



RESEARCH

An independent audit of UK customer-service infrastructure

State of Dual CX

Mode 3/4 Readiness Across 100 UK Brands

100

BRANDS AUDITED

17

SECTORS

0%

READY FOR AN
ARRIVING AI AGENT

Edition 1.0 · Findings collected Q1 2026 · First published March 2026

Author: Maria McCann, Founder, Neos Wave

Licence: Creative Commons Attribution 4.0 International (CC BY 4.0)

DOI: 10.5281/zenodo.20475323 · Open dataset: github.com/neos-wave/state-of-dual-cx

Standard referenced: The Service Handshake v1.1 — DOI 10.5281/zenodo.19046746

ABSTRACT

Summary of findings

Across 100 brands and 17 sectors of the UK economy, customer-service operations have invested heavily and successfully in one kind of artificial intelligence: the AI a human talks to. Seventy-six brands have deployed chatbots, sixty-six proactively disclose that a customer is talking to AI, and eighty-one offer live chat. That investment is real and, for the most part, it works.

This audit examines a different question: what happens when the customer brings their own AI. Consumer AI agents are already navigating support channels, passing identity checks and requesting resolutions on behalf of their users. The finding is uniform. Not one of the 100 brands has published any protocol for receiving an autonomous agent acting for a customer – no declaration of what its own AI is authorised to commit to, what data may be exchanged, how the customer's agent should identify itself, or what happens when something goes wrong.

The headline is a structural gap the report calls the **readiness cliff**: a large majority of brands have capable, mature AI infrastructure on the human-facing side, and none has designed for the agent-facing side. This is not a technology gap – more than half the sample already has the technical prerequisites – but a **design gap**. The report sets out the gap sector by sector, identifies where the stakes are highest, and discusses the implications for organisations, platform vendors and regulators.

A NOTE ON NEUTRALITY AND AUTHORSHIP

This report is observational and vendor-neutral: no brand was contacted, all observations derive from publicly accessible interfaces, and the full dataset is open under CC BY 4.0. It is published by Neos Wave, which also authored The Service Handshake, one proposed open standard for the gap the report identifies. The audit measures the **absence of any published agent-receiving design** – of any origin, proprietary or open – and not the adoption of any single standard. That distinction is preserved throughout.

CONTENTS

1. Methodology & scope – what was and was not assessed; definitions
2. Key findings at a glance
3. The readiness cliff
4. Cross-cutting findings
5. Sector-by-sector analysis
6. Mode 4 candidate ranking
7. Implications
8. Recommendations
9. Appendix: full 100-brand technology table

1. Methodology & Scope

This audit is a light-touch observational assessment of publicly accessible customer-service infrastructure, conducted in Q1 2026. All 100 brands were evaluated against the same dimensions using unauthenticated browser sessions. Where authenticated access was required to reach a support channel, that authentication dependency is recorded as a finding – it is not bypassed. No brand was contacted during the audit.

Definitions: the four interaction modes

Service interactions are now occurring in four modes. **Mode 1** is human to human. **Mode 2** is human to a brand's AI – the mature, familiar interaction the market has built for. **Mode 3** is a consumer's AI contacting a brand's systems. **Mode 4** is a brand's AI and a consumer's AI interacting at the agent layer. Modes 3 and 4 are the subject of this report; they are already occurring, and no equivalent design thinking has been applied to either.

What was assessed

- **Mode 2 maturity** – whether human-to-bot interaction is operational, mature and clearly disclosed. Rated **Green** (fully operational with clear AI disclosure), **Amber** (partial implementation, inconsistent disclosure, or authentication dependency), or **Red** (no digital support capability on public surfaces).
- **Bot presence and disclosure** – whether a bot exists and whether AI identity is disclosed proactively.
- **Chat provider** – the technology stack powering the interface, where identifiable.
- **Developer API availability** – whether a public developer API exists that could support agent-initiated contact.
- **Live chat accessibility** – whether live chat is reachable without authentication.
- **Mode 4 candidacy** – see the caveat below.
- **Mode 3/4 readiness** – whether the brand has published any declaration of how it intends to receive consumer AI agents: what they may do, what data they may access, and what the fallback rules are.

IMPORTANT CAVEAT – WHAT "MODE 4 CANDIDATE" DOES AND DOES NOT MEAN

"Mode 4 candidate" is a **capability signal, not a readiness certification**. A brand qualifies if it has the technical prerequisites – a bot, an API, or both – that would make Mode 4 implementation tractable. It is a deliberately low bar, and the headline figure of 57% should be read strictly in that light: it means "more than half the sample has the raw infrastructure," not "more than half is close to ready." Readiness is measured separately, and is zero.

IMPORTANT CAVEAT – WHAT THE 0% MEASURES

The 0% readiness figure records that **no brand has published any agent-receiving protocol of any kind** – proprietary, open, or standards-body. It is not a measure of adoption of The Service Handshake or of any single framework. The finding would stand unchanged if every reference to a named standard were removed: the receiving side of the interaction has simply not been designed anywhere in the sample.

Data provenance: a note on the technology attributions


Every "chat provider" attribution in this report is inferred from public interface analysis – not confirmed by the named organisation. No company in the audit has publicly declared the system behind its support surface; the attributions are deduced from how each interface behaves and loads. They are therefore indicative, and individual entries may be inaccurate. This caveat applies only to the technology-stack column. The mode ratings, bot presence, AI disclosure, live-chat accessibility and API findings rest on directly observable interface behaviour and carry the normal confidence of an observational audit.

What was not assessed

- Authenticated support experiences (except where the login dependency itself is a finding).
- Internal system architecture not visible from public interfaces.
- Quality of existing bot interactions (containment, CSAT, resolution rates).
- Back-end API capability beyond public documentation.
- Regulatory compliance beyond what is observationally apparent.
- Future readiness based on vendor roadmaps or press releases.
- Tone, sentiment or accessibility compliance of deployed bots.

