

AI SERVICE GOVERNANCE

Start with Your Services

A Leader's 7-Step Playbook for Adopting AI in business, nonprofits, and public institutions.

You can't digitize chaos. So start with the service.

Contents

Before You Adopt AI	3
Introduction	3
About the Author	4
Step 1. Start with Services, Not Technology	4
Step 2. List Your Services	5
Step 3. Choose a Pilot Service	6
Step 4. Set the AI Agent's Autonomy Level	7
★ Step 5. Create the AI Service Passport 8 The centerpiece tool: a three-part AI Service Passport you build for each service (manifesto, storefront, prompt).	
Step 6 · At Scale. The Rules for AI Services	11
Step 7. Launch and Monitor the Pilot	12
To Sum Up	13

Before You Adopt AI

Every leader adopting AI wants the same four things:

- less chaos
- less risk
- more transparency
- better service

So, before you reach for another shiny AI tool, answer these questions about the AI you're about to put in front of your customers:

- What should I hand to AI first — and what shouldn't I touch yet?
- How do I know the AI does what it promised?
- When must it hand the customer to a human?
- What decisions is it not allowed to make?
- Who's accountable when it gets something wrong?
- What stays public, and what stays in the private prompt?
- How do I measure its quality and keep improving it?
- What happens when it fails?
- What data is it learning from?
- What risk is the company carrying?

If you can't answer these, you don't really have an AI solution — you have an unmanaged risk behind a friendly interface. That's the chaos.

And here's the uncomfortable part: no tool fixes it. The fix isn't more technology — it's order. You take the service and you write it down: what you promise, where the limits are, who's accountable, what happens when something breaks. Passport, standard, rules — the label doesn't matter. What matters is what it is: your process, lifted out of chaos and set down as order.

The 7 steps that follow are simply how you get there.

Introduction

95% of companies see no measurable return on Artificial Intelligence (AI) (MIT study, 2025). The reason is almost never the model itself — it's that the organization has no strategy for adopting AI.

This playbook is your step-by-step roadmap: where to begin and what, concretely, to do. By the end, you'll be able to:

- tell your AI services apart from your AI functions;
- draw up a list of services to bring AI into;
- pick a pilot service to start with;
- choose one of the 5 levels of AI agent autonomy;
- build a three-part AI service passport;

- write the rules for developing and delivering AI services;
- and learn how to monitor the quality of your AI services.

None of this is a rigid prescription — the rules around AI are still taking shape. What I'm offering is my own system for standardizing AI services. It grew out of my work moving government services online during administrative reform. I spent years studying that field, as both a scholar and a practitioner, and I'm convinced that the way governments approach digital services work just as well for adopting AI — in business and in the non-profit world alike. So, this playbook fits business leaders, non-profits, and public institutions equally. Don't treat it as dogma; treat it as a working template, a place to start as you find your footing with AI.

About the Author

Sergei Ponomarev — PhD in Political Science, founder of aibusiness.vc

- In 2014, defended a dissertation in political science on the adoption of e-government.
- In 2015, completed a fellowship on government openness at the Kennan Institute in Washington.
- Spent 14 years as an associate professor of political science.
- Spent 7 years leading applied research at the GRANI think tank — running nationwide monitoring of public-service quality, independent assessments, and mystery-shopping studies. Author of the Information Openness Standard for government bodies.
- Spent 5 years as an aide to a municipal deputy.
- Spent 4 years running trainings and workshops for entrepreneurs.
- Has spent the last 18 months building AI agents hands-on.
- linkedin.com/in/sergei-ponomarev

Book a free 20-minute consultation: info@aibusiness.vc

Step 1. Start with Services, Not Technology

The first question on the list: what to hand to AI first, and what to leave alone. And the answer isn't "the technology." Put AI where it pays back fastest, in the services your clients actually buy.

These days, most CEOs are under constant pressure from AI enthusiasts pitching automated secretaries, recruiters, salespeople, consultants — you name it. The catalog grows by the day. And every so often a giant like Anthropic or OpenAI ships a new model that wipes out hundreds of AI startups overnight. It's hard to find your bearings in that chaos — especially if you don't have the technical background to judge whether a given AI solution is any good.

The problem isn't the IT solutions themselves. The problem is chasing the technology: one-off, scattered AI rollouts don't fit together, and they only add to the chaos inside your organization.

Remember the golden rule: you can't digitize chaos.

So, you always start with the process. Break it into procedures, each procedure into tasks, and put a resource and an owner behind every task. And yes — in writing. There's no way around it.

