

Your Next Customer is a Bot

A series exploring the impact of agentic commerce and consumer AI in customer service.

The Four Conversations Nobody Is Having Together

Why commerce protocols, regulation, contact centre design and dispute handling now need one operating response

Author Martin Hill-Wilson

Publication Originally published 2026 | Updated 16 September 2026




Follow for more content at: dualcx.com

Agentic commerce is developing through several connected systems. Technical standards are defining how agents discover products, exchange information and initiate transactions. Regulators are clarifying how existing duties apply when organisations use AI. Contact centres are introducing their own AI. Payment firms and merchants are working out how to record authority and consent.

These activities often have different owners, budgets and measures of progress. The customer encounters their combined effect.

When an AI-mediated purchase goes wrong, the service operation may need to establish what the customer wanted, what authority the agent had, what information it used, what action it took and which organisation can provide a remedy. The readiness question is whether the contact centre can reconstruct that sequence reliably.

Four conversations therefore need to be considered together.



Where the market stands

Agentic commerce has moved into live testing and limited production, but progress is uneven. AI is already being used for product discovery, comparison, review analysis and basket creation. Customers remain more cautious when an agent must choose the final item, authorise payment, accept a continuing commitment or resolve a problem.

Three broad models help explain where control and evidence may sit, although individual journeys can combine elements of more than one. Retailers such as Amazon and Walmart are embedding assistants within customer relationships they already control. Google is developing an open protocol and persistent cart across its services. OpenAI has moved away from completing checkout inside ChatGPT and towards product discovery, merchant applications and referral journeys.

Market model	Where control sits	Service implication
Retailer-owned assistant	Within an established retail, account, payment and fulfilment relationship.	The organisation may hold more of the evidence, but it must still connect agent actions to service records.
Open protocol journey	Across agents, merchant systems, commerce protocols and payment providers.	Evidence and authority can be distributed across several participants.
Browser or general-purpose agent	With an agent acting across sites and services that it does not control.	The merchant may receive an order without a complete or agreed account of the customer's mandate.

These models distribute control differently, but they encounter the same operational limits. Product and inventory data must be current. The agent's authority must be clear. Payment requires evidence of consent. Customers need a way to pause, reverse or challenge an action. Someone must take responsibility when the journey crosses several organisations.

This is why discovery is advancing faster than unattended purchasing. A poor recommendation can be ignored. A purchase creates financial, contractual and service consequences that have to be explained and, where necessary, reversed.

The immediate issue is not the volume of autonomous transactions. It is whether organisations can support the first cases that cross from assistance into delegated action.

1 Commerce protocols

The first conversation concerns the standards that allow agents, merchants, platforms and payment providers to exchange data and initiate actions.

Universal Commerce Protocol, Agentic Commerce Protocol and the associated payment, identity and orchestration mechanisms address different parts of that exchange. They do not yet form one settled architecture.

Google describes UCP as an open standard covering discovery, buying and post-purchase support. Its implementation is being introduced in phases and remains limited to participating merchants and eligible markets. OpenAI expanded ACP towards product discovery in March 2026, allowing merchants to provide product and promotion data while continuing to use their existing commerce and fulfilment systems.

Different journeys may therefore use different combinations of discovery, basket, checkout, payment and merchant systems. Some will remain within an AI service for longer. Others will begin with an external agent and return to the merchant for checkout or support.

Each transition can separate the transaction from its context. A merchant may receive a valid order without receiving a usable record of the customer's purpose, constraints or priorities. Protocols can standardise parts of the exchange, but they do not determine what the service operation needs when the outcome is contested.

2 Regulatory duties

The second conversation concerns the duties that apply when AI contributes to a customer interaction or outcome.

The Competition and Markets Authority confirmed in March 2026 that existing UK consumer law continues to apply when businesses use AI agents. Organisations remain responsible for agents they deploy and should test their operation, monitor outcomes, maintain appropriate human oversight and correct problems promptly.

Data protection requirements depend on the role played by AI and the effect of the decision. The presence of AI does not make every decision subject to the same rule. Where a decision is solely automated and has legal or similarly significant effects, the organisation must consider the applicable safeguards, including information, human intervention and challenge.

These duties have an operational consequence. An organisation cannot demonstrate effective oversight merely by showing that a policy exists. It needs evidence that staff can recognise a contested outcome, obtain the relevant records, intervene with sufficient authority and explain what happens next.

The contact centre is one place where those arrangements become visible. If an adviser cannot access the evidence or secure a review, the weakness affects the customer's ability to obtain a remedy.

3 Contact centre design

The third conversation concerns the use of AI within customer operations.

Much of this work focuses on AI deployed by the organisation: automated service, adviser assistance, conversation summaries, quality monitoring and workflow orchestration. These uses deserve attention, but they cover only part of the change.

Contact centres will also receive work created by AI services outside the organisation. A customer's agent may have compared products, interpreted terms, assembled a basket, selected a payment route or attempted to change a booking before the organisation hears from the customer.

THE ADVISER MAY SEE THE RESULT WITHOUT SEEING



The customer's original goal.



The constraints and priorities given to the agent.



The options presented and information supplied by the merchant.



The agent's authority limits.

✓ Any confirmation or approval.

⋮ The ordered sequence of actions.

✂ The point at which the outcome departed from the customer's purpose.

