



Adaptive Thinking / Paper No.003

When your customer stops using your app

Preparing for the Agent-Ready Enterprise

For years, digital transformation meant bringing the customer into your interface. The next transformation may require bringing your business into theirs.

ADAPTIVE LABS September 2026

01 THE APP ERA

We built an app for everything

Digital transformation has long followed a relatively simple direction: bring the customer into your interface. Build a website. Build a portal. Build an app. Make it easier to search, buy, book, pay, manage and communicate.

Banks built banking apps. Retailers built shopping apps. Airlines built travel apps. Mobility companies built driver apps. Restaurants built booking apps. Each decision made sense.

Together, we created a world in which almost every organisation has built its own digital front door. The problem is that someone still has to walk through all of them: the customer.

Your calendar tells you where to be. Navigation tells you how to get there. Another app handles parking, another dinner, another banking, another shopping, another the car. Each application may be excellent. But the human still orchestrates the whole.

The problem is not that we have too many apps. The problem is that the human still has to orchestrate them.

02 THE SHIFT

From destination to intent

Customers rarely wake up wanting to use an application. They want an outcome: get to the airport, find somewhere for dinner, service the car, buy a suit for a conference, reorganise tomorrow.

The application has simply been the mechanism through which intention becomes action. That mechanism is beginning to change.

TODAY



EMERGING



An agent can potentially operate across boundaries that applications were designed around. It can use context, interact with multiple services, ask for clarification, perform authorised actions and continue adapting when circumstances change.

The agent does not need to replace every application. It may replace the need for the human to orchestrate them.

03 A CHANGE IN CONSEQUENCE MANAGEMENT

From assistant to agent

Imagine you are returning to your hotel after a meeting. You have dinner booked with someone at the hotel restaurant. Traffic gets worse. Your ETA moves by 35 minutes.

An assistant could tell you: "Traffic is heavy. You may be late." Useful. But a context-aware agent can understand the consequences of the delay, your agenda, your preferences and the authority you have already delegated.

● YOUR AGENT

"Traffic has delayed your arrival by 35 minutes. I informed the hotel and Anna that you'll be late. An outdoor table became available at 20:00, so I moved your reservation there. The change added €12, which I paid using your authorised card. Nothing else in your evening agenda is affected."

You did not call the hotel, reopen the restaurant booking, check the calendar, send the message or separately approve a small payment that remained inside a boundary you had already defined. One change in reality created several consequences. The agent managed them.

An agent does not simply execute your instructions. It manages consequences within the authority you have given it.

04 EVERYDAY LIFE

A day with a connected agent

The hotel scenario is not exceptional. Much of everyday life consists of small coordination tasks, and today almost each of them requires a separate app, a separate search and a separate decision.

Now imagine an ordinary day in which a personal agent knows your calendar, your preferences and the limits you have set:

	WHAT THE AGENT DOES	WHAT THE BUSINESS MUST OFFER
07:40	Before leaving, you take a photo of the fridge and the pantry. The agent builds a shopping list, compares online stores and books a delivery for 19:00, when your calendar shows you at home.	<i>Live availability, prices and delivery slots an agent can read and book.</i>
08:20	On the way to work, the agent reminds you that your wife's birthday is tomorrow. It proposes flowers delivered to her office before 16:00 and a table at 20:00 at the restaurant you liked last time – and waits for your yes.	<i>A florist and a restaurant that accept agent requests, confirm availability and honour the details.</i>
13:30	A routine check-up is due. The agent finds a free slot at your clinic that fits between meetings, books it and adds travel time to your calendar.	<i>A booking system with open, bookable slots – not just a phone number.</i>
17:50	Two of tomorrow's meetings overlap. The agent proposes a new time to the other participant's agent and confirms once both calendars agree.	<i>Calendars and systems that can negotiate with each other, within agreed permissions.</i>

None of these moments is spectacular. That is the point. Each one removes a small piece of orchestration from the human – and each one ends at a business that the agent must be able to find, understand and act with.