You can't digitize a living employee — forget it. For all our flaws, a person brings common sense, empathy, context, a feel for human nature, a sense of the culture. AI has none of that. All it has is instructions.

Whatever isn't in the instructions doesn't exist in its synthetic life.

So first you map the process, then you hand it to AI. Skip that, and the agents start making up rules and misleading people.

So that's settled. The next question is which processes to hand over. From a management-theory point of view, any organization's processes fall into two uneven groups:

- Services. The things clients pay for. Eighty percent of your effort and resources should go into producing your services and goods.
- Functions. The work the organization does for itself: logistics, marketing, accounting, legal, IT, and so on.

Most CEOs bring AI into their functions first — it feels like the easier win. But here's the catch:

- Endless tuning. AI tools are generic; your functions are unique to your company.
- Slow feedback. You won't know whether your "content factory" actually works for months.
- A black box. Most CEOs can't tell how well-built an AI agent really is under the hood.
- Quiet sabotage. The rollout comes from the top, people resent it, and staff start quietly undermining it.

I'd flip the order.

Bring AI into your services first.

Here's why that's simpler and safer:

- Money. Services are where you grow. You can only cut a function down to zero — but you can grow a service without limit.
- Feedback. The employee who runs the service knows it cold, so they'll spot right away whether the AI is up to the job.
- Accountability. Because you understand the service, you can keep the AI agent on a tight leash and fine-tune it easily.
- An outside check. If something breaks, you'll hear about it immediately — from your clients.

Action 1. Split your organization's processes into services and functions.

Step 2. List Your Services

You can't choose what to give AI until you can see every service in one place. This turns "what first?" into a concrete shortlist.

A service is not the same thing as a function. The simple test: a service always has an outside customer or applicant — someone who asks you for it and pays.

A service is always provided on request.

With a service, the organization commits to something and hands the applicant a result. Picking and setting up an insurance policy, booking an appointment with a specialist, giving advice — those are services.

A special case is proactive services, where the organization makes the first move — say, a reminder that your insurance is due for renewal. Those happen only with the customer's prior consent.

A function is an internal operation that creates no direct obligation to a particular customer: churning out content, setting prices, internal analytics, optimization.

Here are the differences side by side:

Attribute	Service	Function
External paying client	Yes	No (the organization is its own client)
What it governs	The relationship with the customer	Internal processes
Examples	Policy selection, taking an application, advice	Content factory, pricing, analytics
Where to start with AI	Start here	Leave for later

Action 2. Make a list of every service your organization offers. Set the functions aside.

Step 3. Choose a Pilot Service

Now narrow "what first?" to a single service, and keep "what risk are we carrying?" small on purpose, by picking something low-stakes.

Your first pilot isn't your most important or highest-volume service — it's a practice run, a cheap and safe way to learn. A good candidate is:

A. Simple and self-contained

- Few steps, one owner, no outside dependencies.
- The result shows up in the same session — easy to check.
- Routine and repeatable: the same request, over and over.

B. Low-stakes and reversible

- Low risk — nothing touching hiring, credit, health, or money owed.
- Easy to undo if the AI slips up, with no lasting damage.
- No online payments, and as little personal data as possible.

C. Ideal candidates

- Answering questions about your services — what you offer, how to get it, what it costs. The safest start.
- Booking appointments — take the client's details and reserve a slot.
- Helping people choose — guiding a client to the right service or plan.

Look for a simple, self-contained, reversible, low-stakes service.

Steer clear of, on a first pilot: online payments; taking in and checking documents; bundled services; high-stakes ones (hiring, loans, medical care, legal rights); decisions that are hard to undo; and anything riddled with "special cases."

Action 3. Pick one pilot service to bring AI into.

Step 4. Set the AI Agent's Autonomy Level

This one choice answers two of the hard questions at once: how far the AI may act on its own, and, by the same line, what it's not allowed to do.

You can hand a service to AI with more or less autonomy. The more your AI agent does on its own, the higher the stakes — and the sharper your service description and the sturdier your technical setup need to be. Here's my matrix of AI autonomy in delivering services.

Essentially, it's a scale of how far you pull employees out of the service: each level hands the AI a little more autonomy. One caveat — it's a reference matrix, not a ladder you're obliged to climb. Adapt the levels to your case.