SECTION 2

2. Key Findings at a Glance



*Capability signal only – has a bot, an API, or both. Not a readiness measure (see Section 1).

Companies are investing heavily in AI assistants, and for the most part the investment is working. When a customer contacts a brand today, there is a good chance they will interact with an AI that is named, branded and capable of resolving their query. **None of these brands has designed for what happens when the customer brings their own AI.**

The asymmetry shows up in every metric:

- 76 brands have deployed AI assistants. None has declared what those assistants are authorised to commit to when contacted by another AI.
- 66 brands tell customers they are talking to AI. None asks whether the customer contacting them is also an AI.
- 42 brands have developer APIs. None has published a specification for how a customer's AI agent should use them.
- 81 brands have live chat. None has a receiving protocol for when an agent opens that window on a customer's behalf.

Metric	Count	%
Brands audited	100	–
Sectors covered	17	–
Mode 2 maturity: Green	73	73%
Mode 2 maturity: Amber	25	25%
Mode 2 maturity: Red	2	2%
Mode 4 candidates (capability signal)	57	57%
Brands with bots	76	76%
AI disclosure visible	66	66%
Developer APIs present	42	42%
Live chat deployed	81	81%

Metric	Count	%
Published any agent-receiving protocol	0	0%

3. The Readiness Cliff

The most important comparison in this report is between two numbers: **73% Mode 2 capability and 0% agent-receiving readiness**. That gap is the readiness cliff.

The Mode 2 investment story is substantial. Named bot characters — Fin, Zea, Agentforce, Freddy, TOBi, Cora, Arti, Sandi, JAX, Rita, Ezra, Billie, Theo, Rafeeq, Molli — represent genuine brand investment in AI-mediated service identity. The sector-wide rollout of conversational AI over the past four years has changed consumer expectations at scale.

The agent-receiving story is a blank. No brand across any of the 17 sectors has published how a consumer AI should identify itself on arrival, what goals it may pursue, what data it may access, what it may not commit the brand to, or what the fallback is when agent and brand AI reach a conflict or a request beyond the agent's authority.

THE LIABILITY IS ALREADY ESTABLISHED — MOFFATT V. AIR CANADA (2024 BCCRT 149)

An airline was held liable for its chatbot's unauthorised discount promise, on the principle that an organisation is responsible for the representations of its automated systems — regardless of whether those systems had authority to make them. One hundred brands operate AI systems making commitments across 81 live-chat channels. Not one has published a protocol for what those systems are, and are not, authorised to commit to when the other party is itself an AI. The exposure is not hypothetical; it is precedented.

The cliff is structural, not incidental. Mode 2 was designed for a human-to-brand-AI interaction that is now mature. Modes 3 and 4 require a different design decision — the publication of a receiving protocol — and that decision has not been taken at any of the 100 brands audited.

4. Cross-Cutting Findings

1 • AI disclosure is standard; AI reception is not

Sixty-six of 100 brands proactively disclose AI at first contact. The market has moved decisively on transparency – naming bots, labelling interactions, in some cases publishing AI policies. But this is Mode 2 disclosure, designed for humans who need to know they are talking to a machine. It does not address what a brand's AI does when the party initiating contact is itself an AI. No brand has published the receiving side of that disclosure.

2 • The "Dadbot" scenario has no receiving infrastructure

THE DADBOT SCENARIO

A shorthand, used throughout this report, for the case where one person's AI agent contacts an organisation on behalf of another, more vulnerable person – for example, an adult child's agent managing an elderly parent's energy account, insurance policy or banking. It is the interaction pattern where delegated authority, identity verification and safe fallback matter most, and where mishandling carries the greatest human cost.

The highest-stakes pattern identified is a family member's AI agent managing an energy contract, insurance policy or bank account on behalf of a vulnerable relative. The sectors most relevant – Utilities, Banking, Insurance, Healthcare, Charity – show the weakest agent-receiving readiness against their Mode 2 investment. British Gas, Aviva, NatWest, Bupa and Age UK can all handle the scenario technologically. None has a declared protocol for what an agent is authorised to do, what data it may provide, or what the fallback is when its authority is unclear.

3 • The protocol stack has a missing layer

Shopify's published agent-commerce tooling is clear evidence that the protocol layer for agent-to-system commerce is already being built by market leaders. The missing layer is the support-surface declaration: what happens when a consumer agent that has just completed a transaction needs to contact support? Forty-two brands have developer APIs; none has a published declaration. The API-declaration gap is the protocol stack's missing middle.

4 • CX platform vendors have not designed for their own product's future

The vendors whose products define the agent-era design challenge for their customers have not designed agent-receiving infrastructure for their own surfaces. Zendesk sells the tools UK retailers use to receive interactions; its own support surface has no declaration. Intercom markets Fin as the future of AI-mediated service; it has no protocol for receiving a consumer agent that arrives to resolve a billing dispute with Intercom. Salesforce markets Agentforce as the enterprise standard; its own Agentforce-

powered surface has no declaration. The vendors leading the market toward agent-era CX are the most conspicuously absent from the readiness picture.

5 · Fintech and banking are API-ready but declaration-absent

Open Banking, PSD2 and competitive dynamics have produced API infrastructure technically capable of supporting Mode 4 interaction. The authorisation layer — OAuth, token scoping, delegated access — is present. The semantic layer — what the AI assistant may commit to, and what a consumer agent may request — is entirely absent. The institutions with the most complete technical infrastructure are furthest from having designed what the interaction should be.

6 · Travel has built the bot layer without the agent layer

Seven of eight travel brands are Mode 4 candidates — the highest rate of any consumer-facing sector — and not one has published a receiving protocol. The Moffatt precedent is a travel-sector precedent, so the liability exposure here is the most directly precedented in the audit.

