a real difference between being able to explain, in general terms, how your team uses AI, and being able to produce the specific record of what happened in a specific case. The first is a conversation. The second is what actually holds up when a client, auditor, or new hire asks “why was this approved?”

This matters most exactly when you least want to be figuring it out for the first time — mid-dispute, mid-audit, or mid-complaint. If a customer disputes a quote an AI tool helped generate, or an employee raises a concern about how a decision affecting them was made, or a regulator asks how a specific record was handled, “we generally use AI to help with that” is not an answer anyone accepts. What they want is the specific prompt, the specific output, who reviewed it, and when. If that trail doesn't exist, it doesn't matter how careful you actually were — you can't prove it.

- Wherever AI touches a real decision, capture what was submitted and what the outcome was — not just that a conversation happened.
- Store it somewhere searchable and dated, not buried in a chat history that scrolls away or a personal account nobody else can access.
- Log who ran it and when, not just the content — “who” is usually the first question anyone asks.
- Set a retention period that matches how long you'd need to defend the decision, not just how long the tool happens to keep it.
- Test it occasionally: pick a past decision and time how long it takes someone to reconstruct it.

THE TEST

Pick one AI-assisted decision from the last month — a quote, a hiring note, a customer email. Time how long it takes someone on your team to show exactly what happened, prompt to outcome. If it takes more than a couple of minutes, or isn't possible at all, that's your gap, and it's usually bigger than people expect.

04 Favor Narrow, Well-Scoped Tools Over One Big Assistant

The businesses getting the most out of AI right now are rarely the ones using it the most broadly.

It's tempting to hand one general-purpose assistant the keys to everything — support, HR questions, data lookups, drafting. In practice, a tool that's scoped to do one job well, with clear limits on what it can see and touch, tends to outperform a generalist both on quality and on how much you can trust it unsupervised.

There's a name for what happens next: assistant sprawl. One bot ends up connected to the helpdesk, the HR system, the shared drive, and the CRM — not through any single bad decision, but because it was easier to grant broad access once than to scope it four separate times. It works fine until the day it doesn't: a bad instruction, a malicious message it was asked to summarize, a permission nobody remembered granting. Now the blast radius is everything the assistant could reach, not just the one job it was supposed to be doing. Narrow tools fail narrow. General assistants fail wide.

- Match each AI tool to a specific job, not a general category of “help.”
- Keep the data and permissions for each use case separate from the others, even when it's tempting to reuse the same connection.
- Give each tool its own credentials and its own logs, so an incident in one doesn't force you to review everything.
- Define what happens when the tool doesn't know the answer — a clear handoff to a person beats a confident guess every time.
- Judge success by how reliably a tool handles its one job, not by how many things it can technically attempt.

TRY THIS

List every AI use case in your business right now. For each one, ask: is this scoped to a specific job, or is it a general assistant being asked to do too much? Start narrowing the second group first — that's usually where the exposure is concentrated.

05 Treat Access Control as Day-One Work, Not Cleanup

The businesses we see retrofit security are almost always doing it right after an incident, not before one.

Permissions and data boundaries are cheap to design in from the start and expensive to add after the fact — both in engineering effort and in the trust you lose if something goes wrong first. Treat AI access the same way you'd treat access to your financial systems: default to the minimum, and expand deliberately.

Ask what a given AI tool can actually touch, and the honest answer is almost always more than the job required — full mailbox access when it only needed to read subject lines, edit rights when it only needed to summarize, a standing connection to shared drives when the job only ever touched one folder. Nobody granted that access maliciously. It's just the path of least resistance when you're trying to get a tool working quickly. The cost shows up later, when that one over-permissioned account becomes the single point of failure for everything it was ever connected to.

- Give each AI tool or bot only the access it needs for its specific job — nothing broader “just in case.”
- Use read-only access wherever the job doesn't actually require writing or deleting.
- Time-box access where you can, especially for trials and pilots that tend to quietly become permanent.
- Review AI permissions on the same cycle as your other IT security reviews, not as a separate afterthought.
- Assume every new AI tool starts with zero access, and earn its way up from there.

TRY THIS

Pick one AI tool your team uses today and check what it can actually access — not what you think you set it up to access. If the honest answer is “more than it needs,” that's worth fixing before you add the next tool.

toodlebox

When You Know It's Time to Start Exploring an AI Solution

None of this requires urgency for its own sake. But there's a real difference between "we should look into this eventually" and "we're past the point of handling this ourselves" — and it usually shows up as one of the signals below.

Working through the five ideas above tells you where you stand. The list below tells you what that standing usually means. None of these signals is an emergency on its own — but once two or three are true at the same time, the gap usually isn't closing by itself.

- **You can't produce the tool list without guessing.** If Section 01's exercise turned up tools you didn't know about — or you're not fully confident you found all of them — that's usually a sign the footprint has outgrown a manual, ask-around approach.
- **Someone outside your business has already asked.** A client security questionnaire, a cyber insurance renewal, a vendor risk review — if any of these have asked about your AI practices and the honest answer was "we're figuring that out," the market is moving faster than your documentation.
- **You've had a "wait, who approved that?" moment.** Even a small one. If an AI-assisted decision has already slipped past a boundary you didn't know you needed, Section 02's whiteboard exercise is necessary but probably not sufficient anymore.
- **You couldn't produce the record when someone asked.** If Section 03's two-minute test took longer than two minutes — or wasn't possible — for a decision that actually mattered, that's not a documentation gap. It's an exposure.
- **One tool is doing five jobs, or five tools are doing one job badly.** Either version of Section 04's problem tends to compound rather than resolve on its own; sprawl gets harder to unwind the longer it runs.
- **Nobody can tell you, without checking, what your AI tools can access.** If Section 05's access check turned up "more than it needs" and there's no plan to fix it, that's usually the signal that matters most — it's the one with the highest cost if something goes wrong first.

If two or more of these are true for your business right now, you're past the point where a whiteboard and an afternoon will fully close the gap. That's usually the moment it makes sense to bring in a second set of eyes — not because you've done anything wrong, but because the fixes now involve coordinating across tools, teams, and vendors in a way that's hard to do in the margins of everyone's day job.

toodlebox

How we think about this

None of the five ideas above require a particular vendor — they're just what we've seen separate businesses that adopt AI well from ones that end up cleaning up after it. You can work through all five with a whiteboard and an afternoon, and you'll be ahead of most of the market for doing it.

It's also, not coincidentally, the exact problem we spend our time on. Toodlebox exists because most businesses on Microsoft 365 want the productivity AI offers without losing visibility into what it's doing — and building that themselves, tool by tool, policy by policy, is a slow and risky way to get there.

From here, there are two reasonable next steps — depending on how much of this you want to tackle yourselves.

OPTION 1 — RUN IT YOURSELF

Everything in this guide is built to work without us. The inventory exercise, the boundary list, the two-minute test, the access check, and the six-signal readiness list are yours to use — no cost, no call required. Get your leadership team in a room for an afternoon and start there.

OPTION 2 — BRING IN A SECOND SET OF EYES

If you'd rather have us build it with you — especially if two or more signals from the readiness list are already true — we run a fixed-scope AI Assessment: a paid engagement where we build the tool inventory, draft the boundary policy, test whether your AI-assisted decisions are provable, and audit what your tools can actually access. You leave with a governance plan you can act on, whether or not you hire us to help build it.

Contact Us

Prefer to just grab a time on the calendar? **Book a time with us** — mention this guide and we'll pick up right where it left off.

toodlebox.com | sales@toodlebox.com | 630-468-9308

We keep response times short — you'll hear back from our team as soon as possible.