

When the Audience Becomes an Agent

Why strategic communications must prepare for machine-mediated action



ARTEFACT91

Campaign Intelligence + Strategic Communications

artefact91.com

In this briefing

Why strategic communications must prepare for machine-mediated action

·	Executive Brief
01	From Answer Environment to Action Environment
02	The Current Signal Is Not One Product. It Is Infrastructure.
03	The Artefact91 Action Environment Model
04	When Product Data Becomes Reputation Data
05	Agentic Trust Is Different from Message Trust
06	The Communications Risk Is Not Limited to Purchases
07	Emerging Research Suggests Agents Can Inherit Decision Biases
08	The Agentic Reputation Test
09	Communications Teams Need a Machine-Actionable Fact Layer
10	The Brand Experience May Occur Without the Brand Interface
11	Measurement Must Follow the Agentic Decision Chain
12	A 90-Day Agentic Readiness Agenda
13	Questions for Leaders
·	Artefact91 Point of View
·	Sources & Research Base

HOW TO READ IT

A briefing, not a report: the executive brief carries the argument, the numbered sections carry the evidence and the method, and every external figure is attributed to its source.

ARTEFACT91 BRIEFING 15

Executive Brief

AI-mediated communications is moving into a new phase. The first phase changed discovery: search engines and assistants began summarizing organizations before people visited their websites. The next phase changes agency itself: AI systems are beginning to compare options, make recommendations, book services, assemble carts and—within defined permissions—complete transactions on a person's behalf.

That shift is now concrete enough to matter strategically. Mastercard, Flybits and Rogers Bank announced a Canadian deployment in which an AI assistant selected an eligible product within a cardholder's stated budget and completed the purchase using predefined controls. Mastercard's accompanying Canadian research found 41% of respondents were open to AI assistants for online shopping. Global Payments' latest cross-market research reports that consumers expect AI agents to make 15% of their purchases within five years, up from 9% a year earlier, while half of U.S. respondents remain concerned about payment security. Reuters separately reported that major banks are warning that agentic shopping is advancing faster than consumer-protection and regulatory frameworks.

The evidence still argues against hype. Gartner found only a minority of U.S. consumers willing to let AI make purchase decisions outright, and Visa found that only 23% trusted generative AI to handle payments on their behalf. Adoption, trust and autonomy are not moving at the same speed.

For communications leaders, however, the strategic change is already visible. An institution is no longer communicating only with people who read, watch and decide. It is increasingly communicating into systems that may retrieve information, interpret rules, compare choices and take an authorized action.

The action environment

The communications question is no longer only: "What answer does AI give about us?" It is becoming: "What can an AI system safely and accurately do with the information it finds about us?"

SECTION 01

From Answer Environment to Action Environment

Artefact91's earlier work on AI search and mediated answers focused on a major reputational shift: audiences may encounter a synthesized answer before they encounter an organization's own source material.

Agentic systems add a second layer.

ANSWER ENVIRONMENT — A machine retrieves, ranks and synthesizes information that shapes perception.

ACTION ENVIRONMENT — A machine retrieves, interprets and uses information to shape or execute a decision.

The distinction matters because errors can now travel farther downstream. A stale price is no longer only a factual error in an answer. It can change a purchase. An ambiguous cancellation policy can become a failed booking. An outdated eligibility rule can cause an agent to discard a relevant option. An incorrect location, deadline, service condition or organizational description can influence action before a human ever reaches the institution's website.

This does not mean autonomous agents are about to replace human judgment across every sector. It means

institutions should prepare for an environment in which machine interpretation increasingly sits between intent and action.

SECTION 02

The Current Signal Is Not One Product. It Is Infrastructure.

The consequential development is not any single shopping assistant. It is the simultaneous construction of standards, payment controls, merchant integrations and consumer interfaces that make agent-mediated transactions possible.

