

The Service Handshake

An open interaction standard for service interactions where one or more parties are represented by AI agents

Maria McCann, Founder, Neos Wave

SH-1.2 | Published 1 October 2026

Version 1.1: 10.5281/zenodo.19046746.

Version 1.2: 10.5281/zenodo.23181969.

Licence: Creative Commons Attribution 4.0 International (CC BY 4.0). Schemas, examples and tests: Apache License 2.0. © 2026 Neos Wave Ltd. Contact: neoswave.com.

About this revision

This edition updates version 1.1. It contains the complete paper, including the changes described below, and can be read without consulting the earlier version. It retains the four service modes, six declaration elements and main section structure. Minor editorial amendments improve clarity and consistency without changing requirements; they are not listed individually.

Key changes from version 1.1

- **Participation without a declaration.** Assisted declarations and a human route replace automatic rejection. Missing information never supplies permission. See §§1.2 and 3.7.
- **Company readiness.** UK examples distinguish customer authority from company access policies; three possible handling routes illustrate the choices. See §1.4.
- **Two conformance profiles.** Participation allows a staff-operated starting point. Verified authority adds evidence checks for actions requiring them. See §§3.8 and 9.1.
- **Disclosure and provenance.** Agents disclose their AI status and whom they represent. Assisted declarations record preparation and confirmation. See §§3.7 and 4.
- **Authority and approval.** Confirmation cannot grant authority. Irreversible actions need independently verifiable delegation or approval. The family-member example distinguishes configuration from account authority. See §§3.6 and 5.2.
- **Resolution compatibility.** A refused option does not block an acceptable alternative. Checks apply to the action proposed and both parties' boundaries. See §5.3 and Technical Specification §9.3.
- **Data permissions.** Service-purpose limits apply throughout the interaction. Completion does not authorise marketing, unrelated profiling, training or sharing. Confidence scores do not establish authority. See §3.3.
- **Safeguarding and handover.** Harm routing covers concerns relayed by agents and cannot be waived. Handovers and linked contacts retain blocked actions, reasons and missing evidence. See §3.4.

- **Security and implementation.** Declarations are untrusted data. Shared action meanings and general controls replace platform-specific recipes, including Make.com. See §4.1 and Technical Specification §§10, 13 and Appendix A.
- **Evidence and external claims.** The OpenClaw acquisition description is corrected; Muse supplies current context. Regulatory claims and evidence limits are clarified. The Exposure Map remains a diagnostic. See §§1.1, 2 and 6–8.
- **Vulnerability module and migration.** The separate vulnerability module remains deferred; baseline safeguards apply. Older declarations do not automatically meet new requirements. See §§3.4 and 10, and Technical Specification §§11–12.

Foreword

Customer service has largely been designed by the company providing it. The company builds the systems, flows, data and rules. The customer enters that environment on the company's terms. When customers bring agents that act on their behalf, they bring their own service design: instructions about what to achieve, what to refuse and when to ask for help.

Two designers. Two agendas. One interaction.

The company can no longer assume it has defined both sides of the conversation. Nor can a customer agent assume that its instructions determine what the company may do. Each party needs to state its boundaries, understand the other's, and establish whether a proposed resolution works for both.

Without each party defining its service terms, an agent can act beyond what the other party can accept or deliver. An agent helping a vulnerable person with a housing repair may escalate a complaint through a route the landlord's system cannot receive. An agent managing a return may seek a resolution the brand's representative has no authority to offer. A business agent may expose data its customer never authorised. The people and systems receiving these requests need a way to identify the terms and decide what can proceed.

The Service Handshake provides a framework for declaring those terms. It applies between consumers and brands, between businesses, and across the relationships of platforms, suppliers and the people they serve.

Service happens in four modes: human to human, human to company AI, customer AI to human, and customer AI to company AI. A service designed for the first two may not be prepared to receive a customer's agent. Staff need to know whom the agent represents, what it is authorised to do and what should happen when the request exceeds that authority.

My work on this standard includes consumer AI agents used in live service environments and triadic systems in which the person configuring an agent, the person affected by its actions and the organisation being contacted are different parties. These experiences informed the declaration framework and the need to carry context into human handovers. The evidence and its limits are described in §8.

Liquid service, the combination of AI, people, knowledge and process, requires each party to define the terms on which it participates. A declaration makes those terms available for

assessment. The declaration works only when staff procedures and technical controls enforce those terms.

Humans are not an escape hatch. A handover is part of the service design, and the person taking over needs the context, authority information and restrictions required to handle the matter. Moving a request to a person does not make missing permission appear. A customer should also be able to ask for help without installing software or preparing a permanent declaration. Version 1.2 allows the party using the standard to help record the other party's terms for that interaction.

We are making the standard open so that service designers, organisations and developers can use it, test it and contribute improvements.

Maria McCann, Co-Founder, Neos Wave

1 Why this standard exists

1.1 Evidence of the need

Personal agents entering company service. Meta's September 2026 launch of Muse makes the design problem immediate. Meta describes a personal agent that can browse, complete forms and negotiate for its user, with a separate Sentinel system controlling outbound actions and requesting permission when needed. Those controls govern what happens within the agent's operating environment. A company receiving the request still needs to establish what the person authorised and whether the evidence covers the proposed service action. Internal approval does not, by itself, provide evidence the receiving company can validate. [9, 10]

OpenClaw illustrates the same wider interest in personal agents. Its creator, Peter Steinberger, announced in February 2026 that he was joining OpenAI and that OpenClaw would move to an independent foundation. [2] Whether an agent comes from a large platform or an open project, the service question is the same: whose instructions does it carry, what may it agree to, and how can the other party check?

The Service Handshake gives the receiving service a way to ask those questions without requiring the customer to arrive with a completed declaration.