7 · Government digital services are more ready than their reputation suggests

HMRC's Developer Hub, Making Tax Digital API and Digital Engagement Platform combine to create a government service surface more Mode 4 capable than many commercial brands. AI disclosure across the GOV.UK estate is more consistent than most commercial sectors. The gap between government's actual infrastructure and its reputation for digital laggardness is sharper than expected.

8 · Insurance is the vulnerability blind spot

Insurance's 2/6 Mode 4 candidate rate is the lowest of any regulated consumer sector, yet its interactions are disproportionately concentrated in moments of acute vulnerability — claim time, bereavement, accident, hardship. The sector has invested least in the infrastructure that would enable agent-receiving design, while carrying some of the highest stakes for getting it wrong.

SECTION 5

5. Sector-by-Sector Analysis

Each sector is assessed on the same dimensions. Mode 4 candidacy throughout denotes technical capability only (a bot, an API, or both); no brand in any sector has published an agent-receiving protocol.

CX Platforms (6 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Freshdesk (Freshworks)	Yes	Freddy AI	Yes	Green	Yes
Help Scout	Yes	AI Assistant	Yes	Green	Yes
Intercom	Yes	Fin	Yes	Green	Yes
Salesforce	Yes	Agentforce	Yes	Green	Yes
Tidio	Yes	Lyro AI Agent	Yes	Green	Yes
Zendesk	Yes	Zea	Yes	Green	Yes

The most complete sector on every Mode 2 dimension: 6/6 Green, bots, AI disclosed, APIs, live chat and Mode 4 candidacy. No sector comes closer to the technical prerequisites – and none's absence of an actual declaration is more consequential. Intercom's Fin is the most explicitly agent-era-framed bot in the audit; Zendesk's Zea identifies as an AI agent from the first message; Salesforce has renamed its core AI product Agentforce. Yet none has published what these systems are authorised to do when a consumer agent, rather than a human, initiates contact. The design paradox is sharpest here: these companies have designed the future of CX for thousands of customers and have not designed the receiving layer for their own surfaces. The sector that most needs to adopt a receiving protocol first is the one selling agentic CX to everyone else.

UK Retail (15 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Amazon UK	Yes	Amazon Virtual Assistant	Yes	Green	Yes
ASOS	No	ASOS AI Assistant	Yes	Green	No
Argos	No	Argos Virtual Assistant	Yes	Green	Yes
Asda	No	Unknown	No	Amber	No
Boots	No	Boots AI Assistant	Yes	Green	Yes
Currys	No	Currys Virtual Assistant	Yes	Green	Yes
John Lewis & Partners	No	N/A	No	Green	No
Marks & Spencer	No	Rafeeq	No	Green	Yes

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Morrisons	No	N/A	No	Amber	No
Next	No	NEXT Bot	Yes	Green	Yes
PrettyLittleThing	No	Madison	No	Amber	No
Sainsbury's	No	Virtual Assistant	Yes	Green	No
Tesco	No	Little Helper	Yes	Amber	No
Very.co.uk	No	Very Assistant	Yes	Green	Yes
Waitrose & Partners	No	N/A	No	Amber	No

The largest and most heterogeneous sector: 9 Green, 6 Amber, 7 Mode 4 candidates. Amazon UK has the most sophisticated infrastructure by some distance and is the clearest retail Mode 4 candidate – and simultaneously the most opaque, behind a near-total authentication wall. M&S's Rafeeq does not clearly disclose its AI identity, a transparency gap even at the entry point. Tesco (WhatsApp-first) and PrettyLittleThing (social-messaging-first) route support through Meta-controlled channels, creating an agent surface outside the brand's own API governance. The John Lewis / Waitrose pair – premium brands in the same group – has the least agentic infrastructure in its sub-segments, making the agent-receiving gap not merely absent but structurally unprepared.

UK Retail – Home/DIY (6 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
A0.com	No	A0 Assistant	Yes	Green	Yes
B&Q (diy.com)	No	B&Q DIY Assistant	Yes	Green	Yes
Halfords	No	N/A	No	Green	No
IKEA UK	Yes	Billie	Yes	Green	Yes
Screwfix	No	Screwfix Virtual Assistant	Yes	Green	Yes
Wickes	No	Wickes Virtual Assistant	Yes	Green	No

A clean Mode 2 picture: 6/6 Green. IKEA's Billie is the standout – the only Home/DIY brand with the full stack (bot, API, live chat, disclosure), handling roughly 47% of enquiries without escalation, and using AI to reposition store staff toward design consulting rather than simply replacing them. None of this translates into agent-receiving readiness; Billie has no published protocol. B&Q's dual-bot design (product advice vs Zendesk support) implies a clearer internal sense of what each bot is authorised to do – a structure a declaration would formalise.

UK Utilities (10 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
British Gas	No	Help Finder Bot / Cosmo	Yes	Green	Yes
E.ON UK / E.ON Next	No	N/A	No	Amber	No
EDF Energy	No	N/A	No	Amber	No
Ecotricity	No	N/A	No	Amber	No
Octopus Energy	Yes	N/A	No	Amber	Yes
Octopus Energy US	No	N/A	No	Amber	No
OVO Energy	Yes	OVO Digital Assistant	Yes	Green	No
ScottishPower	No	ScottishPower Chatbot	Yes	Green	Yes
SSE	No	Nero	No	Green	Yes
Thames Water	No	Virtual Assistant	Yes	Green	No

The highest-stakes sector for agent-receiving failure and the most uneven Mode 2 baseline – four of ten have no public bot at all. The Dadbot scenario is most acute here. British Gas's Cosmo is the technical leader, handling billing, meter readings, direct-debit changes and vulnerability flags – all with direct agent-receiving implications that are currently undocumented. Octopus is the sector paradox: a widely cited CX innovator (its Kraken platform is licensed out) whose public support surface has no bot, no live chat and no disclosure, with Mode 4 candidacy resting entirely on API infrastructure. This "back-end built, front-end absent" pattern recurs across the audit.

