

THERE IS NO SPOON



CREATE AGENTIC AI PRODUCT STRATEGIES

Your competitors are quietly automating decisions while you're still building chatbots.

Last week, a fintech startup used an AI agent to cut their sales cycle by 40%. Not with a fancy dashboard or a clever chatbot, but with a system that qualifies leads, drafts contracts, and books meetings while their team sleeps.

Meanwhile, most companies are stuck in *'pilot purgatory'* burning cash on AI experiments that never escape the demo stage.

This book shows you how to break free.

Whether you're a founder, product manager, or operator, this book gives you a step-by-step method to:

- Spot high-impact use cases for AI Agents.
- Design agentic AI systems with memory, reasoning, and guardrails.
- Build and deploy agents that deliver measurable ROI in weeks, not years.

This isn't another trend. It's the next evolution of operational advantage. And if you don't build systems that think, you'll be outpaced by those who do.

.

Tim Daines

There is no Spoon

Create Agentic AI Product Strategies
Learn how to design autonomous AI
Agents

First Edition

Contents

A New Class of Product Thinking	6
The £1.2M Wake-Up Call	7
Beyond Chatbots: What Is an AI Agent?	9
My Journey: From AI to Gen AI to Agentic AI Systems	12
Why Agent Product Strategy Matters Now	16
Why listen to Tim	19
Why Most AI Strategies Fail	20
Mistake #1: Mistaking a prompt for a product	20
Mistake #2: Building hype, not solving friction	21
Mistake #3: No ownership, no direction	21
Mistake #4: Beautiful UI, no utility	21
Bringing It All Together - Purpose, Not Just AI	22
When To Use There Is No Spoon	24
How it all works	26
Day 1 Find the Agent opportunity	26
Day 2 Define and design the Agentic AI Strategic approach	27
How to use this guide	27
Your Next 48 Hours	29
Step-By-Step Guide To Agentic AI Product Strategies	31
When to use Agentic AI Product Strategy	32
Agent Opportunity Checklist	33
Overview of the workshop	33
Key outputs	34
Materials Needed	34
The Agent brief	35
Day 1	39
Introduction to Agentic AI Product Strategy	40
Here's how to do it:	40
Clarify the Objectives	40
Set Ground Rules	40
Introduce the Agenda	41
Preparation Reminders for Facilitators	41
Stage 0: Choosing a decision maker	42
Stage 1: Where does Work Break Down?	43
Resources	43
How To Deliver it	43
Facilitator Tip	45
Milestone	45
Stage 2: Voting on the Agent Opportunity	46
Resources	46
How To Deliver it	46
Milestone	48

Stage 3: How will we know the Agent is Working?	49
Resources	49
How To Deliver It	49
Milestone	51
Stage 4: Vote and Prioritise	52
Resources	52
How To Deliver It	52
Milestone	54
Stage 5: Where does decision-making need help?	55
Resources	55
How To Deliver It	55
Facilitator Tip	56
Stage 6: When should the Agent Act?	58
Resources	58
How To Deliver It	58
Milestone	61
Day 2	65
Playback - Day 1: Laying the Foundations for Agentic AI Product Strategy	66
What we Covered	66
Stage 7: Choose the Best Agent	67
Resources	67
How To Deliver it	69
Milestone	70
Stage 8: Agentic AI Product Roadmap	72
Resources	72
How To Deliver It	72
Milestone	75
Stage 9: Wrap Up & Build Readiness	77
Resources	77
How To Deliver it	77
Milestone	79
Post-Workshop Follow-Up	80
Ongoing Guidance & Optional Support	81
Final Thoughts	82
Mopping Up & Templates	84
You Did It — Now Don't Let It Sit on the Shelf	85
Templates	86
Thanks To Those Who Got This Here	89
Acknowledgements	90
References	90

A New Class of Product Thinking

The £1.2M Wake-Up Call

Six months ago, I found myself back in a familiar scene. A glass-walled boardroom, floor-to-ceiling screens aglow, every seat around the table filled with sharp minds. At the front, a proud product team was about to unveil their crowning achievement, what they called an *"AI insights assistant."* It was the kind of project that gets people talking. A slick, branded interface, a GPT-powered backend, and a name that looked great on investor slides, Insights+.

They'd spent close to £1.2 million on this thing. Designers had obsessed over every pixel. Engineers had wrangled LLM prompts late into the night. The launch event had been polished to perfection. Today, they were certain, was going to be the moment it all paid off. The demo ran smoothly. Dashboards flickered. Buttons glowed. The chatbot greeted us by name, made a friendly quip, and pulled up pages of business data on command. For a moment, everyone in the room was nodding. Impressed by the theatre of innovation.

Then the new CEO, who had been quietly watching from the end of the table, leaned forward. She didn't smile. She didn't congratulate the team. Instead, she cut straight to the heart of it:

"What's the ROI?"

The room went still. You could feel the air shift. Laptops sat open, but no one typed. Eyes darted around, looking for an answer. No one spoke. Tumbleweed. Minutes dragged.

The silence stretched. The CEO's next words cut deeper. *"So we spent £1.2 million to build a search bar?"*

Nobody corrected her. Because that's exactly what they'd done.

The proud smiles faded. It became clear. Despite the code, the effort, the fanfare, no one could say, in plain numbers, what value this *"AI assistant"* had delivered to the business.

This isn't an isolated story. Over and over, I've watched smart people pour time, energy, and cash into building clever technology, only to end up with digital showpieces that don't actually move the business forward. The trap is seductive. It feels like progress, but when the tough questions come, there's nothing to show except a glowing UI and a stack of invoices.

Three miles away, a 12-person startup faced the same challenge: drowning in contract renewals. But instead of hiring consultants or building platforms, they did something radical. They gave the problem to an AI agent.

Not a chatbot. Not an assistant. An agent with a job description:

- Monitor renewal dates.
- Draft contracts using historical data.
- Send follow-ups without being asked.
- Escalate only when genuinely stuck.

Cost: £500/month. Time to implement: 3 days.

The results were immediate, and unmissable:

- 22 hours saved, every single month.
- 17% increase in renewal velocity.
- No more midnight panic for the CEO at month-end.

They didn't just automate a task. They embedded intelligence into the backbone of their business, a system that quietly got on with the job, delivered measurable value, and let humans focus on the stuff that mattered. That's the difference that keeps me up at night.

The first team built technology, something that looked impressive, sounded innovative, but, at the end of the day, changed nothing fundamental about how the business ran.

The second team built a system that owns outcomes, a structure that delivered results, not just demos.

This is the gap I've spent the last 20 years trying to close. It doesn't matter how talented your engineers are, how much budget you can burn, or how many AI solutions you plug in. If you're not building for outcomes, if you can't answer that CEO's question in two sentences or less, you're building for show, not for value. I've seen this pattern across Fortune 500s and startups, in every industry you can name. Teams obsess over what they're building, but lose sight of **why** they're building it, and **who** is supposed to benefit. Too many products, not enough impact.

That's what There Is No Spoon is here to fix. It's a two-day workshop where you flip the script, you stop chasing tech for its own sake, and start building agentic systems to create AI Agents that relentlessly, reliably move the needle. No matter your team size, your tech stack, or your industry.

Ready? Let's get to work.

Beyond Chatbots: What Is an AI Agent?

Before we go any further, let's get crystal clear about what we mean when we talk about "AI agents" and "agentic AI", because if you've been anywhere near tech headlines, you've probably heard the term thrown around, often with little explanation.

Here's a 30-second test to know if you have an AI agent or just another chatbot:

Chatbot: *"What's the status of the Johnson renewal?" "The Johnson renewal is due in 14 days with a contract value of £45,000."*

AI Agent: *[System prompt needed] "I've drafted the Johnson renewal (£45,000, due in 14 days), applied their standard 5% increase based on usage growth, and scheduled it to send Tuesday at 9am when they typically open emails. Their procurement team changed last month"*

AI Agent: *I've updated the contact to Sarah Chen. Want to review before it sends?"*

See the difference? One answers questions. The other owns outcomes.

An AI agent is not just a chatbot with better manners. It's not simply a search box with a language model stitched in, or a productivity feature wrapped in a sleek interface. Most of what's called "AI" today is just a faster way to answer questions, generate content, or surface information, like a digital assistant that only wakes up when you ask, *"Hey, can you help with this?"* But an AI agent is something fundamentally different.

Think of it as an autonomous digital teammate, one you don't have to supervise, handhold, or constantly prod into action. It doesn't just wait for your instructions; it takes initiative. It observes what's happening in your workflow, remembers key context, makes decisions, and acts on your behalf, even when you're not watching.

Let's break that down:

- **It has a real job.** Not just a *"skill"* or a button to press. An AI agent owns a business outcome. For example: handling renewals, qualifying leads, onboarding customers, resolving support tickets, or even monitoring compliance.
- **It remembers context.** Unlike most chatbots, which treat every conversation as brand new, a true agent tracks patterns, recalls previous actions, and learns from feedback. If you've ever had to remind a bot what you talked about five minutes ago, you'll appreciate this.
- **It reasons and takes action.** A chatbot can give you a to-do list. An agent will actually go and do it. It might draft contracts, send reminders, process approvals, update records, or escalate issues, without pausing for your every *"okay, now do this next."*
- **It operates across tools and systems.** AI Agents don't live in a single app, walled off from the rest of your business. They move between your CRM, your email, your internal databases, and your analytics dashboards, pulling information from everywhere, taking action wherever it's needed, and tying everything together behind the scenes.

Imagine this. You're not in the office. Your agent sees a new contract is due. It:

- Checks the renewal history.
- Pulls in the latest usage data.
- Drafts a personalised email for your client.
- Updates your CRM.
- Pings the sales team for approval, and
- Schedules a follow-up call.

You get a summary in your inbox, no manual effort, no need to coordinate 10 different tools, no stress. That's not science fiction. That's what AI Agents are already enabling for teams that make the leap.

