// A STRATEGIC POINT OF VIEW · APRIL 2026

Agent-Ready

A strategic point of view on preparing for the agentic economy

Kristian Arvidsson

CHIEF AI OFFICER

EXECUTIVE BRIEFING · 20 MINUTES

ARVIDSSON.TECH

// BEFORE YOU START

What you will get from this paper

This paper is a strategic point of view on the agentic economy: what is changing, where most organizations are getting stuck, and how leaders can prepare for two realities at once: serving AI agents and deploying them internally. It is not a market forecast. It is a practical framework built from implementation work, operating experience, and a clear thesis about where business is heading.

The gap between what vendors promise and what actually works is enormous. This paper cuts through it. No pitches. No hype cycles. Just what I have learned from building agent systems, advising on implementation, and seeing firsthand what delivers results.

The strategic question is not whether AI agents will matter. It is where they will sit in the flow of work and commerce, and what that means for your business model. Some agents will act inside your company. Others will increasingly act outside it on behalf of customers, partners, and suppliers. The companies that win will not be the ones with the loudest AI story. They will be the ones that become easier to verify, easier to reach, better governed, and richer in usable context.

01	Introduction	03
02	Adoption Without Transformation	05
03	How to Make Your Organization Ready for the Agentic Economy	06
04	A Readiness Framework	12
05	What to do Monday morning	13
06	A final note · About the Author	14
07	Sources	15

// 01 INTRODUCTION

What is the agentic economy?

The internet is splitting into two operating modes: one built for humans, and one increasingly navigated by software acting on behalf of humans.

A parallel web is emerging in which AI agents increasingly discover, evaluate, and execute tasks on behalf of people and organizations. They do not behave like human users. They optimize for structured information, machine-readable trust, and fast execution. They book flights, evaluate vendors, compare prices, sign contracts, and move money. They operate 24/7, they do not get tired, and they do not care about your brand colors.

KEY INSIGHT

The agentic economy is the shift from humans doing business through software to software doing business on behalf of humans. Autonomous actors that negotiate, purchase, and execute.

This is not science fiction. Anthropic introduced the Model Context Protocol (MCP) in November 2024 as an open standard for connecting AI tools to data sources and services. [1] Adoption has been rapid, and MCP is a strong signal that machine-callable interfaces are becoming first-class infrastructure.

Why should executives care?

Twenty years ago, if your business was not on the internet, you were invisible. It did not matter how good your product was. Customers moved online, and the companies that did not follow became irrelevant.

KEY INSIGHT

The same shift is happening again. Only faster.

When AI agents become the primary way your customers discover, evaluate, and transact with services, every assumption about how your business operates gets challenged. Your website might be invisible to agents. Your sales process might be incompatible with autonomous purchasing. Your data might be inaccessible to the systems that need it.

SIGNALS TO PAY ATTENTION TO

Market forecasts vary widely, but most point in the same direction: rapid growth in agentic AI over the rest of the decade. [3] What matters more than the exact number is the operating shift: experimentation is giving way to deployment, and deployment is forcing questions about trust, governance, and interface design. [7, 8, 6]

Having a website became table stakes in the last platform shift. My bet is that agent-readiness becomes table stakes in the next one. I would present that as a strategic forecast, not a certainty. But the cost of ignoring it is high enough that leaders should prepare now.

What the early movers stand to gain

The threat is real. But so is the upside.

When agents become customers, they do not just replace human buyers. They expand the market. An agent operates 24/7 across every time zone. It evaluates more options, transacts faster, and handles complexity that a human buyer would skip entirely. The companies that are ready for agent-driven commerce will access demand that literally did not exist before.

On the other side, companies that successfully deploy agents internally are seeing step changes in output. Not 10% improvements. Multiples. BCG argues that agentic AI could unlock up to $200 billion in net new demand in technology services over the next five years. [6] Whether or not that exact number holds, the directional point is clear: agentic systems are creating new value pools, not just incremental productivity gains. That value goes to the companies that move first.

Here is the thing. Both sides compound. When your business is trustworthy and discoverable to agents, you capture revenue from agent-driven commerce. When your internal agents are well-governed and loaded with context, your cost of execution drops. Revenue up, costs down. That is not an incremental improvement. That is a structural advantage.

