

Orion + Zapier

# The Five Confrontations of Agentic Systems

What AI Exposes About Your Organization, and  
What to Do Before It's Too Late



---

# Table of Contents

**03**

Executive Summary

**04**

Introduction

**06**

CONFRONTATION 1:  
The Speed Trap

**08**

CONFRONTATION 2:  
The Clarity Test

**11**

CONFRONTATION 3:  
The Human Shift

**14**

CONFRONTATION 4:  
The Delivery Reckoning

**17**

CONFRONTATION 5:  
The Customer Verdict

**20**

Conclusion: Built to  
Last

**21**

About Us

# Executive Summary

This ebook gives leadership teams a five-part diagnostic for agentic readiness. Each confrontation identifies a specific organizational pressure that agentic systems expose—from execution speed to cross-functional clarity to customer-facing trust—and ends with a concrete starting question designed to surface where your organization stands today.

It's built for teams that have moved past "should we adopt AI" and are now asking "why isn't it working at scale." The pattern behind that question is consistent: pilots succeed in isolation, then stall when they hit the operating realities of the broader business. The five confrontations name those realities: speed without structural capacity, clarity that doesn't cross departmental lines, workforce adoption without genuine capability building, delivery that ships without measuring impact, and customer experiences that overpromise what the organization can't yet sustain.

Each confrontation builds on the last. Together, they form a framework for leaders who are ready to stop experimenting and start building something that lasts.

## How to use this book:



### Read it as a team

The diagnostic questions are designed for leadership teams, not individuals



### Follow the sequence

Each confrontation builds on the last — they're designed to compound



### Sit with the discomfort

The diagnostic questions are meant to surface gaps



# Introduction

---

For the past two years, the AI conversation has centred on capability. That framing is increasingly incomplete.

As intelligent systems become agentic—able not only to advise but to initiate actions, coordinate tasks, and execute workflows—the more consequential question shifts from *what can AI do?* to *what is AI about to force us to address?*

Agentic systems compress the distance between intent and action. That sounds like an efficiency story, and at the surface level it is. But underneath, it creates pressure that radiates across the entire organization. Speed exposes misalignment. Automation reveals unclear ownership. Faster delivery surfaces broken governance. And customers—whose expectations are now being calibrated by every AI-powered experience they encounter—begin to see through the gap between what an organization promises and what it actually delivers.

This ebook names five confrontations that define the transition: the speed trap, the clarity test, the human shift, the delivery reckoning, and the customer verdict. They aren't a sequence to resolve — they're pressures that arrive in parallel and compound against each other.

And they point to the same uncomfortable truth: agentic AI won't reward the organizations that adopt it fastest. It will reward the organizations that face the operational realities of the agentic era head on.

Each of the chapters that follow ends with a diagnostic question, a place to start examining where your own organization stands. These aren't theoretical. They're designed to surface the specific gaps that will

determine whether agentic systems accelerate your business or simply amplify what's already broken. Read them as a leadership team and answer them honestly— the discomfort is the point.



Agentic AI won't reward the organizations that adopt it fastest. It will reward the organizations that face the operational realities head on.



#### DIAGNOSTIC QUESTION

If your organization deployed an agentic system tomorrow — what would it accelerate, and what would it break?



## CONFRONTATION 1:

# The Speed Trap

---

Agentic AI compresses the distance between intent and action. Work that once moved through stages of analysis, coordination, review, and execution can increasingly happen in one continuous flow, with people supervising rather than manually carrying each step.

That creates a new baseline for speed, both inside organizations and across markets. But here's the finding that should give every executive pause:

**AI does not create speed out of thin air. It amplifies the speed an organization is structurally capable of sustaining.**

This is why some companies see immediate gains while others hit friction. Organizations that are structurally capable of moving fast—clear ownership, disciplined workflows, defined decision rights—absorb the acceleration. The rest discover that the technology works, but the business still can't move quickly.

The competitive consequence is what we call **competitive compression**. When the cost of acting on an idea drops, more organizations operate at higher speed. Smaller teams coordinate complex work more easily. New entrants move faster than their size should allow. The baseline expectation for responsiveness rises across an entire industry category.

A company that takes weeks or months to update pricing, correct product content, or respond to

customer friction is no longer competing against last year's benchmark. It's competing against businesses that sense, decide, and act in hours or days.