Here's the mindset shift:

- With chatbots and “*AI features*”, you're still the orchestrator.
- With AI agents, you're the architect, you design the system (with your domain experts and technical team), set the guardrails, and let the AI agent handle the day-to-day grind.

This is a step-change in how work gets done. Instead of just making old tasks a bit faster or prettier, you're embedding intelligence directly into your operations. You're turning workflows into living systems that quietly deliver results while you focus on the big picture.

So what about *agentic AI*?

Throughout this book, when I say “*agentic AI*”, I don't just mean a smart bot with initiative. I'm talking about building *systems*, complete, outcome-owning architectures where the AI agent is just the visible interface. Behind the scenes, an agentic AI system is a living, operational fabric that integrates memory, tools, governance, and reasoning. It's not just one agent running tasks, it's a cohesive digital infrastructure designed to observe, decide, act, and improve over time.

An agentic AI system isn't built like traditional automation. It has a runtime brain that coordinates tasks, a memory core that accumulates experience, tools that behave like skilled specialists, and a control layer that enforces business guardrails.

Together, they create a system that doesn't need constant human nudging, it proactively collaborates, adapts to changing contexts, and learns what works.

You interact with the agent. But what you're really engaging with is an intelligent system of systems, quietly orchestrating outcomes behind the scenes.

This is what separates productivity gimmicks from business transformation. Agentic AI systems don't just make things faster; they make things *better*, through continuous context, cross-platform integration, and embedded decision logic. They turn your workflows into intelligent operations and free your team to focus on what only humans can do.

The companies that master this won't just move faster. They'll redefine the very idea of how work gets done, who does it, and what's possible when your business has a second brain that never sleeps.

My Journey: From AI to Gen AI to Agentic AI Systems

If you look around most boardrooms and leadership offsites these days, you'll see the same story play out. Leaders are spending big on AI consultants, tools, training, and talent, hoping to squeeze a little more productivity from their teams.

Faster emails. Automated reports. Maybe a chatbot that sounds a bit more human, or a dashboard that promises "*insights at your fingertips*." These upgrades feel like progress, and sometimes, they are.

But let's be honest, for many companies, these are incremental changes. The real, game-changing value is being unlocked elsewhere, often quietly, behind the scenes, by teams thinking in a fundamentally different way.

Here's what's actually happening. The frontier of competitive advantage is moving. Not towards the next AI or the shiniest chatbot, but toward the organisations that are embedding intelligence so deeply into the fabric of how work happens that outcomes improve on their own. This isn't about adding "*AI features*" to legacy systems. It's about redesigning the system itself, so that AI isn't just a bolt-on tool, it's the operational backbone.

I learned this distinction the hard way, across 20 years and £100M+ in AI investments:

At an Oil and Gas Company: Built a data catalog that saved £75M/year. Not because it was smart, but because it acted autonomously, flagging issues before humans noticed.

At a Hospital: During COVID, our AI didn't just predict bed shortages, it reallocated resources in real-time, saving lives when every minute counted.

At a Tech firm: Watched hundreds of teams with amazing tech fail because they built tools, not systems. They had ingredients but no recipe.

Each failure taught the same lesson. AI without ownership is just expensive consulting.

When I joined the Fortune 500 Oil & Gas company, we faced 1,390 separate data sources living in silos. The typical solution? Build better dashboards. What we built instead? A unified, AI-enabled system that didn't just display data, it owned data quality. The system spotted anomalies, triggered corrections, and escalated issues before they became problems. That shift from *"AI that informs"* to *"AI that acts"* delivered £75 million in yearly savings.

At NHS England during COVID-19, the stakes were life and death. We needed to optimize bed allocation across a hospital with 8000 staff and 1000+ patients. Most AI would have stopped at predictions, *"You'll likely run out of beds in Ward 3 tomorrow."* But predictions don't save lives, actions do.

Our system went further. It automatically adjusted discharge schedules, reallocated beds in real-time based on clinical priorities, and learned from each decision. Result, £2.9 million saved annually, but more importantly, the right bed available when seconds counted. That's when I understood, the breakthrough isn't in intelligence, it's in autonomy.

I carried that lesson to Sensyne Health, working on clinical trial automation for rare disease target compound identification. There, the technical challenge was immense, scaling from 10,000 to 22.5 million patient records.

But again, it wasn't the complexity of the models that drove impact; it was how we reimaged the workflow in the pharmaceutical and hospital research labs around them. The AI system didn't just find eligible patients; it validated, tracked, learned, and reported all with minimal human intervention. It owned the outcome, from start to finish. That's what moved the needle for pharmaceutical therapists.

Then came the tech company. This time, the problem was more subtle but just as damaging. The tech had never been better, cloud compute, GenAI with off-the-shelf models. Yet most teams floundered.

Why?

They lacked decision architecture. They lacked memory based systems. They lacked the discipline to design for *systemic* adoption.

Everywhere I looked, I saw the same friction:

- GenAI built to demo, but not to operate.
- GenAI features that required constant maintenance, never quite trusted to act alone.
- Teams that launched "GenAI" projects, then quietly retired them when they failed to scale.

These experiences led me to co-found Templonix with founder Chris Tyson, who saw what I was seeing, businesses didn't need better AI models, they needed AI that could act like a trusted employee. Not waiting for instructions, but owning outcomes. The framework we built helps teams create AI agents in under an hour, not as experiments, but as operational teammates that follow guardrails and deliver results.

Now at Synaptix Solutions, I train teams to make this shift. Because the real competitive advantage isn't in having AI, it's in having AI that works while you sleep.

I learned all of this the hard way. The real breakthrough isn't in the code. It's in the system you build around it, the memory you architect, the guardrails you enforce, the outcomes you measure, and the trust you nurture.

That's the heart of agentic AI. It's what this playbook is about. And it's what will separate the teams who survive the hype from the ones who actually transform their business.

It's tempting to think the teams making the most progress are those with the most impressive tech stacks or the biggest AI research budgets. In practice, that's rarely the case. The true winners are those who have the clearest sense of what outcome they want to own, and the discipline to build the AI systems that deliver that outcome, day in and day out.

They're not mesmerised by the latest model or API. They're not distracted by the noise around *"AI everywhere."*

They focus on what matters:

- Where does work break down?
- What would it look like if an intelligent system owned that outcome, end-to-end?
- How do we measure real-world impact, not just AI activity?

Here's the good news for every team, regardless of size or technical horsepower. You don't need a moonshot budget, a PhD army, or a magic new tool. What you need is clarity of **outcome**, the discipline to design for **real-world** workflows, and the humility to measure what actually **moves** your business forward.

But let's be clear. The window is wide open right now, but it's closing, fast. A new class of competitor is already emerging. Teams quietly deploying agentic AI systems that operate in the background, learning, improving, and delivering results autonomously, with less friction and more leverage every month. They're not waiting for permission from IT. They're not waiting for the perfect AI model. They're quietly taking work off the table, work your team still does by hand. If you don't move soon, you'll find yourself competing with businesses whose operations never sleep, whose workflows never drop the ball, and whose teams are freed up for higher-level thinking.

Agentic AI isn't the next incremental step. It's the next competitive leap. The only question is, *"Will you leap, or watch from the sidelines?"*

Why Agent Product Strategy Matters Now

While you're reading this, your competition is deploying agents that work 24/7.

They're not hiring. They're not training. They're embedding intelligence directly into operations:

- Sales agents that never miss a follow-up.
- Finance agents that spot fraud in milliseconds.
- Support agents that resolve tickets while you sleep.

By 2025, companies will split into two camps: those with autonomous operations, and those still coordinating humans. The gap between them will be insurmountable.

Back in 2018, I was helping Fortune 500 execs roll out digital transformation initiatives. AI was on every agenda. We built dashboards, ran pilots, and launched *"AI-powered"* features that looked great in slides. But behind the scenes? Not much changed.

Most solutions sat on the edge of existing workflows. They surfaced data. Nudged users. Flagged risks. But every insight still needed human judgment. Every automation stalled at the handoff. These systems didn't transform work, they decorated it.

That's when the penny dropped: we weren't building systems. We were bolting AI onto the side of old ones and calling it innovation.

Fast forward to today, and the game has changed. The shift isn't about tools or prompts. It's about ownership. Intelligence isn't something you add, it's something you design for. Agentic AI marks a new phase: not just assistance, but autonomy. Not just outputs, but outcomes. And that requires product strategy.

Teams ask, *"Where can AI help?"* But that's no longer the question. The real question is: *"Where can AI take responsibility?"*

That means defining:

- What the agent owns.
- What memory it needs.
- How it behaves without constant input.
- Where to intervene, and where to let go.
- What success looks like when there's no interface.

Most agent projects fail not because the tech is wrong, but because the product strategy is missing. Without clear goals, guardrails, and governance, agents don't create clarity, they create chaos.

No LLM will design your escalation path. No chatbot will realign your team's incentives. No API will define the boundary between autonomy and control. Only people can do that, and that's where product strategy comes in.

The conversation shifts from *"Let's try AI"* to *"Let's redesign how this work gets done."* It speaks the language of capability, cost, responsibility, and results.

And it sells. Executives don't want more tools. They want fewer handoffs, faster decisions, and scalable operations without bloating headcount. That's what well-designed agents deliver, when built with intent. This is the next strategic advantage. Helping teams design agent systems that run themselves, adapt in real time, and compound in value.

That's what agent product strategy enables. This book gives you exactly that. If you're tired of living in *"pilot purgatory"*, building proofs-of-concept that gather dust after the first demo, or if you're exhausted by initiatives that over-promise and under-deliver, you're not alone. Many teams feel stuck, knowing there must be a better way, but not seeing a path forward. I've distilled 20 years of building AI products into a two-day workshop to design AI agents that deliver ROI: **There is no spoon.**

There is now spoon isn't because I'm an obsessed fan of The Matrix.

