

Agentic Commerce for Retail Leaders: July 2026 Quick-Start Guide

Why being discovered, understood, trusted, and transactable by AI assistants is now a leadership challenge, and what retail teams should do first.

Author: Jeff Douglas, Digital Commerce Leader | Agentic Commerce and AI-Driven Retail

Why this matters now

I spent more than 20 years building a digital commerce channel from zero into a DC360 Top 500 retailer in one of the hardest categories to sell online. After leaving that role in May 2026, I had the space to go deeper on a question I had been tracking closely: how AI is changing retail discovery, shopping behavior, and the path to purchase.

The conclusion became hard to ignore. Being discovered, understood, trusted, and eventually transactable by AI assistants is no longer just a marketing issue. It is a leadership challenge.

Shoppers are already handing more research, comparison, and purchase preparation to AI assistants. At the same time, the infrastructure is arriving quickly. In roughly six weeks, Google introduced Universal Cart, Mastercard debuted Agent Pay for Machines, and Salesforce shipped what it called its biggest Agentforce Commerce release yet. These are not isolated announcements. They point to the same shift. AI is becoming a new decision layer between the shopper and the retailer.

The question for retailers is no longer only whether their site ranks. It is whether an assistant can access their catalog, understand their products, verify their claims, trust their policies, and guide the shopper toward the right purchase without creating operational risk.

That creates a new kind of visibility problem. A product page may look polished to a human while the information an AI assistant needs is incomplete, inconsistent, or buried across product data, feeds, policies, reviews, variant logic, and checkout rules. A person can often work through that friction. An agent may simply move on.

The visibility problem is already here

Public Adobe Analytics data shows AI-driven traffic to U.S. retail sites grew 235% year over year in the first five months of 2026. During Prime Day 2026, that traffic jumped 89% year over year and converted 40% better than non-AI channels like search, email, and social. Human-approved checkout inside or adjacent to AI surfaces is real and expanding, though the models vary by platform and remain early. Fully autonomous purchasing is further along in protocol infrastructure than many retailers realize, but consumer trust in in-chat checkout is still limited.

That means the urgent work is not autonomous transaction governance. It is discovery and catalog readiness. If an AI assistant narrows the field before the customer ever lands on your site, you either appear in that consideration set or you disappear from the journey.

I have lived through enough ecommerce shifts to know the difference between a trend and an operating problem. Mobile, site speed, personalization, retail media, and platform migrations all looked like digital projects at first. In reality, they became leadership challenges because they changed how customers found products, trusted information, and completed purchases.

Agentic commerce is not one thing

A lot of confusion in this market comes from treating all AI commerce as the same thing. It is more useful to separate it into layers:

- **AI-assisted discovery:** The assistant helps the shopper research, compare, shortlist, and understand tradeoffs.
- **AI-to-cart or AI-to-PDP handoff:** The assistant sends the shopper to the exact item, variant, configuration, or cart state on a retailer-owned surface.
- **AI-mediated checkout:** The shopper approves a purchase inside an AI surface or embedded checkout flow, with the merchant still owning the order, payment rules, fulfillment, and service relationship.
- **Autonomous agent purchasing:** The shopper pre-authorizes an agent to complete a purchase later within defined limits. This is the most sensitive layer because it raises consent, fraud, liability, auditability, identity, and payments questions.

Most retailers should not build their roadmap around the last layer first. The practical work starts with the first three: discovery, handoff, and trustworthy checkout readiness.

Executive takeaway:

AI commerce is not just checkout inside ChatGPT. The immediate work is making product data discoverable, trusted, and actionable across AI-assisted discovery, checkout handoff, and emerging agentic payment models. Consumer adoption of in-chat checkout remains early. Forrester's March 2026 commentary on OpenAI's Instant Checkout pullback reinforces that the market is still experimental, which means most retailers should prioritize discovery readiness, clean product data, and checkout handoff before assuming broad consumer adoption of in-assistant purchasing.