The defining divide won't be between organizations that have AI and those that don't. It will be between organizations that built the execution capacity to absorb whatever comes next, and those that kept chasing the tool instead of building the machine.

## Where to start

---

Pick one function you know well. Think about the last time the team identified something that needed to change, such as a pricing error, a product data issue, or a campaign that needed updating. How long between the moment someone noticed and the moment it was resolved? That gap is your current execution speed. Everything in this ebook is about closing it.

But before you try to close these smaller gaps, pay attention to where some of the larger delays actually live. It's rarely inside the function itself. It's almost always in the handoff to another team, another system, or another approval chain. That's the subject of the next confrontation.



## CONFRONTATION 2:

# The Clarity Test

---

The first confrontation asked whether you can move fast within a function. This one asks a harder question: can anyone in your organization see how work moves across them?

Most AI initiatives start inside a single team. Marketing automates content workflows. Sales deploys an AI assistant. Finance builds a reconciliation bot. Within each domain, the results can be real. But the business doesn't run inside departments—it runs in the flows between them. A customer inquiry that starts in marketing, moves to sales, triggers fulfillment, and lands in support doesn't care about your org chart. It cares about whether context survives the journey.

In most organizations, it doesn't.

Information that exists in one system never reaches the team that needs it. Sales doesn't know what marketing knows about a customer's behavior. Support starts blind to what fulfillment already flagged. The friction was always there, but humans absorbed it, filling gaps with tribal knowledge, workarounds, and Slack messages.

**Agentic systems don't absorb cross-functional friction. They amplify it.**

This is where the problem compounds. When individual teams deploy agents independently, each builds its own context, its own logic, its own version of how the business works. Marketing's agent and sales's agent may operate from conflicting assumptions without anyone knowing. The result isn't just inefficiency, it's actively contradictory automation

at scale. A broken process plus AI equals a faster broken process. But disconnected processes plus AI equals an organization that's confidently wrong in multiple directions at once.

The organizations seeing meaningful results aren't just mapping workflows within a function. They're looking at the spaces between functions: the handoffs, the duplicated effort, the information gaps that are invisible to anyone working inside a single department. They're asking which end-to-end flows should be redesigned, whether certain teams still need to operate the way they do, and how to sequence the journey from siloed operations toward something more connected.

That cross-functional view barely exists in most organizations. And it's the prerequisite for everything that follows.

## How cross-functional friction gets absorbed today



### Tribal knowledge

Institutional memory held by people, not systems



### Workarounds

Unofficial processes that fill the gaps automation can't



### Asynchronous Communication Tools

Coordination that should be designed into the flow

## Where to start

---

Map a flow that crosses at least three teams or departments, from customer inquiry to resolution, from campaign brief to live promotion, from order to delivery. Where does context get lost? Where does the same work happen twice? Where does one team's decision depend on information another team has but doesn't share?

When you run this exercise, expect the findings to be uncomfortable but not surprising. The friction you'll surface is almost certainly something your teams already know about— they've been absorbing it with workarounds, tribal knowledge, and Slack threads for years. Naming it is the first step. One method that works well for this kind of mapping is EventStorming, which is specifically designed to surface how work actually moves across teams versus how the org chart says it should.

What you'll also discover is that the people who've been absorbing that friction are the same ones you'll need most in the next confrontation, because they carry the institutional judgment that no agent can replicate.

[Learn the EventStorming method →](#)



## CONFRONTATION 3:

# The Human Shift

---

The question worth asking isn't whether AI will take jobs or create them, it's what happens to human expertise when execution becomes automated.

The answer: expertise gets reconfigured around judgment rather than execution. The people who will matter most are the ones who know what good looks like, ask the right questions, and recognize when AI is wrong. That kind of judgment compounds over time in ways the tools themselves can't replicate. It's not a redistribution of expertise. It's an expansion of what expertise means.

## Two ways to use AI – only one builds judgment



### Directing AI

- ✓ You set the objective and the constraints
- ✓ You evaluate the output critically
- ✓ Judgment compounds over time



### Deferring to AI

- ✗ AI sets the direction; you ship the output
- ✗ You lose the opportunity to develop judgment
- ✗ Speed increases; capability atrophies

The risk is mistaking speed for progress. Organizations that treat AI as a cost-cutting lever get compliance rather than transformation. Employees go through the motions faster, metrics tick up, and the underlying work doesn't improve. As Ethan Mollick warned in Co-Intelligence:

**“If people rely on AI too passively, they lose the opportunity to develop the judgment that truly matters.”**

Zapier found this out firsthand. When CEO Wade Foster set out to build an AI fluency rubric, the first obstacle wasn't employee resistance, it was that leaders couldn't define what “good” looked like because they hadn't used the tools themselves. Fluency at the team level was impossible until leadership developed it first. As Foster put it:

“AI adoption alone isn't transformation. Real transformation happens when AI reshapes how teams think, collaborate, and create value.”