The spoon is infact the AI agent, a simulated object. It's code runs inside the agentic AI technical architecture. Thus, where its possible to bend the spoon, because it isn't a physical object, the AI agent is the same. Code can be manipulated with other code, because of the non-deterministic nature of how an agentic AI system works i.e., what goes in is known, but what comes out is based on what the AI agent believes is true, not by human design.

The approach isn't about hacking your old processes, or slapping an AI label on what you already do as well. You'll learn how to identify processes ripe for autonomous operation, design AI agents that can navigate complex multi-step workflows, and prove value even to sceptical executives. It's about building something new. An intelligent system you can trust to deliver value while you sleep. It means measuring impact, not just activity. It means building for outcomes, not applause.

And if you've read this far, you're probably ready for that leap. Ready to stop dabbling with AI around the edges and start designing the core systems that run your business, quietly, and reliably.

Most importantly, you'll discover that the most powerful agentic AI systems aren't the ones with the most sophisticated architecture. They're the ones designed with such clarity of purpose that they can achieve outcomes without human intervention.

This book isn't for prompt hobbyists or AI tourists. It's for builders, people trying to ship, scale, and survive. Not more tools. Not another chatbot. But real systems that remember what matters, act on goals, reason across steps, and deliver outcomes, on their own.

So, what does it take to get there? You don't need a data science army. You don't need to be an LLM whisperer. You need to think like an architect.

Here's how we'll do it, together:

- **First, I'll show you how to spot the AI agent opportunity already hidden in your workflow.** It's often not where you expect, and usually not where the hype says to look.

- **Next, we'll walk step-by-step through designing, testing, and deploying agents—without needing a machine learning team or moonshot budget.**
You'll see what “minimum viable agent” really looks like in practice.
- **Finally, you'll learn how to measure, refine, and scale agentic systems—so they keep running, learning, and compounding value for your business.**
The goal isn't just to launch, but to keep getting better.

This book isn't theory. It's a practical, lived-in guide for builders and operators who are ready to create real change.

Ready to build systems that work, so you can focus on the work that matters?

Why listen to Tim

The era of static workflows is over. If you're like most builders today, you've already tried AI. You've added a chatbot to your website. Maybe tested GPT for writing. Perhaps automated a few internal tasks.

And yet... you're still underwhelmed. Because despite the hype, most businesses are still operating like it's 2010. Workflows are linear. Data is siloed. Decisions are delayed. And AI? It's sitting there waiting to be prompted.

That's the exact problem I've spent two decades trying to fix.

I've been in rooms where AI was sold as a miracle ...and stayed in those rooms when it failed to deliver. Not in the lab, in live systems. Healthcare. Energy. Education. Government. Finance.

From hospitals trying to predict surgery risk...to public sector teams navigating national-scale digitisation to enterprise product teams drowning in procurement cycles and spreadsheet chaos.

I've helped ship AI-powered tools that don't just demo well, they *work inside the messy system*, where regulation, human behaviour, and legacy tech collide.

This book isn't based on trends. It's the field manual I wish I had five years ago, when LLMs were new, expectations were sky-high, and nobody could explain how to go from chatbot to agentic capability.

You don't need a new tech stack. You don't need to be an AI expert. You just need a clear problem, and a willingness to rewire how your team works.

Why Most AI Strategies Fail

AI is widely misunderstood, and most companies are getting it wrong. I've seen it happen across government departments, global enterprises, health systems, and scaling startups.

The intent is good. The ambition is real. But too often, they end up in *pilot purgatory* — stuck in cycles of demos, showcases, and “innovation theatre” without delivering anything of value.

The numbers speak for themselves:

- **85%** of AI projects fail to deliver business value (*Gartner*)
- **91%** struggle to scale beyond pilot phases (*MIT Sloan*)
- **68%** of executives say AI fell short of expectations (*PwC*)
- Most GenAI experiments never make it past internal demos

So, what's going wrong?

Mistake #1: Mistaking a prompt for a product

I once watched a team proudly demo a ChatGPT wrapper, calling it an “*AI assistant*.” It could generate responses and crack a joke, but had no memory, no actions, no logic, no task completion. When asked, “*Can it help us close renewals faster?*” the room went silent. It wasn't an assistant. It was a proof of concept, mislabelled and misunderstood.

Mistake #2: Building hype, not solving friction

One fintech launched a GenAI sales coach trained on playbooks and blog posts. It gave advice, matched tone, sounded smart, but it didn't know the status of deals, had no pipeline data, and couldn't integrate with Salesforce. The real bottleneck? Reps drowning in follow-ups. The agent missed the mark. Adoption collapsed in a month.

Mistake #3: No ownership, no direction

In one public sector lab, the brief was *"Explore LLMs for citizen services."* Sounds bold. But when asked what task the AI should own, no one had an answer. No workflow. No outcome. Just tech for tech's sake. The project drifted... then quietly vanished.

Mistake #4: Beautiful UI, no utility

At a large healthcare provider, a team spent six months crafting a slick GenAI dashboard. Auto-summaries, smart search, it looked great. But it couldn't flag clinical risk, escalate decisions, or automate handoffs. It visualised information. That's it. No action. No leverage.

These aren't rare failures. They're the standard playbook of AI 1.0. The common thread? **Most teams never made the leap from tools to systems.**

They treated AI like a plugin, not a system component. They built experiments instead of capability. They asked *"how do we use AI?"*, not *"what should AI own?"*

That's why this book exists. Because most teams don't need more models. They need better questions.

Questions like:

- What specific job should this agent be responsible for?
- What memory does it need to operate effectively?
- What decisions can it make on its own, and when should it escalate?
- How will we measure success, and know when it's working?

Real product strategy starts here, not with a model, but with a mission.

Bringing It All Together - Purpose, Not Just AI

Ben Thompson, tech analyst and founder of *Stratechery*, once said:

“Most companies aren’t failing at AI. They’re failing at integration.”

I agree, but I’d go further, **they’re failing at product strategy.**

Not strategy in a slide deck. Not a roadmap full of GenAI pilots. But operational, execution-ready, decision-led product strategy, the kind that earns its place in the business.

It’s not about picking the “*right*” model or experimenting with prompts. It’s about embedding intelligence into operations, with constraints, metrics, and clarity of ownership. Without that, even the most advanced model is just an expensive toy.

So what does good look like?

✗ “We want to use AI in sales ops”

✓ “We’ll deploy a contract renewal agent by Q1, using pricing history, usage data, and legal rules to auto-draft proposals and reduce renewal cycles by 30%.”

That’s not a use case, that’s an **agentic product strategy**. It has a job, a trigger, an outcome, and clear value. It also makes it easy to say *no* to distractions.

This is what most teams miss:

- They chase use cases without understanding workflows.
- They build front-ends with no logic behind them.
- They skip the system architecture, and wonder why nothing scales.

Agentic product strategy fixes this. It starts with defining the job.

Then it scopes the logic, maps the memory, sets the rules, and delivers a system, not a showpiece.

That's what we'll do in this book. Not theory. Not theatre. Just clear, grounded methods to design, build, and deploy agents that create real operational value, the kind that compounds over time.

Here's the core thesis of this book, in plain English:

"The future of a business isn't about who has the flashiest AI model or the most technical firepower. It's about who can embed intelligence so seamlessly into their daily operations that value is created, even when no one's looking. It's about outcomes that run on autopilot, freeing your users to focus on higher strategy, revenue, and customer growth."

Most leaders I know are stuck building yesterday's solutions to tomorrow's problems. They're hiring ChatGPT specialists to craft more effective emails. Training teams on prompt engineering to generate content faster. Building AI features that make existing workflows marginally more efficient. Meanwhile, their competitors are embedding intelligence directly into their business operations. Not as a feature or a tool, but as the foundation of how work gets done. And so, for the last couple of years, we've been dazzled by demos and distracted by technical one-upmanship:

"Which LLM is the best?"

"How human does the chatbot sound?"

"What's the biggest language model you can fine-tune on your laptop?"

But all that surface innovation misses the point. The companies that will pull ahead aren't the ones showing off the latest shiny object. They're the ones quietly, relentlessly building **systems that think, remember, and act in pursuit of real business goals.**

That's when I understood the fundamental shift happening. **We're moving from using AI as a productivity enhancer to embedding it as business logic. This book exists because I've seen the future, and it's not about prompt engineering.**

It's about designing intelligence that operates independently within your business.

Systems that can:

- **Remember what matters, across every interaction.** Imagine a workflow where key details, preferences, exceptions, and past decisions, are never lost or forgotten, even as teams change or priorities shift.
- **Create Agents that act on multi-step goals, without waiting for a prompt.** Instead of passively responding to your commands, these systems anticipate, trigger, and follow through, end-to-end, across processes and platforms.
- **Enable workflows to adapt, learn, and improve continuously, without you micromanaging them.** They notice patterns, spot bottlenecks, and course-correct (and will escalate if not sure to a human). Over time, they get smarter, faster, and more aligned to your business.

The companies building these systems now will have an insurmountable advantage over those still optimising human workflows with AI assistance. And here's the opportunity. **You don't need a PhD in machine learning or a massive AI team to build effective agentic AI systems.** You need clarity about what outcomes matter, systematic thinking about workflow design, and frameworks for measuring autonomous performance.

When To Use There Is No Spoon

Broadly speaking, **There Is No Spoon** is useful when you're ready to move past GenAI demos and actually embed intelligence into the enterprise layer of the way your business operates.

It's ideal when:

- Your AI Agent experiments are piling up, but nothing's landed in production.
- Teams are still working around AI Agent , not with it.
- You've invested in tools, but can't trace them to a measurable outcome.

- You're stuck in "*prompt engineering*" mode and need a clearer path to systems that actually do the work.

Sometimes, it's painfully obvious that a new approach is needed. Other times, it creeps up. Here's a quick "*vibe check*" if you're not sure:

Deploy **There Is No Spoon** when you notice:

- AI features that look impressive but don't reduce actual effort.
- Lots of talk about "*use cases*" with no clear owners or outcomes.
- Repeated workflows are still being done manually, even after pilots.
- Dashboards with insight, but no change in decision speed or quality.
- Low confidence from leadership despite "*innovation*" activity.
- A general feeling of an AI Agent being a burden, not a benefit.

