

Sergei Ponomarev

Your Company's AI Through the Customer's Eyes

Five requests from your most demanding customer:

1. First, ask me what I need.
2. Then prove that your AI benefits me.
3. Next, let's agree on the rules.
4. Talk to me like a human being.
5. Check yourselves through my eyes.

aibusiness.vc

2026

Contents

Introduction. I Am Your Most Demanding Customer	3
Chapter 1. Ask the Customer First.....	5
1.1. My Needs. There Are More of Them Than You Think	5
1.2. Levels of Collaboration. From "Respondent" to "Co-author"	7
1.3. How to Ask So I Don't Lie.....	9
Chapter 2. Why Would Your Customer Need Your AI?	12
2.1. My Three-Number Test	12
2.2. The AI Guillotine: A Presumption of Corporate Inefficiency	13
2.3. Time Is Relative: Your Instant and My Eternity	14
Chapter 3. Agreeing on the Rules.....	17
3.1. The Service Passport: The Rules Your AI Works By	17
3.2. Is Any of This Actually Required by Law?	18
3.3. Let Me Show You: A Sample AI Service Passport	20
Chapter 4. Talk to Me Like a Human Being	22
4.1. Speak Plainly.....	22
4.2. AI as an Occasion to Care About Me	23
4.3. People, Not "Personnel Costs"	25
Chapter 5. How to Evaluate AI	28
5.1. Build an Evaluation System, Not a Collection of Likes.....	28
5.2. The Mystery Shopper: Hire Me	29
5.3. Living Rules: Revise Them Together With Me.....	30
Conclusion. Free Up My Chair	33
Appendix. The Demanding Customer's Final Checklist	34
About the Author	37

In Word: press Ctrl+A, then F9 to build the table of contents.

Introduction. I Am Your Most Demanding Customer

Hello. I am your worst nightmare: your most demanding customer. Pleased to meet you.

I am the person who pays you money, the one your company works for.

Lately, though, I have noticed that many businesses are adopting AI without asking my opinion at all.

They are busy choosing an AI model, hiring a vendor, calculating the cost, estimating how much they can save on staff, and so on. Fair questions, no argument. But the list almost never includes the question they should start with: what will I, your customer, think about all this?

However advanced, modern, economical and marvelous your AI may be, if I, a mere meatbag, do not like it, you will have problems. According to MIT, 95% of generative AI pilot projects end with zero results. And the AI in those projects worked just fine. Something else failed: the customers did not accept it.

Jeff Bezos understood this better than anyone. He built Amazon around customer obsession and kept one empty chair at the table in important meetings. In that chair sat "the demanding customer," the most important employee (!!!) in the room. Decisions were always made with an eye on that chair.

You are in luck. If you are reading this guide, it means I have come to you, bringing the customer's view of AI adoption in your company. I will help you look at the task through your customers' eyes. Use AI as a tool and an occasion to show you care about them, and they will repay you in hard cash.

And to make it easier for you, I will take this heavy role upon myself: I will be your collective, universal and, fair warning, maximally demanding customer. If I like it, everyone will. I am going to tell you unpleasant things, for which I apologize in advance! I will explain what I need, what irritates me, what I am willing to pay for, and what will make me leave and never come back. You can, of course, ignore me. That is your right. Just don't act surprised later.

What gives me the right to think I can pull off this role? My rather unusual professional experience at the intersection of business, academic research and civic policy.

My name is Sergei Ponomarev, and I hold a PhD in political science. For more than 20 years I have been studying how digital services are created and delivered in practice. In 2014 I defended a dissertation on e-government adoption, completed a fellowship on government openness at the Kennan Institute in Washington, DC, and taught at a university for 14 years.

For 7 years I led applied research at a think tank: I ran a nationwide monitoring program for government and municipal services, conducted independent quality assessments and mystery shopping audits to uncover the real problems people face. I developed an Information Openness Standard for public authorities. I worked with several human rights organizations and projects focused on consumer rights.

For 10 years I have been running trainings and seminars, teaching and advising entrepreneurs on service quality.

For the past year and a half I have been building AI agents hands-on and running the portal aibusiness.vc.

In other words, I spent two decades making sure that the services people receive are open, understandable and convenient. First in the public sector, now in business.

Why this matters right now. A large part of what I am going to ask of you in this guide will soon stop being a polite request.

On August 2, 2026, the transparency requirements of the EU AI Act¹ come into force: businesses in scope will need to say honestly when a customer is talking to AI and to label AI-generated content. Fines run up to 15 million euros or 3% of turnover. In the US, similar laws are already in effect in California, Utah, Maine and Colorado². So consider this a friendly warning: everything I ask for, you will have to do sooner or later anyway. Better to do it before your competitors do, and with a smile.

How this guide is organized. Five chapters, five of my requests:

- First, ask me what I need.
- Then prove that your AI benefits me.
- Next, let's agree on the rules.
- Talk to me like a human being.
- Check yourselves through my eyes.

And at the end I will hand you a checklist of all my requirements for your AI.

Read with a pencil in hand, argue with me, adapt things to your company. I am demanding but fair: every request will be explained and backed by a method for meeting it.

Portal: aibusiness.vc

For collaboration: info@aibusiness.vc

Let's begin. And we begin, naturally, with me.

¹ Regulation (EU) 2024/1689 (EU AI Act), Article 50. Full text: eur-lex.europa.eu/eli/reg/2024/1689

² State AI disclosure laws: California SB 243 (2025), Utah Artificial Intelligence Policy Act SB 149 (2024), Maine chatbot disclosure law (2025), Colorado AI Act SB 24-205 (2024).

Chapter 1. Ask the Customer First

You would not believe how rarely this happens.

Companies spend millions on AI adoption and almost never do the cheapest and most useful thing: they do not ask the customer what they actually need.

And it is not as if you had to go far to find us. We call you, write to you, browse your websites, leave reviews. We are practically shouting about our needs every single day. You just have to know how to listen.

In this chapter I will explain what customer needs are made of and how best to work with us at the AI adoption stage. And yes, how to ask so that I do not lie to you. I lie quite often, out of politeness. We will talk about that too.

1.1. My Needs. There Are More of Them Than You Think

Let me start with an uncomfortable truth about our relationship. When I come to you for a service, we are not equals. You made the rules. You are more experienced than I am, because you sell hundreds of services every day. You designed the interface. You decided what your bot can do. Nobody ever asked me whether I want to play by your rules. Sound familiar?

So yes, we look at your business differently. You judge a service by the result: order delivered, case closed, box ticked. I judge it by the process. If the order was completed but I spent half an hour begging your bot to transfer me to a live agent, I consider the deal a failure and your company terrible.

To avoid that, memorize my basic needs. They fall into three large groups. This classification is not from a textbook: it comes from years of service quality research involving thousands of real customers. People everywhere want the same things. Only the scenery and the operators change.

First group: functional needs. Where, when and how.

The place of service. Often the place of service is the company's office. Everyone has a rough idea of what a good office looks like. You do, right? I would genuinely love to hear it.

With AI, the assistant also needs its own "place," which is usually a channel and an interface. Where does your assistant live: on the website, in the app, in a messenger? I am comfortable where I already am, not where it was cheapest for you to put it. What does the AI assistant's "office" look like: is the font large, is the button visible, how many menus do I have to click through to ask a simple question? A cluttered chat with three carousels and ten buttons is like a store where the goods are piled in a heap and the salesperson hides in the back room.

Time. When the service is available and how long it takes. Here AI has a natural trump card: it does not take lunch breaks, does not get sick and does not go on vacation. Can I solve my problem at night on a Sunday? Wonderful. But opening hours are not the whole story.

There is also the duration of the service: how long the process itself takes. A bot that replies instantly but runs me in circles for forty minutes is available around the clock. Around the clock, and badly.

The manner of delivery. How many actions I have to perform. Every extra button, every "please register first," every "fill out this 12-field form" filters out people like me. We are not lazy. We simply remember that your competitor needed three clicks.

You have surely heard of the "one-stop shop" principle. As a customer, I want to solve all my problems in one place. If the bot takes my order but cannot answer my questions, cannot track the order status, and cannot accept a complaint because I have to "go" to a separate form on the website, then you do not have a smart assistant. You have a convoluted quest game.

And yes, staff: a live human must remain in the loop. Not hidden behind seven buttons, but reachable in one tap. I will have a separate conversation with you about this in Chapter 3.

Second group: informational needs. What is going on?

Accessibility. Before I use your smart bot, I first need to understand what it can do, what it cannot do, and what authority it has: can it process a refund or only chat about one? How long do I wait for a human if the bot fails? An accessible service means simplicity and clarity for the customer, without fog. Uncertainty irritates me more than rejection: a refusal I can understand, fog I will never forgive.

Confidentiality. What happens to what I tell you? Where does the chat log go, who reads it, does your AI train on it, who do you share it with? As long as you stay silent, I assume the worst. That is the kind of person I am, and the law, by the way, is on my side.

Reliability and safety. Can I trust what your AI says? If the bot quoted one price and a different one appeared at checkout, that is not a "neural network hallucination." That is misleading the customer. Your AI's promise equals your company's promise. Not ready to answer for its words? Then how can it represent you?

Contact and feedback. Can I reach you, do my problems get solved, what happens to complaints? Note: handling problems and complaints is my informational need, not your "negativity management." I need information: was I heard, what is being done, when will it be fixed. Silence after a complaint infuriates me more than the problem itself.

Third group: emotional needs. Do not drive me crazy!

This is where all those synthetic bots fall apart. I do not demand smiles from AI, but politeness, empathy and understanding are mandatory.