KEY INSIGHT

The companies that prepare for both sides, serving agents and employing agents, will operate at a speed and scale that unprepared competitors simply cannot match.

My expectation: the gap will not close over time. It will widen.

// 02 THE STATE OF PLAY

Adoption Without Transformation

Adoption is moving faster than redesign. That is the real story.

84%

of companies have not redesigned jobs around AI capabilities.

DELOITTE, 2026 [5]

3×

Employees are roughly three times more likely to be using generative AI at work than leaders estimate.

MCKINSEY, 2025 [7]

<10%

of enterprises have scaled agentic systems to deliver tangible value.

MCKINSEY, 2026 [7]

Deloitte's 2026 State of AI in the Enterprise reports that 84% of companies have not redesigned jobs around AI capabilities. [5] McKinsey's workplace research found employees are already using generative AI far more than leaders expect. [7] And McKinsey's more recent work on agents suggests that while many enterprises are experimenting, fewer than 10% have scaled agentic systems to deliver tangible value. [7]

What are the biggest challenges?

Automating the symptom.

The same pattern plays out at organization after organization. A team builds a proof of concept. The demo looks great. Leadership gets excited. The project moves to production. And then it quietly stalls.

The demo worked because the problem was scoped and the data was clean. Production failed because nobody could articulate what the AI should do when things got ambiguous. When the input is incomplete. When the model is uncertain. When a decision requires organizational context that lives in people's heads, not in any system.

Most AI initiatives are measured by deployment speed, adoption rates, or cost savings. None of those tell you whether the AI is solving the right problem. A team deploys an AI chatbot to fix customer service in two weeks and calls it a success. But the real issue is unclear product descriptions causing the inquiries. They automated the symptom, not the cause. The metric looked great. The underlying problem got worse.

// 02 THE STATE OF PLAY

KEY INSIGHT

The organizations that fail with AI give it freedom. The ones that succeed give it structure.

The data problem.

Fewer than half of organizations say their data quality and accessibility are adequate for AI adoption. That is the real bottleneck. Your AI strategy is only as strong as your context layer: the quality, accessibility, and governance of the data your systems rely on. If an AI agent needed to understand your business tomorrow, could it? Or would it drown in scattered spreadsheets, outdated wikis, and knowledge that only exists in the heads of employees who have been there long enough to remember why things are the way they are?

The accountability vacuum.

When an AI agent makes a commitment on behalf of your company, who is accountable? Most organizations have no governance framework for autonomous AI actions. This is not just a risk management issue. It is a business viability issue. My read: "The AI did it" will not survive a courtroom, and it will not survive your board either.

// 03 GETTING READY

How to Make Your Organization Ready for the Agentic Economy

There are two sides to this shift. Your business will interact with agents in two fundamentally different ways, and each requires its own preparation.
On one side, agents will become your customers. Other companies' and consumers' agents will evaluate, discover, and transact with your business. On the other side, you will employ agents. They will work inside your organization, acting on your behalf, carrying your name.

SIDE ONE

When agents become your customers

1. Make your business trustworthy to agents
2. Make your business discoverable by agents
3. Make your business reachable by agents

SIDE TWO

What you need to employ agents

1. Give your agents your context
2. Govern what your agents are allowed to do
3. Judgment

Most organizations are thinking about neither. The ones that prepare for both will lead.

// 03 SIDE ONE

When Agents Become Your Customers

Your next customer might not be a person. It might be a person's agent.

When that agent lands on your business, it will not browse your website, admire your branding, or read your "About Us" page. It will look for machine-readable signals it can verify in milliseconds. If an agent cannot verify you programmatically, you are at a disadvantage immediately. In some categories that may mean it routes around you altogether. In others it may escalate to a human. Either way, the burden shifts toward machine-readable trust and away from brand theater.

This is the new front door to your business. And most companies have not built it yet.

Make your business trustworthy to agents

Right now, your customers evaluate your business with their eyes. They read reviews, talk to sales, make a gut call. AI agents do none of that.

