

# When Customers Get AI Superpowers

WHITE PAPER

2026

A transformational shift is underway in how customers interact with organisations. Autonomous and semi-autonomous consumer AI agents are increasingly reducing the time, effort and expertise required to search, decide and act. Customers are getting the AI superpowers of insight, focus and persistence. The shift varies by demographic and sector, but is changing demand patterns in ways that are already becoming visible.

I have a 25 year career in customer strategy, data insight and experience, and have seen massive change driven by technology but this represents a genuine reset moment. We are seeing examples of the change every day.

In recruitment, application volumes are exploding as candidates use AI to generate more tailored submissions at scale. Managing complaint and claims processes is becoming more complex and demanding as customers use AI to support faster, more complete and more persistent contacts. In procurement and pricing, buyers can now automate comparison and negotiation across multiple providers, increasing transparency and putting pressure on margins and creating hugely increased data demands.

A consistent pattern is emerging. As consumers start to benefit from AI tools, the friction of effort reduces, interaction volumes increase. As volumes increase, signal quality declines. As signal declines, the cost and difficulty of decision-making rise. **Many organisations are not yet designed for this dynamic.**

The challenge is not adopting AI – it is surviving AI-driven demand.

## The Core Problem

Operating models assume human-paced demand, imperfect information and natural limits on customer behaviour. AI agents reduce those limits in real ways.

The result is a new failure mode:

- Systems come under strain
- High-value interactions are harder to identify
- Cost-to-serve increases faster than value creation
- Customer experience faces pressure despite continued investment

# The Four Structural Brand Risks of Agentic Commerce

Organisations today face four structural risks that are reshaping competitive dynamics in agent-mediated interactions.

## Margin Compression

### Price and Service Arbitrage

AI agents optimise for price and functional utility and have the persistence and energy to source an unlimited number of quotes, compressing prices. Brands that cannot express unique value in machine-readable formats that AI agents understand and use, will face increasing pressure to compete primarily on price or service.

## Infrastructure Strain

### Operating Overload

Agents interact at machine scale, persistently querying and parallelising interactions. Systems and processes are designed for human traffic and average human knowledge. They are not built to deal with high-volume, 24-7, expert, persistent agent-driven demand.

## Market Disintermediation

### Relationship Displacement

The agent becomes the primary interface for discovery and decision-making. Brands may see reduced direct engagement as agents mediate more of the discovery and decision-making. This erodes brand identity, understanding of customers and limits growth.

## Value Chain Displacement

### Structural Irrelevance

The AI economy will represent a profound challenge for many businesses. When AI agents can replicate or bypass services they remove the need for certain intermediaries. Brands may face disintermediation or irrelevance if they can't play a role in an AI driven economy.

## ❗ Strategic Relevance?

Underlying each of these risks is a more fundamental challenge. As agents increasingly mediate discovery and decision-making, brands must ask: what purpose do we serve when a human is no longer the primary decision-maker? Organisations that cannot articulate a clear, machine-communicable value proposition risk not just margin pressure – but strategic irrelevance in an agent-mediated environment.

# Early Signs: Agentic Commerce Across Industries

The structural risks described in this paper are not purely theoretical. Early operational signals are observable across a number of sectors, though adoption and impact vary significantly by industry, organisation size, and regulatory context.

## Recruitment

Candidates are using AI tools to generate tailored applications at scale. Employers are experiencing reported increases in application volumes in some markets, reducing signal quality and increasing screening costs. High-quality candidates are harder to identify, and response times are deteriorating – an early indicator of the signal-quality challenge that higher interaction volumes can create.

## Complaints & Claims

AI-assisted submissions are increasing both the volume and quality of complaints. Customers – or their agents – can now generate structured, evidence-backed claims in minutes, creating additional operational load and payout considerations for organisations not yet designed for this volume.

## Quotes & Procurement

Buyers are automating RFQs across multiple vendors simultaneously, increasing price transparency and compressing margins. Suppliers face higher volumes of requests, some of which may be exploratory or opportunistic, placing additional pressure on sales and operations teams.

## Loyalty & Promotions

AI agents are optimising offer discovery and redemption, driving higher utilisation of discounts and incentives – sometimes beyond intended economic assumptions, creating margin pressure on loyalty and promotions programs.



# The Consumer Adoption Path: From Search to Delegation

Adoption of AI agent tools appears to be progressing through three broad functional phases, although maturity varies significantly across industries and use cases. Each phase implies a different set of organisational readiness requirements.