Respect for dignity. Human dignity is hard to describe in words, yet it registers instantly in your gut the moment it is violated. If the bot has not understood the question on the third try, it should offer to call a human, not keep cheerfully answering off-topic. If the bot failed, it is better to apologize like a person, not to "regret any inconvenience caused."

Precision and accuracy. A precise service means getting it right the first time. Not after I have repeated myself twenty times, but the first time. Do not mangle my name, my address, my grandmother's maiden name or my dog's name. Do not lose my order. Do not answer a question I never asked. Accuracy is a form of respect, and I register its absence instantly.

Satisfaction. The final feeling from the interaction: my problem is solved, not "the ticket is closed." You would be surprised how many companies close tickets without solving the problem. The bot has reported success, the statistics look great, and I am left alone with my issue. Know what I do next? I write a review. That review.

Now, lest I be accused of consumer terrorism, an important disclaimer that will save you money. Not all of my preferences need to be researched.

Some of them have long been written into law: the right to full information about a service, to safety, to protection of my economic interests, to a working complaints procedure. That is the baseline, enshrined

in consumer protection laws. You do not need focus groups to discover it. You just need to comply with it. Everything above the baseline is tied to your specific service, your audience and your market. That is what you should ask me about.

Write down this life hack. My preferences are perfectly visible in my behavior, in the traces your AI already records: at which step I abandon the dialog, which question I ask ten times, at what hour I storm into the chat, and after which answer I leave without a word. It is all already in your logs. You just have to look.

And one last thing. Write it somewhere in large letters. I am not interested in your internal regulations. I do not care which model you use, what your pipeline looks like, how your CRM is set up or what KPIs your support team has.

I judge only what is on my side of the counter: my functional, informational and emotional needs. Judge yourselves by them too. That is just how I am. And by the way, so are you, whenever you buy from someone else. Evaluation through the customer's eyes always follows the external procedures, not the internal ones. Your kitchen is your business.

How exactly to watch, ask and involve me, I will explain in the next two sections.

1.2. Levels of Collaboration. From "Respondent" to "Co-author"

Let me reassure you right away: I do not expect you to dance around me over every button. I am an adult, and I understand you run a business, not a hobby club.

So here is the deal. The depth of our collaboration depends on three things:

- how much time you have,
- how much you want me to like the result,
- what is at stake.

A bot that tells me your opening hours is one thing. An AI that decides whether to approve my installment plan is quite another. The higher the stakes, the higher you climb the ladder. And the ladder is simple, just three steps.

Step one: ask me first.

A survey. A short survey of your customer base. Two or three questions in the chat after an interaction. A dozen phone calls. A week of time and a little money.

What you get in return: my priorities instead of your imaginary KPIs. Here is a typical story. A company plans to deploy a voice robot for taking orders, spends six months choosing a vendor, haggles over the price. Then someone finally thinks to ask the customers. It turns out we were basically fine with how orders were taken. What drove us mad was something else: nobody could tell us where the order was and when it would arrive. That is an entirely different project. And, incidentally, a cheaper one.

When this is enough: simple services, low stakes, routine questions. The bot tells people the schedule, order status, delivery terms? Step one, and do not overthink it.

Just, I beg you, ask before the launch, not after. A survey after the launch is not called "preference research." It is called a post-mortem.

Step two: let me try it.

A consumer test. This level is for higher stakes. Show me the prototype before launch. Sit five to seven customers in front of a screen, hand us your bot to tear apart, and stay quiet. Just watch where we stumble.

Let me tell you a secret about why this works better than internal testing. Your testers know the design, how things are supposed to work. They verify the plan, and everything checks out. I do not know how things are supposed to work. I type with typos, start from the middle, ask two questions in one message and drop in a voice note where nobody expected one. I have a talent for breaking other people's scenarios. We all have that talent, and there are millions of us.

What you get in return: a list of failures before the launch instead of after. Feel the difference in money. A failure found on a prototype costs you a prompt edit. A failure I find in the wild costs you a customer, a review and a social media pile-on.

When this is mandatory: the AI stands in the customer's main path. Ordering, booking, personal account, support. Everywhere I walk every day.

Step three: involve me.

Partnership. This level is for serious stakes: AI for the long haul, AI at the core of the business, AI whose decisions touch my money. One-off surveys are not enough here. You need a standing connection.

What it looks like. A panel of regular customers who see every new version first. A beta testers club with early access or a discount. Regular meetings between your team and living customers: we talk, you take notes. A customer council, if you want a formal name.

What you get in return, and pay attention here, because this is the second most valuable benefit in this entire guide. A customer who was involved in creating a service starts to consider it their own. I forgive minor bugs, because I know they get fixed. I do not leave silently, I write to you, because I believe I will be heard. And I defend your product in the comments more fiercely than your PR person, because I am defending something of mine. Calculate what one such brand advocate is worth. And a panel can hold ten of us.

When you cannot do without it: the AI makes decisions about me. Scoring, personalized pricing, approvals and rejections, claims handling. Every case where a mistake hits my wallet, not just your statistics.

How to choose your step

Ask yourself one question: what happens if the AI makes a mistake? If I shrug and rephrase, step one will do. If I get angry and leave, you need step two. If I seriously lose money, time or nerve, climb to step three and do not skimp.

And one small observation to finish. Companies that reach the third step almost never climb down. Not because they are stuck. Because they have tasted it: permanent access to a living customer turns out to be cheaper and more accurate than any analytics bought from outside. Trust a man who spent twenty years on both sides of these studies.

Which step are you on right now? If the honest answer is "none," do not worry. Read on. The next section is about how to ask correctly. Because asking is a skill too: ask the wrong way, and I will lie to you. Here is how to avoid that.

1.3. How to Ask So I Don't Lie

I must warn you about the main trap. If you ask me point-blank, "Will you like our new AI assistant?", I will answer, "Yeah, probably." And I will be lying.

Not out of malice. I feel awkward upsetting a person who is glowing while describing their project. Everyone does this. Even your mom will say the idea is wonderful. Especially your mom. There is a whole book about this called The Mom Test³, and its core idea is simple: people lie to you out of politeness whenever you ask in a way that makes lying easier than disappointing you.

Remember three rules for talking to me.

First: talk to me about my life, not about your idea.

Second: ask about specific things from my past, not my opinions about the future. Otherwise I will sincerely tell you I will use your bot every day, and then never open it once.

Third: talk less, listen more. The more you talk in the meeting, the less you learn about me.

Compare and feel the difference:

- Bad: "Would you use a chatbot for booking?" Good: "Tell me how you booked with us last time. How many attempts? What annoyed you?"
- Bad: "Which assistant features do you need?" Good: "What did you do when you couldn't reach us by phone? Walk me through it."
- Bad: "How do you like our idea?" Good: "What did you overpay for, wait for, or drive somewhere for last month when you didn't want to?"
- Bad: "How much would you pay for such a service?" Good: "How much did this problem cost you last time? In money, hours, nerves?"

Memories of my real suffering are worth more than fantasies about your bright future.

And there are three kinds of garbage you will hear from me that must be thrown away, not carried into a report.

Compliments: "great idea, well done." Pleasant, useless, forget it.

Fluff: "I would definitely," "I always," "I never." Anything that starts with "I would" is worth nothing.

Ideas: "you should also add voice input and a dark theme."

My ideas are not a development backlog. They are a signal to dig: "Why do you want voice input? Tell me about the last time you needed it." Sometimes a real pain hides behind an idea, but it is never where I am pointing.

Now the arsenal. Five tools, from cheap to expensive. Each has its strength, and none replaces the others.

The log. I already mentioned logs. Start here: it is free, and you already have it. My preferences are visible in my traces: at which step I abandon the dialog, which question I ask ten times, at what hour I storm into the chat, after which answer I leave and never return. Behavior does not lie. I may write in a survey that everything is fine, but if eight out of ten people like me abandon checkout at step three, then step three is where your trouble lives.

The method's weakness: logs tell you where, but stay silent about why. For the why, move to the next tools.

³ Rob Fitzpatrick, The Mom Test, 2013.

My complaints and inquiries. Free gold sitting in your desk drawer. Export the last three months of customer messages: what people complain about, what they ask ten times, where conversations break off. The top 10 recurring questions are a ready-made brief for your AI, written by me personally, free of charge and without consultants. By the way, you can have the same AI analyze the export. Here is the prompt:

Here is an export of customer inquiries for the last __ months: [paste].

1. Group the inquiries by topic and count the frequency of each.
2. Write out the top 10 recurring questions verbatim, exactly as customers phrase them.
3. Flag where customers express irritation and what caused it.
4. Highlight questions that took more than a day to answer.
5. Compile a list of the customer's 5 main pains in their own words.

Interviews. Ten conversations of twenty minutes each. Expensive in time, but this is the only tool that answers the question "why." In a conversation I will tell you things you would never think to ask: that before calling you I call your competitor, that my daughter fills out the form on your site because I cannot read the small print, that the word "offer agreement" in your bot scares me. No survey will catch this, because a survey has no field for what you did not anticipate. The rules of the conversation are above: my life, my past, listen.

Surveys. A tool for validating with numbers what you already heard in interviews. First hypotheses, then the questionnaire, not the other way around. Keep it short: five to seven questions, more than that I will not endure. And I beg you, no "rate your satisfaction with our service level on a ten-point scale." Normal questions sound like this: how do you usually contact us, how many times did you have to repeat your question, what should we fix first. The strength: numbers you can lean on when deciding. The weakness: a survey answers only the questions you thought to ask.

Observation. Watching a live customer in their natural habitat is quite fascinating. This is the consumer test from step two. Sit me in front of your prototype and stay silent. Do not prompt me, do not make excuses, do not say "oh, that part isn't ready yet." Just note where I got stuck, what I muttered, and at what moment I reached for my phone to call the old-fashioned way. Painful? Very. More useful than any report? Absolutely. One hour of such observation saves a month of rework.