Mastercard has introduced Agent Connect and expanded its Agent Suite for Merchants so product details, pricing, availability and fulfillment information can be exposed to participating AI experiences while merchants retain control over business rules. Google's Universal Commerce Protocol similarly supports identity linking, cart building, checkout and order management across AI surfaces. Mastercard's Verifiable Intent model is designed to preserve a tamper-resistant record of what a consumer actually authorized.

These are infrastructure developments. They move the conversation from "Can a chatbot recommend a product?" toward "Can systems establish identity, authority, product truth, consumer intent, payment authorization and recourse across an agent-mediated transaction?"

The new interface

For communications teams, that creates a new stakeholder at the interface: the machine acting for a person.

SECTION 03

The Artefact91 Action Environment Model

ARTEFACT91 ACTION ENVIRONMENT MODEL

HUMAN INTENT → AGENT INTERPRETATION → SOURCE SELECTION → RECOMMENDATION →
AUTHORIZATION → ACTION → RECOURSE

HUMAN INTENT — What is the person actually asking the system to achieve? Goals may be precise, ambiguous or incomplete.

AGENT INTERPRETATION — How does the system translate that request into criteria, constraints and priorities?

SOURCE SELECTION — Which sources, feeds, websites, databases, reviews, policies or third-party signals does the agent use?

RECOMMENDATION — How are options ranked, excluded or framed?

AUTHORIZATION — What has the user actually permitted the system to do?

ACTION — What can the system book, buy, submit, send, register, reserve or change?

RECOURSE — If the outcome is wrong, what evidence exists, who is accountable, and how can the user reverse or correct the action?

Traditional communications practice concentrates heavily on source and recommendation. Agentic environments force institutions to think across the full chain.

SECTION 04

When Product Data Becomes Reputation Data

The strongest early evidence comes from commerce, but the lesson extends beyond retail.

Rithum's survey of more than 1,000 U.S. and U.K. online shoppers found that 58% said their trust in a brand decreases when an AI system gives them incorrect product information, and 16% said they would not complete the purchase. Only 5% reported going directly to a retailer's website when verifying an AI recommendation.

That creates an uncomfortable accountability gap. The organization may not have generated the incorrect answer, but the audience may still assign the reputational consequence to the organization.

Information traditionally managed as "operational data" therefore becomes communications infrastructure: price, availability, service conditions, eligibility, deadlines, locations, policies, leadership, accreditations, terms, evidence, exceptions, contact routes and correction mechanisms.

The exposure

If these facts are inconsistent, stale, inaccessible or difficult for machines to interpret, the reputational problem may surface before the communications team knows a decision was made.

SECTION 05

Agentic Trust Is Different from Message Trust

People can distrust a message while still understanding who delivered it. Agent-mediated decisions can blur that boundary.

A user may not know whether a wrong recommendation originated with the AI system, the merchant's product feed, a third-party marketplace, an old webpage, a review, a payment provider, a sponsored placement, or an interpretation error.

The institution should therefore design for traceability. Artefact91 proposes six questions.

IDENTITY — Can the institution tell when it is interacting with an authorized agent?

AUTHORITY — Can the agent demonstrate what the person permitted it to do?

SOURCE — Can the institution identify which information shaped the action?

VERSION — Can it establish which policy, price, rule or statement was current at the time?

REVERSIBILITY — Can the action be paused, reversed or corrected?

RECOURSE — Can a human obtain explanation and remedy without being trapped between platforms?

In an action environment, credibility depends partly on whether the institution can reconstruct what happened.

SECTION 06

The Communications Risk Is Not Limited to Purchases

Commerce is the first highly visible testing ground because payments force difficult questions about identity, consent and liability. The same communications logic can emerge wherever digital systems act on structured public information.

- Travel agents can compare and book.
- Scheduling agents can make appointments.
- Service agents can complete registrations.
- Procurement agents can shortlist suppliers.
- Donation agents could eventually select or route contributions.
- Administrative agents can submit forms or requests.
- Public-service assistants may guide people through eligibility and application pathways.