Air Canada and chatbot information. Version 1.1 discussed *Moffatt v. Air Canada*, a 2024 British Columbia Civil Resolution Tribunal decision concerning misleading bereavement-fare information supplied by the airline's chatbot. The tribunal found negligent misrepresentation and awarded damages. The case illustrates why the information and commitments supplied by a company agent need clear ownership and controls. [3]

For service design, the relevant questions are how the company controls the information its agent supplies, what commitments it permits and how an uncertain request reaches an authorised person. A declaration can record those limits, but the implementation **MUST** apply them.

Shopping agents and consumer authority. Consumer-agent development has also produced mechanisms for recording purchasing authority. Mastercard's March 2026 announcement of Verifiable Intent describes a tamper-resistant record of what a user authorised, intended to

provide evidence that participants can verify. This is relevant evidence infrastructure, not evidence that every service action is covered by a purchasing mandate. [4]

Suitable evidence from these systems can support a service decision when it covers the proposed action. A refund, account change, complaint settlement or cancellation may require different permissions and consequences to be considered. The scope of the proposed action determines what evidence is needed.

The retail observations in §5.1 show the service problem in practice. Section 8 sets out the evidence and the work needed to evaluate the revised standard.

1.2 The four modes of modern service

The four modes identify who is participating. They do not by themselves establish an organisation's readiness or an agent's authority.

Mode	Participants	Application of the standard
1	Human to human	Direct use is optional. Restrictions inherited from an earlier part of the interaction still apply after handover.
2	Human to company AI	Company AI discloses its role and uses its declaration. The person can state requirements in ordinary language.
3	Customer AI to human	Customer AI discloses its role. Staff receive its declaration or help prepare one, retaining the stated limits.
4	Customer AI to company AI	Agents exchange supported declarations or establish assisted participation. Each proposed action is checked before execution.

Missing declarations **MUST NOT** automatically exclude a person from service. Where the counterparty lacks a usable declaration, the conforming service **MUST** offer proportionate assistance to establish the terms needed for the proposed action. An accessible human route **MUST** remain available. Declining assistance **MUST NOT** itself exclude the person from service. Public, non-personal information **MAY** be provided while terms are being established. Access to protected information and consequential actions **MUST** remain blocked where relevant permissions or authority are unresolved. Neither silence nor failure to provide a declaration is consent.

1.3 The design gap

An agent may be able to identify itself, exchange messages and initiate a transaction without having established the terms of a service interaction. The receiving party still needs to know what outcome is sought, which outcomes are unacceptable, what data may be used and what should happen if no permitted resolution is available.

The Service Handshake records those terms and connects them to a decision about a proposed action. It can work with identity, transport and payment systems. It does not assume that those systems lack relevant intent or approval information; it requires the service to establish whether their evidence covers the matter at hand.

1.3a Application contexts

The four modes apply across three commercial contexts. The six declaration elements remain the same. What changes is the relationship between the parties, how terms are established and whose authority is required.

Business to consumer. A consumer or their AI agent may arrive with no prior declaration, or within an existing account relationship. Terms can be stated at contact, referenced from an earlier agreement or recorded with assistance. Receipt of the other party's declaration does not grant permission. A company declaration does not override the consumer's stated boundaries, and the consumer need not adopt the standard for future interactions.

Business to business. There is often a contract, an onboarding process and an existing commercial relationship. The bilateral declaration can be pre-agreed as part of that relationship and used in subsequent agent interactions. Version 1.1 described this as a treaty: an agreed record of the terms within which each party's agent operates.

The parties **MUST** still identify the applicable version and scope. An agent may represent an organisation and be permitted to make commitments within an agreed mandate, but a declaration does not itself make every statement legally enforceable or expand the underlying delegation. Escalation ownership and contractual approval requirements need to be recorded alongside the permitted actions.

Business to business to consumer. This context can involve three sets of terms: an agreement between businesses, the consumer's requirements, and rules for resolving a conflict between them. The resolution hierarchy determines who handles a conflict and how it is escalated. It cannot grant authority over the consumer's affairs or waive applicable safeguards.

Consider a housing association using an agent to manage tenant service contacts on behalf of a local authority. The organisations have agreed what the agent can offer and what **MUST** be referred. A tenant contacts the service about damp and mould. Their requested resolution may fall outside the agent's mandate. The service needs to identify what it can do immediately, what needs approval, who owns the next step and what the tenant will be told. Harm indicators may require a specialist route even if the commercial agreement points elsewhere.

This is why the declaration records both boundaries and fallback rules. A conflict **MUST** have an accountable outcome, and each affected party's authority **MUST** be considered. The invitation and verification requirements in §§3.6–3.8 apply in all three contexts.

1.4 Company readiness for customer agents

When a customer sends an AI agent, the receiving company needs to know whether its terms, access policies and service procedures can accommodate that contact. Permission from the customer and permission to use a company's channel are separate questions. A customer may authorise an agent to seek a lower bill while the company's procedures recognise only an account holder or a registered human representative. Staff then have to decide what can proceed and what needs further evidence.

These questions are already affecting access to services. In September 2026, Amazon blocked Meta's Muse agent from shopping on Amazon.com on customers' behalf. Amazon said it had not

agreed to the access and that third-party agents “should operate openly and respect service provider decisions about whether or not to participate.” This illustrates a distinction central to the Service Handshake: a customer’s instruction to an agent does not, by itself, establish an agreed basis for that agent to interact with the company. [17]

Restrictions can also appear in existing terms. Booking.com’s UK-English customer terms (§A15.2, checked 29 September 2026) restrict automated access and bookings without its prior express written permission. These examples show actual access barriers and stated policies; they do not establish that every restriction is legally enforceable or that companies will follow the same approach. [14]

Existing arrangements for human representatives offer a starting point. Ofcom’s guidance distinguishes third-party bill management, which allows a nominee to discuss and pay bills, from authority to close or change an account. Those arrangements do not automatically extend to software agents. They show why a general statement that someone is authorised to act is insufficient without a defined scope. [13]

Authorisation can provide evidence of whom an agent represents and what it may do. It does not absolve the receiving company of liability. Consumer rights and the company’s own data-protection responsibilities remain applicable. A Handshake declaration cannot itself create delegation, supply a lawful basis for disclosure or override a company’s access conditions. The legal context is discussed in §7. The UK examples and the Amazon.com case illustrate the design questions without limiting the standard to one jurisdiction. [15, 16]

Consider a customer who sends an agent to lower their broadband bill, with instructions not to cancel the service. Where the provider accommodates agent contact, the following are plausible handling routes, not measured outcomes or predictions.