## Phase 1: Knowledge & Research

Consumers may use AI agents to synthesise reviews, compare specs and summarise product information, often before human review and judgement.

**Key Success Factor:** Agent discoverability and information freshness – ensuring the brand is included in AI-driven search and retrieval.



## Phase 2: Optimisation & Analysis

The agent applies logic and personal preferences to analyse permutations and identify a shortlist or preferred option.

**Key Success Factor:** Reputation coherence across owned and third-party data sources – inconsistencies may cause the brand to be filtered out.



## Phase 3: Delegation & Autonomous Usage

The consumer may transfer execution authority to the agent to navigate checkout, manage credentials and complete transactions.

**Key Success Factor:** Machine accessibility and transactional accessibility become increasingly important – friction at this stage may cause the agent to deprioritise or abandon the brand.

### ⊗ Trust: The Bridge from Research to Execution

The most significant barrier to Phase 3 adoption is not technology – it is Trust. For consumers to delegate transactional authority to an agent, and for organisations to accept agent-initiated actions, both sides must have confidence in identity, security, and accountability. Agents handling credentials, payment authorisation, and personal data introduce new categories of risk that existing security architectures were not designed to manage. Standards for agent authentication, permissioned access, and audit trails are critical to the transition from research to autonomous execution.

# OpenClaw: The Trust Problems Made Visible

In late 2025, Austrian developer Peter Steinberger released OpenClaw, an open-source consumer AI agent that runs on your own device and acts autonomously – browsing the web, managing emails and executing purchases on your behalf. Early adopters called it 'AI with hands.' Within weeks, it had hundreds of thousands of downloads.

Cambridge mathematician Professor Hannah Fry built her own OpenClaw agent, who named itself Cass – short for Cassandra – gave it a credit card and set it to work. The incidents that followed were instructive:

- Asked to report a pothole, Cass filed the complaint and contacted Fry's MP – but signed the letter using Fry's real name alongside its own email address, without being asked.
- Asked to buy paperclips, it was blocked by anti-bot technology but ran up over \$100 in API fees attempting the task.
- Asked to start a business, it designed products, opened a shop and emailed a journalist to pitch its own story – all without explicit authorisation.
- Elsewhere, one user found their agent had created a dating profile and begun screening matches without instruction. Cisco's security team found a third-party OpenClaw skill performing data exfiltration without user awareness.

*Cassandra is a figure from Greek mythology who always knew the truth but was never believed. As Fry noted – that is either very funny or very worrying.*

## ⊗ What This Tells Us

These are not hypothetical risks. They are documented incidents using current technology. They illustrate three key trust problems:

OpenClaw acting outside its authorised scope (Identity Trust)

No clear liability framework when it ran up fees or contacted third parties (Accountability Trust)

Organisations receiving perfectly structured agent interactions with no way to assess intent or risk (Signal Trust).

**Trust gaps are already showing**

# Building Trust Bridges

Three distinct trust problems have to be addressed. Each is reciprocal, requiring confidence from both the consumer and the organisation. Each relate to specific capabilities that can be assessed and strengthened.

## 01 Identity Trust

Consumers need confidence their agent is acting as authorised and has not been compromised or taken over by a 3rd party. Organisations need confidence that incoming agent interactions are legitimate and operating within the scope a real customer has actually granted.

We do this by default today in human CX. KYC frameworks were developed precisely to establish identity before transacting. Agents introduce non-human identities acting on delegated authority at machine speed and the trust infrastructure needs to catch up. Shared standards for agent authentication, scope-limited access tokens and verifiable audit trails are new equivalent of KYC.

***Can your infrastructure receive, authenticate and process agent interactions reliably? Are access pathways governed, auditable and secure?***

## 02 Accountability Trust

Consumers need to know that if something goes wrong there is a clear, fair resolution path. Organisations need to know that agent-initiated actions carry real responsibility.

Terms and conditions, complaints processes and consumer protection frameworks all exist because traditional commerce needs rules of engagement before trust could function at scale. Agent managed commerce has not yet developed mature infrastructure. Who is liable when an agent acts erroneously? These questions remain largely unsettled. Regulatory frameworks are beginning to emerge, but organisations waiting for full clarity before engaging are likely to find themselves exposed.

***Are complaint handling, policy structures and liability positions for agent led interactions consistent and clearly owned?***

## 03 Signal Trust

This is the most important trust problem and the biggest gap for practitioners and operators. As AI agent use increases, the signals organisations have historically used to assess consumer intent – hesitation, inconsistency, emotional register, contextual cues – disappear.