Where to start if money is short and time is shorter: logs and complaints today, they are free. Then ten interviews. That is enough to avoid missing badly with your first AI service. Surveys and tests can wait until there is something to validate and something to test.

Chapter 1 Checklist. Check Yourself

No.	Question	Done
1	We know the three groups of customer needs: functional, informational, emotional, and the nine dimensions within them	☐
2	We exported and analyzed complaints and inquiries from the last three months	☐
3	We have the top 10 customer questions, recorded in their own words	☐
4	We looked into the logs: where customers abandon the dialog and after what they leave silently	☐
5	We conducted at least ten interviews by the rules: about the customer's life, about the past, listening	☐

6	We separated facts from compliments, fluff and ideas	☐
7	We chose our step of collaboration: ask, let them try, involve	☐
8	A list of the customer's five main pains, in their own words, hangs in front of the team	☐

Eight items. Closed them all? Congratulations, you now know more about me than ninety percent of your competitors know about their customers. And now let's find out something unpleasant: whether you need AI at all. Don't rush to take offense. The next chapter may save you a great deal of money. We will be counting by my yardstick.

Chapter 2. Why Would Your Customer Need Your AI?

Allow me to ask a question no vendor will ever ask you. Everyone around is adopting AI, and so did you. Well done, congratulations. But what is in it for me?

Seriously. If you deployed AI for internal matters, to cut costs, replace staff, speed up paperwork, that is your business. Manage it as you wish. My opinion is not required.

But the moment you place AI between yourself and me, everything changes. Now, to get your service, I am obliged to talk to your bot. Not "allowed to," mind you, but obliged: you often leave me no other door. And since you made that decision for me, at least answer one question: will this make things better or worse for me? Not for you. For me.

The question is not rhetorical. It has a precise numerical answer, and in this chapter I will give it to you. Then I will show you what to check before any automation, so you do not end up accelerating nonsense. And at the end I will reveal a secret: why your dashboard shows three seconds while I have been fuming for two hours. We will be counting by my yardstick. Fair warning: my yardstick is harsh. But honest.

2.1. My Three-Number Test

Let me start from a distance. For many years the World Bank compared countries using the Doing Business methodology⁴. The idea was brilliantly simple: look at the state through the eyes of an entrepreneur. Take a government service, say, opening a small firm, getting a construction permit or connecting electricity, and count three things.

- How many procedures you have to go through.
- How many days it takes.
- How much it costs.

In one country a new business opened in a couple of steps and a single day. In another it took a dozen offices and a month of waiting. Countries fought for years over their positions in that ranking. Notice what was measured: exactly what a person feels first-hand. Steps, time, money. Nothing else concerns them.

Well, I evaluate your company exactly the same way. Every day, free of charge and without any ranking. I have the same three indicators, and I want you to start counting them before you sign the AI contract with a vendor.

The first indicator: my steps. How many actions I must perform from the moment "I need this" to the moment "I got it." Everything counts: finding where to write to you. Registering. Confirming my email. Filling out the form. Attaching a document. Calling because the form did not work. Repeating my question to an agent because the bot passed nothing along. Every handoff, every "now go to another section" is a step. Steps are good at hiding: inside your chat there seem to be two, but along the entire path from my need to the result there are eleven.

The second indicator: my time. How long the whole story lasts. Not your internal "average handling time," but my time: with the searching, the waiting, the repeat visits and the "your call is very important to us." The service does not begin when I write to your bot. It begins when my need arises, and it ends when the problem is solved and I have forgotten about it. Everything in between goes on my meter.

⁴ The World Bank, Doing Business reports (2003-2021).

The third indicator: my costs. Everything I spend beyond the price of the service itself. Fees and surcharges. A paid support line. Taking a day off to visit your office because "this matter is handled in person only." Printing, signing, scanning. And let's be honest, my nerves belong here too: they do not show up in your books, but they show up beautifully in my reviews of you.

What to do with this? Take one of your services, the most popular one. Walk its path as a customer, from need to result, and write down three columns of numbers: steps, time, costs. That is your "before."

Then take the AI project and honestly forecast the same three numbers "after." If the AI is already running, it is even simpler: walk the path again and compare. You already know where to get the data from Chapter 1: logs, complaints, ten interviews, watching a live customer.

From here, arithmetic decides everything. There are exactly three possible outcomes.

Outcome one. My numbers went down. Seven steps became five. Three hours became ten minutes. "Go to the office" became "solved it with the bot over coffee." Congratulations, this is the case that justifies adopting AI in the first place. Launch boldly. I will say thank you, come back and bring friends. Notice that I never even asked which model you use or how many parameters it has. I do not care. My numbers went down. The rest is detail.

Outcome two. My numbers stayed the same, but yours improved. You cut support costs, unloaded your people, and I felt no difference: I used to solve my problem in a day, and I still do. Fine, that is legitimate. Then one request: do not tell me in your advertising that it was all done "for my convenience." I can count, and that kind of falseness is felt instantly.

Outcome three. Your numbers went down and mine went up. The most common and the saddest outcome. I used to call, and a polite living person solved my problem in five minutes. Now I circle a bot for fifteen minutes while it fails to understand the question, then wait for an agent, then tell the whole story again because the bot passed nothing along. Your report shows savings. My meter shows that I am paying for your savings with my own time. That will not do. I do not need that kind of AI, and I will find a way to let you know: with a review and with my wallet.

Write this down.

AI adoption is justified for the customer when it reduces their steps, time and costs.

At least one of the three numbers, without worsening the others. Everything else is either your internal optimization or my degradation at your expense. The first is acceptable. The second I will not forgive.

And do not panic if the test shows something unpleasant. It is far cheaper to see it on paper now than in the reviews later. The next section explains what to do with the steps the test caught. Spoiler: not all of them deserve acceleration. Some deserve a funeral.

2.2. The AI Guillotine: A Presumption of Corporate Inefficiency

Every company accumulates ballast over the years. Forms nobody reads. Approvals that approve nothing. Reports nobody needs. Steps that exist for a single reason: "that's how we've always done it." You no longer notice this ballast. You are used to it. I notice it every day, because I wade through it to reach your service.

And here is the thought this whole section was written for. AI adoption is the best excuse for a general cleanup your company will get in a decade. You were going to rebuild your processes anyway. So seize the

moment: do not drop a bot on top of what you have now. Upgrade the company. I call this maneuver the AI guillotine: AI adoption as an occasion to improve my service, which is the whole point.

The essence of the guillotine is a flipped lens. Normally a process lives until someone proves it is unnecessary. And nobody rushes to prove it: every process has an owner, the owner has a habit, and the habit has an iron argument, "we've always done it this way."

The AI guillotine flips the burden of proof: from now on, every process is presumed unnecessary until proven otherwise. A presumption of inefficiency for your operations. Every form, every approval, every step walks up to the whiteboard and explains why it exists. Failed to explain? Under the blade. No mercy and no farewell speeches.

Everything is cut in two passes, and the order matters. First remove what should not exist at all: duplicated and obsolete demands on customers. Then honestly divide what remains between humans and AI. In that exact order, because you cannot automate ballast. Or rather, you can, but then you are paying money to make nonsense happen faster. And removing it later will cost twice as much: by then it will have a budget, an implementation report and a defender with a slide deck.

The moral: the AI guillotine is a conversation about the meaning of the work, not about technology, and it must be held with the employees who deal with me directly. And there is no better time than now, while the window of opportunity is open.

So you do not think this is pretty theory, one number from my own practice. I personally took part in such a "cleanup" for a city government. We analyzed more than a hundred services and cut their combined delivery time by more than five hundred days, and six services were eliminated entirely.

Five hundred days sounds modest until you remember that a service is not received by one person. Even if each of those services is used by just a thousand people a year, the city now waits one and a half thousand years less. Every year. Fifteen hundred years of human life that used to be spent waiting and have now been returned to people. Feel the scale.

Now, why do I, your customer, root so passionately for your AI guillotine? I have skin in this game. Three kinds of skin, actually.

The first interest is obvious: my savings. Most of the steps that will not survive the guillotine are my steps. My forms, my waiting, my "please send the document again." You cut the ballast, and my life gets lighter.

The second interest is pragmatic. The money for bad processes and for quality services comes out of the same pocket: mine. Every coin you spend feeding ballast is a coin you did not invest in me: in speed, in people, in the very AI that actually solves my problems. You feed the chaos less, I get better quality.

The third interest is human. The guillotine cuts activities, not people. Free your employees from routine, and they will finally have time for what you cannot delegate to a machine: complex cases, attention, a live conversation, empathy. With me, by the way.

And I know how to be grateful. I repay care in the same coin: I return, I buy, I bring friends. That is how this works. In the next section I will show you why, even after the cleanup, your reports may keep telling you fairy tales: time on your stopwatch and time on mine run at different speeds.

2.3. Time Is Relative: Your Instant and My Eternity

To finish the chapter, a magic trick you will encounter so often that it deserves its own section.

Picture a management meeting. Your team shows a dashboard: after deploying the AI assistant, the average response time is three seconds. Three seconds! Applause, bonuses all around. Now the same story through my eyes.

Say I want to change my order. First I spent 5 minutes hunting for where to even write about it: in my account, in the cart, in the contacts. Found it. Then I spent 10 minutes explaining to the bot what I wanted. The bot cheerfully replied in three seconds about my order, but not about the change. I rephrased. It again replied in three seconds, about my order. I typed "agent." It offered to rate the conversation. After 15 minutes of waiting a live person joined, and I spent another 5 minutes telling them everything from scratch. Total: 40 minutes of my life. On your dashboard, my nightmare looks like a series of brilliant three-second responses.