Prior registration. The company asks the customer to register an approved representative or mandate before protected account discussions. This uses a familiar control, but may require setup and repeat contact. Registration still needs to distinguish discussion, access and commitment. It need not prevent public information or an accessible human route.

Assisted participation. The company offers a Handshake invitation and records the customer’s goal, the no-cancellation restriction and any unresolved authority. It can provide public information while arranging the necessary account checks. Before accepting a new contract, it establishes authority for the actual price, duration, charges and other material terms. The invitation helps the interaction progress; it does not replace those checks.

Reusable verified delegation. The agent presents an existing mandate that the company can independently validate. The mandate identifies permitted actions, limits and duration. Reuse could reduce repeated approval requests, but depends on a trusted evidence source, compatible meanings and adequate expiry and revocation handling. Permission to negotiate does not automatically authorise a new minimum term or cancellation. An out-of-scope offer still needs further approval.

These routes can coexist. An assisted invitation may lead to an established account procedure or recognise a valid mandate already held. The company needs to align its published terms, actual access controls and staff guidance so that a route offered at first contact can work through to resolution. Legal review should address the relevant service and jurisdiction, including applicable rights and duties when a preferred channel cannot be used.

The Service Handshake can help establish arrangements where parties are willing to participate. It cannot guarantee that a company will accept an external agent. Where it is adopted, its intended contribution is to reduce unnecessary repetition while preserving necessary checks. Whether it achieves that needs testing through completion time, repeat contacts, unnecessary refusals and unauthorised actions. The practical question is how much can safely proceed now, what evidence is needed next and who is responsible for obtaining it.

2 Where this standard sits

The standard is independent of any particular protocol or technology provider. Declarations can travel over the transport used by the parties and can reference external identity and delegation evidence.

Layer	Function	Relationship to the Service Handshake
Security and governance	Access controls, monitoring and enforcement.	Enforces applicable policies and restrictions around the operation.
Identity and verification	Establishes who is presenting a request.	Provides trusted identity results where required; identity alone does not establish action authority.
Transport and communication	Carries messages and records between systems.	Carries declarations without deciding their truth or sufficiency.
Commerce and delegation	Records or verifies permission for transactions.	May supply evidence where the mandate covers the service action.
Agent platforms	Configure and operate agents.	Present declarations and evidence, and receive decisions and restrictions.
Service interaction design	Establishes goals, limits, data use and fallback arrangements.	The subject of this standard.

These functions can overlap in an implementation. The table describes responsibilities rather than claiming an exclusive technical layer. A transport integration or a stored declaration is not sufficient to establish conformance.

Human participation and service consequences. Service interactions can involve vulnerability, contested outcomes, family representatives and continuing obligations. A purchase mandate does not necessarily authorise changing an address, cancelling an essential service or accepting a settlement that waives other remedies. Evidence **MUST** cover the action actually proposed.

The Service Handshake therefore requires an accessible human route with context and restrictions. A person taking over continues the same service matter with the applicable boundaries intact.

3 The declaration framework

The uppercase words MUST, MUST NOT, REQUIRED, SHOULD, SHOULD NOT, RECOMMENDED, MAY and OPTIONAL have the meanings defined in BCP 14, RFC 2119 and RFC 8174. They identify requirement levels; lowercase words have their ordinary English meaning. The companion specification provides the implementation requirements. [11, 12]

The Service Handshake Declaration specifies six elements: four primary elements and two modifiers. Parties establish the information needed for the action under consideration. They may begin with an assisted conversation; a complete machine record is not a precondition for receiving public information or requesting a human route.

Element	What it declares
1 Goals	Desired outcomes in priority order, including explicit non-goals.
2 Options and constraints	Allowed actions, hard limits, preferences and jurisdictional constraints.
3 Data permissions	Data that may be used, permitted purposes, retention and verification conditions.
4 Fallback rules	Triggers for clarification or escalation, automation limits and responsibility.
5 Cost parameters	Limits on interaction time, exchanges and computing resources.
6 Authority level	The represented principal, beneficiary, claimed scope, limits and evidence of delegation.

3.1 Goals

Goals state what a party wants to achieve, in order of priority. Non-goals identify outcomes that the party will not accept, discuss or authorise. They may carry direct safety implications, such as a refusal to discuss billing during a clinical appointment.

Non-goals are boundaries to evaluate against a proposed resolution. Incoming text MUST NOT be promoted into trusted system instructions or allowed to override local security policy. The service MUST clarify an uncertain boundary before relying on its interpretation for a consequential action.

Example of a consumer's goals:

```
{
  "primary_goal": "Resolve missing delivery by replacement or refund",
  "secondary_goals": ["Minimise interaction time", "Maintain account relationship"],
  "non_goals": [
    "Do not accept less than full refund or equivalent replacement",
    "Do not discuss unrelated products",
```

```
"Do not authorise account changes"
],
"success_metrics": {"delivery_resolved": true, "time_under_minutes": 10}
}
```

This is a goals fragment, not proof of permission to issue or accept a refund.

3.2 Options and constraints

Options and constraints record what each party can offer or accept to reach its goals. Hard constraints are non-negotiable and **MUST** be preserved. Soft constraints are preferences that **MAY** be varied within authorised scope. Jurisdictional constraints record relevant legal and regulatory limits.