Behind every small convenience sits a business that was ready for the agent.

05 HUMAN INTENT AS THE ANCHOR

Context changes everything

The real shift is not simply from applications to agents. It is from isolated interactions to context-aware orchestration. Calendar, agenda, location, time, preferences, history, relationships, commitments, current activity and permissions can become parts of a changing context.



Within those boundaries, the agent can observe, anticipate, interact, act, confirm and adapt. Outside them, it returns to the human for judgement.



| *Human sets the intent. AI expands the capacity to act on it.*

This is AI as cognitive support extended into coordination and execution.

06 TODAY → EMERGING

The infrastructure is beginning to emerge

This world does not yet exist as one coherent system. But some of its foundations are already being built.

Model Context Protocol¹ is developing common ways for AI applications to connect to tools and data. Agent2Agent² provides an open standard for communication and collaboration between agents. Agentic Commerce Protocol³ and Universal Commerce Protocol⁴ are creating structured ways for agents and merchants to discover products and complete commerce flows. Agent Payments Protocol⁵ addresses secure agent-led payments and delegated authority.

The final architecture remains uncertain. Different standards may coexist. Some ecosystems may remain closed while others become increasingly interoperable. But a direction is visible: AI systems are moving from accessing information toward interacting with capabilities.

The strategic question is no longer only: What can agents do? It is: What must an enterprise become when customers increasingly interact with it through agents?

07 ENTERPRISE RESPONSE

The Agent-Ready Enterprise

An Agent-Ready Enterprise is not simply a company with an AI chatbot, an API or well-organised data. It is an enterprise whose relevant products, services and capabilities can increasingly be discovered, understood and used by authorised agents acting on behalf of humans.

01	DISCOVERABLE	The agent can identify that a relevant product, service or capability exists.
02	UNDERSTANDABLE	It can interpret conditions, availability, constraints, alternatives and consequences.
03	INTERACTIVE	It can communicate changing context, ask questions and receive meaningful responses.
04	ACTIONABLE	It can perform permitted actions rather than merely retrieve information.
05	TRANSACTABLE	It can complete authorised economic transactions with appropriate identity, permissions, payment and confirmation.
06	ADAPTABLE	It can continue interacting when reality changes after the original transaction.

An agent-ready business does not only publish information for machines to consume. It can participate in a conversation between systems around human intent.

08 UX + AGENT EXPERIENCE

When the interface stops being a screen

For most of the digital era, interface design meant creating something a human could see and manipulate: a screen, a menu, a button, a journey. Those interfaces will not simply disappear. Humans will continue to value visual experience, brand, exploration, emotion and judgement.

But another interface is emerging alongside them: one designed for the human's agent. Can an agent discover what the organisation offers? Understand it? Interact? Execute? Understand what happened? Recover when something changes?

The interface of the future may not always be a screen. It may increasingly become a conversation between agents and systems.

What does customer experience mean when part of the customer experience is experienced by an agent?

09 THE CAR AS CONTEXT-RICH NODE

Beyond the in-car assistant

The car is an interesting example – not because automotive is the centre of this transformation, but because a moving vehicle creates unusually rich context: location, destination, ETA, traffic, vehicle state and journey context.

Combine that with a personal agent that understands calendar, agenda, preferences, commitments and permissions, and connect both to an ecosystem of external services. The result can become more interesting than another in-car assistant.

The car can become a context-rich node in the user's wider agent ecosystem.

There is also a simpler reason the car matters. Many of us spend a significant part of every day driving and commuting. At work or at home, attention belongs to whatever we are doing. The drive is different: eyes and hands belong to the road, but the mind is often free.

With a voice-first agent, that time can turn the car into a mobile office – and, more broadly, a mobile life interface. You can go through tomorrow's agenda, move a meeting, dictate a reply, prepare for the next conversation or start a task that will be finished somewhere else – all by voice, without taking attention away from driving.

The car does not need a bigger screen. It needs an agent that turns travel time into thinking time.