First, get the language right

There is a lot of loose vocabulary in this space, and using it loosely is a fast way to lose credibility with senior leaders. Here is the practical distinction I use:

- **SEO:** Search engine optimization gets a human to your site from a search engine. The traffic source is a classic results page, where a person types a query, scans links, and clicks through.
- **AEO and GEO:** Answer engine optimization and generative engine optimization influence how your brand and products show up when AI systems summarize, compare, and explain options. They shape whether you appear in the answer, how accurately you are represented, what sources are cited, and which claims the model has enough evidence to trust.
- **ACO:** Agentic Commerce Optimization is a term coined by Scot Wingo — founder of ReFiBuy — and is gaining traction as the working vocabulary for this layer of the stack. The practical definition: the product-level work that makes your catalog eligible, understandable, trustworthy, and transaction-ready for AI assistants and shopping agents. I use it as an executive framework, not a settled standard, but it is the clearest shorthand available for what retailers actually need to do.

The clean executive rule: SEO wins the click. AEO and GEO influence the answer. ACO prepares the product to be selected, verified, and acted on.

The new goal: own the product card

In traditional search, the goal was to win visibility for a page. In agentic commerce, the goal is broader: own the product card. That means AI shopping engines can correctly recognize the product, map it to the right canonical item, understand its variants and offers, and see your product as accurate, complete, trustworthy, and eligible to recommend.

This matters because agents do not always organize the journey around your PDP. They may organize it around a product card: the recognized product, its attributes, images, reviews, sellers, prices, availability, delivery promises, policies, variants, and related offers. If your identifiers are weak, your variant structure is confusing, or your offer is attached to the wrong product record, you can lose visibility even if your own website looks polished.

Owning the product card means the agent can identify the product correctly, trust the information attached to it, and connect your offer to the shopper's intent. That is the practical bridge from ACO strategy to catalog execution.

In my experience, catalog problems rarely show up as “catalog problems.” They show up as conversion leaks, confused customers, and costly returns. AI assistants will only make that more visible.

The mental shift that matters is this: for the better part of two decades, we optimized to win a click. Now we also have to optimize to win a recommendation we may never see happen. The funnel can start somewhere the retailer does not control.

Move 1: Get your catalog to the agents, cleanly and completely

The principle is simple: your product data has to be available where AI assistants can actually use it, and it has to be consistent enough to trust. A polished PDP is important, but it is only one path. Feeds, APIs, commerce platform catalogs, structured data, and protocol integrations may all shape what an assistant understands, verifies, and recommends. In practice, retailers need more than one path because each source can break, lag, or contradict the others.

What to do

- **Create an agent-access policy**, not just an unblock list. Review robots.txt, bot-management rules, crawler permissions, rate limits, and partner access. Decide which AI crawlers, feeds, APIs, platform catalogs, and partner agents should be allowed, how spoofing is detected, and who approves changes. Opening everything is naive. Blocking everything is self-defeating.
- **Audit structured data on every PDP.** Price, availability, brand, GTIN, MPN, condition, ratings, review count, variants, size, color, material, dimensions, shipping, delivery options, return terms, warranty, seller, merchant of record, and business policy data should be present where relevant and consistent with the feed and checkout.
- **Protect product-card ownership.** Make sure your core products are mapped to the right canonical item across AI shopping engines, marketplaces, feeds, and shopping surfaces. Strong identifiers, clean titles, accurate brand and model data, complete variants, consistent images, and offer-level accuracy determine whether the agent can recognize your product and prefer your offer.
- **Deliver a clean feed or platform catalog in parallel.** An optimized PDP and a direct data path are not redundant. Feeds and catalog integrations give agents fresher, more controlled data. PDPs provide context, authority, and corroboration. Treat feed freshness, suppression reasons, missing attributes, and policy errors like revenue-impacting metrics.

- **Make inventory and price trustworthy.** Product availability, variant eligibility, delivery promise, promotional price, tax/shipping estimates, and final checkout price must reconcile across feed, PDP, cart, and checkout. If the agent sees contradictions, the safer recommendation is often the competitor with cleaner data.
- **Model the variant tree correctly.** In high-consideration categories, one item may be a tree of sizes, colors, finishes, configurations, delivery rules, service options, warranties, and compatibility constraints. If those variants are modeled as disconnected products or hidden behind page logic, an agent can easily recommend the wrong thing.