Level	What the AI agent does	What's left for the employee	Result	What you save (ROI)
1. Information	The AI explains the service: what it is, how to get it, what it costs, who to contact	The employee still delivers the whole service	Information delivered about the service	The employee's time (\$) spent answering questions
2. Data intake (the way in)	The AI takes in and checks the request (a smart form): completeness, correctness, and screens out the wrong fit	The employee gets qualified leads and delivers the service	A qualified lead; the wrong-fit requests screened out	Time (\$) spent taking in, checking, and screening requests
3. Decision prep	The AI does the work itself, drafts a decision against set criteria, and passes it to the employee to sign off	The employee reviews it, signs off / hands it to the applicant	A draft decision to approve	Time (\$) spent analyzing and drafting the decision
4. Full cycle	The AI delivers the service on its own: inform → intake → decide → deliver	The employee is out of the loop	The service delivered to the applicant on its own	The employee's time (\$) the service used to take

Level	What the AI agent does	What's left for the employee	Result	What you save (ROI)
5. AI agent to AI agent	Your AI agent serves the applicant's AI agent	The human applicant acts through their own AI agent	A service delivered to the client's agent (AI2AI)	The recipient's own time (\$)

Inform → Intake → Decision prep → Full cycle → AI2AI

Take the example of running an online course business. Hundreds of people ask the same things every day: which courses, what schedule, who's teaching, the price, any discounts, the start date, how to pay.

- Level 1. The AI agent fields all of these and hands the person off to the school's admin — and that alone clears away a mountain of repetitive work.
- Level 2. The agent doesn't just answer — it gathers the client's details and builds a request: goal, preferences, current level, budget, preferred time. Then it qualifies the lead: a fit or not. The wrong fits it politely turns away or redirects, and the admin gets a qualified lead with a ready-made profile.
- Level 3. On request, the agent picks the course, format, and schedule itself, puts together a tailored offer, and issues an invoice. All the admin does is approve it and send it on.
- Level 4. The agent handles the client start to finish: answered questions, qualified, chose, enrolled, took payment, and now sends reminders before each lesson — the admin is out of the loop.
- Level 5. It's no longer a person reaching your AI agent, but the client's AI agent: it reads your machine-readable service passport, stacks it against other offers on the market, weighs it against the client's goals, budget, and schedule, and signs them up.

Whoever describes their AI services well first will win the market of AI agents.

Level 5 already exists in simple consumer cases: ask your ChatGPT or Claude to find you a product online, and it will. Full AI2AI — agents negotiating with each other by machine-readable rules — is still on the horizon, but we'll get there in a couple of years, once every AI service has a machine-readable passport. That will be a genuine revolution in the market.

Action 4. Decide how much autonomy you want for your pilot AI service.

Step 5. Create the AI Service Passport

This is where most of the list gets answered: what the AI promises, when it must pass to a human, what it may never do, who's accountable, what it logs, and what's public versus private.

One thing to be clear about: there is no single, accepted standard today for describing AI services. No law, no industry norm, no off-the-shelf passports. Everything here is my own take, drawn from studying how the public sector does it. Take it or leave it — your call.

When I said earlier that you describe a service before handing it to AI, I don't mean a few sentences — I mean a specific document that lays out the whole process of delivering it. That document is what I call the AI service passport. Why do you need one?

The AI service passport does three jobs at once: manifesto, storefront, and prompt.

It has at least three sections, each with its own job:

- General Provisions → Manifesto on the use of AI;
- Standard → Consumer storefront;
- Procedures → Part of the system prompt.

General Provisions

From August 2026, EU law starts requiring you to disclose that a customer is talking to AI — with fines of up to 3% of turnover if you don't. So you need to build that disclosure in now, and the passport is one way to do it. Anyone should be able to find out that your organization uses AI when it deals with customers. That's why I call this part your public manifesto on the use of AI.

Open disclosure that you use AI — to anyone, no exceptions.

I'd build the "General Provisions" section from two parts:

- What's specific to this service. What the service covers; who can apply and what's expected of them; categories of applicants (profiling) — who gets a different path, who to recognize, who to send straight to a human.
- The shared requirements from your Rules on AI services (more on the Rules in the next step). This covers your organization and its licenses and certificates for the service; contact details and how to appeal; your principles for delivering AI services; the "this is AI" disclosure; and whatever lawmakers require down the line.

The AI Service Standard

Every website has a list of the company's services. Usually, they're described in a line or two — spelling it all out is tedious, nobody reads long copy, and people are used to just asking.

AI agents change that. More and more people look for products and services not on websites but through ChatGPT or Claude — it's called GEO, optimizing for AI search. And an AI agent reads everything, top to bottom; if there isn't enough, it simply won't recommend you. So the passport gets a second section, the "Standard," which acts as a storefront for your clients' AI agents.

