

# Design for the AI agent, not just the customer

The shift from customer experience (CX) to agent experience (AX): win by being chosen inside the agent's flow.



For two decades we optimized for the customer. Now the customer delegates. The next era of commerce is intent-first, and advantage belongs to enterprises that machines can understand, trust, and select without a human watching.

## Your next buyer is a machine, and it is already shopping

- Buying is shifting from human-led interaction to delegated decisions, as AI agents move from assistants to active buyers across industries.
- Gartner expects AI agents to drive \$15 trillion in B2B purchases by 2028, making AX an execution layer, not an experiment.

- Designing for AX is architectural: catalog metadata must become agent-ready decision intelligence, governed ontologies, and action-grade APIs, not richer feeds.
- Pricing, promotions, and brand shift from persuasion to proof, since agents select only what they can justify and trace.
- Advantage moves from owning the storefront to being chosen upstream, which demands a new, agent-aware measurement and attribution model.

## Why the shift from CX to AX matters now

For the last two decades, enterprises optimized relentlessly for customer experience (CX). They mapped journeys, refined touchpoints, and invested heavily in personalization, interface design, and omnichannel engagement. That model delivered results, until the buying journey itself began to change. A more fundamental shift is now underway, and many organizations are still underestimating its implications.

The customer is no longer the only decision-maker in the buying journey. Increasingly, AI agents are becoming active participants and, in some cases, primary actors. As AI systems grow more capable, commerce is moving from human-led interaction to delegated decision-making. This is no longer a distant possibility. Retailers, financial services firms, healthcare providers, and enterprise platforms are already experimenting with AI agents to automate routine decisions, reduce friction, and improve responsiveness.

**This gives rise to agent experience (AX): the set of product, data, and transaction capabilities that let AI agents discover, evaluate, and act on behalf of customers. The implication is significant. Businesses must now design not only for human customers, but also for the agents representing them.**

Traditional CX rests on a simple assumption. A human customer engages with a brand, evaluates options, and makes a choice. Agentic commerce breaks that model. In an agent-driven world, customers delegate rather than interact. Instead of browsing product pages, comparing reviews, and applying filters, a customer can simply instruct an agent: “Find me the best running shoes under \$120 for daily use.” The agent then executes the decision workflow. It searches across platforms, compares specifications, prices, and reviews, evaluates trade-offs, and either recommends or completes the purchase.

This is the essence of the shift from CX to AX. Competitive advantage increasingly depends on how well your business can be understood, evaluated, and transacted with by machines, not just by people. The broader momentum behind AI investment, agent-based automation, and autonomous workflows suggests AX is becoming an execution layer for commerce, not merely an assistive layer for customer support.

# What agentic commerce actually is

**Agentic commerce** is a model in which AI agents autonomously search, compare, negotiate, and transact within digital ecosystems on behalf of users. Rather than relying on direct human interaction with brands and platforms, these agents execute tasks such as product discovery, purchase decisions, subscription management, and service optimization based on user preferences, policies, and real-time context.

In this model, the primary interface is no longer always the end customer. It is often the customer's agent. Analyst forecasts already point to a meaningful rise in interactions and transactions initiated or influenced by autonomous agents, while users grow more comfortable delegating repeatable tasks to AI. Gartner expects one in three GenAI interactions running autonomously with agents by 2028. That changes what businesses must optimize for: structured product data, interoperable APIs, transparent rules, trust signals, and reliable

## What actually changes when designing for AX

Designing for AX is not about making existing product feeds slightly richer. That was already part of digital commerce before agentic AI. The real shift is architectural. Enterprises must move from catalog metadata to agent-ready decision intelligence: a governed layer of product, customer, policy, context, and outcome signals that lets agents reason, compare, justify, and act with confidence.

In other words, AX requires a fundamental shift in information architecture. Agents do not simply read product descriptions. They interpret intent, evaluate trade-offs, check constraints, apply rules, and select the next best action. To support that, businesses need an intelligence layer that connects what a product is, who it is right for, when it should be recommended, under what conditions it can be fulfilled, and why it is the best choice for a given journey.

## Product data evolves from catalog data to decision intelligence

Traditional product data answers basic catalog questions: name, price, description, specifications, images, availability, and category. Agent-ready product intelligence answers decision questions. Which product best fits this user's intent? What trade-offs matter? What evidence supports the recommendation? What constraints or risks apply? What alternatives should be considered? Delivering that means layering in several things at once:

- **Intent semantics** that map customer goals, use cases, preferences, exclusions, and situational needs to the right products or services.
- **Decision attributes** such as compatibility, suitability, substitutions, bundles, total cost of ownership, serviceability, sustainability, and risk factors.