But there's also a more proactive lens to this. You don't need to wait until frustration sets in. The workshop works just as well as an opportunity-finder, a way to spot latent value and design new capacity into your workflows.

Here are some practical ways teams are using the workshop today:

- You're starting a new quarter and want to align product, ops, and execs around AI Agent investment that actually drives ROI.
- You're working with a client or stakeholder who "*wants to use AI Agent*" but isn't clear on what that means.
- You've identified a process that's eating time, and suspect an agent could take over, but need to validate how.
- You're building a prototype and want to define memory, goal states, triggers, and guardrails before writing any prompts.

How it all works

There Is No Spoon isn't just a workshop. It's a system design sprint grounded in the operational realities of your business, shaped by the lessons we learned from using the [Templonix Framework](#).

Templonix taught us that the problem with most AI projects isn't technical, it's structural. Teams were chasing “*AI features*” before they defined the outcomes. Prompting models before clarifying purpose. Designing interfaces without logic, memory, or measurable success.

Those mistakes led us to five critical design principles:

- **Define Goal Ownership** – Agents must deliver outcomes, not just assist with tasks.
- **Model Memory** – Decisions require continuity across time, not stateless interaction.
- **Map Triggers & Loops** – Agents must know when to act, escalate, and adapt.
- **Design System Boundaries** – Define where agents integrate, and where they don't.
- **Align Metrics to Outcomes** – If it doesn't move the business, it's not worth building.

If you want to compress months into days, there's a faster path. The **There Is No Spoon** workshop takes your team from confusion to clarity in 48 hours. No theory. No slides. Just hands-on system design with immediate results.

Here's what actually happens:

Day 1 Find the Agent opportunity

Morning: We map where work breaks. Not in theory, in the actual experiences of your team. *Where do decisions stall? What gets repeated endlessly? Where does knowledge get trapped?*

Afternoon: We identify which decisions matter. *Who makes them now? What happens when they're wrong? What would change if they happened instantly, correctly, every time?*

Day 2 Define and design the Agentic AI Strategic approach

Morning: We decide which agents will be prioritised first. What must be remembered? What rules never change? What patterns indicate success or failure?

Afternoon: We build the system strategy. *When does the agent need to be built? When does measurement begin? How do further agents get built? Who owns the implementation of the agents?*

Executive Track: While teams design, leaders align. We validate ROI, prioritise use cases, and ensure what gets built actually matters to the business.

You leave with:

- A validated agent use case with clear ownership.
- A memory architecture encoding your team's expertise.
- Guardrails and metrics for safe deployment.
- A shared vocabulary for "AI that works" vs "AI that impresses".
- A 30-day implementation roadmap.

Now that you understand how **There Is No Spoon** works, and why it exists, it's time to get practical. The next sections will walk you step-by-step through preparing, running, and delivering your first workshop.

How to use this guide

This guide is here to help you spot, design, and deploy your first AI Agent and Agentic AI system, without needing a machine learning team or a moonshot idea. The goal is simple, move from AI confusion to operational clarity. One workflow at a time.

It's not a textbook. Not a research paper. A hands-on guide to building something that works, fast.

Here's how to get started:

- **First, understand what you're building.** This isn't just a smart UI or a tool with an LLM bolted on. You're building an intelligent system that can observe, decide, and act, not just respond. Once you make that mental shift from "AI feature" to "AI teammate," everything else starts to fall into place.
- **Then, read with intent.** Skim the concepts, flag the examples that sound like your business, and start imagining where these agents could go. You're not looking for the biggest or sexiest use case. You're looking for the most annoying, repetitive, costly task your team quietly dreads every week.
- **Print the worksheets or make a copy.** Write things down. Map workflows. List actions. Get hands-on. This is about creating systems that work in your real context, not someone else's slide deck. You're not building an AI "experiment." You're redesigning a job.
- **Don't skip the thinking stage.** Everyone wants to jump into prompting and tools. That's how most teams fail. Take the time to answer: What is this agent responsible for? What decisions should it make? What guardrails does it need? What does success look like?
- **Test fast. Learn faster.** Your first agent should be simple, safe, and testable within a week. Whether it's built with Airtable and Zapier or GPT and Slack, the only thing that matters is whether it delivers a useful result with minimal human input. That's your proof of value.
- **Get feedback from real users.** Let the person who used to do this task try it. Ask them where it fails. Where it surprises. Where it actually helps. That's your goldmine. Use it to refine, not just the tool, but your thinking.

This is how you move from *"We're exploring AI"* to *"We've got an agent handling that now."* Start small. Start now. And use the checklist below to make sure you're on the right track.

Your Next 48 Hours

The companies winning with AI aren't smarter than you. They just started.

Your competition is already building agents. Every day you wait, they pull further ahead. But you can catch up faster than you think.

Your move:

1. **Today:** Pick one broken workflow. Just one.
2. **Tomorrow:** Map what an agent would need to own it.
3. **This week:** Build a proof of concept. Ugly is fine. Working is essential.

Stop admiring the problem. Start building the solution. The future isn't about humans versus AI. It's about humans with agents versus humans without them.

Step-By-Step Guide To Agentic AI Product Strategies

When to use Agentic AI Product Strategy

Some of the most valuable opportunities for Agentic AI aren't hidden deep in strategy decks, they're sitting right in front of you, in the day-to-day work your team quietly endures.

That recurring task no one enjoys.

The document that gets updated every week but never quite makes a difference.

The email sent out on repeat with just a few tweaks.

These aren't edge cases.

They're signals, pointing straight to workflows that are ripe for change.

You'll know it when you hear it. *"Why are we still doing this manually?" "Didn't we send that already?" "I feel like I do this same thing every Monday."*

These aren't just employee gripes, they're flashing lights on the dashboard. When a process only works because someone nudges it forward, chases a colleague, or pastes the same data in two places, it's not really working. That's where Agentic AI earns its place, by turning recurring, logic-based effort into autonomous, and goal-driven action.

There's also a bigger picture at play. When teams are stretched thin, when speed matters more than ever, when doing more with less becomes the mandate, that's when you need something that can operate with context, memory, and intent. Not a tool. A teammate. An agent. The kinds of systems that don't just assist, but act. That don't just respond, but decide.

And yet, many teams miss this moment because they start from the wrong place. They begin with the tool, the model, or the buzzword. But agentic design doesn't start with what's possible, it starts with what's painful. You don't build agents to impress. You build them to fix. The best indicator isn't technical readiness, it's operational frustration.

So if you're seeing delays, repetition, decision bottlenecks, or manual patchwork to keep things moving, that's your opening. You don't need to ask whether your business is *"ready for AI."* The question is simpler, *"where would it help if something just handled this... every time?"*

This is how you move from “*We’re exploring AI*” to “*We’ve got an agent handling that now.*” Start small. Start now. And use the checklist below to make sure you're on the right track.

Agent Opportunity Checklist

To validate whether a workflow is ready for an agent, ask these four questions:

Question	Why It Matters
Is it repeatable?	Agents thrive on recurring patterns, if it's done weekly, it's a candidate
Can it be triggered?	You need a clear start signal (like a date, event, or field value)
Is reasoning required?	Agents shine when logic, rules, or judgement are needed, not just lookups.
Does it end in a decision or deliverable?	Agents must complete a useful task, not just analyse.

If you can answer **yes to at least 3 out of 4**, you've got an agent opportunity worth pursuing. Let's find yours.

Overview of the workshop

Day	Exercise	Purpose
1	Friction Mapping	Identify where work breaks down, judgement bottlenecks, rework loops, and data blind spots. There are signals for where AI agents can create real leverage.
1	Decision Mapping	Surface the key decisions inside broken workflows. Who makes them? When do they fail? This uncovers candidate areas for agent ownership.
1	Trigger & Memory Ideation	Brainstorm what the agent needs to remember and when it should act, escalate, or adapt, using a memory map.

Day	Exercise	Purpose
2	Choose the best agent	Prioritise the best Agentic AI strategy with team voting and decider choosing.
2	Agentic AI Product Roadmap	Map out the key agentic product strategies over a timehorizon of 3-12 months and allocate to owners with due dates.
2	Wrap Up & Build readiness	Finalise your agent specification and define next steps, whether that's prototyping, vendor evaluation, or internal build. Leave with clarity and commitment.

Key outputs

- A clear **agent use case** tied to business value.
- A **workflow map** showing friction, decisions, and opportunity.
- A **memory and trigger blueprint** aligned with domain logic.
- A **lightweight agent design** ready for internal or vendor build.
- Cross-functional alignment on **next steps** and success metrics.

Materials Needed

- Miro / MURAL / FigJam board (for remote) OR,
- Whiteboard, sticky notes, and Sharpies (for in-person).
- Access to key workflows, subject matter experts, and sample data.
- Agent Canvas (provided).
- Timer or facilitator clock.
- Decider (with the authority to move forward).

Whether you're a product lead, AI strategist, or operations executive, the workshop helps you cut through AI noise and get to what matters, building agentic systems that do real work.

The Agent brief

To show you exactly how to run **There Is No Spoon** in a real-world setting, we'll walk through a fictional example rooted in very real challenges.

Meet Lina, a mid-career product manager working at a fast food e-commerce company operating across the MENA region.

Lina's company delivers meals, offers loyalty rewards, and handles a high volume of orders daily through web and mobile. The pressure is relentless. Every quarter brings a new feature push. Every metric is a moving target. She's keeping up... but barely.

Here's what Lina said on our first call:

"I've been working in product for 7 years. But I feel stuck. GenAI is changing everything, and I don't know how to catch up.

I keep hearing about agents, how they can automate workflows, personalise journeys, increase LTV... but no one's teaching how to actually design them. My execs are asking, 'Where's our AI roadmap?' and I don't have a confident answer.