09 CONTINUED

One intention, different contexts

On the way to work, the agent might say: "Your wife's birthday is tomorrow. Based on your agenda, I can send flowers to her office before 16:00 and book 20:00 at the restaurant you liked last time. Shall I arrange both?"

Nothing in that sentence is spectacular on its own. What matters is what it connects: a date that might otherwise be remembered too late, a relationship, an agenda that shows when the evening is free, a preference learned from a previous visit, and two services – a florist and a restaurant – that the agent can actually reach and act with. The agent brings them together. The human decides.

There is no reason to open three apps on a car screen. The conversation starts in the car, the confirmation can arrive on the phone, and the evening itself happens somewhere else. The device changes. The intention does not.



The context changes. The intent continues.

Automotive conversational AI is already moving toward richer natural-language interaction. BMW, for example, describes its 2026 integration of Alexa+ as extending in-car interaction beyond fixed commands toward natural conversation. Our point is broader: the vehicle can be one contextual node in an agent ecosystem rather than another isolated digital destination.⁶

10 AI AS COGNITIVE SUPPORT

Delegated agency

If an agent can communicate, transact and adapt, how much should it be allowed to do? The useful answer is not maximum autonomy. It is delegated agency: the human defines the intent and the limits; the agent coordinates and executes within them.

An agent might be allowed to notify someone about a delay, move a reservation within 30 minutes, spend up to €20 to preserve a preference, reorder frequently purchased items or schedule routine vehicle service around the calendar. A major purchase, contractual commitment or high-consequence decision may require explicit human judgement.

The objective is not to remove the human from the loop. It is to stop requiring the human to operate every loop.

This makes identity, permissions, decision rights, transparency, auditability, reversibility and escalation part of the architecture. Emerging payment and authorisation standards already point toward scoped permissions and delegated authority rather than unlimited autonomy.^{4 5}
7

Autonomy does not mean absence of control. It means freedom to act within understood boundaries.

11 WHEN THE BOUNDARY MOVES

The rational decision trap

Now imagine a management meeting. A proposal appears on the screen:

NEW CUSTOMER APP
INVESTMENT: €5 MILLION

The business case is compelling: better UX, higher engagement, more customer data, greater loyalty, cross-selling and lower service costs. Every assumption may be rational. The application may even succeed.

What if we are optimising the interface just as the interface itself begins to move?

The problem is not that investing in an application is necessarily wrong. The problem is assuming that today's system boundary will remain tomorrow's system boundary. A rational organisation can make a rational investment based on a rational business case – and still misunderstand the system in which that investment will eventually operate.

Rational decision. *Wrong system boundary.*

Instead of asking only "How do we build a better app?", another question belongs beside it:

How will customers interact with us if our app is no longer necessarily their primary interface?

12 PREPARE FOR SEVERAL FUTURES

Don't predict the winning architecture

One universal personal agent? Several specialised agents? Platform agents? Automotive agents? Enterprise agents? Agents negotiating with other agents? Open ecosystems? Closed ecosystems? Probably some combination we cannot yet see clearly.

We do not know. And we do not need to know. Replacing today's rigid assumption with another rigid assumption about tomorrow would miss the point.

The better question is which capabilities remain valuable across several plausible futures: accessible and structured information, interoperability, clear capabilities, identity and permissions, executable services, transactions, context, trust and adaptability.

Don't design the organisation around one predicted future. Design it to participate in several possible ones.

13 THE ORGANISATIONAL IMPLICATION

Preparing rather than predicting

Until now, transformation often meant putting more of the enterprise behind a better digital interface. The next transformation may require something different: not simply another layer on the outside, but a change in how the organisation exposes, connects and governs the capabilities inside.

That affects technology architecture, but also product design, operating models, data ownership, decision rights, governance, partnerships and potentially where competitive advantage sits.

Companies have competed for decades to own the customer interface. In an agentic world, part of that interface may increasingly sit outside the company. That does not make the enterprise less important. It makes its ability to participate in a wider system more important.