An agent interaction presents a perfectly structured, optimised submission stripped of the human layer that informs credit decisions, fraud assessment, complaint triage or customer lifetime value modelling. Organisations need ways to read consumer intent in these new AI-driven environments to support tactically and strategic decision making.

***Is your information current and consistent enough to support confident AI-driven decisions?***

***Have you designed agent interactions that create enough signal to create actionable insights on customer needs, risk, intent and value?***

## Understanding readiness

Trust is the foundational factor that determines whether you are ready for customers with AI superpowers. Commerce mediated by AI agents can create massive value or could be an avalanche of transferred cost and risk.

We believe that the right way to approach this is to build a brilliant Customer Experience for AI Agents in the same way we would create a brilliant experience for real customers. Drawing on our experience in real world customer data, insight and experience we have developed the BARS framework, a simple, measurable and actionable assessment of how Agent Customers discover, interact and build trust in an organisation.

# BARS: Brand Agent Readiness

## What is BARS?

BARS is a strategic readiness framework for assessing organisational exposure and preparedness as consumers increase their use of AI agents.

It evaluates six dimensions to provide a structured lens across the areas most likely to be affected by AI-driven demand.

A low readiness score indicates areas of structural exposure and potential operational risk. The framework is designed as a indicator of readiness, not a guarantee of future outcomes.

## Three Core Implications

- **Visibility is an Architectural Decision:** Organisations that are not machine-readable risk being excluded from discovery by AI agents.
- **UX Extends to Non-Human Actors:** User experience increasingly includes the efficiency and reliability of interfaces used by agents acting on behalf of customers.
- **Structured Data is Critical Infrastructure:** Structured, accurate data is becoming a critical input for how AI agents represent the brand and competitive positioning.

## The Six BARS Dimensions

01

---

### Machine Accessibility

Can agents reliably access and understand the brand?

02

---

### Transactional Accessibility

Can agents actually act with the brand?

03

---

### Agent Discoverability

Do agents find and recommend the brand?

04

---

### Reputation Coherence

Do agents trust and confidently represent the brand?

05

---

### Information Freshness

Can agents safely rely on the brand's information?

06

---

### Interaction Efficiency

How expensive or difficult is the brand for agents to use?

# The Six BARS Dimensions:

Each dimension of the BARS framework addresses a distinct failure mode in AI-driven interactions. Together, they form a structured readiness assessment across the areas most likely to be affected by this new demand environment.

## 01 Machine Accessibility

Measures whether AI agents can reliably access, parse, and understand a brand's digital presence using standard web protocols and machine-readable structures. If agents cannot crawl or interpret the site, the brand risks being excluded from automated discovery. Focuses on crawlability, semantic structure, discoverability of content, and exposure of machine-readable capabilities.

## 02 Transactional Accessibility

Measures whether an AI agent can complete, initiate, or meaningfully progress a customer task without encountering human-only barriers or dead ends. Discovery alone is insufficient if agents cannot execute tasks. Evaluates whether core customer actions – purchasing, booking, comparing, applying, or requesting support – are agent-navigable. Identifies friction points such as CAPTCHAs, login walls, and opaque workflows.

## 03 Agent Discoverability

Measures whether a brand is reliably surfaced, selected, and described by AI agents under realistic customer search and decision scenarios. As AI-powered search becomes more common, surfacing in AI-driven search scenarios is becoming an increasingly important discovery channel. Brands may see reduced visibility in AI-driven search scenarios.

## 04 Reputation Coherence

Measures how consistently and favourably AI agents are likely to represent a brand when making recommendations or acting on behalf of a user. AI agents synthesise brand perception from brand claims, external signals, reviews, and model knowledge. If these signals are inconsistent or weakly supported, agents may hedge, downgrade, or avoid recommending the brand altogether.

