

Your Next Customer is a Bot

A series on agentic commerce, customer contact and service changes needed when AI acts for customers.

Muse has made external AI agents a current service issue

Customers can now send agents into browser journeys and contact centres. Service operations need rules for access, authority, handover and case ownership.

Author Martin Hill-Wilson

Publication October 2026 | Pre-reading for the LinkedIn Live



Follow for more content at:

dualcx.com

In my revised article, [Why These Four Conversations Need Orchestrating](#), I argued that commerce protocols, regulation, contact centre design and dispute handling need one operating response.

Meta's Muse is now producing the type of service contact described in that article.

According to [Meta's launch announcement](#), Muse can open a browser, fill in forms, negotiate, shop and arrange travel on a customer's behalf. It can continue working after the customer closes the app and return when it needs approval for an action such as sending an email or making a purchase.

Muse also calls businesses. [Reuters reported](#) that customers have used its phone service to book appointments, check stock and obtain quotations. Tests have included negotiating a broadband bill and contacting an insurer. Market Sentiment's account adds examples such as waiting on hold, cancelling insurance, arranging a return and seeking a flight credit.

These are not demonstrations of a future shopping protocol. They are current service interactions arriving through browsers and telephones that organisations already operate.



Friction no longer limits demand in the same way

The [Market Sentiment article](#) treats Muse mainly as an investment and market disruption story. Its thesis is that personal agents will weaken businesses whose revenues depend partly on customers finding comparison, switching, cancellation, claims or redemption too inconvenient.

That is a plausible investor thesis rather than settled evidence about adoption or market impact. The article was published within two weeks of Muse's launch. Claims about mainstream behaviour, disrupted business models and future winners should therefore be treated as an early interpretation.

Its service observation is more immediate. Much existing demand is suppressed by the effort required to pursue it. Customers abandon queues, decide that a small refund is not worth the time, allow a renewal to proceed or accept a failed return because the next step involves another call.

An agent does not face the same constraint. It can wait, retry, compare offers and continue a task while the customer does something else. That changes which requests become economical to pursue.

The effect is likely to be felt first in work where persistence has value:

Cancellations and supplier changes.

Price comparisons and retention negotiations.

Refunds, credits and compensation claims.

Returns and booking changes.

Loyalty point redemption.

Appointment and availability checks.

Complaints that require repeated contact across several organisations.

The service workload has not been removed. Part of the cost has moved from the customer's time to the organisations receiving the requests.

This will affect demand forecasts. Abandonment becomes a weaker indicator of customer intent when an agent can remain in the queue. Small claims can arrive in greater numbers because the customer no longer has to decide whether each one deserves personal effort. Retention offers may be requested and compared more systematically. A failed attempt may generate automatic retries unless the agent can see that the case is already in progress.

Contact centres will need to distinguish human demand from machine initiated attempts, without losing sight of the customer task that produced both.



'AI mediated' does not identify one actor


Muse has also exposed a problem with the phrase 'AI mediated service'. It can describe a chain of actors rather than one stable participant.

Reuters found that Meta had been testing a human concierge inside Muse. Some calls presented as part of the AI service were quietly passed to contractors who contacted businesses on the customer's behalf. Meta introduced the test because some organisations ended calls when they learnt that the caller was an AI. Internal tests indicated that people completed more calls successfully.

Employees raised concerns that sensitive customer information could pass to contractors without adequate disclosure. A Meta executive accepted that testing calls placed by contractors without proper disclosure had been a mistake, and the company withdrew the feature for the time being. Meta said it would release any future version only after adding suitable privacy protections and disclosures.

This distinction matters. The public phone feature and the human concierge test were not the same thing. The phone service was being introduced to Muse users. The contractor handover was tested with employees and was then withdrawn.

The test still reveals an operating problem that applies beyond Muse. During one customer task, control can move between:



The receiving adviser may hear a person and assume that person is the customer. The customer may believe that the task remains inside a protected AI service. The contractor may have access to information that the customer intended to share with a machine under a different set of privacy assurances.

Identity checks designed for a direct conversation between one customer and one organisation do not describe this arrangement. Nor does a transaction record that names only the account holder and merchant.

The service record needs to show who was acting at each stage, the capacity in which they acted and when control changed hands.



Organisations need an external agent policy

Amazon's response to Muse shows why this cannot wait for an organisation to adopt an agentic commerce protocol.

[GeekWire reported](#) that Amazon blocked Muse from shopping on its site. Amazon said Meta had not agreed access, Muse did not identify itself while browsing, and its use of customer credentials raised privacy and security concerns. Meta disputed the account of its credential handling, saying that Muse could use credentials held in secure storage without seeing passwords or payment details.

The disagreement is important because Muse does not depend on every organisation offering an agent interface. Where an API is available, it can use that route. Where one is not available, it can use a browser. Its calling capability provides another route into an organisation.

Refusing to adopt a protocol does not prevent external agents from arriving. It means that they may arrive through channels designed for people, without agreed identity, authority, data or evidence standards.