I'm not afraid of the tech, but I am afraid of building the wrong thing, or worse, spending six months prototyping something that doesn't work in the real business."

After that call, it was clear **There Is No Spoon** was the right workshop to help her.

We agreed to run a 2-day workshop with her cross-functional squad to help them:

- Identify a workflow where an AI agent could truly own an outcome.
- Translate their domain expertise into a memory system the agent could use.

- Design a loop where the agent could observe, act, and improve, not just prompt and reply.
- Define the performance metrics that would matter to the business

As we walk through this guide, you'll see how we'd run the **There Is No Spoon** workshop with Lina and her team. You'll follow the exercises using her context, see the decisions made in real time, and understand how each step shifts them from "*AI anxiety*" to "*agentic clarity*."

When you're ready to run it yourself, just drop in your context. The process remains the same.

Let's get started.

Day 1

Introduction to Agentic AI Product Strategy

30 minutes

Before the work begins, it's critical to set the stage, not just logistically, but mentally. The introduction frames the purpose of the session, sets expectations, and gives everyone a shared sense of direction.

Here's how to do it:

Welcome and Context

Start by welcoming the group and acknowledging the time and focus they've committed to this session. Reaffirm that this isn't another brainstorming session, it's a structured sprint to identify and architect real AI opportunities that deliver business impact.

Clarify the Objectives

Explain why they're here and what they'll walk away with. For **There Is No Spoon**, the goal is simple. To define a specific, high-impact use case where an AI agent can own an outcome, and build the blueprint to make it real.

You can say something like:

"By the end of Day 2, we'll have a mapped-out agent use case, memory requirements, system boundaries, escalation rules, and a business metric it must move. This isn't a tech demo. It's about strategic clarity."

Set Ground Rules

To keep the energy high and the outcomes sharp, establish a few norms:

- **Be fully present**, no multitasking or open laptops.
- **All ideas are welcome**, even half-formed thoughts.
- **Challenge with respect**, decide with clarity.
- **No "what-about-isms"**, defer edge cases until later.

Optional: Use a physical or virtual “*parking lot*” for great ideas that don’t belong in the current moment.

Introduce the Agenda

Give a high-level overview of the workshop flow so participants know what to expect. For example:

“Today we’ll focus on friction and decisions, where work breaks down, and which decisions actually matter. Tomorrow we’ll shift into memory and agent design, translating domain knowledge into architecture and defining the agent’s job.”

This primes the group for the rhythm of the sprint, **discovery first, system design second**.

Preparation Reminders for Facilitators

Before the session begins, make sure:

- The agenda is clear, adaptive, and aligned to business goals.
- The physical or digital workspace is fully set up.
- All required materials are accessible (sticky notes, markers, digital tools).
- If remote: everyone has their camera on and is in a distraction-free space.

Great workshops don’t happen by accident. They’re designed for clarity, momentum, and progress, before the first word is even spoken.

Now you’re ready to begin.

Stage 0: Choosing a decision maker

Before you dive into the workshop, the most important choice you'll make isn't about AI. It's about alignment.

Every team needs a Decision Maker. The person who can make confident calls when the group gets stuck. This isn't about hierarchy or ego. It's about momentum.

In an agentic AI sprint, decisions shape the agent's autonomy, its triggers, and how it integrates into the business. Waiting for consensus on every detail will grind the sprint to a halt.

The role of the Decision Maker is simple:

- Make the final call when the group can't agree.
- Resolve trade-offs between ambition and feasibility.
- Keep the sprint moving forward with clarity and focus.

Who should it be?

Usually it's:

- A product owner or team lead.
- An operational leader with deep workflow knowledge.
- A stakeholder responsible for the outcome the agent will own.

But it must be someone with:

- Authority to make strategic calls.
- Context on the problem space.
- Trust from the team.

The best Decision Makers act more like editors than dictators, they bring focus, not friction. Once they're chosen, make it official at the start of the session.

Give them the green light to lead when clarity is needed most.

Stage 1: Where does Work Break Down?

Activity Name: Friction Mapping

Duration: 45 minutes

The first step is about finding the pain in the workflows where human time is wasted, decisions slow things down, and nobody has end-to-end visibility.

This isn't about vague inefficiencies. We're looking for **judgement bottlenecks**, **rework loops**, and **data black holes**, the kinds of friction that AI agents can relieve with real business impact.

Your job as a facilitator is to surface these pressure points and translate them into opportunities for intelligent automation.

Resources

- Sticky notes (any colour, rectangular).
- Sharpies (one per participant).
- Whiteboard or large wall space.
- Digital whiteboard if remote (use FigJam or Miro with a "Friction Mapping" frame).

How To Deliver it

1. Set the context (5 minutes)

Say:

"Today, we're going to find the 'grind' in your work, the parts where things get stuck, repeated, or forgotten. These are the tasks that slow you down and pull people into coordination chaos. This is where agents can make a real difference."

2. Explain what we're looking for (5 minutes)

Ask participants to think about:

- Repeated processes that happen every week.
- Tasks that rely on multiple tools or people.
- Moments where delays or errors cost time, money, or trust.

Prompts to say out loud:

- *“Where do we lose time chasing updates?”*
- *“What task could a smart intern do with clear instructions?”*
- *“Where do things break when one person is on holiday?”*
- *“Where are people making judgement calls that could be standardised?”*

3. Quiet writing (15 minutes)

Each person writes 2–3 friction points on sticky notes. Use one idea per note.

Examples:

- *“Sales team chases signed contracts manually every Friday”*
- *“Customer onboarding gets stuck waiting for IT provisioning”*
- *“Nobody knows who owns the monthly KPI report – it always gets missed”*

4. Cluster the patterns (10 minutes)

Group the notes by theme. Use broad categories like:

- Data handoffs.
- Manual approvals.
- Reporting delays.
- Customer follow-ups.

Let the group help sort them, and ask questions to clarify where the friction comes from.

5. Apply the Agent Signal Test (10 minutes)

Now test the top 3–5 patterns using the *Agent Opportunity Test*:

Question	Pass / Fail
Is it repeatable?	
Can it be triggered automatically?	
Does it require logic or decision-making?	
Does it end with a clear output?	

If it gets 3 or more passes, you’ve found a candidate for an AI agent.

Facilitator Tip

Encourage teams to write a draft **opportunity statement** right after the test:

*“We want to automate **[process]** so that **[team]** saves time and **[business outcome]** improves.”*

Feedback Loop from Support

“We want to automate insight extraction from support tickets so that the product team saves time and feature prioritisation improves.”

Delivery Issue Escalation

“We want to automate escalation of failed deliveries to CS so that the customer service team saves time and complaint resolution speed improves.”

Customer Reorder Reminders

“We want to automate repeat order nudges so that the growth team saves time and customer retention improves.”

Milestone

By the end of this exercise, the team will have:

- A list of high-friction workflows.
- A clear understanding of where agents could drive value.
- One or more validated agent opportunities ready to explore in Day 2.

This is where your agentic design journey truly begins. Let's move to decision mapping.

Stage 2: Voting on the Agent Opportunity

Exercise Name: Agent Candidate Voting

Duration: 25 minutes

Now that we've mapped several promising agent opportunities, it's time to focus. This step helps the team prioritise one high-impact agent use case to guide the rest of the time.

We'll do it in two phases:

1. Straw Poll Vote to hear team preferences
2. Decider Vote to lock in the primary candidate

This ensures we combine diverse input with strategic clarity, and avoid endless debate.

Resources

- One green voting dot per person
- One gold star or larger dot for the Decider
- Sticky notes or whiteboard with all agent opportunities visible
- For remote: Use Miro, MURAL or FigJam voting tools with voting names shown

How To Deliver it

1. Straw Poll Voting (5 minutes)

Say:

"Review the three agent opportunities we mapped earlier. Choose the one you believe will have the greatest impact on the business if automated. Think about effort, outcomes, and repeatability."

- Hand out one green voting dot per person.
- Ask them to write their initials on the dot.
- Set a 5-minute silent review timer.
- At the end, everyone places their dot at the same time to reduce bias.

2. Share & Discuss (10 minutes)

Say:

“Let’s go around and quickly hear why each person voted the way they did. This gives us context before the final decision.”


- Call on each person briefly (1 min max).
- Ask them to state the agent use case they picked and why.

3. Final Decider Vote (10 minutes)

Say:

“Now it’s time for the Decider to pick the single agent opportunity that will guide the rest of this sprint.”

- Hand the Decider a larger voting dot (or gold star).
- The Decider may:
 - Choose one agent use case as-is.
 - Or combine aspects of two (if it’s straightforward and doesn’t need debate).



Reminder: The Decider’s role is not to be the smartest person in the room — it’s to keep momentum by making tough calls.

Milestone

The team has aligned on one clear, high-potential agent opportunity.

This now becomes the focus of the remaining sprint activities, from defining memory to setting agent boundaries.

Stage 3: How will we know the Agent is Working?

Exercise Name: Success Signals

Duration: 30 minutes

Now that the team has selected a promising agent opportunity, we need to define what success actually looks like, not in vague aspirations, but in clear, measurable outcomes.

We do this by crafting “**As measured by...**” statements tied to the selected agent brief. This ensures the agent’s success can be tracked, validated, and aligned to real business impact.

Resources

- Rectangular sticky notes (one block per participant, a different colour to the brief).
- Sharpies (one per participant).
- Whiteboard or wall space for sticking notes.
- **For remote workshops:** Create a dedicated “*Success Signals*” section under the chosen brief in your digital whiteboard tool.

How To Deliver It

1. Introduce the idea of success signals (5 minutes)

Say:

*“Now that we’ve picked our agent idea, it’s time to get specific. You’re going to write ‘**As measured by...**’ statements for that opportunity. This will help us know if the agent is delivering real value — or just creating noise.”*

2. Show examples (5 minutes)

Write the chosen agent brief on the board (if not already there), then add a few examples like:

- **Agent Brief:** *“We want to automate escalation of failed deliveries to CS so that the customer service team saves time and complaint resolution speed improves.”*

→ **As measured by:** 40% reduction in average time-to-resolution for failed deliveries.