Think of the AI service Standard as your product page on a marketplace.

The Standard spells out what the service is, what you get, what it costs, the rules it runs by, and the timelines. Again, templates for this already exist — you don't have to invent anything. Your job is to fill the template with your own substance. After all, the person who owns the service is your employee: only they know its every step, every wrinkle, and what clients actually want. That hard-won knowledge is what goes into the AI service Standard.

The Standard comes in two parts:

- The human-readable part. A plain-language description of the service for the client, listing its key features.
- The machine-readable part. The same information, formatted so an AI agent can read it easily.

The AI service Standard is your storefront. We used to walk the shops and read the signs ourselves; soon our AI agents will do the legwork for us. And before another AI agent will recommend you to its owner, it has to study you and your services first. It makes sense — you gather the facts before you make a deal.

The service Standard is what lets two AIs shake hands before a deal.

I get that this can feel like a lot — even overkill. But trust me: work through it once and it all clicks into place. And if you've read this far, congratulations — you're past the hardest part of the playbook. Keep going!

Procedures

You can think of the Passport's third part as the AI agent's job description — and the most valuable part of the system prompt.

The Procedures are the step-by-step recipe for delivering the service.

This section lists every procedure in the service: exactly what the AI agent does, in what order, on what timeline, where the forks are, on what criteria it decides, when it has to escalate and pass the client to a human, what counts as a good outcome, what it logs to learn from, how it rates the interaction, and so on. Spell out the hard limits, too: the decisions the AI must never make on its own, and the lines it may not cross under any circumstances. What you forbid is as much a part of the job description as what you allow. This is the section you'll turn into the agent's technical prompt.

First write the procedures, then drop them into the AI agent's prompt.

Note that everything in the Passport is meant to be public. Don't confuse the AI service Passport with the AI agent's prompt.

- The AI service Passport — a public document for customers, about the rules.
- The AI agent's prompt — an internal, technical document for the developers.

The prompt is the engine the AI runs on — the brains of your digital employee. It may hold proprietary know-how or sensitive data: internal pricing, internal procedures, decision criteria, personal data, and the like.

For now, the law only asks you to warn customers, "you're dealing with AI." What it will eventually require you to disclose about the prompts themselves, nobody yet knows.

Right now, you're under no obligation to show what's under the hood. You could skip the passports entirely — but by now I think you see the point: the fuller and more accessible your information, the more readily clients' AI agents will come to you. So, we'll have to feel our way together to the right balance between being open and protecting your own interests.

Here's what I'd do for now:

- Routine procedures — the simple, industry-standard ones with nothing sensitive in them — go straight into the Passport, published.
- Procedures with anything sensitive are internal-only: they go into the AI agent's prompt, and the Passport mentions them without spelling them out.

That was probably the toughest chapter. But you made it!

If you have questions, book a free consultation: info@aibusiness.vc

Action 5. Build your AI Service Passport: general provisions, standard, procedures.

Step 6 • At Scale. The Rules for AI Services

Skip this on day one. It's what keeps everything sane once you have lots of AI services. At scale, some questions go company-wide: who's accountable, how a customer appeals, what quality you guarantee, what risk you carry. The Rules answer them once for every service.

Not for your first pilot. Reach for this once you have ten-plus AI services.

So much for passports. Nicely done! You could stop here — and for one pilot, you should. But there's a catch waiting at scale. With under a dozen AI services, life is simple. But what about hundreds? And things keep changing: phone numbers, emails, internal rules, plus new requirements from lawmakers. Pushing every change into hundreds of passports is a real slog — even for AI.

That's why corporate documents use what's called a Single Source of Truth — here, the Rules. You put all the shared information in this one document, and everything else just points back to it instead of repeating it.

The Rules are the rules you build every AI service passport by.

So, at scale, you keep two kinds of document:

- A Passport for each AI service — its general provisions, standard, and procedures.
- The Rules for developing and delivering AI services — the one document that sets the shared frame for every Passport.

The Rules hold everything you'd otherwise have to copy into every passport. Set it up right once, and from then on, your AI agents pull what they need automatically. Your data stays consistent and current, and you're done with copy-paste.

Make no mistake — the rules lawmakers put on AI will only grow.

Here's the full map of what a complete setup covers — you won't build all of it at once, but it's worth seeing the whole shape:

- Disclosure policy. Which services you run with AI, and how you tell clients. A direct requirement of Article 50 of the EU AI Act.