- **Evidence signals** including verified reviews, expert ratings, historical performance, return reasons, warranty outcomes, and fulfillment reliability.
- **Contextual signals** such as customer segment, geography, inventory position, delivery promise, loyalty status, and consented behavioral history.

The objective is not just to be found by agents, but to be understood, reasoned over, and confidently selected. A product with complete catalog data may still be invisible to agents if the surrounding decision context is missing.

## A product ontology intelligence layer becomes foundational

Agents need more than attributes. They need meaning. This is where a product ontology intelligence layer becomes critical. It defines the entities, relationships, rules, synonyms, constraints, and business meanings that let agents interpret products in context rather than matching keywords or reading static descriptions. That layer should connect four domains:

- **Product ontology:** products, variants, components, features, use cases, compatibility, dependencies, and substitutions.
- **Customer ontology:** intent, preferences, constraints, life stage, usage context, affordability, accessibility, and loyalty signals.
- **Commerce ontology:** price, promotion logic, availability, delivery promise, returns, warranties, service rules, and payment options.
- **Trust ontology:** source provenance, confidence score, policy compliance, review authenticity, risk indicators, and explainability markers.

Without this semantic foundation, agents may retrieve information but struggle to reason correctly. With it, they move from shallow comparison to grounded recommendation, explaining why a product is suitable, what assumptions were used, and what alternatives exist.

## APIs become action and context channels, not just integration endpoints

For AI agents, the interface is not primarily the website or app. It is the set of discoverable, reliable, governed services through which agents can query, validate, reserve, personalize, and transact. APIs therefore need to expose not only product and inventory data, but also rules, eligibility, policies, decision support, and executable actions. In practice, that means the ability to:

- Query product, policy, inventory, fulfillment, and service data in real time.
- Validate eligibility, compatibility, constraints, consent, and user-specific rules before recommending.
- Simulate scenarios such as delivery options, bundle savings, return implications, or upgrade paths.
- Execute actions such as reserve, purchase, modify, cancel, escalate, or hand off to a human.

If these services are fragmented, slow, undocumented, or inconsistent, the agent experience fails even when the human-facing experience looks polished. AX readiness depends on API reliability, machine-readable contracts, policy orchestration, and controlled access to authoritative data.

## Pricing, promotions, and brand signals become explainable decision inputs

Agentic commerce puts commercial logic under machine scrutiny. Agents compare total value, detect ambiguity, challenge inconsistent pricing, and prefer options they can justify against the user's stated intent. This changes the role of pricing, promotions, and brand from persuasive messaging to decision-grade evidence:

- **Pricing logic** must be transparent, current, and inclusive of taxes, fees, delivery, returns, and service costs.
- **Promotion logic** must be machine-interpretable, with clear eligibility, exclusions, expiry, and stacking rules.
- **Brand trust** must be expressed through measurable signals such as fulfillment reliability, complaint resolution, service history, verified outcomes, and policy adherence.
- **Decision traceability** must explain why a recommendation was made, what data was used, and what confidence level applies.

In an AX world, competitive advantage shifts from being visually persuasive to being semantically clear, operationally reliable, and decision-ready. The brands that win will be those whose data architecture gives agents enough trusted context to recommend them with confidence.

This does not make CX irrelevant. Humans will remain central in emotional, complex, or high-stakes decisions. But as routine choices are delegated, enterprises must design for two audiences at once: people who need intuitive experiences, and agents that need structured meaning, decision context, and executable access.

## The risks that come with delegated decision-making

Agentic commerce introduces a broader operational risk surface, and it deserves the same rigor as the opportunity. Security becomes more complex because AI agents can be targeted, spoofed, manipulated, or exploited through the systems they access. Organizations will need stronger identity controls, transaction safeguards, monitoring, and exception handling to prevent misuse.

Privacy is equally important. Agents often require access to personal preferences, financial information, behavioral patterns, and contextual data to act effectively. That raises the need for explicit consent, data minimization, strong governance, and clear accountability for how agent-accessible data is used.

There are also real ethical considerations. Agents may optimize for outcomes that do not fully align with user intent, organizational policy, or social expectations. To manage that risk, businesses need guardrails for fairness, explainability, escalation, and human oversight, especially where decisions affect access, pricing, or customer treatment.

# Where agentic commerce goes next

Websites and apps will continue to exist for brand storytelling, inspiration, complex journeys, loyalty engagement, and human reassurance. They will no longer be the default starting point for every transaction. The primary battleground shifts from owning the storefront to being selected inside someone else's agent-led decision flow.