The organisations best prepared may not be those that correctly predict which agent, platform or protocol will win. They may be those that recognise early enough that the boundary is moving.

14 THE OTHER SIDE

What the enterprise could lose

None of this comes without cost. When an agent stands between the customer and the enterprise, several things organisations have long relied on may weaken.

RELATIONSHIP	Customers may increasingly talk to their agent rather than to you. Direct contact, first-party data and everyday presence can thin out.
BRAND	When an agent compares options on price, availability and conditions, differentiation that lived in the interface may matter less.
MARGIN	Platforms that operate agents may position themselves as new gatekeepers between customers and businesses – and charge accordingly.
ACCOUNTABILITY	When an agent acts wrongly on a customer's behalf, responsibility between the customer, the agent provider and the enterprise is not yet clearly defined.

These risks do not argue against preparing. They argue for preparing deliberately: an enterprise that shapes how agents interact with it is in a stronger position than one that is simply reached by them.

Being agent-ready is not only about being usable by agents. It is about deciding on what terms.

15 WHERE TO START

Questions worth asking now

None of this requires a perfect answer yet. But a few honest questions are worth asking in almost any organisation today:

- 01 If a customer's agent tried to find, compare and buy our core offer today, where would it fail?
- 02 Which of our capabilities could we safely expose to authorised agents – and which should always stay with a human?
- 03 What would we lose if part of the customer relationship moved outside our interface, and how would we protect it?
- 04 Which of our current digital investments assume that our app remains the primary interface?
- 05 Who in the organisation owns the question of agent readiness?

In the end, this is not really a paper about AI. It is about seeing a system whose boundaries may be changing.

A framework built around today's reality can optimise today's reality extremely well – until reality changes. Adaptive Thinking starts somewhere else: observe reality, change perspective, challenge the boundary, understand the relationships, then design the response.

Reality should inform the framework. The framework should not constrain reality.

The emergence of agents may not make today's applications obsolete. That is the wrong prediction to optimise around. The more useful question is:

What would our organisation need to become if the way customers interact with businesses fundamentally changes?

We do not need a perfect answer today. We need organisations capable of finding better answers as reality evolves.

THE FUTURE DOES NOT NEED TO BE PREDICTED PERFECTLY.
IT NEEDS TO BE PREPARED FOR INTELLIGENTLY.

Reflection expands perspective.
Execution creates change.
Adaptive Thinking connects the two.

ADAPTIVE LABS adaptivelabs.eu



NOTES & SOURCES

Technology signals referenced in the paper

1. Model Context Protocol – open protocol for connecting AI applications to external tools and data sources.
<https://modelcontextprotocol.io>
 2. Agent2Agent (A2A) Protocol / Linux Foundation – open agent-to-agent communication and collaboration standard.
<https://a2a-protocol.org/v1.0.0/>
 3. OpenAI / Stripe, Agentic Commerce Protocol – structured merchant/product integration and commerce flows.
<https://developers.openai.com/commerce>
 4. Shopify / Google, Universal Commerce Protocol – agentic commerce from discovery through checkout and post-purchase operations, including human hand-off.
<https://www.shopify.com/ucp>
 5. Google, Agent Payments Protocol (AP2) – payment-agnostic framework for secure agent-led transactions and delegated authority.
<https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol>
 6. BMW Group, 5 January 2026 – Alexa+ integration in BMW Intelligent Personal Assistant and natural-language in-car interaction.
<https://www.bmwgroup.com/en/news/general/2026/alex.html>
 7. Model Context Protocol, Enterprise-Managed Authorization, 18 June 2026 – centrally managed authorisation for MCP servers.
<https://blog.modelcontextprotocol.io/posts/enterprise-managed-auth/>
-

EDITORIAL NOTE

The paper distinguishes current infrastructure from the proposed organisational synthesis. "Agent-Ready Enterprise" is used descriptively rather than as a claim of originating the broader "agent-ready" terminology. Scenarios such as delegated low-value reservation changes illustrate a plausible next architecture, not a claim that such end-to-end capability is universally available today.