As measured by...

40% reduction in average time-to-resolution for failed deliveries.

As measured by...

20% decrease in customer churn rate

As measured by...

5-minute decrease in average time to extract insights.

As measured by...

25% increase in repeat order rate.

Remind participants that good success signals are **specific, ambitious, and clearly tied to the outcome**, not just activity or effort.

3. Silent writing (10 minutes)

Each person writes **1–2 “As measured by...” statements** for the chosen agent opportunity.

One statement per sticky note.

4. Stick and read-out (10 minutes)

Each participant sticks their metrics to the wall or digital board and reads them out loud, no explanation needed.

Group similar metrics together and spot patterns.

Facilitator Tips

- Push for precision: If someone writes “more insights” or “faster service,” ask:
“What number would prove it’s happening?”
- Highlight leading vs lagging indicators: Sometimes it’s helpful to note both early signs (engagement, usage) and hard impact (cost reduction, revenue lift).

Milestone

By the end of this step, the team should have a **clear set of success signals** for the chosen agent, setting the foundation for future validation, feedback loops, and iteration.

Let me know if you'd like sticky note visuals or metrics for a specific agent example.

Stage 4: Vote and Prioritise

Exercise Name: Pick Success Metrics

Duration: 20 minutes

Now that the team has generated clear “*As measured by...*” statements for their chosen agent opportunity, it’s time to narrow them down.

We’ll do this in two quick steps:

1. A **Heat Map Vote** to highlight team preferences.
2. A **Decider Vote** to make the final selection.

Our goal is to help the team agree on **2–3 success metrics** that will guide how they evaluate and track the agent’s business impact.

Resources

- Red voting dots (3 per person).
- Sharpies or pens for initials.
- One larger/different-colour dot for the Decider.
- **Remote workshops:** Use voting features in Miro, FigJam, or MURAL.

How To Deliver It




1. Heat Map Vote (5 minutes)

Say:

“Each of you will now vote for the 2–3 metrics you think are the most important indicators of success. These are what we’ll prioritise going forward.”

- Hand out **3 red voting dots** per person.

- Ask participants to place their votes silently.
 - They can place all three dots on one metric if they feel strongly.
 - They can vote for their own metric.
 - Once everyone has voted, scan the board to highlight the highest-voted sticky notes.
- Remind everyone: This is just surfacing the front-runners, not the final pick.



2. Decider Vote (10 minutes)

Say:

“Now the Decider will choose the final 2–3 success metrics that we’ll use to track this agent opportunity moving forward.”

- Give the Decider a **larger or standout dot** to mark their picks.
- The Decider can either:
 - Select the top-voted metrics OR choose a different combination if it better reflects the business context.


Encourage clarity, not compromise. The Decider’s job is to **unblock the group**, not manage by committee.

Milestone

By the end of this step, the team will have:

- Chosen 2–3 high-impact success metrics.
- Anchored the agent opportunity to tangible business outcomes.
- Created a basis for progress tracking, iteration, and proof of value.

You can now create a simple visual (Agent + Selected metrics) to keep this visible throughout the rest of the workshop.



Stage 5: Where does decision-making need help?

Exercise Name: Decision Mapping

Duration: 45 minutes

Now that we've aligned on the agent opportunity and what success looks like, we're going to explore how decisions get made in your current customer and business journey.

This is where many agentic AI projects fall apart, not because of the tech, but because **nobody mapped the actual decisions that need augmenting.**

We'll use this stage to locate the human decisions, bottlenecks, or moments of judgement where AI could help, not replace, people.

Resources

- Journey map (created in earlier stages or brought in).
- Red sticky dots (3 per participant).
- Sharpies.
- Post-its or labels for marking decision points.
- Wall or whiteboard space.
- **For remote:** use a collaborative whiteboard tool like Miro with dot-voting features.

How To Deliver It

1. Review the End-to-End Map (10 mins)

Say:

*"Let's walk through your current customer journey or workflow. We're not looking at features or systems, we're looking for **human decisions**: moments where a person has to choose, judge, approve, triage, or escalate."*

Mark these key questions:

- Where is a decision currently manual or slow?
- Where does someone rely on gut feel?
- Where do decisions affect multiple teams?
- Where does delay or inconsistency hurt the customer or the business?

Facilitator Tip

Use a different colour to annotate these decision points (e.g., a red marker or dot).

2. Red Dot Prioritisation (10 mins)

Say:

“You each have 3 red dots. Use them to vote on the decision points you think are most valuable, or painful, to automate or support with AI.”

Instructions:

- Participants can place all 3 dots on one decision or spread them.
- Vote silently and at the same time to avoid anchoring bias.
- If remote, enable private voting, then reveal at once.

Facilitator Tip

Encourage participants to vote based on **impact** and **feasibility**—not just frustration.

3. Quick Group Debrief (10 mins)

Say:

“Let’s review the top voted decision points. These are the high-value targets for AI augmentation.”

Discuss briefly:

- Are these decisions made frequently enough for AI to help?
- Are they made by experts, or general staff?
- Do they use data that could be structured, labelled, or extracted?

Use this discussion to flag **top candidates for agent augmentation**, you'll use these in the prototyping stage.

Goal	Resolve Customer Escalations Efficiently							
Main Stage Workflow	Before			During				After
Stages of activity	1 Ticket Received AI Agent	2 Initial Triage AI Agent	3 Low Priority	4 High Priority AI Agent	5 Escalation L1 AI Agent	6 Escalation L2 AI Agent	7 Escalation L3 AI Agent	8 Resolution AI Agent
Pain points	• Delayed ticket visibility	• Incorrect classification of priority	• Slow resolution time	• Insufficient resource allocation	• Communication gaps with managers	• Lack of urgency from senior managers	• Overloaded directors	• Resolution not confirmed or tracked
Needs	• Automated ticket intake	• Accurate priority tagging	• Faster queue assignment	• Clear escalation protocol	• Manager alerting mechanism	• Senior management buy-in	• Escalation visibility for directors	• Closure confirmation process
Doing	 Customer submits issue AI Agent Opportunity: • Auto-capture ticket details • Extract key information • Suggest solutions from KB	 Support team reviews AI Agent Opportunity: • Auto-classify priority level • Analyze sentiment & urgency • Route to correct queue	 Assigned to low-priority queue	 Escalated to high-priority queue AI Agent Opportunity: • Monitor SLA thresholds • Auto-assign to available agents • Predict resolution time	 Email sent to agent's manager AI Agent Opportunity: • Auto-trigger escalation emails • Include ticket summary • Set response deadlines	 Email sent to Senior Finance Managers AI Agent Opportunity: • Compile escalation history • Auto-notify stakeholders • Track response times	 Email sent to Director of Finance Ops AI Agent Opportunity: • Executive summary generation • Risk assessment alerts • Auto-schedule follow-ups	 Tickets marked closed AI Agent Opportunity: • Auto-verify resolution • Send confirmation to customer • Update knowledge base
Data Sources used	• Support ticket system	• Priority rules engine	• Case resolution logs	• Agent workload tracker	• Manager contact directory	• Org escalation tree	• Director escalation triggers	• Closed ticket audit log
Technology	• CRM/API intake system	• AI classifier for priority	• Queueing algorithm	• Agent assignment tool	• Email automation service	• Escalation tracking dashboard	• Workload distribution platform	• Resolution analytics/reporting

Milestone

By the end of this step, your team will have:

- A set of clearly mapped decision points across the customer or business journey.
- A short list of the most critical decisions worth enhancing with agentic AI.
- A shared understanding of where AI fits *into* existing human workflows, rather than replacing them.

Stage 6: When should the Agent Act?

Exercise Name: Trigger & Memory Ideation

Duration: 45 minutes

Now that you've mapped the journey and identified success metrics, it's time to design how the agent will *think*.

This step focuses on building the agent's internal logic:

- What it should remember.
- When it should act or adapt.
- And where it should escalate or stop.

We'll do this by building a Memory Map, using three coloured sticky notes to define the agent's cognitive architecture.

Resources

- 3 colours of sticky notes (Memory, Trigger, Handoff).
- Sharpies.
- Previous journey map or workflow.
- Wall space or digital board section titled "Agent Mind Map".
- **Optional:** Use the Prompt Stack template from your Agentic AI toolkit.

How To Deliver It

1. Set the Scene (5 minutes)

Say:

"Now we're designing the agent's brain—how it remembers what's important, when it takes action, and when it hands it off to a human. These are the mechanics behind agentic AI."

2. Walk Through the 3 Categories (5 minutes)

Explain each one with examples:

- **Memory**

What information should the agent retain and carry forward?

Examples:

- Last 3 customer complaints.
- Onboarding progress.
- Delivery tracking history.
- Time since last contact.

- **Triggers / Reasoning**

What should prompt the agent to act, escalate, or stop?

Use a REACT-style loop: *Plan* → *Act* → *Observe* → *Repeat*

Examples:

- “If churn risk > 0.6 → notify account manager”.
- “If message is unanswered in 3 days → send follow-up”.
- “If sentiment = negative → reduce automation confidence”.

- **Handoff**

When should the agent involve a human?

Examples:

- Any refund request over £100.
- Failed delivery detected twice.
- High-priority SLA breached.

3. Silent Brainstorm (15 minutes)

Each person writes:

- 3 Memory notes (yellow).
- 3 Trigger notes (orange).
- 2 Handoff notes (blue).


Encourage real examples based on the journey map.


4. Stick and Connect (10 minutes)

Stick all notes on the wall/board in columns by type.

Draw arrows from:

- Memory → Trigger → Handoff. Look for patterns and logic loops.





5. Group Reflection (10 minutes)

Ask:

- Which triggers are most critical to get right?
- Are there too many handoffs or not enough?
- Does the agent remember the right things at the right time?

Milestone

By the end of Stage 6 Day 1, the team has built a functional “*mind*” for the agent.