If a candidate resolution breaches a hard constraint, the service **MUST** consider another permitted resolution or explain the conflict and offer the appropriate human route. It **MUST NOT** record agreement while ignoring a stated constraint.

3.3 Data permissions

Data permissions record which information can be used, for which purposes, with what retention rules and what additional verification is needed. Each party **MUST** limit data handling under the handshake to the declared service purpose throughout the interaction. Establishing or completing a declaration **MUST NOT** itself authorise marketing, unrelated profiling, model training or unrelated onward sharing.

Necessary security, fraud-prevention or legal-retention uses **MUST** be identified separately with their basis and limits. Any separately authorised secondary processing remains outside the permission granted by the handshake and requires its own applicable basis. A declaration cannot waive another party's legal obligations or establish consent that has not been given.

A conforming service **MUST** limit its requests to information needed for the service matter and relevant authority. An initial case reference is not necessarily sufficient for protected account access.

The legacy trust_levels fields may record operational confidence signals. Confidence in an answer or a counterparty is not evidence of delegation, data permission or approval. The technical specification separates those signals from the checks needed before an action.

3.4 Fallback rules

Fallback rules define what happens when an interaction cannot proceed within its declared parameters. Implementations need routes for uncertainty, conflicts and harm indicators, as well as resource limits.

- Confidence or uncertainty conditions determine when autonomous handling **MUST** pause.
- Conflict conditions determine what happens when goals or constraints cannot be reconciled.
- Defined harm indicators trigger the applicable safeguarding or specialist route.

Vulnerability should be considered in relation to the situation as well as recognised customer groups or ongoing support needs. Bereavement, financial pressure, an unfamiliar process or difficulty understanding a proposed commitment can change what support a person needs in a particular interaction. Membership of a group does not establish incapacity, and a person need not belong to an identified group to need support. In Mode 2, the company agent can respond to expressed needs by adapting the pace, explanation or channel and offering human help while preserving the person's choices. Support needs may emerge throughout the interaction; they do not automatically require escalation or restriction. Defined harm concerns still trigger the applicable safeguarding response.

A conforming service **MUST** define observable harm indicators, a proportionate response and an accountable human route for its service context. The baseline in Technical Specification §6.3 provides a starting point. Services **SHOULD** obtain specialist input where the context or level of risk calls for it. Indicators **MUST NOT** depend on unrestricted agent inference or speculative diagnosis. A supporting organisation **MUST** preserve its applicable safeguarding routes even if the counterparty asks to suppress them. These safeguards still apply when a person is represented by an agent. A conforming consumer agent **SHOULD** relay these concerns when its principal expresses them. A company agent **MUST** treat a relayed concern as raised and apply the relevant safeguarding response. The record **MUST** distinguish a relayed report from an independently established fact, and disclose only the information needed for the response. A report does not itself establish a diagnosis or justify an indefinite restriction.

The `maximum_automation_scope` field records what an agent may initiate but not complete without approval. For example, an agent may discuss a provider switch while lacking authority to complete it. Approval **MUST** come from an authorised person or process; the presence of a human is not sufficient on its own.

A human handover **MUST** preserve the proposed action, every unresolved restriction, the reason for each block, the evidence or decision needed next, relevant data-use limits and the accountable role. A staff member taking over **MUST NOT** remove a restriction without the required evidence or authorised decision. Opening another contact or changing channel **MUST NOT** reset an unresolved restriction on the same case or action. Restrictions need a stated scope, review point and route for resolution; they **MUST NOT** become unexplained, indefinite blocks on an account. Sensitive safeguarding information **MUST** be shared only where needed, without speculative diagnoses or unexplained denial of ordinary service.

Version 1.1 announced a dedicated vulnerability declaration module for 1.2. This version strengthens baseline routing and handover requirements but does not deliver that separate sector-specific module. The module remains deferred; the baseline requirements apply to this version.

3.5 Cost parameters

Cost parameters set limits on the time, exchanges, tokens or effort an interaction may consume. Approaching a limit **SHOULD** trigger a useful next step, such as a summary and callback route. These limits help prevent repeated negotiations without progress.

Interaction budgets do not grant permission to incur financial commitments. Financial limits belong with the relevant authority and need the currency, scope and period applicable to the proposed commitment.

3.6 Authority level

Authority level records whom the agent represents, who is affected, what it claims it can commit to and what evidence supports that claim. The principal is the person or organisation represented for the relevant action. The beneficiary receives or is affected by the outcome. The configurator, called the architect in version 1.1, may be a different person. Configuring an agent **MUST NOT** be presumed to grant authority over another person's account.

For example, when a family member configures an agent to act for a parent on the parent's utility account, the parent is the represented principal for that account action. The family member's ability to authorise a change depends on the actual delegation. A power-of-attorney label in a declaration is a claim requiring verification, not proof by itself.

Confirmation by a principal or counterparty agent establishes that the declaration reflects that confirmer's claims. It **MUST NOT** expand authority beyond the supporting evidence. Irreversible actions **MUST** have delegation evidence or principal approval that the relying party can validate without relying solely on the requesting agent's assertion.

Approval may occur within the same application if it produces independently checkable evidence from a trusted approval source. The evidence **MUST** link the approval to the principal, requesting agent or authenticated session, specific action and material terms, and state its validity conditions. The agent saying that its customer approved is insufficient on its own.

3.7 Participation when adoption differs

Make participation easy, keep each party's boundaries intact, and verify authority at the point where an action requires it.

When both parties hold suitable declarations, they exchange or reference them. When only one party uses the standard, it offers to help record the other party's terms for that interaction. This is an assisted declaration. When neither party uses the standard, ordinary service continues under their existing procedures without a claim of SH conformance.

A conforming AI agent **MUST** disclose that it is AI and whom it represents at first contact. If this is unclear, the receiving party **MUST** seek clarification before relying on an assisted declaration. Disclosure is not verification of identity or authority. Uncertain AI status is a reason to clarify relevant representation, not a reason to invent an identity or deny ordinary service. Public information and normal account procedures remain available; unresolved permission blocks only the affected protected action.