Neither of us is lying. We simply hold different stopwatches. Yours starts when I write to the bot and stops when the bot replies. Mine starts when my need arises and stops when the problem is solved and forgotten. Everything in between lands on my bill: the searching, the waiting, the repetitions, the transfers, the "tell me from the beginning."

This is not an AI disease. It is the old disease of all services. A car service shop boasts: "the part swap takes twenty minutes." And for the customer: book a week ahead, leave work early, drive over, wait in line, wait for the procedure, wait for the receipt, drive back. Twenty minutes of work cost the customer half a day of life. Three seconds of your AI mean 40 minutes on my planet. Interstellar in action.

Why this is dangerous in practice. The metric you pay bonuses for starts to govern behavior. Reward response speed, and you will get fast responses. Fast does not mean useful: the bot will lightning-quick answer the wrong question, instantly close dialogs and swiftly ship me off to "the help section." Your stopwatch will set records. Mine will spin like a taxi meter. And then someone at a meeting will wonder why customers are leaving despite such wonderful metrics.

So write these down:

Measure the customer's entire journey: time from need to resolution.

The whole path, from my first action to the solved problem. Not average response speed, but how long the entire story took.

The share of issues resolved on the first attempt. Truly resolved: I got what I wanted and did not come back with the same question. A closed dialog and a solved problem are different things. We have covered this.

Repeat contacts. How many times I came back with the same problem. Every repeat means that last time you reported a resolution that did not happen.

The number of transfers and repetitions. How many times I was passed from hand to hand and how many times I told my story from scratch. Every "tell me from the beginning" subtracts from my love for you.

How do you see your real time? There is one way, and it is free: walk the path yourself, or ask someone from your household. Do not ask your team. They will show you the dashboard. Take your phone, imagine you need to reschedule a delivery, return an item or move an appointment, and walk the entire path as an ordinary customer. With the button hunt, the bot, the waiting. Time it from "I need this" to "done." It is the cheapest audit of your life and, I suspect, the most sobering. Incidentally, in Chapter 5 I will teach you to make this check regular and systematic.

For now, let's wrap up the chapter. You have learned to count AI's benefits by my yardstick, held a general cleanup and now own two stopwatches instead of one. Check yourself.

Chapter 2 Checklist. Check Yourself

No.	Question	Done
1	We counted the customer's three numbers: their steps, their time, their costs. Before adoption and after	☐
2	We honestly identified our outcome: better for the customer, the same for the customer, worse for the customer	☐
3	We ran the guillotine: removed unnecessary steps and processes before automation, not after	☐
4	We divided what remained between humans and AI, after listening to the people who do the work	☐
5	We measure the customer's time from need to resolution, not the bot's response speed	☐
6	We track repeat contacts and "tell me from the beginning"	☐
7	Someone from management walked the customer's path personally, stopwatch in hand	☐

Seven items. If they are all closed, you hold something rare: an honest answer to the question "why does our customer need our AI." Most companies have a vendor's slide deck instead of an answer.

Now that we have established your AI is genuinely useful to me, it is time to agree on the rules of the game. What your bot promises me, what I am entitled to demand of it, and where its authority ends. That is Chapter 3, and it is about the most underrated word in business: the word "standard."

Chapter 3. Agreeing on the Rules

So, we have confirmed that your AI genuinely helps me: it cuts my steps, time and costs, and the ballast has been thrown out. Now, before I start using it, let's agree on the rules of the game.

Because right now, hand on heart, I am playing against the casino. Only you know what your bot can do. You decide how long I wait for a reply. When to call a human is determined by an algorithm I have never seen. And when something goes wrong, you shrug: "technology, what can you do." No. Since you have placed a machine between us, put down in writing what it owes me. I call this the standard, or the service passport, and it is the most underrated word in business.

3.1. The Service Passport: The Rules Your AI Works By

In Chapter 1 we mapped out my needs. Now a trick question: how do you make sure your AI actually meets them? Every day, in every dialog, including the one that happens at three in the morning on a Sunday?

The answer is duller than you would like and more reliable than you would expect: rules. My needs must be sewn into the rules of service delivery, and the rules must be written into a single document: the service passport. Your employee has a job description. And your bot, which talks to thousands of customers a day, lives by gut feeling? Not acceptable. Every AI service gets a passport.

The passport is not technical documentation. Which model you use, which prompts, how the integrations work, you may write it down, but as a rule I am not interested. That is your kitchen. The passport describes not the bot's insides but the rules by which I can receive a proper service. That is: what I am entitled to, what I am not, and what I will sue you over.

Here is what I, your customer, want to see in the passport of every one of your AI services:

- What the service is, who provides it and to whom.
- What I will receive and by when. The result of the service and the time you promise to deliver it.
- What the bot can and cannot do. The full list of actions and, more importantly, the boundaries: it processes returns but does not settle payment disputes. Honest boundaries save my time.
- What data the bot will ask of me. An exhaustive list. Anything beyond the list, the bot has no right to request.
- What it costs. If the service is paid, the price is in the passport, not a surprise at the end.
- On what grounds I can be refused. A closed list. "Denied" without a reason from the list is not a refusal.
- When the bot must call a human. Pre-defined cases: I asked directly; the bot failed to understand twice; the case is atypical; the conversation turned into a conflict. Pre-defined, not "at the algorithm's discretion." This is my emergency exit, and I want to see it on the map.
- What gets recorded. Every bot decision is logged: what it decided, by what criterion, when. A black box, like on an airplane. If we argue, we pull the recording.
- A check before delivery. Before handing me the result, the bot verifies it against the passport. Not sure? It does not deliver. It calls a human. Better to wait a minute than to receive confident nonsense.
- Where to complain and how fast you will answer. The channel, the deadline, the form of the response.

Look closely at this list: it is my needs from Chapter 1, translated into rules. The functional ones became deadlines and boundaries, the informational ones became lists and grounds, the emotional ones became the emergency exit to a human. The circle has closed: you asked what I need and wrote it into the document the machine runs on.

Now three rules that make a passport real rather than decorative.

Rule one: the passport is public. It hangs where I will find it: next to the chat, on the service page. Written in human language, briefly. Not "your request is being processed," but "we are looking into it and will reply within a day." Every piece of officialese in the text reads to me as "we can't be bothered to talk to you like humans." A passport buried in a forty-page user agreement exists legally, but not humanly.

Rule two: word equals deed. A gap between the passport and the bot's actual behavior is not a "documentation inaccuracy." It is misleading the customer, with everything that follows. Wrote "a human joins within two minutes"? Kindly make it two.

Rule three, my favorite: you answer for the AI. The phrase "it was the algorithm's error, we are not to blame" does not work. The bot quoted a price? That is your price. The bot promised a deadline? That is your promise. A company may not hide behind its AI, just as it may not hide behind an intern: you hired it, you answer for it.

And a small glimpse of the future, so you understand where all this is heading. In a year or two it will not be me coming to you for a service, but my personal AI assistant: to compare prices, place an order, ask questions. Your beautiful website means nothing to it. It needs the service passport in machine-readable form: what you do, what you require, what it costs, where the boundaries are. Whoever has a passport, my agent will strike a deal with in seconds. Whoever does not, it simply will not understand, and it will move on to your competitor. So the passport is not bureaucracy. It is your future language for talking to an entire new class of customers in the agent economy.

And for you personally the passport is convenient for one more reason: it is a ready-made technical brief. "Build us a good bot" is not a brief. A service passport handed to a vendor is a brief you can accept work against and hold someone accountable to. As a bonus, disputes with me turn into arithmetic: promised in the passport, delivered, case closed.

3.2. Is Any of This Actually Required by Law?

I am impressed by how stoically you have endured this consumer assault of mine. Surely you have already asked yourself how many more wishes this guy has. Do we really have to do all this? Yes and no. The service passport, at the time of writing, is still my consumer wish. But some of it is already becoming mandatory. And take it from a political scientist: this list of obligations will only grow.

On August 2, 2026, Article 50 of the European AI Act comes into force: the transparency chapter. If your business serves European customers, it may well fall in scope, and part of my requests stops being just requests on that day. The idea of the law is simple and sounds suspiciously like what I have been repeating this entire guide: a person has the right to know when they are dealing with a machine or with something a machine produced.

Here is what the law requires and how it maps onto your service passport:

Legal requirement	The line in your passport
A person must clearly understand they are talking to AI	The bot introduces itself at the start: "You are chatting with an AI assistant." Not "sweet Anna," not "a tired

	agent," a soulless silicon algorithm
AI content must carry machine-readable marking	Texts, images, audio and video created by your AI carry a "made by AI" mark that software can detect
Deepfakes must be disclosed	Synthetic voices and generated faces in ads are honestly labeled as artificial
AI texts for the general public must be flagged	Articles and materials written by AI on matters of public interest disclose it
Emotion recognition: banned at work and in education, elsewhere requires notice	If your AI analyzes my tone and mood, it says so plainly. And the HR tool reading employees' emotions is switched off entirely

Now, why it is time to start taking me seriously.

A transparency violation costs up to fifteen million euros or three percent of turnover. Prohibited practices, like that same emotion recognition in the office, up to thirty-five million or seven percent. Note: the percentage is calculated on global turnover, not European revenue. For a company with fifty million in revenue, one such fine devours the entire year's profit. The giants have the fat to survive it. A mid-sized business does not.

And do not exhale with the thought "this is about big tech." The obligations flow downhill, to those who deploy AI. Put someone else's bot on your website? You are in scope. Running ads with a synthetic voice? You are in scope. The law distinguishes between those who built the model and those who use it, but transparency toward the customer is demanded of both.