In practical terms, the homepage becomes less important, while product intelligence, pricing logic, inventory accuracy, fulfillment reliability, identity controls, and API readiness become more important. The customer may still see a checkout screen, but the real evaluation may already have happened upstream, inside a personal assistant, a platform-native shopping agent, or an enterprise procurement agent.

Commerce platforms and SaaS vendors such as Shopify, Salesforce Commerce Cloud, SAP Commerce, Adobe Commerce, and BigCommerce will therefore evolve from storefront platforms into agentic commerce operating layers. Their relevance will depend on how well they expose merchant catalogs, policies, checkout, payments, fraud controls, loyalty, and post-purchase workflows to external agents through trusted protocols and governed services. Shopify's move toward agentic storefronts and AI-channel commerce is an early signal. The platform of the future is not just a place to build a store. It is a distribution, orchestration, and transaction layer for agents.

Enterprises will face a strategic choice. They can build their own branded agents, plug into third-party agents, or operate a hybrid model. My view is that most large enterprises will need the hybrid. A brand-owned agent is useful for loyalty, service, advice, and repeat engagement. But third-party agents will control a growing share of discovery and comparison, especially when customers ask neutral, cross-brand questions. The enterprise therefore needs to be agent-ready both inside its own ecosystem and outside it.

This also points to the rise of domain-specific personal agents. Consumers may not rely on one universal agent for everything. They may use a health agent for care plans and insurance decisions, a travel agent for itinerary planning and disruptions, a finance agent for subscription optimization and savings, a shopping agent for replenishment and comparison, and a workplace agent for procurement or approvals. Commerce will increasingly happen at the intersection of these agents. A travel agent may book a trip, a finance agent may validate the budget, a loyalty agent may apply rewards, and a merchant agent may fulfill the transaction.

## Why agent-led ROI needs a new measurement model

This is why agentic commerce ROI cannot be measured only through traditional web conversion. Enterprises will need a new measurement model that captures how agents discover, evaluate, recommend, and complete decisions. The core question changes from "Did the customer click and convert?" to "Was our product trusted, selected, and acted on by the agent?" That reframes the metrics that matter:

- **Agent visibility:** share of relevant agent prompts where the brand or product is retrieved, considered, or shortlisted.
- **Agent preference:** rate at which the agent recommends the brand versus alternatives for similar intents and constraints.

- **Agent-to-transaction conversion:** percentage of agent-led recommendations that progress to cart, checkout, purchase, reorder, renewal, or service action.
- **Operational impact:** reduction in assisted-service volume, search friction, returns, failed orders, manual interventions, and cost-to-serve.
- **Customer value impact:** change in average order value, repeat purchase rate, basket quality, subscription retention, loyalty usage, and lifetime value for agent-mediated journeys.

Attribution will also need to change. The last click becomes far less meaningful when the decision happens before the visit. Enterprises will need agent-aware attribution models that recognize agent discovery, recommendation, negotiation, and execution as distinct events. A future conversion path may look like this: intent captured by a personal agent, product shortlisted through a third-party agent, eligibility validated through the merchant's API, checkout completed through an embedded payment layer, and post-purchase service handled by the brand agent. Each event should be measurable.

The strongest enterprise play is therefore not to replace the existing commerce stack, but to build an agent commerce control plane on top of it. This control plane governs product semantics, agent permissions, API access, pricing rules, identity, payments, trust signals, telemetry, experimentation, and attribution. It becomes the layer that decides which agents can access what, under what conditions, with what level of authority, and how the enterprise measures value from those interactions.

**The next era of commerce will not be website-first, app-first, or marketplace-first. It will be intent-first. The winners will be enterprises that make their products, policies, and services legible to agents, measurable across agent-led decisions, and trustworthy enough to be selected when no human is looking at the screen.**

## About Brillio

Brillio is The Enterprise AI Accelerator helping Fortune 1000 companies move from AI ambition to scaled impact, faster. Powered by our AI accelerator platform – Agentic Data and Application Management (ADAM), Brillio is one of the fastest-growing digital technology service providers, delivering transformation across five core workstreams: business-led transformation, customer experience transformation, AI and data engineering, digital engineering, and infrastructure engineering.

With 14 delivery locations across North America, Europe, and Asia and a team of over 6,000 customer-obsessed professionals, Brillio combines deep industry expertise, modern engineering, and accelerators to deliver measurable outcomes. Headquartered in Dallas, Texas, Brillio serves clients globally with a commitment to speed, scale, and measurable impact.



<https://www.brillio.com/>

Contact Us: [info@brillio.com](mailto:info@brillio.com)