This includes:

- What it remembers.
- How it decides.
- When it steps back.

These are now ready to be written into prompt stacks for prototype testing in Claude after creating the product strategy map in stage 8, Day 2.

Day 2

Playback - Day 1: Laying the Foundations for Agentic AI Product Strategy

Welcome the team back and re-anchor them in the strategic journey we've started. Remind them that Day 1 wasn't about theory or exploration, it was about surfacing the real operational pain and decisions where AI agents can drive value.

What we Covered

Friction Mapping: Where Work Breaks Down

We uncovered the judgment bottlenecks, rework loops, and data handoffs that are costing time and trust. These are the moments where intelligent agents can relieve pressure and restore flow.

Decision Mapping: Where Human Judgment Slows Progress

We zoomed in on the key decisions that currently block automation. The team voted on which ones matter most, those are our prime candidates for agent support or delegation.

Trigger & Memory Ideation: Designing the Agent's Mind

We defined what the agent should remember, when it should act, and when it must escalate to a human. Using Memory, Trigger, and Handoff patterns, we began shaping a logic system that mimics sound judgment, not blind automation.

Now that the pain is visible, the decision points are prioritised, and the agent's mind is forming, Day 2 is about turning this into prototypes, planning, and proof.

Let's move forward.

Stage 7: Choose the Best Agent

Exercise Name: Agent Strategy Voting

Duration: 45 minutes

By now, your team has surfaced multiple agent concepts, each designed around real friction, critical decisions, and scalable logic. It's time to move from possibilities to priorities.

This step uses structured, time-boxed voting to identify the most promising agent opportunities—those that are strategically aligned, practically viable, and energise the team.

The goal is not to crown a perfect solution today. It's to choose the best *starting points*—the agents we'll shape into prototypes and system flows.

Resources

- Red voting dots (10 per person) for Heat Map.
- One large green dot per person for Straw Poll (write initials).
- One Decider dot (different colour or shape).
- Wall space, whiteboard, or digital board with all agent concepts displayed.
- Sharpies or pens.

Once voting is complete, it will look a little like this.



How To Deliver it

1. Set the Scene (5 minutes)

Say:

“We’re about to vote on the agent opportunities we think are most worth exploring. The aim here isn’t to debate, it’s to rapidly visualise team energy and insight. This process happens in three rounds:

1. Heat Map.
2. Straw Poll.
3. Decider to Vote.”

Encourage participants to focus on *alignment to the long-term goal, clear triggers and memory logic, and value to the business and customer.*

2. Heat Map Voting (15 minutes)

Give everyone 10 red dots.

Say:

“Walk through each agent concept. Place dots on anything that feels compelling, smart memory logic, a strong handoff, a clear ROI. You can cluster dots or spread them out. Vote on what excites you.”

Tips:

- Dots can go on any part of the template: goal, reasoning, memory triggers, handoffs.
- Participants can vote on their own.
- No talking during this round—pure signal gathering.

Once complete, read out where the energy is clustering.

3. Straw Poll (10 minutes)

Hand each participant one large green dot and a Sharpie to write their initials.

Say:

“Now choose the ONE agent you personally believe has the most potential to deliver real business impact. Stick your green dot next to it and be ready to explain why.”

Go around the room:

- Each person gives a 60-second pitch:

“I picked this one because...”

No discussion, just signal and insight.

4. Decider Vote (10 minutes)

Now the pre-assigned Decider steps in.

Say:

“Based on all the team input, the Decider will now pick 2–3 agent strategies to move forward.”

They place a uniquely marked dot on the final selections. Let them combine or slightly tweak if needed, but keep momentum.

Milestone

By the end of this step, the team will have:

- 2 to 3 prioritised agent strategies.
- Shared conviction on where to focus next.
- A strong foundation to move into an agentic Ai product roadmap

These selected agents now form the **core blueprint** for design, validation, and implementation in the next stage.

Stage 8: Agentic AI Product Roadmap

Exercise Name: Roadmap Creation

Duration: 60 minutes

Now that the team has selected the most promising agent strategies, it's time to translate that momentum into a clear action path.

This step maps each chosen **Agentic AI concept** onto a simple 3–12 month timeline, assigns an accountable owner, and sets the rhythm for delivery check-ins.

The roadmap doesn't require full project plans—it's about momentum, visibility, and early commitment.

Resources

- Whiteboard or large wall area.
- Timeline drawn with sections:
 - Next 3 Months.
 - 3–6 Months.
 - 6–12 Months.
 - 12+ Months.
- Sticky notes for concepts (use the original agent templates or color-coded summaries).
- Additional sticky notes for owners and check-in dates.
- Sharpies.
- For remote teams: A digital timeline board with columns.

How To Deliver It

1. Set the Scene (5 minutes)

Say:

“We’re now moving from idea to delivery. Our task is to map out when we think these Agent strategies will begin, and who will help make them happen.”

Explain:

- You’re creating a **high-level delivery rhythm**, not detailed Gantt charts.
- Focus on “When should we start?” and “Who will lead it?”.

2. Draw the Timeline (5 minutes)

Create a clear, horizontal roadmap across your wall or digital board. Divide it into four phases:

- **Next 3 Months:** Quick wins, urgent friction.
- **3–6 Months:** Larger but achievable agent builds.
- **6–12 Months:** More complex or dependent work.
- **12+ Months:** Horizon concepts or scale-ups

Encourage the team to think about dependencies and business readiness.

3. Place Agent Strategies (30 minutes)

Take each selected agent concept and ask:

- *“When should we start this work?”*
- *“What needs to happen first?”*
- *“Is this a quick build, a phased release, or a bigger systems shift?”*

Stick each concept in the phase that feels most realistic.

Group related agents together where possible (e.g. customer support agents before escalation agents).

Facilitator Tip:

Don't over-sequence the map, this is a shared compass, not a sprint backlog.


4. Quick Review & Adjustments (10 minutes)

Step back and ask:

- “Does anything feel too soon or too far away?”
- “Are we overlooking a prerequisite?”
- “Are there any concepts competing for the same resources?”

Make minor tweaks. Keep the energy brisk.

It should evolve something like this:

**5. Assign Owners and Check-In Points (15 minutes)**

Now lock in momentum by naming ownership.

For each roadmap item:

- Ask: “***Who will lead the next step?***”
 - Write the name on a sticky note and place it next to the concept.

- Then ask: **“When should we check in again?”**
 - Note a loose date (e.g. “early Q4” or “Feb 2026”).

Facilitator Tips:

- One owner per concept is enough.
- If no one owns it, question whether it’s viable.
- Use this moment to surface resource needs or support gaps.

6. Decider Sign-Off (5–15 minutes)

Say:

“Let’s check in with our Decider to make sure the roadmap reflects where we need to go.”

The Decider should:

- Review all roadmap items.
- Confirm timing and accountability.
- Flag any misalignments or resource gaps.

Once approved, this roadmap becomes the team's **execution charter**.

Milestone

The team now has a visible, time-bound Agentic AI Product Roadmap showing:

- Which agent strategies will be activated and when.
- Who is leading each initiative.
- When the next decision pulse will happen.

This becomes your foundation for technical mapping, prototype development, and business alignment going forward.

Stage 9: Wrap Up & Build Readiness

Exercise Name: Workshop Closing

Duration: 10 minutes

This final stage is about sealing clarity, commitment, and momentum.

Over the past two days, the team has turned complex systems, decision bottlenecks, and AI hype into a **clear Agentic AI product strategy**. This closing step confirms alignment, finalises the agent spec, and outlines what happens next—whether that's prototyping, procurement, or internal builds.

A strong close creates ownership and turns strategy into movement.

Resources

- Snapshot of the final roadmap.
- Copies or screenshots of chosen agent strategies.
- Optional: blank sticky notes or a shared document for reflections.

How To Deliver it

1. Recap the Journey (3 minutes)

Say:

“Over the past two days, you’ve built something exceptional. You started by mapping real friction and uncovering the decisions that hold your business back. Then, you designed intelligent agent strategies, prioritised the best ones, and created a roadmap to build them. You didn’t just talk about AI—you architected how it will work for your business.”

Quickly revisit:

- The long-term opportunity you’re aiming to unlock.
- The best-fit agent strategies selected.

-  The roadmap and next 90-day focus.
- Owners and delivery momentum.

Keep this short and focused. Let the team see the shape of their own work.

2. Acknowledge the Work (2 minutes)

Say:

“Building agentic systems isn’t easy, it means changing how we think about products, decisions, and automation. This team has shown deep insight, creativity, and forward-thinking. Thank you for your openness and drive.”

Pause for a moment so it lands. This reinforces shared accomplishment.

3. Lock in the Next Steps (3–5 minutes)

Clarify how momentum continues post-workshop.

Use this checklist as your guide:

-  **How outputs will be shared**
(e.g. Google Drive folder, Notion hub, internal doc)
- **When you’ll regroup**
(e.g. “Let’s do a check-in 2 weeks from now to review builds in motion.”)
- **What happens next**
 - Prototype build.
 - Internal stakeholder alignment.
 - Vendor discovery or data integration prep.

- **How progress will be communicated**

(e.g. Slack thread, executive briefing, async updates)

Optional but powerful:

Appoint a “**Roadmap Keeper**” to ensure priorities stay visible and timelines don’t drift.

4. Optional: One-Sentence Reflection (2 minutes)

If time allows, close with a fast reflective moment:

Say:

“On a sticky note or in the chat, write: What are you most excited to see this agent system do in the real world?”

Let 2–3 people share their reflections. This energises the close.

Milestone

By the end of this step, the team:

- Has total clarity on which Agent strategies will move forward.
- Understands exactly what the next steps look like.
- Knows how and when progress will be reviewed.
- Feels a sense of pride, purpose, and urgency.

Leaves not with a vague ambition, but with a **working Agentic AI product strategy** that’s about to go live.

That’s how we move from friction to function, through strategic AI product design that thinks, remembers, and acts.

Let the build begin.