The only relief: systems already on the market before August 2 got a reprieve until December 2, 2026, and only on one point, the machine-readable marking of content. Everything else, including "admit you are an AI," lands on August 2 with no extensions.

What I advise doing before August. Inventory every point where AI works for you: bots, content generation, synthetic media in ads, tone analysis. Add an explicit introduction to every dialog: one line at the start, cheap now, expensive later when the complaint arrives. Label your generated marketing. And write all of it into the service passport, because AI lives by an iron rule: what is not in the instructions does not exist in its synthetic life. Your human employee intuitively knows that pretending to be someone else is wrong. Your bot knows nothing until you write it down.

Look at the table once more: the law demands exactly what I asked of you in the first chapters out of goodwill. Introduce yourself honestly. Do not pass off machine-made things as living ones. Do not read my emotions in secret. If you build the service out of care for me, you will have met the spirit of Article 50 before you ever opened its text, and sharply reduced your compliance risk. And you gain something money cannot buy: in an era when nobody is sure what is real, trust becomes the hardest currency there is.

The test is simple. Imagine a regulator walks in on August 3 and asks you to show how every AI touchpoint of yours discloses itself. Can you show it in five minutes? If not yet, you know what to do. There is still time. Not much.

3.3. Let Me Show You: A Sample AI Service Passport

Theory is nice, but you are surely already picturing a two-hundred-page tome, an approvals committee and six months of work. Exhale. For a simple service the passport fits on one page and takes one day to write, together with the employee who runs this service by hand today.

Let's take the simplest, most common service, the one you should start with: providing information. A language school, say, puts an AI assistant on course questions. Here is its passport in full, as I see it.

AI SERVICE PASSPORT: Information about the school's courses

1. General provisions. Service: information about courses, schedules, prices and enrollment. AI agent autonomy level: information only. Enrollment and payment are handled by a staff member. Who may apply: any visitor to the site or messengers. Result: an accurate answer and a clear next step. Introduction: the assistant opens with "I am the school's AI assistant." Feedback: a "Complain" button in the chat, a human reply within one business day.

2. The standard: what we promise. Covers: the list of courses and formats; schedules; prices and discounts; how to enroll and what happens next. Channels: website chat, messengers, email. Timing: instant replies in chat, waiting no longer than a minute; email within one business day. What it does not do: does not judge your language level "by eye," does not give pedagogical advice, does not discuss other clients. Grounds for refusal: the question is not about the school (we will point you in the right direction); requests for other people's data. Price: free. Quality: answers strictly from the approved standard, nothing "improvised"; the question is closed in full; handover to a human preserves the context.

3. The AI agent's procedure. Steps: introduce itself as AI, clarify the question, answer strictly by the standard, offer a next step or a human. Must call a staff member if: the client asked for a human; it failed to understand twice; the topic turned to enrollment, payment or refunds; the person showed strong emotions. May never: invent things absent from the standard; quote prices or dates not in the standard; give legal or medical opinions; request data beyond the necessary. Logging: every dialog is recorded: question, answer, whether it was handed to a human, the client's rating.

Machine-readable card (for customers' AI agents):

```
service: Course information
provider: Language school [name]
description: courses, schedule, prices, enrollment
channels: chat, messenger, email
price: free
response_time: instant; email <= 1 business day
languages: en, [your languages]
disclosure: service is provided by an AI assistant
escalation: enrollment, payment, conflict, on client request
contact: [link and phone]
updated: [date], version 1.0
```

Service owner: [name, title]. Version: 1.0, dated [date].

That is the whole passport. One page, and look how many of my requirements it closed at once. I know I am talking to a machine. I know what it can and cannot do. I know the price it quoted is real, because improvisation is forbidden to it in writing. I know my emergency exit: two misunderstandings, and I am handed to a human instead of being run in circles. And I know every conversation is recorded, so a dispute, if one arises, will take a minute.

Notice three things.

First: the procedure section, once filled in, goes into your agent as its system instruction. The passport is not a paper about the bot. It is the bot: the document and the instruction coincide, word equals deed by construction.

Second: the machine-readable card at the end. Remember I said my AI assistant will be visiting you soon? The card is a shop window for it: the fuller the card, the more willingly other people's agents will recommend you to their owners. Consider it free advertising in a channel your competitors have not heard of yet.

Third: the passport has an owner with a first and last name. Not "the team," not "the department," but a person who answers. To me, the demanding one, this matters more than everything else.

Start with one service, the simplest one, like this one. Fill in the passport, launch, watch the logs, adjust. Then take on the next. In a couple of months you will have something almost nobody on the market has: an AI that works by rules and by customer preferences.

Chapter 3 Checklist. Check Yourself

No.	Question	Done
1	Every customer-facing AI service has a passport: general provisions, standard, agent procedure	☐
2	The passport is published where the customer will find it and written in human language	☐
3	The passport fixes: what the bot can and cannot do, deadlines, required data, price, a closed list of grounds for refusal	☐
4	The cases when the bot must call a human are spelled out: customer's request, two misunderstandings, money, conflict	☐
5	The bot introduces itself as AI at the start of the conversation; generated content and deepfakes are labeled (Article 50, deadline August 2, 2026)	☐
6	Employee emotion recognition is switched off; customers are warned about tone analysis	☐
7	Every bot decision is logged: what it decided, by what criterion, when. Before delivering a result the bot checks it against the passport	☐
8	The passport has an owner with a first and last name, and any gap between the passport and the bot's behavior is treated as an emergency, not a "nuance"	☐
9	The machine-readable service card is published for customers' AI agents	☐

Nine items. If they are all closed, congratulations: your AI works by rules that I have seen and can rely on. By my demanding observations, such companies are rare birds, so you have just overtaken not ninety percent of your competitors but ninety-nine.

But I still have a question bigger than the rules of one service. A service is a service, but what about your AI across the company? How much of it is there, where does it live, who answers for it, what happens to my data in all that machinery? That is the next chapter: how to talk about all of this like a human being. To me and, don't be surprised, to your own employees.

Chapter 4. Talk to Me Like a Human Being

You have built the AI agent, written the service passport, and we have agreed on the rules. Now let's talk about how you present all of this. To me, your customer. This stage is critical: present it wrong, and I walk. And by the way, have you forgotten your employees? Don't be surprised: your team also needs the AI news delivered properly, to defuse their anxiety and potential sabotage, so that I, once again, don't end up the casualty.

Let's start with the foundation. There is no point informing anyone if you do it in a language that needs translation. Let me explain about clarity. It is a sore spot.

4.1. Speak Plainly

Open your bot's chat history with customers and look for phrases like "your request is being processed," "the provided data is undergoing verification," "under the current terms of service." Found them? Your AI speaks officialese, a dead language no living person uses. I write to the bot "where is my order," and it answers me like a bureaucrat replying to a complaint. It feels as if I hadn't tried to buy something but had violated something. This is genuinely infuriating.

Officialese used to have an excuse: lawyers wrote the texts, to be safe, and rewriting them was expensive. Your bot has no such excuse. A language model can speak simply, warmly and clearly. It is literally its native talent. If your AI talks like a tax inspector, someone set it up that way or was too lazy to set it up otherwise. It was a choice, not a limitation of the technology. And when I read "your request is being processed," I feel that choice on my skin: you couldn't be bothered.

The rules of human language are simple. Here are the six main ones:

- Short sentences. One thought per sentence. If a phrase has three commas and two "whiches," chop it up.
- Human words. Not "effectuate payment" but "pay." Not "in the event that questions arise" but "if anything is unclear." Not "functionality" but "what it can do."
- Active voice. Not "a refund will be issued" but "we will return your money." In the first version the money returns itself, mysteriously, and there is no one to ask. In the second there is a "we," and that "we" answers.
- Lead with the point. The answer to my question in the first line, details below. Do not make me read three paragraphs of preamble to learn "yes" or "no."
- Concrete numbers. Not "shortly" but "by the end of the day." Not "a fee may apply" but "a 2 percent fee." Fog, I remind you, I do not forgive.
- No internal jargon. Your "tickets," "escalations" and departmental abbreviations stay in the department. I have an order, a question and a desire to get an answer.

For clarity, before and after the translation into plain language:

Before	After
"Your request has been registered and will be reviewed within the established timeframe"	"We got your question, we will reply by 6 pm"
"To effectuate the refund of funds it is necessary to provide payment details"	"To return your money, send us your card number"

"At the present moment the given item is absent from the warehouse"	"This item is out of stock, back on Friday. Want me to hold one?"
---	---

Feel it? In the second version I am talking to someone who cares, even if it is a machine.

The test is simple: read the bot's reply out loud, as if talking to a neighbor. If it cannot be said out loud, rewrite it. And to avoid rewriting by hand, give the job to the bot itself. Here is the prompt:

Here are our AI assistant's standard replies: [paste].
 Rewrite each one by the rules of human language:
 short sentences, active voice ("we will return", not "will be returned"), the point in the first line, concrete dates and numbers, no officialese or internal jargon. Tone: a friendly adult, no baby talk. Output a table: before / after.

Half an hour of work, and I will notice the effect from the first message. And here is the important part: plain language is not "dumbing it down," as some experts think. It is respect in its most practical form: you spend your effort to save mine. And it is in this very language that you will need to write everything covered in the next two sections: the story for me and the story for your team.

4.2. AI as an Occasion to Care About Me

Let me share a secret that will save you marketing budget. I, your customer, am deeply indifferent to your technology. Neural networks, models, agents: those are your toys. But I care to the point of trembling about how you treat me. Which is why one and the same event, "we have adopted AI," can be presented in two ways. The first: quietly install the bot and wait until I stumble upon it and draw my own conclusions. The second: turn the launch into a vivid gesture of care for me. Guess which way most companies choose. And guess which one pays off.