Every service organisation therefore needs a policy that answers three questions.

1 Which routes are accepted?

The organisation should state whether external agents may use an API, website, app, messaging service or telephone channel. The answer can vary by task.

An agent checking opening times presents little risk. An agent cancelling insurance, changing payment details or accepting a continuing commitment requires stronger controls. A single rule for all agent contact will either obstruct low risk requests or expose consequential actions to inadequate checks.

Where access is refused, the organisation should provide a controlled alternative. A block without a usable route does not remove customer demand. It redirects the customer or agent towards another channel, repeated attempts or a complaint.

2 What authority must be proved?

Account access does not establish the scope of the customer's permission. An agent may be able to sign in while lacking authority to change a contract, disclose personal data or accept an offer.

BEFORE COMPLETING A CONSEQUENTIAL ACTION, THE SERVICE OPERATION NEEDS TO ESTABLISH

- | | |
|--|---|
| ☒ The identity of the customer. | ☒ The identity of the agent service and the current actor. |
| ☒ The outcome the customer requested. | ☒ The decisions the agent may make without returning to the customer. |
| ☒ Any limits on price, product, data, duration or acceptable alternatives. | ☒ The point at which the authority expires. |
| ☒ The actions that require fresh customer approval. | |

Memorable information and shared security answers are poor substitutes for delegated authority. They increase access without proving what the customer permitted the agent or contractor to do.

3 What happens when the actor changes?

The [Service Handshake](#) was proposed to preserve goals, constraints, data permissions, authority and fallback arrangements when work passes between customers, AI services and human teams.

Muse shows that the handshake must be renewed whenever control changes. A transfer from AI to contractor is not an internal implementation detail if it changes who can see the data, interpret the instruction or speak for the customer.

AT EACH TRANSFER, THE RECEIVING PARTY NEEDS A COMPACT RECORD OF

- ☒ Who is acting and for whom.
- ☒ What the customer wants to achieve.
- ☒ What has already been attempted.
- ☒ What information has been used or disclosed.
- ☒ What authority remains.
- ☒ Which action needs confirmation.
- ☒ Where the transcript, approvals and outcome will be retained.

The transfer should stop if that information is missing and the next action carries material financial, contractual or privacy consequences.



Measure the customer task, not each contact

External agents will expose a weakness in contact centre measurement. One customer task may produce several calls, browser sessions, retries, handovers and organisations. Counting each event as an independent contact hides the cause and inflates or fragments the operational picture.

An agent that makes five attempts to cancel one policy has created one customer task, five attempts and one eventual outcome. Those measures answer different questions.

1

CUSTOMER TASK

5

ATTEMPTS

1

EVENTUAL OUTCOME

SERVICE TEAMS SHOULD ADD MEASURES THAT SHOW WHETHER AGENT ACTIVITY CAN BE RECOGNISED AND RESOLVED

- ☒ The number of customer tasks involving an external agent.
- ☒ The routes used and whether each one was approved.
- ☒ Automated attempts per customer task.
- ☒ The proportion of contacts where the current actor was identified.
- ☒ The proportion where delegated authority was verified.
- ☒ The number of transfers between AI and people.
- ☒ Evidence completeness when the case reaches an adviser.
- ☒ Time to customer remedy.
- ☒ Time to decide responsibility between participating organisations.
- ☒ Costs recovered where an external service contributed to the failure.

First contact resolution remains useful, but the unit of analysis needs to be the customer's task. Otherwise, an automated sequence of failed attempts can appear as contact volume without revealing the unresolved purpose behind it.



What Muse adds to the four conversations

Muse does not create a fifth conversation. It shows what happens when the existing four remain separate.

The commerce discussion focuses on how agents gain access and complete tasks. Regulation determines the duties attached to identity, data, disclosure and customer authority. Contact centre design determines whether advisers can recognise the request and act on it. Dispute handling determines whether the evidence survives long enough to establish a remedy and allocate the cost.

Muse also shows that the service operation cannot assume there is one external agent throughout the case. A browser agent may become a voice agent. The voice agent may pass the call to a contractor. The contractor may speak to an adviser who has no record of the customer's original instruction. Each participant can complete its immediate action while the service record loses the reason for the task and the limits placed upon it.

THE IMMEDIATE RESPONSE IS OPERATIONAL

- 1 Define which routes external agents may use for each type of task.
- 2 Verify the customer's delegated authority before consequential action.
- 3 Renew the Service Handshake whenever the actor changes.
- 4 Link every attempt, transfer and outcome to the customer task that caused them.

Each accepted route should have an owner, authority rules, confirmation points, an evidence record and a tested route to remedy.

GET READY FOR DUAL CX

Brainfood and Neos Wave offer diagnostic and design services to prepare contact centres for agentic commerce. We also run [LinkedIn Live discussions](#), [podcasts](#) and a [resource centre](#) to help educate, swap experiences and guide the redesign for how the new world of Dual CX works.

www.dualcx.com

Your Next Customer is a Bot

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

© 2026 Brainfood Consulting. All rights reserved.