Move 2: Make product information agent-readable and intent-rich

For years, short product copy often worked. A tight title, a few clipped specs, and lifestyle copy could get a human to click. That is no longer enough. AI assistants read for meaning and match against natural-language shopper intent, such as “I need a lightweight two-person tent for weekend backpacking that can handle rain, packs small, sets up quickly, and still has enough room for two adults and gear.” A four-word title cannot answer that. Complete, structured, accurate product data can.

What to do

- **Fill in the product record, not just the page.** Materials, dimensions, compatibility, warranty terms, certifications, constraints, service requirements, assembly, delivery limitations, and product limitations all matter. The goal is completeness and accuracy, not word count for its own sake.
- **Add conversational attributes.** Older product feeds rarely captured use case, occasion, compatibility, buyer intent, skill level, customer constraints, pairing logic, audience fit, or decision context. Agents need these fields to answer questions like “which option is best for my situation,” “what works with what I already have,” or “what should I choose if I care most about durability, budget, ease of use, or long-term value?”
- **Use direct language early.** Make the product type, ideal customer, core use case, most important constraint, and key differentiator explicit in the first lines and in structured fields. Do not make the model infer the important facts from lifestyle copy.
- **Use customer language as missing data.** Site search terms, zero-result searches, support chats, returns reasons, reviews, delivery complaints, financing questions, and sales associate notes reveal the language real shoppers use. Feed that language back into attributes, FAQs, PDP copy, buying guides, and product taxonomy.
- **Separate truth from persuasion.** Agents need facts they can verify. Marketing language still matters, but the core product record should be literal, specific, and testable.

The practical test:

Ask whether an assistant could answer the five hardest customer questions about the product without guessing, calling customer service, or reading between the lines.

Move 3: Win the evidence and trust layer

An AI assistant should not take your marketing copy at face value. It corroborates. It may compare your PDP, structured data, feed, reviews, third-party sources, price, delivery promises, return terms, Q&A, and public reputation. The agent needs evidence, and that evidence lives both on and off your site.

What to do

- **Build product-level review depth.** Volume, recency, specificity, and usefulness matter. Forty recent reviews with fit, size, use-case, delivery, assembly, durability, and return detail are often stronger evidence than four generic five-star ratings.
- **Use PDP FAQs to tell your authoritative side.** A genuine product-level Q&A gives the agent clean, attributable text and gives you a place to correct stale, incomplete, or wrong third-party claims.
- **Earn presence in credible retrieved sources.** Identify which buying guides, category authorities, review sources, marketplaces, forums, and expert roundups show up when AI systems answer your category questions. This is not old link-building. It is evidence management.
- **Close contradictions.** If your PDP says one delivery promise, your feed says another, reviews complain about a third, and checkout adds a fourth, the agent has a trust problem. It will usually resolve uncertainty by recommending something easier to verify.
- **Make policies machine-readable.** Returns, warranties, delivery, financing, installation, service areas, white-glove options, and exclusions should not be buried in PDFs or vague policy pages. These are buying-decision inputs.

Retailers often underestimate how quickly small inconsistencies become trust problems. A customer might forgive one confusing delivery message. An assistant comparing five options may simply avoid the uncertainty altogether.

Move 4: Prepare for AI-to-cart, AI-surface checkout, and payment trust

This is the area most teams either overhype or ignore. Human-approved AI checkout is no longer hypothetical—Microsoft Copilot Checkout launched in January 2026 with PayPal, Shopify, and Stripe; Google’s Universal Cart started rolling out in the US in May 2026. At the same time, consumer trust and adoption are still developing. Forrester’s March 2026 commentary on OpenAI’s Instant Checkout pullback reinforces that the market is still experimental. The right posture is practical readiness, not science fiction.

What to do

- **Treat AI-to-cart as a conversion path.** If an assistant sends a shopper to your site, land them on the exact variant, preserve price and availability, keep cart state where possible, surface shipping and returns quickly, and avoid forcing the customer to restart the journey.
- **Know which checkout model you are in.** The purchase may happen on the merchant site, in an in-app browser, through an embedded AI-surface checkout, through a native AI checkout, or through a future pre-authorized agent payment. These are different operating models and should not be described as one thing.
- **Define payment and authorization governance now.** Who owns consent records, fraud risk, chargebacks, loyalty identification, tax, returns, customer service, refunds, substitutions, backorders, and order changes when an agent initiates or completes part of the transaction. Do not wait until the first broken agentic order to answer that.