Here is what care sounds like. Not "we have implemented an innovative AI solution to optimize customer service." That is about you and your pride. But like this: "Dear customers, you are the most important people to us, so we built an assistant that answers you within a minute at any time of day. The live team hasn't gone anywhere: one click, and you are talking to a person. Our staff member [name] is responsible for the assistant. Write to them if anything is off. And send us ideas for improvement: we really do read them." The same event. The difference is who stands at the center: you or me. When it is me, I feel it, and I remember.

There are more channels for this story than you might think. Pick by your audience:

- A letter to customers: "we adopted AI with care for you," two paragraphs, warm and without pathos.
- A social media post in the same tone.
- A short announcement page on the website.
- A notice right in the chat: "An AI assistant is replying. Need a human? Tap here." This one, by the way, is also a legal requirement. Remember Chapter 3.
- A client cheat sheet: how to use the assistant, three steps.
- An FAQ with honest answers: is it really AI? is my data safe? how do I reach a human?
- And if you have a physical location, hang a sign at the entrance: "AI works here: it helps faster, a human is always nearby."

Visitors smile, take photos and, most importantly, are not startled when a machine answers them. Each of these formats takes an evening of work. Together they turn a technical event into a story about how you thought of me.

By the way, a gift, so you cannot say I only make demands. Here are eight ready-made "AI works here" stickers, drawn for my client communication kit. Take them: put one on your website, in your chat widget, on the office door, in your newsletter footer. No licenses, no small print. The full editable kit with all the templates lives at aibusiness.vc.



Now further. I received your letter and came to your website. What should I find there?

One document. Public, prominent, titled, say, "How We Use AI."

Note: not a page for a specific bot, but the story of your entire AI household at once. Because the passport we wrote in Chapter 3 is sewn into a specific service, and you may have many such services: a chat bot, an

in-app assistant, a voice line. Before using any of them, I want the full picture: which AI services you have at all, what they can broadly do, what happens to my data, who in the company answers for all of it, and where to complain if something goes wrong. Plus a last-updated date, so I can see the document is alive, not a monument to the year 2025.

And here I want to share an engineering trick that will earn you gratitude not just from me but from your own developer. It is called the "single source of truth," and we used it when drafting service regulations. At first, contact details were sewn into every regulation of every service, but whenever an email or phone number changed, dozens of documents had to be rewritten. So the frequently changing information is better moved into one shared document that all the regulations simply reference. The change is made once, in one place.

The same story awaits your bots. Today you have one bot, and you have calmly sewn into its prompt the name of the responsible person, the support email and the working hours. In two years you will have twenty bots, maybe a hundred. And then the employee quits, the email changes, the office moves. What, rewrite a hundred prompts? Why. All the information that can change: names, contacts, hours, links, prices, lives in one place, in that same public AI document. The bots do not store it. They pull it from there. Change it in one place, and it changes everywhere. Your document for customers and your reference for bots turn out to be the same file, and the company's word, outside and inside, matches by construction.

And one more reason to build a single source of truth, the most far-reaching one. Soon it will not only be me visiting your site but my personal AI agent: I will task it with finding the best deal, asking questions and preparing the purchase. Your beautiful design means nothing to it. It does not look at pictures. The first thing it will look for is who it can talk to here in its own synthetic tongue. It will want to know which chat bots and assistants you have, who handles what, where the machine-readable service cards are, where its requests should go. For it, the single AI document is not an "about us" page but the entry point for negotiations. Whoever has one, my agent will close a deal with in seconds. Whoever does not, it will honestly skip and move on to your competitor. So you are writing this document not only for today's me, but for tomorrow's stream of customers brought by their agents.

A beautiful structure has emerged. Appreciate it: the passport describes one specific service, the single source of truth describes the company. I learn about your AI before meeting it, communicate with it by clear rules, and always know the information is accurate and current. One audience remains, the one almost everyone forgets during AI adoption, and wrongly so: the people who work for you, your employees. It is your business, of course, but if something goes wrong with the AI, it is them I will be switched to, and then I will yell at. So they too deserve a warning. And herein lies the main value of this guide.

4.3. People, Not "Personnel Costs"

Your people serve me, and I feel their fear through the counter as clearly as I feel officialese in the chat.

Here is a scene playing out in thousands of companies right now. A bot was deployed. Nobody really explained anything to the staff, but everyone has read the news about "AI will replace millions of jobs." So I call, and a live manager confides: "Don't bother with that bot, it's useless. Let me sort you out like a human being." Recognize it? That is not a bad employee. That is a frightened employee defending their job the only way they can: by quietly sinking your bot. And I am the one standing between two fires, and it is not cozy here.

Now let's talk honestly about the root of that fear, which is to say, about you, dear executive. I know what crossed your mind when you heard about AI. Payroll is the fattest line of expenses, and the thought of "replacing everyone with synthetic bots" is tempting to the point of trembling. All the more so because the report looks superb to investors: costs cut, well done, bonus for the board.

So allow your demanding customer to say something rude about this strategy. Rude, but arithmetically precise. You can only cut costs down to zero. You can grow endlessly. The agenda of "fire everyone to save money" is a false agenda, and here is the question that kills it: what will you do with the money you saved? Put it on deposit? Saving is not a strategy. It is a one-time trick. Growth is a strategy.

Look at the real trap of thinking here. Most executives see AI as a way to digitize their current level: do the same things, but cheaper. Few ask the next question: what level could we move up to thanks to AI? What new services could we offer customers that used to be beyond our reach: not enough hands, time, money?

Remember the AI guillotine from Chapter 2: we cut the fat not for the sake of saving as such. We were clearing space and freeing resources for growth. AI gives you processes that used to be fundamentally inaccessible. It is a unique chance to start doing what you have never done. Scary? Yes. But trading that chance for "we laid off three operators" is like buying an airplane to haul potatoes to the market next door.

And now about your people. Listen carefully, because I am speaking as the customer. Your employees know me better than any bot. They have worked with me for years. They know my moods, my emotions, my quirks that are written down nowhere. They remember what I look for, what I complain about, what my business is, when my season comes and what I can be offered. The bot has none of this knowledge. It is not in the CRM. It is in heads and fingertips. In the atmosphere that forms at the moment of a deal.

Fire the people, and you throw away the most valuable knowledge in the company: the knowledge of me, your most important person. The right scheme is exactly the reverse: routine to the AI, the customer to the human. Let the machine fill out forms and answer the routine, while a living person guides me, watches over me, talks to me warmly and offers what a bot would never think of.

And one more thing that changed quietly and radically: every one of your employees can now be a builder. Assembling an AI agent today is easier than mastering a spreadsheet with macros. Someone on your customer team has surely been sitting on an idea for five years: "we really should also do this, but there's no one to do it." Well, now there is someone: her, with AI in her hands. Ask your team not "who do we cut" but "what new things can we now offer customers." The difference in the question is the difference in the company's fate.

How to tell all this to the team, short and to the point:

First: formulate the principle and repeat it. We are cutting activities, not employees. Routine goes to the machine, people move to more interesting work closer to the customer. Saying it once is not enough. People believe deeds: the first employee freed from routine who receives a new role, not a termination letter, will convince better than a hundred town halls.

Second: train them. What the bot may do, what it may not, which data must never be uploaded, how to check AI answers, how to hand a customer from the bot to a human so the customer does not retell everything. An hour at the start, a short refresher every quarter, onboarding for newcomers. And by the way, this is no longer just my advice: since February 2025, European law directly obliges companies to

ensure adequate AI literacy⁵ of staff working with AI. While everyone was discussing the 2026 deadline, this obligation had already arrived.

Third: open an idea channel. Let anyone propose: "let's have the AI do this." Half the ideas will turn out to be gold, because your people see what neither you nor consultants see: the customer's real life.

Fourth: appoint an owner for the AI household. The company needs at least one person for whom this is written into the job description, not into their enthusiasm: collecting employees' ideas and turning the best into working bots, writing briefs for vendors and accepting their work, maintaining the service passports and that single source of truth, training the team, following the news of models and laws, which, as you have noticed, update more often than price lists. In a small company it is one employee with dedicated hours, in a large one a whole department. Call it the AI competence center. And I, the customer, benefit directly: when something goes wrong, I know the company has a specific person who will learn about it and fix it, not "the thing sort of runs itself."

Chapter 4 Checklist. Check Yourself

No.	Question	Done
1	The bot's replies pass the "read it aloud to a neighbor" test: no officialese, the point in the first line, concrete deadlines	☐
2	Customers were told about the AI adoption as an act of care: what it gives them, a human nearby, a named person responsible	☐
3	Clear channels were used: a letter, a post, a notice in the chat, a cheat sheet, an FAQ	☐
4	The website has one public document about all the company's AI: services, boundaries, data, responsible person, where to complain, last updated	☐
5	Changeable information (contacts, names, prices) lives in a single source of truth, and the bots pull it from there instead of storing it in prompts	☐
6	The team has been told the principle "we cut activities, not employees," and it is backed by deeds	☐
7	Employees are trained: the rules, forbidden data, checking answers, handing a customer to a human (a legal requirement since February 2025)	☐
8	An idea channel is running: what else can we offer customers with AI	☐
9	A head of the AI competence center is appointed: an employee responsible for ideas, bots, passports, training and news	☐

In sum: it takes plain language, an honest story and respect for two groups of people: those who buy from you and those who work for you. And now the final chord of this guide: let's check how all this beauty actually works. Through my eyes. Chapter five: complaints, mystery shoppers and rules that live.