- Client rights. What a recipient of AI services can demand: to know it's AI; an explanation of the AI's decision; a switch to a human; safety; a quick appeal; and so on.
- Service principles. Written once here — every passport inherits them.
- Your liability. What you stand behind. Remember: "the AI did it" doesn't get you off the hook with the client.
- Quality criteria — set out both from the customer's side and from the organization's.
- Register of AI services, and how passports get approved. Which services run on AI; how each passport is written, signed off, and revised.
- Appeals against AI decisions — the form, the deadlines, the root-cause review, and updating the passport after a complaint.
- Machine-readable disclosure — a service card (about 13 required fields) for the client's AI agent.
- Improvement policy — how the service learns, gathers feedback, and gets better.
- Independent review — the ground rules for an outside audit of what the AI agent "promised" versus what it "did."

You won't need most of this for your first service — but every item earns its place, and together they're your shield against client lawsuits and the AI's own missteps.

Action 6 (later). Once you pass ~10 services, write the Rules for developing and delivering AI services.

Step 7. Launch and Monitor the Pilot

Here you answer the question that worries every leader most: does the AI actually do what it promised — and what happens when it doesn't?

Quick recap: together we've learned to separate services from functions, list your services, and pick a pilot; we've covered the levels of AI agent autonomy; and we've sorted out the passport (and, for later, the Rules). With those documents in hand, it's easy — on your own or with developers — to make them the backbone of the AI agent's prompt. Thanks to them, the agent stops being a black box and turns transparent, accountable, and under your control.

All that's left is to put the AI agent live and start watching how it works. Let me point out what to watch for and how best to gather information.

Monitoring that changes nothing is not monitoring.

What to watch. What you're monitoring is how closely practice matches the passport. You're hunting for the gap between what the passport promises and how the AI agent actually behaves. Log every deviation — bad and good alike. Pay special attention to:

- the usual triggers for escalation — where the agent handed off to a human, and where it didn't but should have;
- your customer-profiling categories;

- what quality looks like from the customer's side;
- what quality looks like from the organization's side;
- and build a library of edge cases.

And decide in advance what happens when it fails. When the agent hits something it can't handle, it should stop, hand the case to a human, and log it — never improvise. A clean failure is a handoff, not a confident wrong answer. The loop: find the gaps, make the fixes and check again.

How to gather it. Plenty of proven methods:

- Test purchase (mystery shopping) — experiencing the service as a client would (partial or full, one-off or repeated).
- Customer surveys and built-in feedback — how the client actually felt; the agent can even ask ("how did I do?").
- Peer panels and shared experience — companies in nearby fields trade practices and compare AI service passports.
- Independent assessment — a specialized outside audit.

Action 7. During the pilot, measure what was promised against what's delivered, find the gaps, and turn them into improvements.

To Sum Up

Go back to the ten questions you couldn't answer on page one. Run through them now and you can answer everyone. That's the whole difference: between owning an AI service and hosting an unmanaged risk.

Most people see AI as a technology. I see it as a new social mechanism, a new kind of social interaction. In a few years, everyone will have their own AI agent that can deal on our behalf not just with companies but with government — getting permits, licenses, certificates. Sure, some judgment calls and in-person steps will stay. But hand AI even a third of that bureaucracy, and at scale it adds up to billions of hours saved for people and businesses. That's what we should all be aiming for.

This market still has to be built and learned together, by the people who use AI services and the people who provide them. And it all starts small: with one honestly described AI service. It's a vast new ocean of opportunity, and hundreds of professions that don't exist yet.

So, where to begin bringing AI into your organization today:

- Split every process into functions and services.
- List your organization's services. Set the functions aside.
- Pick one service: simple and self-contained, low-stakes and reversible.
- Choose the AI agent's autonomy level: inform → intake → decision prep → full cycle → AI2AI.
- Build the AI Service Passport: general provisions, standard, procedures.
- At scale, write the Rules for developing and delivering AI services.
- Launch the pilot and start monitoring your AI service.

Don't put it off, August 2 and the AI Act rollout are right around the corner. Thanks for reading this playbook all the way through! I'd love to hear your thoughts and suggestions.

Book a free 20-minute consultation: info@aibusiness.vc

We'll sit down together and map your services, find where AI belongs and where it doesn't, and turn that into a roadmap you can act on. I'm not selling you a template. What I offer is the whole process: putting your services in order so you can put AI inside them safely, transparently, and on your terms.

May AI be with you.