- **Map failure modes.** Wrong variant, stale price, unavailable item, incompatible accessory, delivery restriction, service-area mismatch, backorder, promo mismatch, tax/shipping change, financing change, and return-policy surprise are not edge cases in complex commerce. They are the exact cases an agent must avoid or escalate.
- **Measure the handoff.** Track AI-originated sessions, feed-driven referrals, AI channel attribution, app-driven journeys, conversion rate, abandoned cart causes, post-purchase service contacts, returns, and customer satisfaction after AI-influenced purchases. The channel will be hard to manage if analytics remains click-centric only.

Know the plumbing: the protocol stack

Leaders do not need to personally implement every protocol, but they do need to understand what each layer is trying to solve. These layers can stack. They do not all compete for the same job.

- **ACP (Agentic Commerce Protocol):** OpenAI's protocol, developed with Stripe and leading merchants, is open infrastructure for AI-native commerce in ChatGPT. It supports agent-ready commerce flows, including richer product discovery and merchant-provided product data. OpenAI has moved away from its initial standalone Instant Checkout model and is prioritizing better product discovery and merchant-owned checkout experiences. The important merchant principle is that ACP is designed to keep merchants involved as merchant of record, with responsibility for fulfillment, returns, support, and post-purchase communication wherever ACP-powered checkout occurs.
- **UCP (Universal Commerce Protocol):** Google's broader commerce-journey layer is designed to give AI agents, merchants, commerce platforms, and marketplaces a shared language across discovery, cart, checkout, and post-purchase actions. The practical merchant reality is straightforward: agents need reliable access to product data, pricing, delivery promises, checkout options, and order status without requiring custom, one-off integrations for every emerging AI surface.
- **AP2 and payment-trust protocols:** Google's Agent Payments Protocol, Visa's Trusted Agent Protocol, Mastercard's Agent Pay, and related FIDO work are aimed at the trust layer behind consumer agentic commerce: agent identity, user authorization, cryptographically signed mandates, payment credentials, spending limits, fraud controls, auditability, and liability. Mastercard's June 2026 AP4M extends this into a different frontier — continuous, machine-to-machine payments between agents, including microtransactions — which matters less for retail checkout today and more as a signal of where autonomous commerce infrastructure is heading. This infrastructure is moving quickly, but it does not mean fully autonomous consumer retail purchasing is mainstream today.
- **MCP (Model Context Protocol):** MCP is an open standard created by Anthropic and donated to the Agentic AI Foundation under the Linux Foundation. It provides a secure way for AI systems to connect with external tools and data sources. In commerce, MCP can help an agent query catalog, inventory, order, service, CRM, or ERP systems, but it does not by itself solve commerce semantics, payment authorization, fraud controls, customer consent, or merchant operational policy.

The practical takeaway: supporting more than one path is how you stay present across ChatGPT, Gemini, Copilot, Claude, brand-owned agents, and whatever comes next. The goal is not protocol fandom. The goal is reliable product understanding, safe action, and measurable commerce outcomes.

What the major platforms can and cannot currently do

The major AI platforms are moving toward the same broad destination, but they are not interchangeable. Each platform is developing a different mix of discovery, product data ingestion, checkout handoff, embedded checkout, agentic workflow, and payment authorization.