⁵ EU AI Act, Article 4 (AI literacy), in force since February 2, 2025.

Chapter 5. How to Evaluate AI

So, you did everything right: asked me, counted the benefits, wrote the AI service passport, presented it to everyone like a human being. The bot works, the dashboard is green, the team is happy. One last question remains, the trickiest one: how do you know it works?

Just do not show me the dashboard again. We established back in Chapter 2 that your stopwatch and mine run differently, and that "dialog closed" does not mean "problem solved." Internal metrics tell you how the machine labors. The truth about service quality lives on my side of the counter. This chapter is about building an evaluation that looks through my eyes: without a budget for sociologists, without reports for the drawer, and without self-deception.

5.1. Build an Evaluation System, Not a Collection of Likes

Let me start with the main rule, which will save you a great deal of time and money. Before collecting any number, any review, any rating, ask one question: what decision will we make based on this data? If there is no answer, do not collect it. Data without a decision is not analytics. It is stamp collecting. Ask yourself: what happens if we do not have this information? If the honest answer is "nothing," then there is no reason to gather it.

Hence the second rule, twin brother of the first: it only makes sense to evaluate what you control. If a parameter cannot be changed, measuring it is entertainment. You control the bot's response speed? Measure it. My mood on a Monday? Leave that to the astrologers.

Now, what exactly to evaluate. Everything is already in place. We built it in the previous chapters. It just needs assembling. The evaluation of an AI service stands on two pillars:

The first pillar: are the promises kept? You have a service passport, and a passport is a ruler. Promised a reply within a minute: what share of dialogs made it? Promised a human in one click: how many clicks did it actually take? Promised not to ask what you already know: does it re-ask? Every line of the passport becomes a verifiable indicator. That, by the way, is why I insisted so hard on measurable promises: you cannot put a ruler against the slogan "we strive for the best."

The second pillar: did things get better for me? Remember my test from Chapter 2: steps, time, costs? It works after launch too. Plus my feeling: is the problem solved, will I come back, will I recommend you. The first pillar without the second is dangerous: you can flawlessly keep promises I do not need.

How to collect data without going broke. A small business does not need sociology with hundred-person focus groups. There is an approach called lean data⁶, built for verifying impact without big budgets.

Four principles: ask the customers themselves, not experts; collect exactly as much as the decision requires; use simple free tools; ask briefly and often, not lengthily and once a year. A five-question survey sent on time beats an annual study lying in a drawer.

And what to ask me, so as not to annoy me or waste my time. Here is a five-question backbone as a gift. Adjust it to your needs:

- Did you manage to solve your problem with the assistant? If not, what got in the way?
- How much time did the whole thing take, from "I wanted" to "I got"?

⁶ Acumen, The Lean Data Field Guide: acumen.org.

- Compare it with how things were before: simpler, the same, harder?
- Would you recommend us to friends? Why?
- What should we fix first?

Pay attention to the third question. It is my favorite. Evaluation always needs a comparison: with how things were, with a competitor, with expectations. "Did you like it?" is a question about nothing. People answer "yes" out of politeness. We covered that in Chapter 1. "Did it get simpler than before?" is a question with an honest answer.

What remains is to assemble a system out of this, meaning add a calendar and an owner.

- Continuously, by itself: logs (where dialogs get abandoned, where people call for a human, where the bot re-asks) and a one-click micro-rating after each dialog.
- Quarterly: a short survey of the base using those same five questions, plus a check of every promise against the passport.
- Twice a year: the exercise described in the next section. Patience.

Every number, I repeat, must end in a decision: you saw that a third of dialogs break at step three, you fix step three and look again a month later. You saw it and sighed: consider that you never saw it.

That, essentially, is the whole infrastructure: the passport as a ruler, five questions as a thermometer, a calendar as discipline, and one rule on top: no decision, no measurement. And now the tastiest part. There is an evaluation method that I, your demanding customer, love most tenderly of all. So much that I am ready to take part in it personally. Next section: the mystery shopper.

5.2. The Mystery Shopper: Hire Me

A confession of love. Of all the evaluation methods I have seen in twenty years, this is my favorite. It is cheap, honest and hits dead center. It is called the test purchase, or mystery shopping, and if you take only one method from this whole chapter, take this one.

The essence is indecently simple: walk through your own service as an ordinary customer. Not as the executive who knows where every button is and what was meant. As a person off the street: anonymously, with a real need, from "I need this" to "I got it," writing down everything that happens along the way. The view not from behind the counter, but from my side of it. The difference, as you will discover, is enormous.

Why this method seems made for AI services. First, testing a bot requires neither agents nor a budget: no actors to hire, no branches to drive to. A laptop and half an hour will do. Second, you can repeat it every week if you like. The bot will not recognize the inspector and start trying harder. Third, and most important, a live walkthrough catches what logs will never show: dead ends with no exit, a tone that sets your teeth on edge, and that feeling of "I am nobody to them," which no metric records but which decides everything.

How to run a proper test purchase, five steps:

First: invent the legends. At least three scenarios. The typical case: an ordinary question from an ordinary customer. The atypical one: a situation not covered in the FAQ. And the conflict one: you are unhappy, you write sharply, you demand a human. It is on the second and third scenarios that bots crumble, and it is exactly these that nobody checks before launch.

Second: prepare a scorecard. What to record: how many steps, how much time, where you stumbled, the bot's verbatim replies, whether it introduced itself as AI, how many times you had to rephrase, whether a

human was called and whether the context was passed along. And crucially: next to every item, the corresponding line from the service passport. Remember I said the passport is a ruler? The test purchase is the act of laying the ruler down. Promised "a human within two minutes"? Start the timer.

Third: walk the path for real. Behave like a living person, not a tester: type with typos, start from the middle, ask two questions at once, drop in a voice note. Do not prompt yourself and do not forgive the bot with "well, we haven't finished that part yet." The customer does not know what you have not finished, and does not want to know.

Fourth: record the evidence. Screenshots, screen recording, times by the clock. Not for a court, for the debrief: "the bot was rude" sounds like an opinion, while a screenshot of the rudeness is a fact nobody can argue with at the meeting.

Fifth, the decisive one: consequences. Write this in large letters: a test purchase without follow-up actions is a wasted half hour. Every finding becomes a decision: what we fix, who fixes it, by when. And a month later, a repeat purchase along the same scenario: did it get fixed, or did we just have a nice chat. Purchase, debrief, fix, verify. A cycle, not an event.

Who should do the shopping? In ascending order. The best first shopper is you, the executive: half an hour in the role of your own customer sobers you up better than a quarterly report, as generations of directors can confirm. Next, an employee from another department who does not know "how it is supposed to work." Then an acquaintance who is not in the loop at all. And for regular full checks you need an outside person with a proven method and a fresh eye. You could invite me, the author of this guide. I have, shall we say, some experience.

And finally, the thought this whole section was written for. Every customer of yours is already a mystery shopper. The customer walks your service every day, notices every dead end and every rudeness, keeps a mental scorecard. The only difference is that you never receive the customer's report. Or rather, you do, but in the most expensive form there is: in reviews, in the silence of the departed, and in revenue that "somehow dipped." A test purchase is a way to get the same report earlier, cheaper and without the losses.

Run a test purchase regularly, every six months, and after every major update. The vendor silently updated the model? Your bot may have changed without asking you.

And now the final section of this guide. About the fact that evaluation is not a finish line but a circle: technology changes, I change, and your beautiful rules go stale faster than you think.

5.3. Living Rules: Revise Them Together With Me

Let me open the final section with bad news. Everything you are about to build from this guide will go stale.

Look at who is working against your beautiful documents.

Technology: the vendor updated the model, and the bot started talking differently, asking neither you nor me.

The law: we have seen how requirements arrive in waves, and August 2026 is not the last one.

And finally, me. My expectations grow mercilessly: what delighted me yesterday is the norm today and an irritant tomorrow. A bot answering within a minute felt like a miracle until your competitor started

answering in ten seconds. I did not get worse. I simply got used to it. Such is the nature of the most demanding customer, and there is no cure.

The conclusion: AI documents come in exactly two kinds, living and dead. A living document gets revised, it has a last-updated date, and I, by the way, look at that date. A dead one was written, framed and admired. A two-year-old service passport proudly marked version 1.0 tells me one thing: nobody has looked in here since launch. And a bot living by a dead passport lies to me on schedule.

How to keep things alive without bureaucracy. Two mechanisms.

The first: triggers, events after which a revision is mandatory. The vendor updated the model. You launched a new service. A spike of complaints on one topic. The law changed. If any of these happened, the owner of the AI household sits down and checks the passports against reality.

The second: a calendar. Once a quarter, a quick "promises versus facts" check based on test purchases and logs. Once every six months, a major revision: what to remove, what to add, what to change.

Note that the material for revision accumulates by itself: complaints, logs, surveys, test purchases. Speaking of complaints, write down the rule that is worth the most: every justified complaint ends not only with an apology to me but with a change to the rules. An apology treats the case. A rule change treats the cause. A company that apologizes for the same thing for the second year running is apologizing in vain.

And now the main thing. Revise not for me, but with me.

Remember the ladder from Chapter 1: ask, let them try, involve? Back then I was talking about the launch. Well, the ladder does not end at launch. It becomes permanent. A panel of a dozen regular customers who see each new version of the bot first. A beta club with early access. A short conversation with living customers before every major passport revision: what infuriates, what is missing, what should go. Call it whatever you like, a customer council if you want formality. The essence is the same: I take part in editing the rules by which I am served.

What is in it for you, besides my gratitude? This: I am the cheapest and most motivated expert available. A consultant takes the money and leaves. I stay and keep using the service. I notice breakages before your metrics do, because I live inside your service. And I am, damn it, personally invested in everything working: I am the one who has to use it. Give me a legitimate channel of influence, and I will stop expressing my opinions exclusively through reviews.