The answer is not access to a model's private reasoning. The service operation needs observable records: mandate, constraints, information supplied, options presented, approvals, actions and transaction events.

The Service Handshake provides a complementary design pattern for this problem. It describes how goals, constraints, data permissions, authority and fallback arrangements could pass between customers, AI services and human teams when work changes hands.

Without this context, the adviser must reconstruct the case from the customer's account and a transaction record that shows what happened but not why.

4 Dispute handling

The fourth conversation concerns what happens when an agent-mediated outcome is challenged.

Existing complaint, refund, chargeback and redress routes still apply. The gap lies in the evidence and coordination needed to use them when several agents, platforms or organisations have influenced the result.

A conventional payment record may show that a transaction was authenticated. It may not establish whether the product, amount or continuing commitment matched the customer's current mandate.

Emerging agentic payment models include records of agent identity, mandate and transaction intent. These technical records do not by themselves decide legal responsibility or establish whether the customer understood what the agent would do.

AGENTIC DISPUTES MAY THEREFORE REQUIRE EVIDENCE ABOUT

🕒 What the customer authorised and how long that authority remained valid.

🏠 What the agent was permitted to infer or change.

⇒ Which options it presented or excluded.

📄 What information each organisation supplied.

☑ Whether the customer approved the final action.

🔗 How an error moved between agents and systems.

ACCRA uses four provisional operating categories to help organisations examine these cases:

CATEGORY**Intent mismatch**

The recorded instruction is followed, but the result does not serve the customer's underlying purpose.

CATEGORY**Scope drift**

The agent changes or infers requirements beyond its authority.

CATEGORY**Stale mandate**

The agent acts on an instruction or permission that is no longer current.

CATEGORY**Cascading agent error**

An error passes between agents or systems and changes the outcome.

These categories support investigation and pattern analysis. They are not legal classifications or payment reason codes.

The same caution applies to Type A and Type B failures. Type A indicates that an AI deployed by the organisation contributed to the failure. Type B indicates that an external customer agent or platform contributed. The distinction helps identify where evidence may sit and how costs might later be recovered. It should not delay immediate customer support.



What changes when they meet

When these conversations are considered separately, each can appear manageable. The protocol team can demonstrate technical integration. Legal and risk teams can produce an AI policy. The contact centre can deploy adviser assistance. Payment and complaints teams can continue using established dispute processes.

A customer journey can still fail between them.

The operating problem also changes with the model. A retailer-owned assistant may sit close to the account, catalogue, payment and fulfilment records. An open protocol journey may distribute those records across several participants. A browser agent may act across sites that did not design their systems or service processes for it. The contact centre cannot assume that every agent-mediated transaction will produce the same evidence or route to resolution.

An external agent may interpret the customer's instruction, use merchant information supplied through one protocol, complete payment through another system and hand the resulting problem to a contact centre that cannot see the mandate or change the outcome.

No single component has necessarily failed in isolation. The weakness lies in the connections between them.

This produces service demand that does not fit neatly into an ordinary complaint or an interaction with the organisation's own AI. The case combines delegated action, incomplete evidence, several organisations and uncertainty about who can act.

The customer does not need a technical account of that fragmentation. The customer needs one party to understand the immediate problem, preserve the evidence and move the case towards a remedy.



Readiness questions

Agentic commerce should therefore be treated as a question of organisational readiness rather than as a protocol implementation or contact centre technology project.

The ACCRA framework examines readiness across seven dimensions:

Data and visibility

Platforms and integration

Workflow participation

Handoff and resolution

Workforce capability

Governance and accountability

Dispute and resolution handling

Leadership teams can begin with seven questions:

- 1 Can the organisation identify transactions in which an AI agent participated and determine whether that agent used an agreed route?

- 2 Can the service team retrieve the mandate, approval and event records needed to investigate?

- 3 Can systems preserve context when work passes from an AI service to a person?

- 4 Can advisers distinguish merchant, agent, payment and fulfilment failures?

- 5 Do they have authority to provide a provisional remedy where delay would harm the customer?

- 6 Can governance teams connect repeated service cases to changes in products, policies and controls?

- 7 Is there a defined route for evidence exchange, cost recovery and disputes involving another organisation?

A low score in one area can prevent the others from working. Detailed logs have limited value if advisers cannot access them. Skilled advisers cannot reconstruct evidence that was never captured. A sound escalation route will still fail if nobody has authority to act.



A practical review

Select one customer journey in which an external AI could discover, compare, configure or purchase a product or service.

Bring together service, commerce, operations, payments, data, legal and risk. Trace the journey from the customer's instruction to possible complaint or recovery.

FOR EACH STAGE, RECORD

- | | |
|---|--|
| ☒ Which organisation or system acts. | ☒ What evidence it creates and passes to the next party. |
| ☒ What authority is being exercised. | ☒ How the customer can stop or change the action. |
| ☒ What happens when information conflicts. | ☒ Which team provides the immediate remedy. |
| ☒ Which team investigates responsibility and recovers cost. | ☒ Which records the contact centre can actually see. |

Use the gaps to select specific changes from the ACCRA Library of Interventions. Each change should have an owner, a test and evidence that it improved the handling of the next case.

The immediate planning task is to make agent-mediated transactions visible, reviewable and recoverable before they become routine service demand.

Your Next Customer is a Bot

A series from Brainfood and Neos Wave. Follow for more content at dualcx.com