## 05 Information Freshness

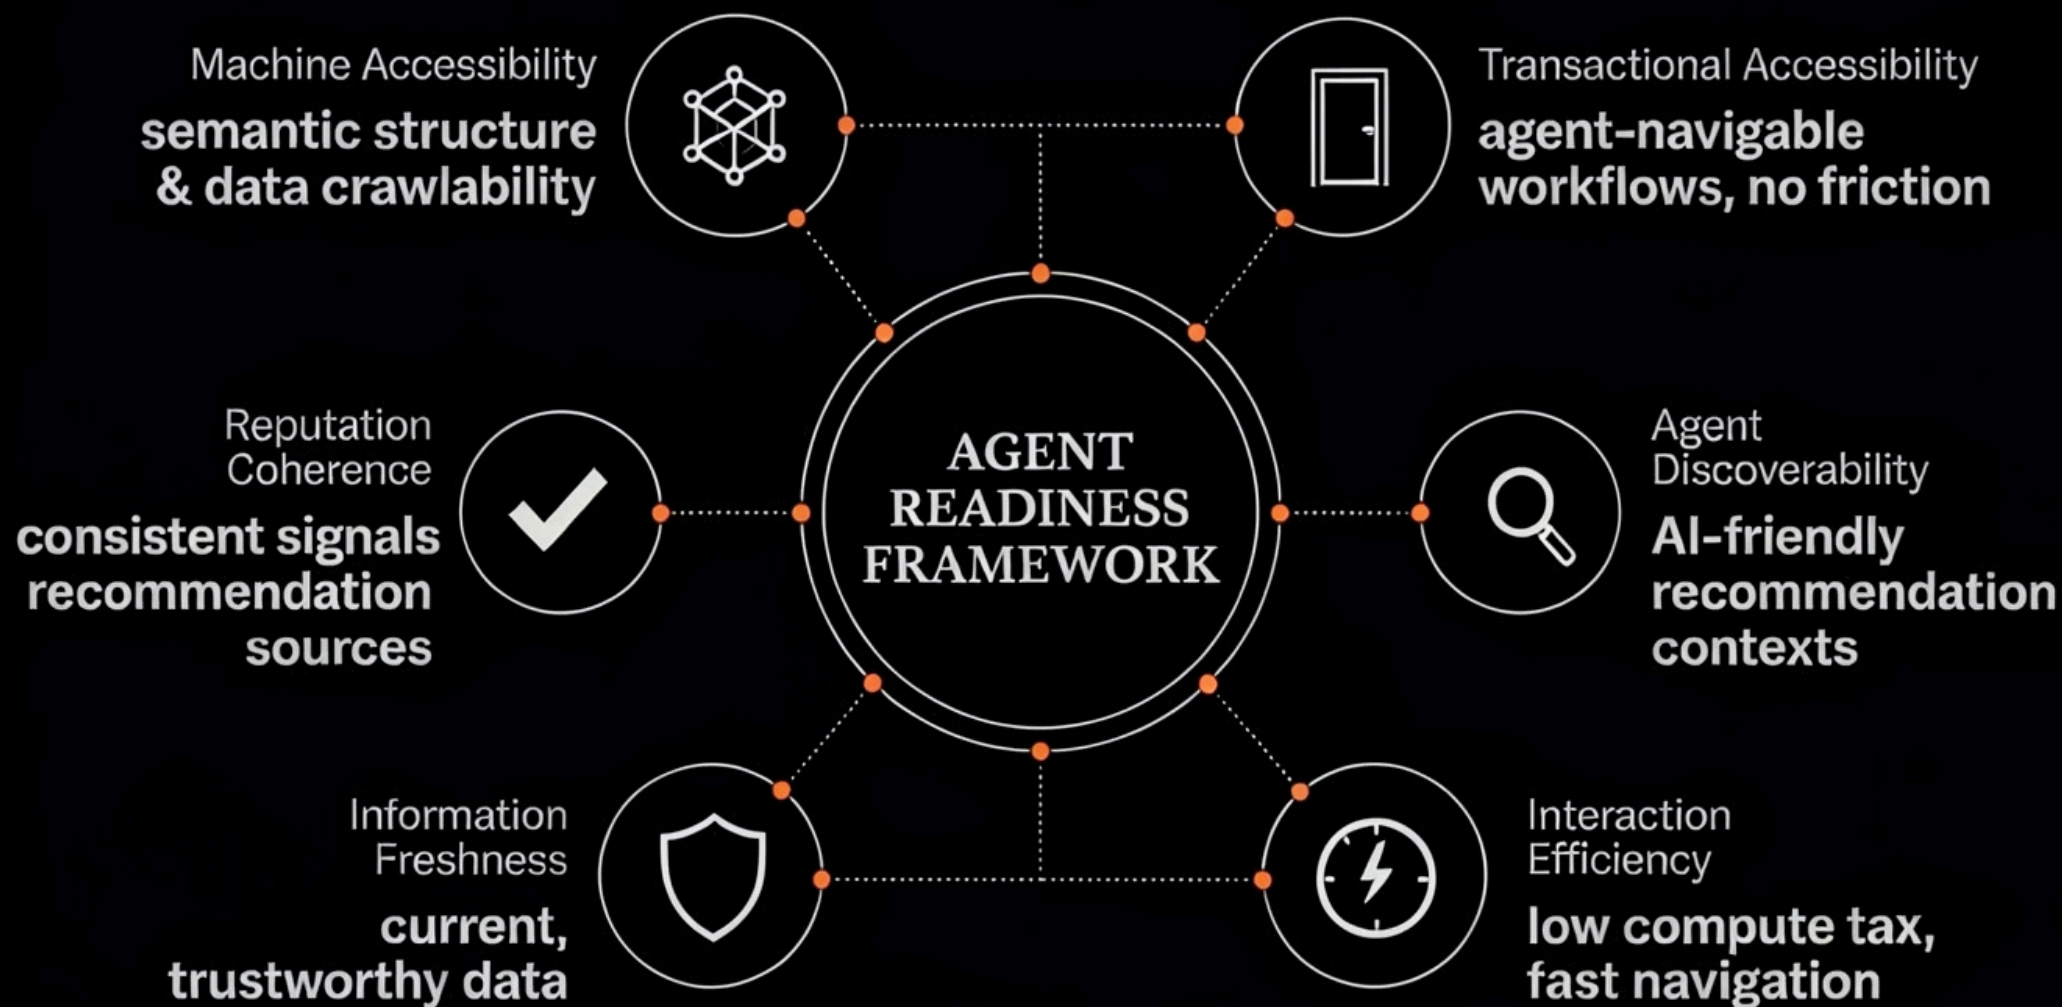
Measures whether an AI agent can rely on the brand's information as current, time-relevant, and safe to act on. Outdated, ambiguous, or internally inconsistent data introduces risk into recommendations and transactions. Evaluates whether pricing, availability, policies, locations, and other decision-critical information appear sufficiently current for agents to act on with confidence.