Widespread autonomous execution is not established across all of these domains. The strategic point is narrower: once a system moves from explaining options to acting on them, institutional information becomes part of the decision architecture.

That is a communications issue because public language, policies, FAQs, notices and evidence increasingly become inputs into machine-mediated choices.

SECTION 07

Emerging Research Suggests Agents Can Inherit Decision Biases

New academic work should be treated as early evidence rather than settled fact.

A September 2026 preprint by Davood Wadi and Yu Ma tested eight commercially deployed large language models acting as shopping agents. Under unconstrained conditions, common pricing cues rarely produced poor choices. But when information acquisition carried a cost and the goal prompt was vague, agents sometimes skipped diagnostic information and selected suboptimal options in ways that resembled human heuristics.

The important communications implication is not that agents are “irrational.” It is that information architecture matters. If critical facts are difficult to retrieve, costly to inspect, inconsistently structured or hidden behind weak

interfaces, the agent's decision process can change.

The implication

Clarity, accessibility and structured evidence are not only usability issues. They may shape machine-mediated decisions.

SECTION 08

The Agentic Reputation Test

Artefact91 recommends testing high-consequence customer and stakeholder journeys against seven questions.

CAN IT FIND US? — Is authoritative information machine-retrievable across the relevant platforms?

CAN IT IDENTIFY US? — Are organizational identity, locations, leadership, services and credentials unambiguous?

CAN IT UNDERSTAND THE RULE? — Are prices, eligibility criteria, service limits, deadlines, exceptions and conditions stated clearly enough to be interpreted correctly?

CAN IT DISTINGUISH CURRENT FROM OLD? — Are superseded policies and outdated pages clearly retired, redirected or labelled?

CAN IT VERIFY THE CLAIM? — Can important claims be traced to evidence rather than promotional language?

CAN IT ACT SAFELY? — Are authorization, limits and confirmation appropriate to the consequence of the action?

CAN THE PERSON RECOVER? — Is there a clear path to human support, explanation, correction and remedy?

An institution is not agent-ready because an AI assistant can crawl its website. It is agent-ready when the full journey from intent to recourse can be governed.

SECTION 09

Communications Teams Need a Machine-Actionable Fact Layer

The emerging environment requires more than “AI-friendly content.” Institutions should build an authoritative fact layer.

CURRENT — Outdated information is actively retired.

CONSISTENT — Critical facts match across websites, feeds, profiles, media materials and third-party systems where the institution has control.

SPECIFIC — Policies and service conditions use concrete language rather than ambiguous promotional wording.

STRUCTURED — Important information can be retrieved and interpreted reliably.

EVIDENCED — Claims can be traced to authoritative supporting material.

VERSIONED — Material changes to high-consequence information can be reconstructed.

CORRECTABLE — There is a defined route for fixing errors in owned and relevant third-party environments.

This is the operational version of Artefact91's concept of retrievable proof.

SECTION 10

The Brand Experience May Occur Without the Brand Interface

One of the most important changes is the separation of institutional experience from institutional interface.

A consumer can discover, compare and potentially purchase without meaningfully visiting the seller's digital property. A traveller may move from request to itinerary through an assistant. A stakeholder may encounter a summarized policy without reading the underlying document.

That changes what "owned media" means. The website remains important, but it becomes one node in a larger evidence and transaction system.

The question

Not only whether the organization has a clear digital journey, but whether its facts remain clear when the journey is assembled somewhere else.

SECTION 11

Measurement Must Follow the Agentic Decision Chain

Traditional digital measurement emphasizes visits, impressions, engagement and conversion. Agent-mediated environments require additional indicators.

AGENT DISCOVERABILITY — Does the institution appear in relevant machine-mediated choices?

FACT CONSISTENCY — Do different systems retrieve the same critical information?

EXCLUSION RATE — Is the institution being incorrectly filtered out because of missing or ambiguous data?