The supporting party **MUST** offer to record the desired outcome, unacceptable outcomes, relevant data permissions, claimed authority and fallback preferences to the extent needed for the proposed action. Staff can do this during the ordinary service conversation; a separate form or software installation is not required. It presents a summary for correction and confirmation. It **MUST** preserve unknowns and refusals and **MUST NOT** fill gaps with permissive defaults. Missing information **MUST NOT** be treated as permission. Assisted declarations **MUST** record `declaration_provenance` as assisted.

Suggested invitation: “I can record your boundaries for this conversation. Please tell me whether you are an AI acting for someone, whom you represent, the outcome you want and anything you must not agree to. We will check approval separately if a proposed action needs it.”

The invitation applies in either direction. A consumer agent can ask a company to state its terms, but cannot certify the company's policy or authority on its behalf. A consumer can decline an invitation and use the human route. Ordinary access and authority checks remain in force.

3.8 Requirements profiles and proportionate checks

Participation sets the minimum requirements: AI disclosure, an invitation where needed, no inferred consent, preservation of boundaries, prevention of unauthorised actions, an accessible human route, applicable safeguarding and service-purpose data use. A small contact centre can implement this through conversation and a structured case record. The consumer does not need to write JSON, and participation does not initially require cryptographic infrastructure.

Verified authority includes Participation and adds the evidence checks needed for protected or consequential actions. The service checks the trusted issuer or approval source, principal, scope, audience, validity and status; relates the evidence to the proposed action; and records its decision. Each party sets its own financial limits. There is no universal monetary threshold in the standard.

Action category	Minimum treatment
Public information	May proceed before a declaration is complete if no personal data or commitment is involved.
Protected information	Establish relevant identity, access authority and data permission before disclosure.
Service action	Establish compatible terms and authority for the exact action before execution.

Additional requirements depend on reversibility, financial exposure, recurring or cumulative commitments, account security, essential-service consequences and possible harm. Reversibility does not grant authority. Irreversible actions need independently verifiable delegation or principal approval. A Participation implementation that cannot perform the required checks **MUST** hold the action or route it to a process that can. Existing staff or account procedures can provide the required approval when they establish the approver's authority, show the actual terms and produce a record the service can validate. A login code alone establishes neither delegation nor approval of a particular action. Technical Specification §8.3 gives a refund example.

4 Machine readable schema

The declaration is designed for service design, legal and operations teams as well as agents and software systems. JSON is the defined machine-exchange format. A staff-operated Participation workflow may record equivalent information in a form or case note.

The top-level fields retained from version 1.1 include `protocol_version`, `principal_id`, `agent_id`, `beneficiary_id`, `identity_tokens`, `issuer` and the `valid_from` and `valid_to` timestamps. The six declaration elements remain nested objects. SH-1.2 adds declaration identifiers and revisions, disclosure, provenance, confirmation and unresolved fields.

The companion Technical Specification provides field references and worked JSON examples. The schema files are `sh-1.2.schema.json` and `sh-1.2.action-assessment.schema.json`, published with this edition. Structural validation does not establish the truth of a claim, the authenticity of evidence or permission to execute an action.

Identity tokens reference external verification mechanisms. The service **MUST** obtain trusted verification results when needed; an incoming status of `valid` is still a claim. An action assessment is stored separately from the declaration and records the relying party's decision and reasons.

Storage and implementation. Declarations can be stored in a CRM, ticketing system, orchestration service or external configuration store. Staff **MUST** be able to see relevant restrictions where they make decisions. Protected operations **MUST** check permissions at execution; displaying the declaration or placing it in an agent's prompt does not enforce it.

The standard does not prescribe a technology stack. Organisations **SHOULD** version their templates, assign review responsibilities and retain appropriate change records. The technical specification describes workflow controls, security and migration from 1.1. Appendix A defines a small common service vocabulary so unfamiliar systems can identify complementary actions without first agreeing every local name.

Companies adopting the Model Context Protocol (MCP) to let consumer AI agents use their services can apply the Service Handshake within that approach. Existing checks and approvals can be used where they are sufficient. The Handshake sets out what the interaction needs to preserve: each party's boundaries, agreed uses of information, safeguards and a route to human help. Using MCP does not, by itself, establish that these requirements have been met.

4.1 Security considerations

A counterparty declaration is untrusted data across all six elements. Implementations **MUST** validate its structure and interpret its terms through trusted local controls. Text supplied as a goal, non-goal or evidence reference **MUST NOT** change system instructions, tool access or security policy. An incoming claim that permission has been verified is still a claim.

The service **MUST** enforce the resulting decision at the protected operation and carry unresolved restrictions across linked contacts. Internal fraud rules and approval thresholds **MAY** remain private, provided the actual proposed terms and required next steps are communicated. Keeping an internal limit private **MUST NOT** conceal material conditions of an offer or override an applicable obligation. Technical Specification §§10 and 13 define the implementation controls.

5 Worked examples

5.1 Mode 3 retail delivery resolution

In the three live retail calls described in version 1.1, a consumer AI agent contacted service staff about a missing delivery and completed the interaction without staff identifying it as AI. The calls informed the invitation and disclosure requirements in this revision. Section 8 explains their scope and limitations. [1]

The following declarations illustrate how that type of interaction should be handled under the revised standard. The AI agent discloses its role at first contact. Staff use an existing declaration or offer to record one.

Element	Consumer terms	Company terms
Goals	Resolve the missing delivery by full refund or equivalent replacement.	Resolve the delivery issue within policy.
Options and constraints	Replacement within 48 hours; no voucher settlement.	Refund to the original payment method, replacement or referral to an authorised role.
Data permissions	Order information for this case; no marketing.	Relevant order and account information for resolution, subject to access checks.
Fallback rules	Request human review if unresolved after three exchanges.	Route uncertainty or harm indicators to the responsible team.
Cost parameters	Maximum ten minutes before offering another route.	Summarise and arrange follow-up when the time limit approaches.
Authority level	Claimed permission to accept a refund or replacement; no account changes.	Authority to offer the specified remedy within company limits.