That starts with answering the question employees are already asking (even if they aren't saying it out loud): What's in it for me?

It deserves a real answer. You're not just asking people to learn a new tool. You're asking them to take on higher-order work that pushes them outside their comfort zone – work that may significantly change what their role means – while the tools themselves shift every few months. Leaders who haven't done that work themselves can't credibly ask their teams to do it. And a one-hour town hall isn't going to cut it.

In practice: work with teams to define what judgment looks like in each role. Create safety for people to say “this isn't working yet.” And measure the quality of decisions, not just the volume of output. Because if speed is the only metric, you'll reward compliance and call it transformation.

## Where to start

---

Ask yourself this question: are you, your peers, and your direct reports learning to direct AI, or just deferring to it?

Zapier's AI Fluency Rubric is a useful lens — three levels from capable through transformative, assessed across mindset, strategy, building, and accountability, with expectations defined role by role. Start there with your leadership team, then use it to open the same conversation across the organization.

That distinction — genuine capability versus faster output — is what determines whether the next confrontation accelerates your business or exposes what's already broken.

[Explore Zapier's Rubric →](#)



## CONFRONTATION 4:

# The Delivery Reckoning

---

When code, tests, and workflows can be generated in hours instead of weeks, execution stops being the constraint. What remains—alignment, governance, validation, and the ability to connect delivery to business outcomes—are the parts most companies have always quietly struggled with, but never had to treat with urgency. Now they do.

In traditional delivery, teams could absorb ambiguity. Unclear priorities, competing stakeholders, loosely defined success criteria— these things slowed delivery down, but the slowness itself created a buffer. By the time something was built, alignment usually caught up.

Agentic systems eliminate that buffer. When a team can build a working onboarding flow in days rather than weeks, misalignment doesn't disappear, it compounds. Rework arrives faster, misdirected effort scales faster, and organizational debt accumulates faster.

**AI is the accelerator, but alignment is the steering wheel. At high speeds, being off by a few degrees takes you somewhere completely different than you intended without time to course-correct.**

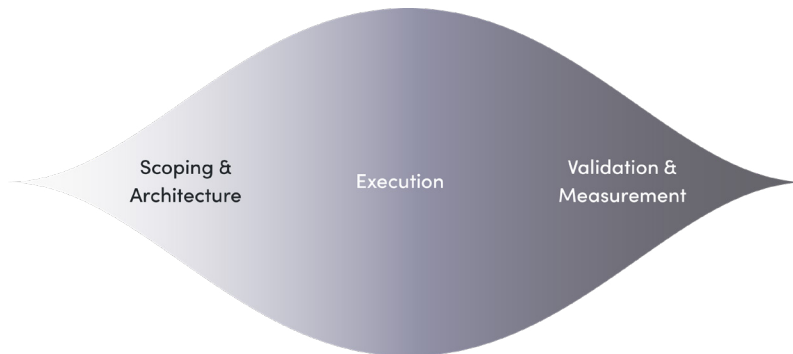
The delivery lifecycle is changing shape. The middle—which is to say production and execution, the part teams historically spent most of their time on—is compressing. What’s expanding is the front and back. More rigorous work up front: decision ownership, architecture, outcome definition. More rigorous work at the back: validation, measurement, evidence of impact.

Governance needs redesigning too. Approval chains, steering committees, and scheduled release


windows assumed that change is expensive and should be tightly controlled. That assumption doesn’t hold anymore. The fix isn’t eliminating governance, it’s separating high-risk changes from low-risk ones, pushing decision rights closer to delivery teams with guardrails rather than layers of approval, and treating governance as something that enables speed rather than manages it from a distance.

The hardest part is measurement. As delivery accelerates, the ability to connect that delivery to outcomes often doesn’t improve. More features shipped, more experiments running, but when someone asks whether any of it moved the needle, the honest answer is often: we’re not sure. Treating measurement as a precondition, not a follow-up, is how that changes.