- **ChatGPT / OpenAI** is one of the most important consumer AI surfaces for product discovery and comparison. OpenAI now supports richer shopping experiences in ChatGPT, merchant-provided product data, and Agentic Commerce Protocol-based commerce flows. The important distinction for retailers is this: being discoverable in ChatGPT, being accurately represented in product comparisons, and being eligible for an approved checkout or app integration are related but separate readiness questions. The broad retailer priority is still catalog quality, feed integrity, product accuracy, and a clean path from AI discovery to merchant-owned conversion.
- **Google** is building a broad commerce layer across Search, AI Mode, Gemini, Shopping, Google Pay, and merchant integrations. UCP and Universal Cart matter because they point toward a more standardized path from product discovery to cart, checkout, and post-purchase actions while keeping the retailer as merchant of record in eligible flows. Retailers should view Google as both an AI answer surface and a commerce infrastructure path, but not assume every feature is live for every merchant, category, or region.
- **Microsoft Copilot** is one of the clearest examples of AI-surface checkout. Copilot Checkout allows eligible shoppers to move from product discovery to purchase inside Copilot with participating commerce and payment partners while merchants remain merchant of record. That model can reduce friction, but it raises operating questions around attribution, customer data, loyalty identification, returns, service ownership, and the post-purchase relationship.
- **Claude / Anthropic** should not currently be treated as a major top-of-funnel consumer shopping destination. Its importance is different. Claude, MCP, and Anthropic's agent research matter most for internal merchant workflows, tool use, catalog operations, customer-service workflows, data cleanup, and experiments in agentic negotiation and decision-making.
- **Shopify** is turning AI-channel readiness into platform plumbing through Shopify Catalog, Agentic Storefronts, ChatGPT discovery, Copilot Checkout, Google AI Mode, and Gemini participation. The nuance is eligibility. Shopify can help merchants distribute structured product data and participate in emerging AI shopping surfaces, but it does not remove the merchant's responsibility for clean product records, accurate policies, trustworthy inventory, variant logic, and operational governance.
- **Salesforce** is positioning Commerce Cloud and Agentforce Commerce as the connection layer between AI-driven shopping and the enterprise systems that govern pricing, promotions, customer context, fulfillment, order management, service, and B2B/B2C workflows. For Salesforce merchants, the opportunity is connection depth. The risk is assuming the platform integration alone solves governance, attribution, product data quality, or cross-functional ownership.
- **Amazon and Walmart are taking different control strategies, and both matter as reference points.**
 - ❖ **Amazon is taking a controlled, first-party approach.** It has brought Rufus's product expertise into Alexa for Shopping across Amazon's app, website, and Echo Show surfaces, while continuing to build the experience inside its own ecosystem. Amazon says Rufus helped more than 300 million customers in 2025, and Alexa for Shopping now connects product research, shopping history, personalization, price alerts, and shopping actions across Amazon-owned surfaces. Amazon has also fought unauthorized third-party agent access, winning a preliminary injunction — currently stayed pending appeal — blocking Perplexity's Comet agent from accessing password-protected areas of Amazon to shop on behalf of customers. At the same time, Amazon joined the Universal Commerce Protocol's Tech Council in April 2026, and it purchased sponsored placements inside ChatGPT during Prime Day. The practical lesson is not simply 'block agents.' It is that Amazon wants agentic shopping

to happen on its own terms — defending its store while keeping a seat at every table where the standards are being written.

- ❖ **Walmart is taking a more distributed, but still merchant-controlled, approach.** After OpenAI moved away from its standalone Instant Checkout model, Walmart brought its own shopping assistant, Sparky, into ChatGPT as a dedicated in-app experience with account linking, loyalty, and Walmart-owned checkout, with Gemini integration following. Walmart’s own stated reason is instructive: in-chat checkout converted roughly three times worse than a click-through to walmart.com, while ChatGPT drove new customers at about twice the rate of search. Most retailers are not Amazon and cannot win with a fully defended perimeter. The more practical lesson is closer to Walmart’s: participate deliberately where AI discovery happens, but keep control of product data, cart logic, checkout, service, and governance.

Two cautions matter. First, the infrastructure is still being built before the economics are fully proven. Platform fees, ranking rules, attribution standards, customer data norms, retail media models, and ownership of the customer relationship are still evolving. Second, platform plumbing does not remove the retailer’s work. It moves the work upstream into catalog readiness, product evidence, policy clarity, payment governance, measurement, and operational trust. In an agentic commerce environment, bad data does not stay hidden. It travels faster.

Why this can’t wait until Black Friday

Last holiday season, before most of the infrastructure in this guide existed, Salesforce reported that AI and agents influenced 20% of global online retail sales — \$262 billion — through personalized recommendations and deeper customer engagement. Salesforce also reported that retailers deploying their own AI agents saw a 59% higher growth rate, averaging 6.2% YoY sales growth versus 3.9% for retailers without them.