If the agent does not use SH, staff can prepare an assisted declaration and ask it to confirm the summary. That confirmation does not prove delegation. The retailer can explain its public returns policy while it checks the authority needed for protected order information or a refund.

A financial limit, such as GBP 100 for a refund, is an example of company policy rather than a standard-wide threshold. It MUST be expressed with the relevant currency and scope. Both parties' authority and the conditions attached to the remedy MUST be satisfied.

5.2 Mode 3 family-configured agent

A family member configures an assistant to help an elderly parent with utility contracts. The agent contacts a provider to discuss a broadband renewal. The configurator, the parent and the provider have distinct roles and interests.

The declaration identifies the represented account holder, the family representative and the claimed delegation. It records permitted billing queries, any limits on discussing a renewal, and restrictions on cancellation or switching. It does not treat the configurator as the person entitled to approve every action.

In this example the declaration permits discussion only and prohibits cancellation or switching. A verified delegation does not override those hard constraints. If the account holder later wants to consider a switch, an authorised person **MUST** first revise the applicable declaration and scope. The service then assesses a new proposal and obtains any required approval. Until that happens, the current agent cannot complete the switch. A missing delegation remains unresolved after transfer to staff.

The provider's safeguarding route can still trigger. An assisted declaration or approval from the other party cannot disable defined harm indicators. The handover records the concern and necessary response with the minimum sensitive detail required.

5.3 Mode 4 brand AI meets consumer AI

A company agent receives a consumer agent's request. Both parties provide declarations, either directly or with assistance. The service assesses a proposed resolution against the applicable goals, constraints, authority and data permissions.

Suppose the company offers refunds, replacements and vouchers. The consumer refuses vouchers but accepts a refund to the original payment method. The unused voucher option does not make the interaction incompatible. The refund may proceed if its terms and the required authority checks are satisfied. Technical Specification §9.3 defines this decision process.

Problem to address	Required handling
Repeated exchanges without progress	Apply declared resource limits and provide a next step.
Inconsistent escalation	Use defined triggers and an accountable role.
Unclear decision history	Record the proposed action, declarations, checks and reasons.
A proposed action exceeds scope	Hold or deny that action and identify what is needed next.
Human staff take over	Transfer restrictions and reasons; obtain any missing authority before proceeding.

For example, a cancellation beyond the customer agent's authority remains on hold when a staff member takes over. If the request changes or relevant evidence expires, the service reassesses it. An escalation changes who handles the matter; it does not itself grant permission.

5.4 Mode 2 appointment enquiry

A person contacts a clinic's AI receptionist to arrange a consultation. The agent identifies itself as AI and explains that it can help with appointments, while treatment advice comes from the clinical team. It offers an available time and explains the fee and relevant conditions. The person does not need to prepare a declaration; their preferences and the service's boundaries are established through the conversation. The agent confirms the booking only when the

person has agreed, the action falls within its authority and the booking system confirms completion. Otherwise, it makes clear that the request remains pending.

The person then asks which treatment would suit them. The agent explains the limit of its role and offers a human route, carrying forward the appointment details and the person's question. If the person says they feel overwhelmed or do not understand the terms, it slows down, explains clearly and offers help without assuming incapacity. This illustrative example shows how the Handshake can support an ordinary service conversation while keeping authority, agreement and responsibility clear.

6 The DualCX Exposure Map

The Exposure Map is a diagnostic framework for assessing readiness across the four modes. It maps those modes against goals alignment, data handling and fallback design. It identifies areas to investigate rather than assigning a readiness score merely from the type of interaction.

Mode	Goals alignment	Data protocol	Fallback design
1 Human to human	Do scripts and procedures preserve the person's requirements?	Are access, purposes and retention controlled?	Are escalation ownership and inherited restrictions visible?
2 Human to company AI	Does the agent operate within declared company limits?	Does the service enforce permissions throughout the interaction?	Can staff receive the context and blocked actions?
3 Customer AI to human	Can staff receive or help establish the agent's declaration?	Can staff distinguish disclosure, identity and authority?	Can unresolved requests reach the right authorised person?
4 Customer AI to company AI	Are proposed resolutions checked against both parties' limits?	Are data permissions and authority enforced at execution?	Are holds, conflicts and safeguarding passed on with reasons?

Each cell should be supported by observed procedures, configuration or test results. A declaration field, a policy document or a green dashboard indicator does not by itself demonstrate that the control works. The diagnostic can be used independently of any commercial implementation; a service seeking SH conformance MUST meet the relevant requirements.

Tests should include successful resolutions, justified holds and unnecessary blocks. Check whether another channel can bypass approval and whether a human handover preserves the restrictions. The historical benchmark reported in 1.1 is discussed below rather than assumed to describe every organisation.

7 Regulatory context

The Service Handshake records declared purposes, authority claims and decisions for review. Organisations still need to establish and meet their legal and sector-specific duties. Section 1.4

illustrates how customer mandates, company access conditions and approval procedures can differ. An authorisation record does not exempt the company from its own obligations. [15, 16] Adopting the standard does not itself establish compliance, and an incoming declaration cannot waive a person's rights.