OLD



NEW



## Where to start

---

Whether your team is shipping daily, weekly or monthly, one question applies: when your team picks what to build next, can they state the expected outcome, the metric, and the baseline? If the answer is usually no, then speeding up delivery won't close the gap between output and impact, it'll widen it.

As a follow-up question, consider what would need to become true to make this change, regardless of how quickly your team delivers.

That follow-up question—what would need to become true—is worth sitting with. Because the answer almost always points beyond the delivery team itself. It points to whether the organization has the clarity, the people, and the operational honesty to connect what it ships to what the customer actually experiences. And that's where the final confrontation waits.



## CONFRONTATION 5:

# The Customer Verdict

---

This is where the previous four confrontations converge. When customers can ask for outcomes in natural language, the distance between request and fulfillment becomes part of the brand itself. Customers don't experience your technology. They experience what your organization can reliably do.

Most agents are built to feel human: warm, responsive, confident. That makes sense as a starting point, but it creates a problem that's difficult to design around. No human is as fast, as consistent, or as available as a well-built agent.

**When the experience glitches or responds in a way that feels slightly off, customers feel it acutely. It's the same discomfort that makes near-realistic CGI faces feel uncanny.**

The instinct is usually to solve this by making agents more human. But that approach may be chasing the wrong target. The more honest solution is to stop pretending— not in a cold way, but in a way that gives customers accurate expectations about what they’re interacting with and what it can reliably do. Organizations that set those expectations clearly and then meet them consistently will build more trust than those that sell a human-feeling experience and periodically fail to deliver it.

Agentic systems also make it dangerously easy to overpromise. You can now create demos, proto-

types, and experiences that imply capabilities your organization hasn’t actually built yet, selling outcomes that depend on operational readiness that doesn’t exist. The gap between what the product promises and what it delivers has never been easier to create, and customers have never been faster to notice.

The brands that earn durable trust won’t be the ones with the most sophisticated demos. They’ll be the ones that defined accountability, built guardrails, and set realistic expectations, so that what the product promises is what actually gets delivered.

## Why agentic AI makes overpromising dangerously easy



### Demos are easy to build

Prototypes imply capabilities the org hasn’t built yet



### The invisible gap

Operational readiness lags behind what’s being sold



### Customers notice fast

Every AI they’ve used sets the bar

## Where to start

---

Look at what your customer-facing agents actually do today, or what you would like them to be able to do. Where is the experience already reliable enough to earn trust? Where is there a gap between what the agentic experience implies and what the organization's systems can consistently deliver? Close the gap or reset the expectation, either works. Leaving it unaddressed does not.

Getting this right doesn't hinge on having the most sophisticated agents. It relies on the work described in every chapter before this one: building the execution speed, cross-functional clarity, workforce judgment, and delivery discipline that lets successful organizations promise only what they can reliably deliver.



# Conclusion

---

In [Built to Last](#), Jim Collins drew a distinction between time telling and clock building. Time telling is having the right answer at the right moment— a charismatic leader, a breakthrough product, a well-timed market play. Clock building is constructing an organization that generates good answers long after any single leader, product, or strategy has run its course.

Chasing today's AI model is time telling; building the organization that absorbs whatever model comes next is clock building.

The five confrontations in this ebook are clock-building work: execution capacity, operational clarity, workforce trust, delivery discipline, customer honesty. None of these expire when the technology shifts. They're what makes every future tool more effective.

The organizations that did this work while competitors were still debating which model to use won't need to start over when the next wave arrives. They'll have built the thing that lasts.

Ready to assess where your organization stands? Orium's Agentic Strategy Canvas and Studio sessions help leadership teams turn diagnosis into action.

[Explore the Agentic Strategy Canvas →](#)

[Talk to our team →](#)



## **Orium**

**Orium** builds agentic systems that help enterprise brands move faster and stay ahead. By bringing together data, content, and commerce, we create experiences designed for collaboration between people and intelligent agents— helping our clients work more efficiently, move faster, and gain a lasting competitive edge.

## **—zapier**

Zapier is the most connected AI orchestration platform, connecting more than 9,000 apps to help businesses of all sizes move faster and work smarter. From startups to enterprises, millions of users rely on Zapier to automate repetitive tasks, streamline operations, and unlock the full potential of AI. Zapier is a fully remote company committed to helping people focus on the work that matters most. Learn more at [zapier.com](https://zapier.com).



**Orium**