UK Telecoms (5 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
BT	Yes	BT Digital Assistant	Yes	Green	Yes
Sky UK	Yes	Sky Virtual Assistant	Yes	Green	Yes
Three UK	No	Three Digital Assistant	Yes	Green	Yes
Virgin Media	No	Terri BOT	Yes	Green	Yes
Vodafone UK	Yes	TOBi / SuperTOBi	Yes	Green	Yes

The highest-performing non-platform sector: 5/5 Green, named bots, disclosure, live chat and Mode 4 candidacy. Vodafone's SuperTOBi – moved from IBM Watson/LUIS to Azure OpenAI generative AI – is the most advanced consumer-facing telecoms deployment in the audit, yet has no declaration for agent-initiated contact: the most capable bot in the sector has the least protection against agentic contact it was not designed for. Sky's announced shift of 2,000 call-centre roles toward digital support makes agent-receiving design a near-term operational necessity, not an optional future. Telecoms also carries

a compounding exposure: it operates the communication infrastructure consumer agents will use to reach other brands, creating multi-agent chains no brand has a protocol for.

SaaS (8 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
HubSpot	Yes	HubBot	Yes	Green	Yes
monday.com	Yes	Amy / Tim	Yes	Green	Yes
Pylon	Yes	N/A	No	Green	Yes
Shopify	Yes	Virtual Help Centre Assistant	Yes	Green	Yes
Slack	Yes	Agentforce	Yes	Green	Yes
Stripe	Yes	Assistant	Yes	Green	Yes
Xero	Yes	JAX	Yes	Green	Yes
Zoom	Yes	ZVA	Yes	Green	Yes

The most technically complete sector: 8/8 Green, 8/8 Mode 4 candidates, 8/8 with developer APIs – and 0/8 on readiness. Shopify's agent-commerce tooling (UCP/MCP) is the single most significant structural development in the audit: it is designed for agent-to-system interaction, yet Shopify has no receiving declaration for its own support surface. Xero's JAX – built on Anthropic Claude and described as an agentic platform – handles most first-contact support, yet has no protocol for when the party initiating the session is another AI: an agentic front-end with no specification for encountering another agent. monday.com's dual-provider design (Qualified for sales, Ada for support) already separates sales-context from support-context authority – exactly the scoping a declaration would formalise.

UK Fintech (3 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Monzo	Yes	AI Assistant	Yes	Green	Yes
Revolut	Yes	Rita	Yes	Green	Yes
Starling Bank	Yes	N/A	No	Green	Yes

The cleanest Mode 4 candidate profile in the audit – 3/3 Green, 3/3 APIs – and 3/3 gate all interactive support behind authentication, so there is no public chat surface to receive an agent at all. Monzo's bespoke in-app assistant plus OAuth API is the closest technical profile to readiness, but with no declaration of what the assistant may do or what a consumer agent may initiate. Starling's API is a reference implementation for open banking; a machine-readable declaration could be published alongside the API documentation with no conversational interface required at all.

UK Banking (8 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Barclays	Yes	Barclays Digital Assistant	Yes	Green	Yes
Halifax	Yes	Virtual Assistant	Yes	Green	No
HSBC UK	Yes	Virtual Agent	Yes	Green	Yes
Lloyds Bank	No	Virtual Assistant	Yes	Green	No
NatWest	Yes	Cora / Cora+	Yes	Green	Yes
Nationwide	Yes	Arti	Yes	Green	Yes
Santander UK	No	Sandi	Yes	Green	Yes
TSB Bank	No	TSB Smart Agent	Yes	Green	Yes

The only sector to score full marks on both bot deployment and transparency: 8/8 Green, 8/8 bots, 8/8 disclosure; the deepest named-bot investment of any sector. NatWest's Cora+ (IBM Watsonx, with OpenAI collaboration for agentic features) is the most sophisticated bank bot in the audit – the bank has already navigated the governance of deploying generative AI in a regulated context, but has not extended that framework to what Cora should do when it meets a consumer agent. The gap is design, not technology. Nationwide's Arti discloses that it summarises conversations for human colleagues on escalation – a nascent form of context handoff, the closest thing in banking to thinking about agent-to-agent continuity.

UK Insurance (6 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Admiral	Yes	A.V.A.	Yes	Green	Yes
Aviva	No	Aviva Chatbot	Yes	Green	Yes
AXA UK	No	N/A	No	Amber	No
Compare the Market	No	AutoSergei	No	Amber	No
Direct Line	No	Direct Line VA	Yes	Green	No
Vitality	No	N/A	No	Amber	No

The vulnerability blind spot: 3/6 Green, 2/6 Mode 4 candidates – the weakest regulated consumer sector in the study, despite interactions that cluster at moments of acute vulnerability (claims, bereavement, accident, hardship). Admiral's A.V.A. is the only sector entry with the full infrastructure stack and the sole credible Mode 4 candidate by technical measure. The critical use case is an agent managing a household's insurance and initiating a claim on behalf of an incapacitated user – needing a protocol for what the bot may accept, what it may not commit to, and what the human handoff looks like. None exists.

UK Travel (8 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Airbnb	Yes	AI Assistant	Yes	Green	Yes
Booking.com	Yes	Helpbot AI	Yes	Green	Yes
British Airways	Yes	Chatbot	Yes	Green	Yes
easyJet	No	Virtual Advisor	Yes	Green	Yes
Ryanair	No	Molli	No	Amber	No
Trainline	Yes	Juno / Cleo	No	Amber	Yes
TUI UK	Yes	Theo	Yes	Green	Yes
Uber	Yes	Unknown	Yes	Green	Yes

The highest Mode 4 candidate rate of any consumer-facing sector (7/8), and the most directly precedented liability. British Airways is the closest parallel to *Moffatt v. Air Canada*: a consumer agent interacting with BA's chatbot about a cancellation or delay-compensation claim creates exposure no declaration currently mitigates. Booking.com's Helpbot AI is the most globally scaled bot in the sector. Ryanair's Molli is the outlier – present but undisclosed, and it requires name and email before interaction, a pre-authentication step that is itself a friction point for an agent that cannot supply human personal data.