A customer's agent looks for trust signals it can parse. API uptime. Security certifications. Payment processor reputation. Verified identity. Structured data. If it cannot verify you, it routes to a competitor. Not because they are better. Because they are verifiable.

Stripe processed $1.9 trillion in transactions last year. [4] "Powered by Stripe" already functions as a trust signal for humans. For agents, payment and identity infrastructure may become a machine-readable trust signal with outsized weight in supplier selection. Verified rails will matter more when the evaluator is software, not a person.

01

WHAT TO DO NOW

- [ ] Audit your business the way an agent would: probe what it can programmatically verify without talking to a human.
- [ ] Invest in machine-readable trust signals: verified payment processing, API reliability metrics, security certifications, structured business data.
- [ ] If your competitor has these signals and you do not, their customers' agents will choose them by default.

// 03 SIDE ONE

Make your business discoverable by agents

Building is free. Discovery is not.

When AI makes it trivial to ship new software, most of what gets built will never be seen by a single customer. When supply becomes infinite, curation becomes the scarcest resource. Google, Apple, Amazon, and TikTok get stronger in an AI-saturated world because they decide who gets attention.

But for agents, the discovery problem is entirely new. If a customer's agent needs to find a supplier, a service, or a partner, how does it find you? There is no agent-native app store yet. My bet: the company that builds it will own the commerce layer for the next decade. Until then, your discoverability depends on what you build today.

02

WHAT TO DO NOW

- [ ] Make your services accessible to agents, not just humans. Do you have an API? Structured data? An MCP server?
- [ ] Consider your position in the distribution chain. Are you a gatekeeper, a participant, or invisible?
- [ ] Distribution strategy is no longer a marketing line item. In the agentic economy, it is a survival question.

Make your business reachable by agents

Agents do not pick up the phone. They do not fill out contact forms. They do not wait on hold. If your business cannot be reached the way agents communicate, you are not reachable at all.

Today, agents reach businesses through the same channels humans do. Email is the obvious first one. AgentMail's March 2026 fundraise is one signal that founders see agent-to-business communication as an emerging infrastructure category, with email as an obvious starting point. [2] That does not make email the long-term winner, but it does suggest that agent-reachable business channels are becoming a real design problem. Tomorrow it will also be MCP servers, agent-callable APIs, and protocols that have not been named yet.

The companies winning here are the ones that already treat their inboxes, support channels, and APIs as machine-callable, not just human-facing. Auto-responses an agent can parse. Structured replies. APIs that return clean data instead of HTML pages built for human eyes.

03

WHAT TO DO NOW

- [ ] Audit how an agent would actually reach your business today. Is there a path that does not require a human?
- [ ] Make sure your email, support, and API layers return structured, parseable responses. Agents cannot read screenshots.
- [ ] Treat machine-callable communication channels as first-class infrastructure, not an afterthought.

// 03 SIDE TWO

What You Need to Successfully Employ Agents

The other side of the coin. Your organization will deploy agents that draft documents, communicate with customers, evaluate vendors, and potentially move money. They will act on your behalf. They will carry your name.

Getting this right is the difference between a competitive advantage and an expensive liability.

Give your agents your context

The model is a commodity. Claude, GPT, Gemini. Pick any of them. What makes them powerful is your context. Your company data. Your customer relationships. Your internal processes.

KEY INSIGHT

An agent without your context is a chatbot. An agent with your context is a junior employee who never sleeps.

This is what most AI strategies get wrong. They start with the technology. They should start with the data. If you deploy an agent tomorrow and it cannot access your company's knowledge, it will make things up. It will give customers wrong information. It will make decisions based on generic training data instead of your specific business reality.

Three signs your data is ready for agents: you own your data and its pipelines without relying on third parties to access it. You can share data across the organization so everyone, human and agent, operates from the same information. And you use your data to test new opportunities, not just report on past performance. If you cannot check all three, your agents will be flying blind.

Fewer than half of organizations say their data quality and accessibility are adequate for AI adoption. That is the real bottleneck. Your AI strategy is only as strong as your context layer: the quality, accessibility, and governance of the data your systems rely on.