RECOMMENDATION ACCURACY — Are products, services, eligibility rules or institutional claims being represented correctly?

ACTION FAILURE — How often does the journey fail between recommendation and authorized action?

CORRECTION LATENCY — How long does it take to change a materially wrong machine-mediated representation?

RECOURSE QUALITY — Can users understand and resolve failures?

The objective is not to optimize blindly for machines. It is to understand whether machine mediation is changing

human outcomes.

SECTION 12

A 90-Day Agentic Readiness Agenda

DAYS 1–30 — MAP THE ACTION ENVIRONMENT — Identify the five highest-consequence journeys in which external AI systems could already shape discovery or action. Audit what major assistants and search environments say. Identify authoritative data sources, third-party dependencies, outdated information and unclear rules.

DAYS 31–60 — STRENGTHEN THE FACT LAYER — Resolve contradictions. Improve structured and machine-retrievable information. Establish owners for prices, policies, eligibility, deadlines, service terms and corrections. Define which actions require explicit human confirmation.

DAYS 61–90 — TEST ACTION, NOT JUST ANSWERS — Run scenarios in which an agent recommends the wrong option, uses an old policy, selects an unintended service, misreads a condition, completes an unauthorized action or routes a user toward a weak source. Test explanation, escalation, reversal and recourse.

The goal is not to predict which agentic platform will dominate. It is to make the institution more robust across platforms.

SECTION 13

Questions for Leaders

Which stakeholder decisions could an AI agent already influence without visiting our owned channels?

Which five facts would be most damaging if an agent got them wrong?

Can we prove which version of a policy or offer was authoritative at the moment an action occurred?

Where does our customer or stakeholder journey depend on ambiguous language?

If an AI system acts incorrectly using our information, who owns the correction?

Can a person reach a human and reverse the outcome?

ARTEFACT91 BRIEFING 15

Artefact91 Point of View

AI search changed who explains the institution. Agentic AI begins to change who acts on the institution's information. That is the consequential shift.

Strategic communications has spent years improving messages for human attention. The next challenge is to protect meaning, authority and trust when the audience's intent may be interpreted and executed by a machine.

The institution does not need to communicate less humanly. It needs to make its facts, rules, evidence and accountability more legible across a system in which people increasingly delegate parts of the decision process.

The answer environment is becoming an action environment.

ARTEFACT91 ACTION READINESS SEQUENCE

INTENT → INTERPRET → VERIFY → AUTHORIZE → ACT → TRACE → REMEDY → ADAPT

ARTEFACT91 BRIEFING 15

Sources & Research Base

- Mastercard — Agent Connect and Agent Suite for Merchants.
- Mastercard Canada — Mastercard, Flybits and Rogers establish benchmark for secure, consumer-controlled agentic commerce in Canada.
- Visa — Visa Trust Index for agentic commerce.
- Global Payments — Agentic Commerce Report 2026.
- Reuters — Banks warn AI shopping bots raise scam, fraud and data-privacy risks.
- Gartner — Consumers want AI shopping help, but not AI purchase decisions.
- Rithum — The New Discovery Engine: consumer trust and AI product discovery.
- Wadi, D. & Ma, Y. — Shopping by algorithm: How agentic AI deploys human heuristics as a surrogate consumer.
- Google — Universal Commerce Protocol.
- Mastercard — Verifiable Intent.

Method Note

This briefing synthesizes current public research, industry infrastructure announcements and emerging academic evidence with Artefact91 analysis. Survey findings are reported with their original scopes and should not be treated as directly interchangeable across markets or methodologies. Vendor announcements are evidence of product direction and infrastructure development, not proof of universal adoption. The Wadi and Ma research is a preprint and is used as emerging evidence rather than settled empirical consensus. Artefact91 frameworks are analytical interpretations intended to support institutional decision-making and should be adapted to jurisdiction, sector, risk level and technical context.