UK Delivery / Logistics (5 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Deliveroo	Yes	Rider Support AI Agent	Yes	Green	Yes
DPD UK	No	DPD Chatbot	No	Green	Yes
Evri (formerly Hermes)	No	Ezra	Yes	Green	No
Just Eat	No	N/A	No	Amber	No
Royal Mail	No	N/A	No	Amber	No

High agent-relevance: tracking, rescheduling and returns are among the most frequent early consumer-agent use cases. Deliveroo's rider support runs on Sierra, an AI-native platform with agent-era thinking built into its product design – even without a published declaration for either surface. Royal Mail is the surprising gap: the national postal service has no live chat, no bot and no disclosure, a material operational risk as consumer-agent adoption grows.

UK Media (6 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
BT Sport / TNT Sports	Yes	BT Digital Assistant	Yes	Green	No
Disney+	No	Virtual Assistant	No	Amber	No
Netflix	No	N/A	No	Green	No
NOW (formerly NOW TV)	No	NOWBot	No	Green	No
Spotify	No	AI-Powered Tool	Yes	Green	No
The Guardian	No	N/A	No	Amber	No

The only sector with zero Mode 4 candidates: no support-context API exists at any of these brands, so there is no technical pathway to Mode 4. Netflix has sophisticated triage and AI-powered search but no disclosure, no named bot and no support API. The sector's agent-receiving urgency is genuinely lower – interactions are subscription, billing and technical rather than financial or legal – but as subscription management is delegated to household agents, media will face the same contact pattern as every other sector without any current infrastructure for it.

UK Healthcare (2 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Bupa UK	Yes	BV (Bupa Virtual Assistant)	Yes	Green	Yes
NHS	Yes	N/A	No	Green	Yes

A small sector with disproportionate significance: health-adjacent agents (medication reminders, appointment booking, condition monitoring, claims) are among the earliest consumer-agent use cases in active development. Bupa's BV is disclosed and API-backed; the healthcare context adds a data-sensitivity dimension no current declaration addresses – what a consumer agent may disclose to BV on a patient's behalf, and what BV may commit to. The NHS API platform is substantive, but its webchat has no named bot, no disclosure and no protocol for out-of-hours agent contact or requests requiring clinical authority.

UK Government (3 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
DWP	No	N/A	Yes	Amber	No
GOV.UK	No	HMRC Digital Assistant	Yes	Green	No
HMRC	Yes	HMRC Digital Assistant	Yes	Green	Yes

More capable than its reputation suggests. HMRC's Developer Hub, Making Tax Digital API and Digital Engagement Platform make it technically more Mode 4 capable than many commercial brands – and MTD's API infrastructure is explicitly designed for software-agent interaction. An agent managing a self-employed person's tax obligations is a direct extension of current policy, not a speculative future; what is missing is the declaration of what the assistant may commit to and what an agent may submit. DWP's transparency-without-infrastructure position (AI acknowledged, no public bot or API) is unique in the audit.

UK Property (3 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
OpenRent	No	N/A	No	Amber	No
Rightmove	Yes	N/A	No	Amber	No
Zoopla	No	N/A	No	Amber	No

The weakest sector on Mode 2 metrics: 0/3 Green, 0/3 candidates, 0/3 disclosure – no consumer-facing conversational AI at all. This is significant given that Rightmove and Zoopla mediate some of the most consequential financial decisions consumers make. A consumer agent searching listings or managing a rental application would meet a static search interface built for human browsers, not agent-mediated transactions. The gap here is structural: this is a sector that needs to build Mode 2 before it can contemplate Mode 4.

UK Charity / Advocacy (4 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Age UK	No	Virtual Assistant	Yes	Green	No
Council Tax Support	No	N/A	No	Red	No
Macmillan Cancer Support	No	Macmillan Support Bot	Yes	Amber	No
Shelter UK	No	N/A	No	Amber	No

1/4 Green, 0/4 candidates – partly a resource constraint, partly a deliberate human-first philosophy reflecting the vulnerability of the populations served. Age UK serves precisely the demographic most likely to have interactions mediated by a trusted family member's AI – the Dadbot scenario – and has no protocol for receiving them. Macmillan is notable for having published an AI policy (one of very few), giving it governance infrastructure closer to a declaration than most commercial brands, even though that policy does not yet cover the agent-receiving question.

Public / Regulatory (2 brands)

Brand	Dev API	Bot	AI disclosed	Mode 2	Mode 4 cand.
Citizens Advice	No	Citizens Advice Bot	Yes	Green	No
Ofgem	No	N/A	No	Red	No

The sharpest binary in the audit. Citizens Advice runs a triaging bot that identifies itself and offers email fallback, serving a population with complex, multi-agency needs increasingly managed by family members or agents. Ofgem – the regulator whose remit covers the energy sector this audit finds most critical – has no digital engagement channel and no AI infrastructure, and so cannot itself model the agent-receiving standard it might eventually require of the entities it regulates.

6. Mode 4 Candidate Ranking

Fifty-seven of the 100 brands meet the threshold for Mode 4 candidacy: sufficient technical infrastructure – bot, API, or both – to make implementation tractable. **Candidacy is a capability signal, not a readiness assessment.** Candidates are ranked by infrastructure sophistication across four dimensions: API maturity, bot capability, live-chat accessibility and AI transparency.

Tier 1 – Full-stack infrastructure (API + named bot + live chat + AI disclosure)

Intercom leads: REST API, webhooks, SDKs, the Fin agent with explicit AI identity, and live chat without login. Intercom has built the enabling infrastructure for Mode 4 for its customers – which makes its own absence of a receiving declaration the most precise instance of the gap in the study. The CX platform vendors (Salesforce, Zendesk, Freshworks, Tidio, Help Scout) cluster at this tier alongside the most complete SaaS and telecoms brands (Zoom, HubSpot, Shopify, BT, Vodafone, Sky).