And so the circle of this guide closes. Appreciate the symmetry. We began with "ask me first." We end with "and never stop asking." Between those two phrases fit everything: the benefit test, the guillotine, the passport, the honest story, the trained team, the living evaluation. AI adoption turned out to be not a project with an end date but a relationship. And relationships, as we know, stay alive exactly as long as the people in them keep talking.

Chapter 5 Checklist. Check Yourself

No.	Question	Done
1	Every metric we collect is tied to a decision: we know what we will do when we see the number	☐
2	We evaluate only what we control	☐
3	The service passport is used as a ruler: every promise is turned into a verifiable indicator	☐

4	Once a quarter customers receive a short survey: was the problem solved, how long did it take, is it simpler than before, what should we fix	☐
5	Logs and the after-dialog micro-rating run continuously, and someone actually reviews them	☐
6	A test purchase is run every six months and after every model update, on three scenarios: typical, atypical, conflict	☐
7	Every test purchase and every justified complaint ends in a change: to the prompt, the passport, the rules	☐
8	All AI documents have a last-updated date, triggers and a revision calendar	☐
9	Customers are involved in revisions: a panel, a beta club or a regular conversation before updates	☐

Nine items. If they are all closed, your AI does not merely work. It grows up together with your customers. That is the rarest quality of all, and unlike models and prompts, competitors cannot copy it.

Conclusion. Free Up My Chair

Well, that is all, dear executive. Your most demanding customer has had his say.

Let's count what you now have. You know my needs, all three groups, and you know how to ask so that I do not lie. You can count AI's benefits by my yardstick: steps, time, costs, and you have run the guillotine, sending under the blade everything that stole our time for years. Your AI services have passports, and your promises have numbers. You talk to me like a human being, you presented the launch as an act of care, and you did not forget to motivate your own employees. And you built an evaluation that looks through my eyes, all the way to test purchases and living rules that grow up together with me.

Not bad for one guide. And notice: almost none of it required a "digital transformation" budget. It required something else, and frankly, something more expensive than money: the willingness to look at your business from my side of the counter.

Remember Bezos's empty chair from the introduction? Let me reveal the last secret. You do not need to buy that chair. It is already standing in your meeting room. It is just usually covered with folders of vendor presentations. Every time your team argues about models, prompts and budgets, look at that chair and ask the question our whole conversation started with: what will the customer think about all this? Now you know what I would answer. I sit in that chair permanently, even when it looks empty to you.

One confession before we part. I have been harsh with you, snide in places, and I never properly apologized. I will not apologize now either. Because every one of my jabs is actually an invitation: I badly want you to succeed. I am the one who has to use what you build.

There will soon be millions of companies adopting AI. Companies adopting AI in the customer's interest are still a handful. Join the short line. It moves faster.

In the appendix you will find the final checklist: all my requirements from the five chapters, gathered into a single list. Print it, walk through it, count the check marks. Then walk through it again every six months: as you will recall, I grow more demanding by the day.

Thank you for reading to the end and not taking offense. That, by the way, is the best sign that your company is going to be just fine.

Your demanding customer.

P.S. I will be glad to read your feedback: info@aibusiness.vc

Appendix. The Demanding Customer's Final Checklist

All my requirements from the five chapters, gathered into one list. How to use it: print it, walk through it honestly, count the check marks. A check mark is earned only by "done and working," not by "planned for next quarter." Then come back to the list every six months: as you will recall, I grow more demanding by the day.

There are exactly forty-two items. Science fiction fans know: that is the answer to the ultimate question of life, the universe and everything. In our case the question was: "do customers love your AI?"

Chapter 1. Ask the Customer First

No.	Requirement	Done
1	We know the three groups of customer needs: functional, informational, emotional, and the nine dimensions within them	☐
2	We exported and analyzed complaints and inquiries from the last three months	☐
3	We have the top 10 customer questions, recorded in their own words	☐
4	We looked into the logs: where customers abandon the dialog and after what they leave silently	☐
5	We conducted at least ten interviews by the rules: about the customer's life, about the past, listening	☐
6	We separated facts from compliments, fluff and ideas	☐
7	We chose our step of collaboration: ask, let them try, involve	☐
8	A list of the customer's five main pains, in their own words, hangs in front of the team	☐

Chapter 2. Why Would Your Customer Need Your AI?

No.	Requirement	Done
9	We counted the customer's three numbers: their steps, their time, their costs. Before adoption and after	☐
10	We honestly identified our outcome: better for the customer, the same, worse	☐
11	We ran the AI guillotine: removed unnecessary steps and processes before automation, not after	☐
12	We divided what remained between humans and AI, after listening to the people who do the work	☐
13	We measure the customer's time from need to resolution, not the bot's response speed	☐
14	We track repeat contacts and "tell me from the beginning"	☐
15	Someone from management walked the customer's path personally, stopwatch in hand	☐

Chapter 3. Agreeing on the Rules

No.	Requirement	Done
-----	-------------	------

16	Every customer-facing AI service has a passport: general provisions, standard, agent procedure	☐
17	The passport is published where the customer will find it and written in human language	☐
18	The passport fixes: what the bot can and cannot do, deadlines, required data, price, a closed list of grounds for refusal	☐
19	The cases when the bot must call a human are spelled out: customer's request, two misunderstandings, money, conflict	☐
20	The bot introduces itself as AI at the start; generated content and deepfakes are labeled (Article 50, deadline August 2, 2026)	☐
21	Employee emotion recognition is switched off; customers are warned about tone analysis	☐
22	Every bot decision is logged: what it decided, by what criterion, when. Before delivery the bot checks the result against the passport	☐
23	The passport has an owner with a first and last name, and a gap between the passport and the bot's behavior is treated as an emergency	☐
24	The machine-readable service card is published for customers' AI agents	☐

Chapter 4. Talk to Me Like a Human Being

No.	Requirement	Done
25	The bot's replies pass the "read it aloud to a neighbor" test: no officialese, the point in the first line, concrete deadlines	☐
26	Customers were told about the AI adoption as an act of care: what it gives them, a human nearby, a named person responsible	☐
27	Clear channels were used: a letter, a post, a notice in the chat, a cheat sheet, an FAQ	☐
28	The website has one public document about all the company's AI: services, boundaries, data, responsible person, where to complain, last updated	☐
29	Changeable information (contacts, names, prices) lives in a single source of truth, and the bots pull it from there	☐
30	The team has been told the principle "we cut activities, not employees," and it is backed by deeds	☐
31	Employees are trained: the rules, forbidden data, checking answers, handing a customer to a human (a legal requirement since February 2025)	☐
32	An idea channel is running: what else can we offer customers with AI	☐
33	An owner of the AI household is appointed: an employee or competence center responsible for ideas, bots, passports, training and news	☐

Chapter 5. How to Evaluate AI

No.	Requirement	Done
34	Every metric we collect is tied to a decision: we know what we will do when we see the number	☐

35	We evaluate only what we control	☐
36	The service passport is used as a ruler: every promise is turned into a verifiable indicator	☐
37	Once a quarter customers receive a short survey: was the problem solved, how long did it take, is it simpler than before, what should we fix	☐
38	Logs and the after-dialog micro-rating run continuously, and someone actually reviews them	☐
39	A test purchase is run every six months and after every model update, on three scenarios: typical, atypical, conflict	☐
40	Every test purchase and every justified complaint ends in a change: to the prompt, the passport, the rules	☐
41	All AI documents have a last-updated date, triggers and a revision calendar	☐
42	Customers are involved in revisions: a panel, a beta club or a regular conversation before updates	☐

Count Your Check Marks

0-14. The red zone. Your AI lives its own life, and I live mine, and our relationship rests on my patience. Which, I warn you, is not infinite: I am already eyeing your competitors, and my draft review is ready. The good news: you now know all forty-two steps to fixing it, and the first ones cost zero money.

15-28. The middle. Not bad at all, seriously. It shows that someone thought about me. But I find the holes faster than you patch them, and my loyalty so far rests on habit, not love. Pick the five cheapest of your unclosed items and close them within a month: you will feel the effect before you feel the cost.

29-41. A strong position. I recommend you to friends, and that, as you know, cannot be bought with advertising. Close the remaining items not for the check marks but for the last frontier: they are exactly what separates you from the best.

42 out of 42. A customer champion. I take off my hat and do what a demanding customer almost never does: praise without reservations. I defend your AI in the comments more fiercely than your PR person. One request: do not relax. In six months I will come back with this list. And it may well have grown longer.

How to Work Through the Checklist with Your Team

Do not fill it in alone in your office. Print a copy for everyone at your next team meeting: let each person count the check marks on their own, honestly and without peeking. Then compare. The spread of scores is itself a diagnosis: if the owner counts 30 and the support team counts 12, two different AIs live in your company, mine and yours. Discuss the gaps, pick the five cheapest unclosed items and agree who closes them and by when. Forty-five minutes, and the checklist stops being reading and becomes a plan.

About the Author

Sergei Ponomarev holds a PhD in political science and has spent more than 20 years making services open, understandable and convenient for the people who use them: first in the public sector, now in business.

He ran a nationwide monitoring program for government and municipal services, conducted hundreds of independent quality assessments and mystery shopping audits, developed an Information Openness Standard for public authorities, completed a fellowship at the Kennan Institute in Washington, DC, and taught at a university for 14 years. For the past year and a half he has been building AI agents hands-on and running the portal aibusiness.vc.

Website: aibusiness.vc

Email: info@aibusiness.vc



Scan to open aibusiness.vc