01

WHAT TO DO NOW

- [ ] Map where your critical business knowledge lives. Is it structured, accessible, and correctly permissioned?
- [ ] Invest in making organizational knowledge machine-readable. The intent gap closes when your context is explicit.
- [ ] Treat your data layer as strategic infrastructure, not an IT project. Your agents are only as good as the context you give them.

// 03 SIDE TWO

Govern what your agents are allowed to do

Your agents will file documents, draft contracts, communicate with customers, and potentially move money. When one of them makes a mistake, and it will, who is accountable?

Not the model provider. Not the tool vendor. You.

"The AI did it" is not going to survive a courtroom. It is not going to survive your board either. Every agent that acts on behalf of your company carries your name, your reputation, and your liability.

The better AI gets, the more serious this becomes. A hallucination in an internal chat is annoying. A hallucination in a customer-facing email is a reputation problem. A hallucination in a financial report is a lawsuit.

Governance is moving from compliance concern to operating requirement. As enterprises push from experimentation into real deployment, the question is no longer whether AI needs oversight. It is where accountability sits when software acts on behalf of the company.

02

WHAT TO DO NOW

- [ ] Define what your agents are allowed to do without human approval. Write it down. If it is not documented, it is not governed.
- [ ] Assign accountability for every category of agent action. If an agent sends an email, signs a commitment, or makes a purchase on your behalf, someone specific owns that outcome.
- [ ] Build governance before you scale your agents, not after. Retrofitting accountability after an incident costs lawsuits, lost trust, and months of remediation. Building it in from the start barely costs anything.

// 03 SIDE TWO

Judgment

Judgment is the ability to evaluate AI output, direct it with intent, and decide what should exist. When execution gets cheaper, deciding what is worth building becomes the differentiator.

When producing things is free, what you choose to produce becomes the entire game. Product decisions. Design sensibility. Editorial judgment about what is worth building and what is not.

When production becomes free, the temptation is to build everything. Ship it all. But "fast" without direction is just expensive noise. The hardest part is not getting AI to produce output. It is deciding which output is worth keeping.

Your agents will have the same problem at scale. They can produce endlessly. The question is whether what they produce is any good. That requires people with domain expertise and judgment. People who can look at what an agent did and know whether it is right or wrong.

On the agentic web, this shows up as orchestration quality. The winning agent systems will not have the best underlying model. They will be the ones where someone with deep domain expertise carefully tuned the workflows, chose the right tools, and made a thousand small editorial decisions about how the agent should behave.

03

WHAT TO DO NOW

- [ ] Invest in people who can evaluate AI output, not just people who can prompt it (that is commodity work).
- [ ] When execution gets cheaper, differentiation shifts upstream: better judgment about what to build, what to automate, what to approve, and what to ignore.
- [ ] The executive who cannot evaluate what AI produces cannot lead in the agentic economy. Stay close to the material.

// 04 ASSESS YOURSELF

A Readiness Framework

Ten questions to assess your position in the agentic economy. Score yourself honestly.

WHEN AGENTS ARE YOUR CUSTOMERS

01	If an AI agent evaluated your business for trustworthiness today, would it pass?
02	Do you have third-party verification that agents could use to validate your services automatically?
03	If a customer's agent looked for your services today, would it find you?
04	Do you have an agent-accessible interface (API, structured data, MCP server)?
05	Can an agent communicate with your business and get a parseable response without a human in the loop?

WHEN YOU EMPLOY AGENTS

06	Is your critical business knowledge structured and machine-accessible?
07	Do you have people who can evaluate AI-generated output and know whether it is right or wrong?
08	When execution is free, do you have a distinct point of view on what should exist?
09	Do you have a governance framework that documents who is accountable when an AI agent acts autonomously on behalf of your company?
10	Do you have at least one AI agent in production, operating autonomously on a real business task?

8-10
YES ANSWERS

You are ahead of most. Formalize and operationalize what you have.

4-7
YES ANSWERS

You have gaps. Prioritize the questions where you answered no.

0-3
YES ANSWERS

The agentic economy is arriving whether you are ready or not. Start with context and governance. Everything else builds on those.

// 05 NEXT STEPS

What to do Monday morning

You do not need to transform your organization overnight. But you do need to start.