Article 50(1) of the EU AI Act places the design obligation for informing people about direct AI interaction on providers, subject to its conditions and exceptions. SH's first-contact disclosure rule is an operational requirement of this standard, including agent-to-agent participation; it does not transfer the provider's legal duty to receiving staff. Applicable legal duties **MUST** be assessed for the actual system and use. [5]

Under the GDPR, purpose limitation, data minimisation, storage limitation and an appropriate legal basis remain separate obligations. Recording a permission in a handshake does not by itself establish valid consent or another legal basis. [6]

Regulatory area	Relevance to service design
EU AI Act	Identify the provider and deployer roles and applicable transparency duties. Do not assume that all customer-service AI is high risk or that a declaration fulfils every duty. [5]
GDPR and applicable UK data protection law	Establish the basis, purpose, access and retention arrangements for actual processing. A declaration records relevant boundaries but does not replace that assessment. [6, 7]
UK Data Use and Access Act and PECR	Keep service-purpose processing separate from electronic marketing permissions and other applicable requirements. [7]
UK consumer protection	Ensure service representations and resolutions comply with applicable consumer obligations. Authority limits cannot remove obligations owed by the company. [8]

Evaluations of this revision **SHOULD** document disclosure, permissions, method and safeguards alongside the service outcome.

8 Evidence base and maturity

The standard draws on my service-design work and the observations presented in version 1.1. The three retail calls involved services without brand-side declarations. They demonstrated a gap in handling customer agents, but did not test bilateral declarations or establish a general success rate. The worked examples here explain how the proposed controls address that gap; they are not evidence that those controls have been validated in live operation.

Mode	Evidence described in version 1.1	Limit on interpretation
1 Human to human	Established service procedures and operations.	Existing practice does not demonstrate SH conformance.

2 Human to company AI	Existing contact-centre AI deployments.	Deployment alone does not establish effective declaration or authority controls.
3 Customer AI to human	Three live retail calls described in version 1.1.	Small observational sample; bilateral SH-1.2 controls were not tested.
4 Customer AI to company AI	Schema implementation and testing described as in development.	Structural validation is not evidence of safe execution or successful live bilateral operation.

In version 1.1, I reported Exposure Map assessments across 85 organisations, with 100% scored green for Mode 1, 52% for Mode 2 and 0% for Modes 3 and 4. These are the original observations, not current market estimates. The sample selection, scoring rubric and underlying dataset are not included in this paper, so the figures cannot be independently reproduced from it. They do not establish conformance with SH-1.2. [1]

The worked examples in this revision illustrate the proposed rules. Evaluation should test whether people understand disclosure, assisted declarations accurately preserve boundaries, compatible remedies remain available, authority is verified and restrictions survive handover. Reports should state the sample, setting, permissions, method, failures and limitations.

9 Adopting this standard

9.1 A practical starting point

The standard is open to service designers, CX and operations leaders, technology teams and organisations handling contacts from customer agents. A useful starting point is one service matter with an accountable role.

1. Choose a use case, such as a missing delivery, billing query or appointment enquiry.
2. Record the company's goals, permitted actions, data purposes, authority limits and fallback rules.
3. Prepare an invitation and a case form for customers or agents without a declaration.
4. Train staff to distinguish stated claims, confirmation and verified authority. Show blocked actions and reasons where decisions are made.
5. Connect protected actions to the account and approval procedures that control them. Hold or refer actions whose checks cannot be completed.
6. Test disclosed agent contacts, refusals, missing delegation, safeguarding and human handovers as well as ordinary successful cases.
7. Measure repeat contacts, resolution time, escalation quality, staff workload, unresolved authority, unnecessary refusals, actions allowed without sufficient permission and relayed-concern rates.

For agent contacts, record the proportion of interactions containing a relayed concern, the proportion of those routed to a human or specialist, and the response given. Review changes over time and, where proportionate, by service matter and agent provider. Investigate whether patterns reflect customer need, inaccurate relaying or use of safeguarding as a shortcut to

human help. Rates alone do not establish misuse. Use only the data needed for this service-quality review; it MUST NOT become customer profiling. Monitoring MUST NOT suppress or delay the required response to an individual concern. Keep ordinary human-help routes available so safeguarding is not the only way to reach a person.

A small contact centre can begin with the Participation profile using a form, invitation script, action checklist and handover template. Machine exchange requires the defined JSON representation. Verified authority requires the additional evidence and execution checks for the action types claimed. A public conformance claim MUST identify the version, profile, use cases and actions covered. A badge MAY link to that scope statement; it MUST NOT imply assurance beyond it.

Organisations building an MCP service can begin with one supported action, map its inputs and outcomes to the declaration framework, and connect the required checks to the operation that performs it. Existing information and verified permissions should be reused where sufficient, with assistance or further approval sought only for unmet requirements. Evaluation should test whether boundaries survive changed proposals, retries and human handovers, as well as whether an ordinary request succeeds.

Changes to local templates and action vocabularies should be versioned and reviewed. Proposed changes to the common schema should be submitted for review rather than treated as local conformance automatically.

9.2 Benefits by stakeholder

The following are intended uses and benefits to test, rather than measured outcomes of SH-1.2.

Stakeholder	Practical use	What to evaluate
CX and operations leaders	Define handling for contacts from customer agents.	Resolution quality, staff workload and avoidable referrals.
Service designers	Specify goals, boundaries and human routes across all four modes.	Whether the process preserves the person's requirements.
Legal and risk teams	Review records of purpose, authority and decisions.	Whether actual practices meet applicable obligations.
Technology teams	Implement a common declaration and action-assessment format.	Enforcement, interoperability and failure handling.
BPOs and outsourcers	Record the scope delegated by client organisations.	Whether staff and agents act within that scope.
Consumers and their agents	State desired outcomes, refusals and fallback preferences.	Accessibility, retained boundaries and ability to obtain help.

An open question for service

Much of service design begins with an assumption: the organisation creates the environment, and the customer navigates it. The organisation defines the channels, the available choices and the route to resolution. A customer bringing an agent introduces another set of instructions, another view of an acceptable outcome and another way of conducting the interaction.

If this becomes ordinary practice, adding an agent-handling procedure will address only part of the change. We will need to reconsider who designs a service, how its terms are established and how responsibility is maintained when neither party controls the whole interaction.