That was one holiday cycle ago, with a fraction of today’s protocol stack live. Universal Cart, Copilot Checkout, Agentforce Commerce, and ChatGPT’s expanded discovery layer have all shipped since then. This holiday season is likely to be the first meaningful stress test for AI-mediated discovery, AI-originated traffic, and emerging checkout handoff models at scale. Many retailers are not yet ready for that shift.

The work in this guide is not a 2027 roadmap, and it is not an October project. Catalog audits, structured data cleanup, feed integration, and policy clarity take months to execute and longer to stress-test under real volume. The work is not glamorous, but it is commercially important. Retailers who do it before peak are more likely to remain visible, trusted, and actionable when AI-assisted shopping journeys accelerate. Being “ready by Black Friday” is already too late. The real deadline is Labor Day, not Thanksgiving. Anyone who has worked through a holiday season knows that the fixes you wish you had made in October usually needed to start in July.

If your team does not have the bandwidth to run this internally on that timeline, this is exactly the kind of work platform partners, data agencies, and specialized readiness consultancies exist to accelerate. Compressing a multi-month catalog and structured-data overhaul into weeks is possible, but it usually requires outside hands, not just outside advice.

The checklist below is not a maturity model. Treat it as your punch list, and work backward from Labor Day.

Minimum readiness checklist

A leadership-ready audit should pressure-test the following items before claiming agentic readiness:

- Products have accurate IDs, GTINs or MPNs where appropriate, brand data, taxonomy, variant relationships, titles, descriptions, images, and availability.

- Feed, PDP, structured data, cart, and checkout agree on price, sale price, availability, fulfillment promise, return terms, and shipping options.
- Top products have product-card coverage and ownership checks: they are mapped to the correct canonical product, attached to the correct variants and offers, and represented with accurate title, image, price, availability, review, policy, and fulfillment data.
- Product variants are modeled as variants, not as disconnected products or hidden template logic.
- Policy data is readable: returns, warranties, delivery zones, white-glove service, installation, financing, exclusions, and merchant of record details.
- AI crawlers, commerce partners, platform feeds, and API access are governed intentionally rather than handled through accidental bot blocking or blanket allow rules.
- The organization can identify AI-originated journeys, attribute orders where the platform supports it, and measure returns, service contacts, and margin after AI-influenced purchases.
- There is an owner with authority across merchandising, ecommerce, IT, analytics, CX, payments, legal, and operations.

Where this goes next: Context Capture

Once the catalog is clean and complete, the next edge is feeding agents the one thing they still lack: how real shoppers actually talk about your products. The richest source of context is already inside the business, often unused.

- The search and prompt language people use to find, and fail to find, your products.
- Zero-result searches that reveal exactly what your catalog does not know how to answer.
- The reasons behind cart abandonment, returns, service contacts, warranty issues, and delivery complaints.
- The projects, occasions, constraints, objections, and use cases buried in reviews and sales conversations.
- The language sales associates use when they translate a messy customer need into the right product configuration.

Capturing even a fraction of this and feeding it back into product data is how you move from keeping up to getting ahead. Most merchants collect some of it. Very few operationalize it into catalog structure, PDP content, feeds, FAQs, buying guides, and AI-surface answers.

Who owns this

This is where good ideas usually stall inside retailers. Not because the goal is wrong, but because the work sits between teams. In most retail organizations, agentic readiness cuts across functions that rarely move in sync. Merchandising owns attributes and catalog depth. IT or engineering owns feeds, APIs, crawler access, and checkout integration. Brand, CX, and service own reviews, Q&A, policy clarity, and reputation. Payments and risk own authorization, fraud, chargebacks, and compliance. Analytics owns attribution and intent data. Legal has a voice in consent, customer data, platform terms, and liability.

Each function can do its part well and the brand can still lose, because no single team owns the outcome. A senior leader must own agentic readiness end to end, with enough understanding to pull teams together and connect the work to the P&L. This is a leadership mandate before it is a project plan.

The bottom line

AI is becoming another decision layer in product discovery and consideration. It does not replace every channel, and it is not the customer. But when it narrows the field before the shopper reaches the cart, it can shape which products become visible, trusted, and actionable.

If an assistant cannot reach your catalog, understand your products, verify your reputation, reconcile price and availability, or trust the path to purchase, it is more likely to favor a competitor that is easier to understand and safer to recommend.