Tier 2 – Strong infrastructure with one gap

Monzo is representative: a bespoke AI Assistant, a mature OAuth API and explicit disclosure, with the authentication wall as the single limiting factor on public-surface access. The regulated API-rich brands (NatWest, Nationwide, Barclays, HSBC, Starling, Revolut, Stripe, Xero, Airbnb, Booking.com, British Airways, TUI, HMRC, Bupa) populate this tier – strong on infrastructure, gated or undisclosed on one dimension.

The full ranked list and the per-dimension scoring are available in the open dataset at github.com/neos-wave/state-of-dual-cx.

7. Implications

1 · The demand signal is structural, not speculative

Readiness is at zero across every sector – including those with the most sophisticated infrastructure in UK consumer services. These are not organisations unaware of the agent trajectory: Shopify is publishing agent tooling, Vodafone runs SuperTOBi on Azure OpenAI, Xero runs JAX on Anthropic Claude, NatWest runs Cora+ on Watsonx. They have built Mode 2 comprehensively and have not yet designed the agent-receiving layer. The demand signal therefore arises from a design gap visible across the most capable brands in every sector – not from ignorance or inaction.

2 · Adoption topology favours platform partners

The CX platform sector – most capable, least ready, and powering the support surfaces of every other sector – defines the adoption topology. Zendesk appears in this audit as a vendor, as the infrastructure behind Deliveroo, Next, B&Q, Zoopla and Trainline, and as the layer beneath dozens of brands outside it. The same is true of Salesforce (John Lewis, BA, Sainsbury's, Slack, The Guardian, AXA), LivePerson (HSBC, Sky, Nationwide, Thames Water), Genesys (Admiral, Three, Virgin Media, Direct Line, OVO) and Sprinklr (BT, Netflix, Macmillan). An agent-receiving declaration templated at the platform layer would propagate to hundreds of brands without each designing independently. The 57 candidates are the natural first wave for direct engagement; the platforms beneath them are the leverage point for scale.

3 · Regulatory timing is actionable

Moffatt v. Air Canada (2024 BCCRT 149) established liability for AI commitments regardless of authority. The UK Consumer Duty requires good outcomes in automated interactions; Ofgem's vulnerability frameworks extend in principle to AI-mediated contact; FCA AI/ML guidance is in development; GDPR rights apply to automated decision-making that intersects directly with agent interaction design. For the four regulated sectors here – Banking, Insurance, Healthcare, Government – an agent-receiving declaration is increasingly a compliance prerequisite, not merely good practice. The window for proactive declaration, before a format is mandated, is open now and will not stay open indefinitely.

8. Recommendations

For brands

Immediately: audit what your named AI assistant is actually authorised to commit to – not what it is technically capable of saying. The gap between those two things is your Moffatt exposure.

Near term: publish a receiving declaration – human- and machine-readable – specifying what goals an inbound consumer AI may pursue, what data it may access and provide, what your AI is authorised to commit to, what is out of scope, and the fallback when confidence or authority is insufficient.

Regulated sectors (Banking, Insurance, Healthcare, Government): frame the declaration as a compliance document, not an innovation initiative. The Consumer Duty, FCA guidance and GDPR rights all create obligations a declaration directly addresses; its absence is increasingly a regulatory-risk position.

Vulnerable-consumer sectors (Utilities, Banking, Insurance, Charity): design explicitly for the Dadbot scenario – what a family member's delegated agent may do, how it must identify itself, what verification is required, and the human escalation path when vulnerability signals are present.

CX platform vendors: the most urgent recommendation in the audit. Publish your own receiving declaration for your own surface, then build declaration tooling into your platform as a native feature. The platform that ships agent-receiving readiness as a default capability will define the standard for the market it powers.

For regulators

FCA and PRA: the API-declaration gap in banking and insurance is a regulatory gap, not merely a design one. Guidance on AI-assistant authority scoping would give institutions the mandate to act on infrastructure they already have.

Ofgem: the energy sector's Mode 2 gaps and the Dadbot scenario together form a specific regulatory risk; vulnerability frameworks should explicitly address AI-mediated interactions, and the regulator's own digital absence undermines its capacity to model the standard it may need to set.

ICO: guidance on automated decision-making under GDPR should explicitly address the data-rights implications of AI-to-AI interactions, where a consumer agent provides personal data to a brand AI with no human intermediary.

Scaling this research

This 100-brand, 17-sector audit establishes a baseline methodology extensible to further sectors and geographies – financial services beyond banking and fintech, UK local government, retail banking in the US and EU, healthcare beyond Bupa and the NHS, and the global travel sector. The receiving-declaration

gap is expected to be universal across geographies until a major market event – a regulatory mandate, a high-profile Moffatt-equivalent ruling, or platform adoption – creates the first wave.

THE STANDARD REFERENCED IN THIS REPORT

The Service Handshake v1.1 is one open, vendor-neutral standard (CC BY 4.0) that defines the receiving declaration this report finds missing: how an AI agent identifies itself, what it is authorised to do, and what happens when things go wrong. It is published independently of this audit and available at DOI 10.5281/zenodo.19046746. It is referenced here as an example of the class of solution the findings call for, not as the only possible answer.