Post-Workshop Follow-Up

To ensure the Agentic AI strategy turns into tangible business outcomes, the following actions will guide implementation and alignment across the next phases:

1. Workshop Outcomes Summary (within 48 hours)

All participants will receive a packaged summary including:

- **Finalised Agent Strategy Deck:** including key concepts, priorities, and memory logic.
- **Decision Maps:** visual maps of friction points and agent handoff zones.
- **Agentic Roadmap:** strategy timeline with owners and next checkpoints.
- **Implementation Tracker:** initial actions, due dates, and ownership for each agent concept.
- **Calendar Invites:** for follow-ups and status check-ins

2. First Check-In (Week 2)

A 30-minute sync to:

- Review what's been kicked off and by whom.
- Unblock early challenges (e.g. data access, team capacity).
- Clarify any confusion around next steps.
- Share early signals or wins from prototypes or vendor exploration.

3. Quarterly Reviews (Every 3 Months)

Create momentum and keep direction aligned by reviewing:

- Progress against the roadmap (by strategy and owner)
- Any changes in data availability, tech stack, or product constraints

- Emerging opportunities or agent capabilities that could be added.
- Strategic pivots needed based on market or internal shifts.

These reviews act as **guardrails for iteration** rather than static plans.

4. Six-Month Refresh Workshop

Run a 3–4 hour refresh session to:

- Refine long-term outcomes based on new learnings or executive direction.
- Reassess chosen agents and explore potential 2.0 upgrades.
- Evaluate performance across Memory, Reasoning, and Handoff dimensions.
- Identify capability gaps (e.g. new tooling, process adoption, org design).
- Add new agent concepts or adjacent use cases (e.g. onboarding → renewals).

Optional: use this session to **review vendor build vs. internal ownership trade-offs**.

Ongoing Guidance & Optional Support

Depending on your preferred delivery model, Synaptix can support via:

- Fortnightly sprint advisory (design reviews, prompt stack QA, scaling roadmaps).
- Train-the-trainer sessions to build internal capability.
- Memory architecture or vector store integration planning.
- Vendor technical due diligence.

We're here to ensure your Agentic AI systems don't just launch, they scale, sustain, and succeed.

Next Step: Keep a shared workspace active (Notion, Miro, etc.) to track deliverables, owners, and cross-functional dependencies. This is where the build begins. Let's turn intelligent product strategy into intelligent execution.

Final Thoughts

You've now built something most teams haven't even started thinking about.

In just two days, you've gone from questions and friction to a working blueprint for agentic AI inside your business. You've designed agents that don't just automate tasks, they observe, remember, escalate, and reason. You've turned pain points into triggers, built memory systems that reflect real behaviour, and mapped decision points that shape outcomes.

Most teams are still asking how AI fits. You've answered that by showing exactly where it drives value, what data it needs, and how to scale it safely.

This wasn't about trends or tech theatre. It was about fixing the real inefficiencies slowing your business down, rework loops, manual handoffs, judgment bottlenecks. And now you've turned those into systems that can learn, act, and scale.

What you've built isn't a plan on paper. It's a launchpad. You've created agent roles, logic flows, safety rails, and roadmaps. You've aligned decision-makers. You've chosen where to start and who will lead. You're ready to move.

While most teams are still on the sidelines waiting for AI to *"just work,"* you've put yourself in the driver's seat. You've moved faster, with more structure, and more intention.

The edge doesn't come from being first. It comes from being ready. Now go build it.

Mopping Up & Templates

You Did It. Now Don't Let It Sit on the Shelf

You've just completed your first Agentic AI Sprint, and that's no small thing.

Over two focused days, you didn't just talk about AI. You designed it. You mapped real friction in your workflows. You uncovered where decisions break down. You imagined agents that could reason, remember, and act. And then you made tough calls about which ones matter most to build.

Now, you have something few teams ever get: a shared, strategic view of where AI belongs in your business, and a roadmap to make it happen.

But this is just the beginning.

The real value starts when you turn strategy into systems. When the first agent goes live. When you test it, learn from it, and iterate. When your team begins to think differently, not about adding AI, but about *running the business on it*.

You're now in a different category.

You're not chasing the AI hype.

You're shaping your workflows, your customer experience, and your internal operations around intelligent automation. And you're doing it with purpose, not panic.

This sprint gave you the framework. The next step is execution.

Pick one agent and build it. Run another sprint. Teach your team. Share the workshop. And when you're ready to scale, we're here to help, with tools, templates, cohorts, and enterprise workshops.

One last reminder. The companies that win in this wave won't be the ones with the most AI tools. They'll be the ones that know *how to think in systems*. You've just taken the first step.

Now, go finish what you started. Let's build.

Templates

Customer Journey Map Template

How to Use This Template:

1. Replace placeholder text with your specific journey information

2. Add or remove stage boxes based on your journey length (typically 5-10 stages)

3. Use bullet points in data cells for easy scanning

4. Keep descriptions concise but specific

5. Consider adding emojis or icons to the "Doing" section for visual appeal

Goal

[Define the overarching goal of this journey - What are you trying to achieve?]

Main Stage Workflow

Before

During

After

Defining Your Stages:

Before:

Awareness, research, decision-making phases

During:

Active engagement, transaction, or service delivery

After:

Follow-up, support, retention activities

Stages of activity

1

[Stage Name]

2

[Stage Name]

3

[Stage Name]

4

[Stage Name]

5

[Stage Name]

6

[Stage Name]

Stage Naming Tips:

Use action-oriented names (e.g. "Submit Request" not just "Request")

Keep names short (2-3 words max)

Ensure stages flow logically from one to the next

Pain points

List frustrations:

What's difficult?

What causes delays?

What confuses users?

List frustrations:

What's difficult?

What causes delays?

What confuses users?

List frustrations:

What's difficult?

What causes delays?

What confuses users?

List frustrations:

What's difficult?

What causes delays?

What confuses users?

List frustrations:

What's difficult?

What causes delays?

What confuses users?

List frustrations:

What's difficult?

What causes delays?

What confuses users?

Needs

What users need:

Requirements

Expectations

Desired outcomes

What users need:

Requirements

Expectations

Desired outcomes

What users need:

Requirements

Expectations

Desired outcomes

What users need:

Requirements

Expectations

Desired outcomes

What users need:

Requirements

Expectations

Desired outcomes

What users need:

Requirements

Expectations

Desired outcomes

Doing

User actions:

What they do

How they do it

Tools they use

Consider adding icons

User actions:

What they do

How they do it

Tools they use

Consider adding icons

User actions:

What they do

How they do it

Tools they use

Consider adding icons

User actions:

What they do

How they do it

Tools they use

Consider adding icons

User actions:

What they do

How they do it

Tools they use

Consider adding icons

User actions:

What they do

How they do it

Tools they use

Consider adding icons

Data Sources used

List data/systems:

Databases

Documents

Analytics

External sources

List data/systems:

Databases

Documents

Analytics

External sources

List data/systems:

Databases

Documents

Analytics

External sources

List data/systems:

Databases

Documents

Analytics

External sources

List data/systems:

Databases

Documents

Analytics

External sources

List data/systems:

Databases

Documents

Analytics

External sources

Technology

Tech/tools used:

Software

Platforms

APIs

Hardware

Tech/tools used:

Software

Platforms

APIs

Hardware

Tech/tools used:

Software

Platforms

APIs

Hardware

Tech/tools used:

Software

Platforms

APIs

Hardware

Tech/tools used:

Software

Platforms

APIs

Hardware

Tech/tools used:

Software

Platforms

APIs

Hardware

Journey Map Best Practices:

Be specific:

Use real examples and actual quotes from users

Focus on emotions:

Pain points should capture frustration, needs should express desires

Stay factual:

Base your map on research, not assumptions

Keep it visual:

Consider adding icons, emojis, or screenshots

Update regularly:

Journey maps should evolve as processes change

Optional Enhancements:

Add an "Opportunities" row to highlight improvement areas

Include "Emotions" or "Feelings" row with emotion indicators

Add "Touchpoints" to show channels of interaction

Include metrics like time spent or satisfaction scores

MONTH 1	MONTH 2	MONTH 6	9 - 12 MONTHS	12+ MONTHS

**Thanks
To Those
Who Got
This Here**

Acknowledgements

While I've been the face of this sprint and guide, it wouldn't exist without the help, ideas, and grit of others.

Chris Tyson

Co-founder of Templonix and one of the sharpest AI system thinkers I've had the pleasure of working with. Chris helped translate an agent technology architecture into something product teams could actually use. Without his technical brain and ability to turn complex logic into teachable flows, this two-day workshop would've stayed a deck on a shelf. You can learn more about his work at templonix.com, it's where the framework powering this sprint was born.

Cameron Fehr

A long-time friend and a natural educator. Cameron's background in building educational systems helped me break down agentic product strategy into something that *clicks* for teams in just two days. His step-by-step thinking, classroom wisdom, and talent for sequencing learning moments had a massive influence on how the sprint flows.

To both of you, thank you for believing in the mission of helping business teams build smarter systems.

References


Knapp, J., Zeratsky, J., & Kowitz, B. (2016). *Sprint: How to Solve Big Problems and Test New Ideas in Just Five Days* The OG design sprint framework. We evolved and adapted many of its facilitation principles for an AI-first world.

Rumelt, R. (2011). *Good Strategy / Bad Strategy: The Difference and Why It Matters*. A must-read for anyone tired of vague, buzzword-filled plans. Rumelt's thinking on leverage, clarity, and real choices deeply influenced how we structure agentic strategy.

Alex Hormozi *\$100M Leads: How to Get Strangers to Want to Buy Your Stuff*

Hormozi's "one channel, one avatar, one goal" principle is baked into how we think about agent capability scoping, simple focus wins every time.

Strategy Sprint by AJ&Smart Their approach to rapid strategy formulation inspired how we shaped the two-day arc of this sprint, moving from chaos to clarity with velocity.



TIM DAINES

SYNAPTIX SOLUTIONS