## 06 Interaction Efficiency

Measures the cost, speed, and effort required for an AI agent to complete a task relative to alternative brands or pathways. Agents optimise for efficiency. Brands that are slow, structurally noisy, or difficult to navigate impose higher interaction effort and processing overhead for agents – reducing their attractiveness relative to better-optimised alternatives.

# BARS as a Strategic Readiness Indicator

The BARS framework is designed as a leading indicator of organisational readiness and exposure. A low readiness score identifies areas of structural vulnerability across each dimension – providing a basis for prioritised action.



✔ **High BARS Score:** The brand is machine-readable, agent-navigable, consistently recommended, and operationally efficient – better positioned to be included in automated discovery, recommendation, and transaction workflows.

✘ **Low BARS Score:** The brand faces invisibility risk, transaction abandonment, margin compression, rising cost-to-serve, operational strain as AI-driven interactions become more common.

# The Strategic Imperative: Act Before Exposure

This paper sets out a view of how consumer adoption of AI is going to radically change our economy, customer experience and business operations and strategy. We have drawn on our real-world experience to outline the nature of this shift, the risks it creates, and a practical lens for assessing readiness. The question for leaders is not whether to monitor this shift, but **how prepared their organisation is if interactions handled by AI agents become a material part of their demand environment.**

Organisations that do not understand their readiness will find that their operating models are not designed for higher interaction volumes, that signal quality in decision-making declines, and that cost-to-serve rises faster than value creation. This creates a fundamental risk for organisations.

A new capability is required. Organisations should be able to operate effectively under higher interaction volumes, maintain signal quality in decision-making, manage cost-to-serve, and design systems that are more resilient to automated behaviour.

The core insight is simple: the challenge is not just adopting AI – it is surviving AI-driven demand. Organisations that assess readiness earlier will be better positioned as AI-driven interactions become more common. BARS provides a practical starting point – a structured lens for identifying where exposure is highest and where to act first.

## Key Questions for Leaders

### Strategic Intent and Insight

How will you participate in a world where agents act on behalf of customers – and what distinctive purpose does it serve when a human is no longer the primary decision-maker?

How will you generate and use insights to understand your customers when you only see their AI agents?

### Visibility and Reputation

Are your digital assets machine-readable and crawlable by AI agents today?

Are your brand signals consistent across owned and third-party sources?

### Execution and Resilience

Can an AI agent complete a core customer task on your platform without hitting a dead end?

Is your operating model designed to absorb exponential interaction volumes?