APPENDIX

9. Full 100-Brand Technology Table

Mode ratings, bot presence, AI disclosure, live chat and API columns rest on directly observable interface behaviour. **All chat-provider attributions are inferred from public interface analysis, not confirmed by the named organisations – they are indicative and may be inaccurate.**

Brand	Sector	Mode 2	Chat provider (inferred)	Bot	AI disc.	Live chat	Dev API	Mode 4 cand.
Freshdesk	CX Platform	Green	Freshchat	Freddy AI	Yes	Yes	Yes	Yes
Help Scout	CX Platform	Green	Help Scout Beacon	AI Assistant	Yes	Yes	Yes	Yes
Intercom	CX Platform	Green	Intercom	Fin	Yes	Yes	Yes	Yes
Salesforce	CX Platform	Green	Agentforce	Agentforce	Yes	Yes	Yes	Yes
Tidio	CX Platform	Green	Tidio	Lyro AI Agent	Yes	Yes	Yes	Yes
Zendesk	CX Platform	Green	Zendesk	Zea	Yes	Yes	Yes	Yes
Amazon UK	UK Retail	Green	Amazon (custom)	Amazon Virtual Assistant	Yes	Login req.	Yes	Yes
ASOS	UK Retail	Green	Microsoft Azure OpenAI	ASOS AI Assistant	Yes	Yes	No	No
Argos	UK Retail	Green	Unknown	Argos Virtual Assistant	Yes	Yes	No	Yes
Asda	UK Retail	Amber	Salesforce	Unknown	No	No	No	No
Boots	UK Retail	Green	Synthetic	Boots AI Assistant	Yes	Yes	No	Yes
Currys	UK Retail	Green	Emplifi / Vyntelligence	Currys Virtual Assistant	Yes	Yes	No	Yes
John Lewis	UK Retail	Green	Salesforce	N/A	No	Yes	No	No
Marks & Spencer	UK Retail	Green	Google Cloud CCAI	Rafeeq	No	Yes	No	Yes
Morrisons	UK Retail	Amber	N/A	N/A	No	No	No	No
Next	UK Retail	Green	Zendesk	NEXT Bot	Yes	Yes	No	Yes
PrettyLittleThing	UK Retail	Amber	Mission Labs	Madison	No	Yes	No	No
Sainsbury's	UK Retail	Green	Salesforce	Virtual Assistant	Yes	Yes	No	No
Tesco	UK Retail	Amber	WhatsApp / Meta	Little Helper	Yes	No	No	No
Very.co.uk	UK Retail	Green	IBM Watson / Salesforce	Very Assistant	Yes	Login req.	No	Yes

Brand	Sector	Mode 2	Chat provider (inferred)	Bot	AI disc.	Live chat	Dev API	Mode 4 cand.
Waitrose	UK Retail	Amber	WhatsApp	N/A	No	No	No	No
AO.com	Home/DIY	Green	Engageware / custom	AO Assistant	Yes	Yes	No	Yes
B&Q	Home/DIY	Green	Zendesk	B&Q DIY Assistant	Yes	Yes	No	Yes
Halfords	Home/DIY	Green	Salesforce	N/A	No	Yes	No	No
IKEA UK	Home/DIY	Green	Custom (Ingka Group)	Billie	Yes	Yes	Yes	Yes
Screwfix	Home/DIY	Green	NICE CXone	Screwfix VA	Yes	Yes	No	Yes
Wickes	Home/DIY	Green	Rocket.chat	Wickes VA	Yes	Yes	No	No
British Gas	UK Utilities	Green	Custom (Vite)	Help Finder Bot	Yes	Yes	No	Yes
E.ON UK	UK Utilities	Amber	N/A	N/A	No	No	No	No
EDF Energy	UK Utilities	Amber	N/A	N/A	No	No	No	No
Ecotricity	UK Utilities	Amber	N/A	N/A	No	No	No	No
Octopus Energy	UK Utilities	Amber	N/A	N/A	No	No	Yes	Yes
Octopus Energy US	UK Utilities	Amber	N/A	N/A	No	No	No	No
OVO Energy	UK Utilities	Green	Genesys Cloud	OVO Digital Assistant	Yes	Yes	Yes	No
ScottishPower	UK Utilities	Green	Zendesk	ScottishPower Chatbot	Yes	Yes	No	Yes
SSE	UK Utilities	Green	NICE CXone / Genesys	Nero	No	Yes	No	Yes
Thames Water	UK Utilities	Green	LivePerson	Virtual Assistant	Yes	Yes	No	No
BT	UK Telecoms	Green	Sprinklr	BT Digital Assistant	Yes	Yes	Yes	Yes
Sky UK	UK Telecoms	Green	LivePerson	Sky Virtual Assistant	Yes	Yes	Yes	Yes
Three UK	UK Telecoms	Green	Genesys Cloud	Three Digital Assistant	Yes	Yes	No	Yes
Virgin Media	UK Telecoms	Green	Amazon Connect	Terri BOT	Yes	Yes	No	Yes
Vodafone UK	UK Telecoms	Green	Microsoft Azure OpenAI	TOBi / SuperTOBi	Yes	Yes	Yes	Yes
HubSpot	SaaS	Green	HubSpot	HubBot	Yes	Yes	Yes	Yes
monday.com	SaaS	Green	Qualified / Ada	Amy / Tim	Yes	Yes	Yes	Yes
Pylon	SaaS	Green	Pylon	N/A	No	No	Yes	Yes