1 **Audit your context layer.** Where does your company's knowledge live? Could an AI agent access it? This is the foundation everything else depends on.

2 **Define your governance boundaries.** What are your AI agents allowed to do autonomously? Who is accountable? Write it down before you scale.

3 **Test your trust signals.** Evaluate your business the way an agent would. What machine-readable indicators of trustworthiness exist? What is missing?

4 **Ask the strategic question.** What do you own that gets more valuable as AI gets better? If the answer is nothing, your positioning needs to change.

The companies that treated the internet as a fad in 1998 spent the next decade catching up. The companies that dismissed mobile in 2008 spent five years rebuilding everything. The window to build defensible positions gets shorter each time.

KEY INSIGHT

AI made execution free. What it didn't make free: trust, context, accountability, distribution, and judgment. That is where companies will compete next.

The question is not whether to engage with the agentic economy. It is whether you will be positioned to lead in it.

WATCH OUT FOR THESE WARNING SIGNS

- An agent looking for your services today would not find you, or could not trust you
- Your knowledge lives in spreadsheets and people's heads, not in systems
- Nobody is accountable when an AI acts on your behalf
- Your AI initiatives are scoped to "use AI", not to a specific business outcome
- You are not getting real value from the agents you have deployed

// 06 CLOSING

A final note

LLMs do not come with instruction manuals. Not for how to use them, and certainly not for how to extract as much value as possible. Old software did. You bought a product, read the documentation, followed the steps. There was a right way to use it. That era is over.

There is no universal deployment playbook for agents. Principles transfer. Patterns transfer. But operating designs still have to be tested against your workflows, your data, your customers, and your risk profile.

This means treating your AI strategy the way emergency rooms treat patients: treat the patient, not the protocol. The playbook matters, but the moment it stops matching reality, you have to adapt. The organizations that succeed will be the ones willing to step outside the usual methods, form a thesis, test it, and adjust.

There is no single right answer. There is only what works for you, right now, with what you have. Start there.

// ABOUT THE AUTHOR

Kristian Arvidsson

I am a Chief AI Officer who still pushes code every week. I build agent systems, evaluate AI tools, and advise on implementation. Before that, I founded an API-first e-commerce backend from the ground up.

> *Innovation happens beyond the documentation.*

This paper comes from that hands-on perspective. The gap between what vendors promise and what actually works is enormous, and the only way I can tell the difference is by staying close to the material. I do not write about AI from a distance. I build with it every day.

Connect with me on LinkedIn.

// SOURCES

Sources

1 Anthropic, "Introducing the Model Context Protocol" (November 2024). Open standard for connecting AI tools to data sources and services. modelcontextprotocol.io

2 "AgentMail raises $6M to build email infrastructure for AI agents." TechCrunch, March 2026.

3 Precedence Research, "Agentic AI Market Size, Share, and Trends." Market forecasts across research firms consistently point to rapid growth in agentic AI through the end of the decade. Exact multi-year figures vary by methodology and are revised frequently; cited here for directional magnitude only.

4 Stripe, "2025 Annual Letter" (February 2026). $1.9 trillion in total payment volume processed in 2025, up 34% year over year.

5 Deloitte, "State of AI in the Enterprise" (2026). 84% of companies report they have not redesigned jobs around AI capabilities. deloitte.com/cz-sk/en/services/consulting/research/the-state-of-ai-in-the-enterprise.html

6 BCG, "The $200 Billion AI Opportunity in Tech Services" (2026). Estimates agentic AI could unlock up to $200 billion in net new demand in technology services over the next five years. bcg.com/publications/2026/the-200-billion-dollar-ai-opportunity-in-tech-services

7 McKinsey, "Superagency in the Workplace: Empowering People to Unlock AI's Full Potential at Work" (2025). Employees are roughly three times more likely to be using generative AI at work than leaders estimate. Related 2026 McKinsey analysis indicates that while many enterprises are experimenting with agents, fewer than 10% have scaled agentic systems to measurable business value.

8 PwC, "2026 AI Business Predictions." AI moves from experiments to enterprise-scale impact; trust and governance positioned as prerequisites for deployment.