The Service Handshake is an early framework for examining those questions. Its declarations and safeguards give practitioners something to implement, challenge and test. They do not settle how competing service designs should meet, which arrangements people will trust or whether the resulting services will be easier to use.

Readiness therefore involves more than systems and procedures. It involves questioning assumptions about who initiates, who negotiates, who decides and whose definition of a successful outcome counts. A company may be technically capable of receiving a customer's agent while still expecting the customer to follow a journey designed entirely on its terms.

We are at the beginning of this work. The emerging model of service has yet to be built, and familiar assumptions may become less reliable before workable alternatives are established. This framework is offered as a contribution to that process, and an invitation to help shape what follows.

10 About this document

This document is licensed under the Creative Commons Attribution 4.0 International licence (CC BY 4.0). You may share, adapt and use it for any purpose, including commercially, provided you credit Maria McCann and Neos Wave Ltd, link to the licence (<https://creativecommons.org/licenses/by/4.0/>) and indicate any changes.

The JSON schemas, examples and tests published with the Technical Specification are licensed separately under the Apache License 2.0 (<https://www.apache.org/licenses/LICENSE-2.0>).

Name and marks. The Service Handshake name, logo and badges are not licensed under CC BY 4.0 or the Apache License 2.0. Adaptations are welcome under those licences, but must be clearly identified as modified and must not be presented as the Service Handshake, as an official edition, or as endorsed by Neos Wave. Conformance claims follow §9.1.

Citation. McCann, M. (2026). *The Service Handshake* (Version 1.2). Neos Wave Ltd.
<https://doi.org/10.5281/zenodo.23181969>

This is a living standard. Contributions should report pilot methods, observed failures and feedback from service designers, operations, legal, risk and development teams. Revised

versions should identify the important changes without requiring readers to reconstruct them from text edits.

Version 1.0 and version 1.1 were published in March 2026. This edition retains the 1.1 structure and incorporates the changes listed at the front. The earlier publication and its citation remain available. SH numbers identify editions of the standard; a minor-number increment does not promise schema or behavioural compatibility. Implementations use the declared schema version and the migration rules in Technical Specification §11. Reading an older declaration does not automatically establish that it meets SH-1.2 requirements; the companion specification defines version handling.

About the author

Maria McCann is the co-founder of Neos Wave and has spent over twenty years designing and operating customer experience at scale, including senior roles at Spotify and ASOS. Her work includes consumer-side agent deployments, triadic AI systems and live contact-centre tests.

Neos Wave develops service transformation work related to this standard. The open standard is available independently of a commercial implementation. Any claim of conformance identifies the version, profile, use cases and actions covered, including for Neos Wave implementations.

References

- [1] McCann, Maria. The Service Handshake v1.1 and companion Technical Specification. Neos Wave, March 2026. <https://doi.org/10.5281/zenodo.19046746>.
- [2] Steinberger, Peter. OpenClaw, OpenAI and the future. 14 February 2026. <https://steipete.me/posts/2026/openclaw>.
- [3] Moffatt v. Air Canada, 2024 BCCRT 149. British Columbia Civil Resolution Tribunal, 14 February 2024. <https://www.canlii.org/en/bc/bccrt/doc/2024/2024bccrt149/2024bccrt149.html>. Case retained from version 1.1; see also the case account at <https://www.dentonsdata.com/airline-ordered-to-compensate-a-b-c-man-because-its-chatbot-provided-inaccurate-information/>.
- [4] Mastercard. When AI starts buying for you, trust becomes the product. 5 March 2026. <https://www.mastercard.com/content/mccom/eu/en/news-and-trends/stories/2026/verifiable-intent.html>.
- [5] European Commission. Transparency obligations under Article 50 of the AI Act. <https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act>. Consulted 28 September 2026.
- [6] Regulation EU 2016/679, Articles 5 and 6. <https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng>.
- [7] Information Commissioner's Office. UK organisations stand to benefit from new data protection laws. June 2025. <https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2025/06/uk-organisations-stand-to-benefit-from-new-data-protection-laws/>.

- [8] Competition and Markets Authority. Direct consumer enforcement guidance CMA200.
<https://www.gov.uk/government/publications/direct-consumer-enforcement-guidance-cma200>.
- [9] Meta. Introducing Muse. 8 September 2026.
<https://about.fb.com/news/2026/09/introducing-muse-personal-ai-agent/>.
- [10] Meta AI Research. How We Built Safety Into Muse.
<https://research.meta.ai/blog/security-and-safety-for-ai-agents-our-approach-with-muse>.
Consulted 28 September 2026.
- [11] Bradner, S. RFC 2119, Key words for use in RFCs to Indicate Requirement Levels.
<https://www.rfc-editor.org/rfc/rfc2119>.
- [12] Leiba, B. RFC 8174, Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words.
<https://www.rfc-editor.org/rfc/rfc8174>.
- [13] Ofcom. Managing telecom accounts on someone's behalf powers of attorney and third party bill management.
<https://www.ofcom.org.uk/phones-and-broadband/vulnerable-customers/power-of-attorney>.
Consulted 29 September 2026.
- [14] Booking.com. Customer terms of service, UK-English version, §A15.2.
<https://www.booking.com/content/terms.en-gb.html>. Consulted 29 September 2026.
- [15] Competition and Markets Authority. Writing a fair contract for customers.
<https://www.gov.uk/guidance/writing-a-fair-contract-for-customers>. Consulted 29 September 2026.
- [16] Information Commissioner's Office. A guide to controllers and processors.
<https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/controllers-and-processors/controllers-and-processors-a-guide/>. Consulted 29 September 2026.
- [17] Bishop, Todd. "Amazon blocks Meta's Muse AI assistant in new standoff over agentic shopping." *GeekWire*, 21 September 2026.
<https://www.geekwire.com/2026/amazon-blocks-metas-muse-ai-assistant-in-new-standoff-over-agentic-shopping/>.