Brand	Sector	Mode 2	Chat provider (inferred)	Bot	AI disc.	Live chat	Dev API	Mode 4 cand.
Shopify	SaaS	Green	Custom / UCP+MCP	Virtual Help Centre Asst	Yes	Login req.	Yes	Yes
Slack	SaaS	Green	Salesforce Agentforce	Agentforce	Yes	Yes	Yes	Yes
Stripe	SaaS	Green	Custom (RAG)	Assistant	Yes	Login req.	Yes	Yes
Xero	SaaS	Green	Custom / Anthropic	JAX	Yes	Login req.	Yes	Yes
Zoom	SaaS	Green	Zoom Virtual Agent	ZVA	Yes	Yes	Yes	Yes
Monzo	UK Fintech	Green	Custom (proprietary)	AI Assistant	Yes	Login req.	Yes	Yes
Revolut	UK Fintech	Green	Custom (proprietary)	Rita	Yes	Login req.	Yes	Yes
Starling Bank	UK Fintech	Green	Custom (proprietary)	N/A	No	Login req.	Yes	Yes
Barclays	UK Banking	Green	Custom / proprietary	Barclays Digital Asst	Yes	Yes	Yes	Yes
Halifax	UK Banking	Green	Nuance (Lloyds Group)	Virtual Assistant	Yes	Login req.	Yes	No
HSBC UK	UK Banking	Green	LivePerson	Virtual Agent	Yes	Yes	Yes	Yes
Lloyds Bank	UK Banking	Green	Custom (authenticated)	Virtual Assistant	Yes	Login req.	No	No
NatWest	UK Banking	Green	IBM Watsonx	Cora / Cora+	Yes	Yes	Yes	Yes
Nationwide	UK Banking	Green	LivePerson	Arti	Yes	Yes	Yes	Yes
Santander UK	UK Banking	Green	Enterprise (proprietary)	Sandi	Yes	Yes	No	Yes
TSB Bank	UK Banking	Green	IBM Watson	TSB Smart Agent	Yes	Login req.	No	Yes
Admiral	UK Insurance	Green	Genesys PureCloud	A.V.A.	Yes	Yes	Yes	Yes
Aviva	UK Insurance	Green	Nuance	Aviva Chatbot	Yes	Yes	No	Yes
AXA UK	UK Insurance	Amber	Salesforce	N/A	No	Yes	No	No
Compare the Market	UK Insurance	Amber	Salesforce	AutoSergei	No	Yes	No	No
Direct Line	UK Insurance	Green	Genesys PureCloud	Direct Line VA	Yes	Yes	No	No
Vitality	UK Insurance	Amber	Stonly	N/A	No	No	No	No
Airbnb	UK Travel	Green	Custom (proprietary)	AI Assistant	Yes	Login req.	Yes	Yes
Booking.com	UK Travel	Green	Custom (proprietary)	Helpbot AI	Yes	Yes	Yes	Yes
British Airways	UK Travel	Green	Salesforce	Chatbot	Yes	Yes	Yes	Yes

Brand	Sector	Mode 2	Chat provider (inferred)	Bot	AI disc.	Live chat	Dev API	Mode 4 cand.
easyJet	UK Travel	Green	Hubtype	Virtual Advisor	Yes	Yes	No	Yes
Ryanair	UK Travel	Amber	Custom / unknown	Molli	No	Yes	No	No
Trainline	UK Travel	Amber	Zendesk	Juno / Cleo	No	No	Yes	Yes
TUI UK	UK Travel	Green	Smartagent	Theo	Yes	Yes	Yes	Yes
Uber	UK Travel	Green	Custom (proprietary)	N/A	Yes	Login req.	Yes	Yes
Deliveroo	Delivery/Logistics	Green	Zendesk / Sierra	Rider Support AI Agent	Yes	Yes	Yes	Yes
DPD UK	Delivery/Logistics	Green	Custom (proprietary)	DPD Chatbot	No	Yes	No	Yes
Evri	Delivery/Logistics	Green	SmartAgent / Amazon Connect	Ezra	Yes	Yes	No	No
Just Eat	Delivery/Logistics	Amber	Zendesk / Custom	N/A	No	Login req.	No	No
Royal Mail	Delivery/Logistics	Amber	N/A	N/A	No	No	No	No
BT Sport / TNT Sports	UK Media	Green	Sprinklr	BT Digital Assistant	Yes	Yes	Yes	No
Disney+	UK Media	Amber	Custom (Disney)	Virtual Assistant	No	Yes	No	No
Netflix	UK Media	Green	Sprinklr	N/A	No	Yes	No	No
NOW (NOW TV)	UK Media	Green	Custom (proprietary)	NOWBot	No	Yes	No	No
Spotify	UK Media	Green	Custom (proprietary)	AI-Powered Tool	Yes	Login req.	No	No
The Guardian	UK Media	Amber	Salesforce	N/A	No	Yes	No	No
Bupa UK	UK Healthcare	Green	Microsoft Dynamics 365	BV	Yes	Yes	Yes	Yes
NHS	UK Healthcare	Green	Unknown / Custom	N/A	No	Yes	Yes	Yes
DWP	UK Government	Amber	N/A	N/A	Yes	No	No	No
GOV.UK	UK Government	Green	Custom (DEP)	HMRC Digital Assistant	Yes	Yes	No	No
HMRC	UK Government	Green	HMRC DEP	HMRC Digital Assistant	Yes	Yes	Yes	Yes
OpenRent	UK Property	Amber	Custom (contact form)	N/A	No	No	No	No
Rightmove	UK Property	Amber	Freshdesk	N/A	No	No	Yes	No
Zoopla	UK Property	Amber	Zendesk	N/A	No	No	No	No
Age UK	Charity/Advocacy	Green	Custom	Virtual Assistant	Yes	Yes	No	No

Brand	Sector	Mode 2	Chat provider (inferred)	Bot	AI disc.	Live chat	Dev API	Mode 4 cand.
Council Tax Support	Charity/Advocacy	Red	N/A	N/A	No	No	No	No
Macmillan Cancer Support	Charity/Advocacy	Amber	Sprinklr	Macmillan Support Bot	Yes	Yes	No	No
Shelter UK	Charity/Advocacy	Amber	Twilio / Ciptex	N/A	No	Yes	No	No
Citizens Advice	Public/Regulatory	Green	Amazon Connect	Citizens Advice Bot	Yes	Yes	No	No
Ofgem	Public/Regulatory	Red	N/A	N/A	No	No	No	No



State of Dual CX · Edition 1.0 · Published by Neos Wave Limited (Company No. 12002488), London. Findings collected Q1 2026. Open dataset and methodology: github.com/neos-wave/state-of-dual-cx. The Service Handshake v1.1 is available under CC BY 4.0 at DOI 10.5281/zenodo.19046746. All observations are based on publicly accessible information; no brands were contacted during the audit. This report is licensed CC BY 4.0.