That matters most in the hardest categories, especially spec-heavy and return-sensitive goods where mismatches are expensive and customer confidence is hard-earned. Discovery is already shifting. Checkout is evolving across platforms, protocols, and payment models. The retailers that win will be the ones that make the journey easier for both the customer and the assistant acting on the customer's behalf.

Underneath all of this, the job has not really changed. I have always been focused on one question. How do we make it easier for the customer to do business with us? Every replatform, page-speed improvement, search decision, merchandising decision, and operational tradeoff comes back to the same goal: removing friction between a person and the thing they want to buy.

Agentic commerce is the newest version of that work.

That is why this problem matters to me. I am not approaching agentic commerce as a trend to comment on. I am looking at it as an operator who has spent more than 20 years building ecommerce systems where product data, customer trust, conversion, delivery, service, and platform execution all have to work together.

This is the problem I think is most worth solving in retail right now.

— Jeff Douglas

Digital Commerce Leader | Agentic Commerce and AI-Driven Retail

About Jeff Douglas



Jeff Douglas is a digital commerce executive with more than 20 years building and scaling consumer-facing ecommerce for a Berkshire Hathaway omnichannel retailer. He started in sales in 1996, took over internet operations in 2003, launched the company's ecommerce business, and grew it from \$0 into a major home-furnishings operation ranked in the Internet Retailer / Digital Commerce 360 Top 500. As Director of Ecommerce, he owned the operational P&L: revenue, conversion, average order value, and customer lifetime value. He also led Product and UX, Digital Marketing, Site Merchandising, and the enterprise migration to Salesforce Commerce Cloud.

He is an active voice in digital commerce and agentic retail, most recently serving as a national panelist on "The New Era of Agentic Commerce" in 2026. He has been quoted in *Furniture Today* and *Total Retail*, and has spoken at IRCE and for the HFA-Home Furnishings Association.

Contact: jeff.douglas@outlook.com • linkedin.com/in/jeff-douglas-72a0094

Sources Reviewed

Public sources reviewed for this guide:

- [OpenAI: Powering Product Discovery in ChatGPT and merchant product feeds](#)
- [OpenAI: Instant Checkout and the Agentic Commerce Protocol \(original launch\)](#)
- [Google: UCP launch announcement at NRF 2026](#)
- [Google: Universal Cart and Agent Payments Protocol at I/O 2026](#)
- [Salesforce: Agentforce Commerce announcement \(June 2026\)](#)
- [Shopify: Agentic Commerce on Shopify — How It Works](#)
- [Anthropic: Project Deal agentic commerce experiment \(anthropic.com, April 2026\)](#)
- [Forrester: What It Means That The Leader In “Agentic Commerce” Just Pulled Back \(March 2026\)](#)
- [EMARKETER: FAQ on agentic commerce — how brands should act now](#)
- [Visa: Trusted Agent Protocol and Intelligent Commerce milestones \(Oct–Dec 2025\)](#)
- [Forrester: The State of Agentic Commerce, Q2 2026 \(May 2026\)](#)
- [GeekWire: Amazon, Walmart, and the agentic commerce race](#)
- [Adobe Analytics: 2026 Prime Day Insights \(June 2026\)](#)
- [OpenAI: Agentic Commerce Protocol \(ACP\) developer documentation](#)
- [Google: Universal Commerce Protocol guide](#)
- [Google: Under the Hood — Universal Commerce Protocol \(developer blog\)](#)
- [Microsoft: Copilot Checkout announcement \(January 2026\)](#)
- [Shopify: Millions of merchants can sell in AI chats](#)
- [Shopify: Agentic Plan launch \(April 2026\)](#)
- [FIDO Alliance: trusted AI agent interaction standards](#)
- [Digital Commerce 360: OpenAI agentic commerce updates \(March 2026\)](#)
- [EMARKETER: FAQ on GEO and AEO — where AI search and SEO overlap in 2026](#)
- [Mastercard: Agent Pay for Machines launch \(June 2026\)](#)
- [Forrester: Agentic Payments in B2C Commerce — Where We Are Now \(April 2026\)](#)
- [Scot Wingo / Retailgentic: Agentic Commerce Optimization Update](#)