

The New Discovery Layer

Preparing Retail for Agentic AI



how do these jeans fit?



Rise

Low

Mid

High

Inseam: 29"

Runs small

True to size

Runs large

Table of Contents

3	Part I: Introduction
4	Part II: The Evolution of Product Discovery
6	Part III: Agents in Practice
10	Part IV: The Data & Infrastructure Layer
14	Part V: Risks & Missteps
17	Part VI: Readiness to Results
19	Conclusion

We're witnessing the early stages of a radical shift in digital commerce: the rise of AI shopping agents. It's a transformation that redefines discovery, giving shoppers a new way to find products and make purchasing decisions. But unlike previous shifts that primarily changed where people shopped, the agentic era changes who — or rather, what — does the shopping.

For the first time, consumers can delegate part or all of the discovery process to an AI agent: an intelligent system that perceives input, reasons over options, and acts to achieve a goal. Instead of clicking through product pages, comparing features, and reading reviews, shoppers can describe their needs through a prompt and let their agentic AI assistant do the research, comparison, and possibly even the transaction. It's a fundamental change in the relationship between retailers and customers, introducing a new intermediary that acts on behalf of the shopper.

This isn't science fiction. Early versions of AI-assisted shopping already exist, and consumer comfort with conversational experiences is growing fast — according to the [2025 State of Ecommerce survey](#), 64% of shoppers have used generative AI (GenAI) to shop (up from 51% in 2024), and 58% say they are comfortable with AI-assisted discovery.

Shoppers increasingly expect AI to cut through the noise and surface options that truly fit. For many retail leaders, that pace of change can still raise uneasy questions about what role agents will play in ecommerce.

The concern is understandable. Retailers have spent decades perfecting digital storefronts — optimizing conversion funnels, personalizing product pages, and crafting brand experiences that guide shoppers from discovery to checkout. It's natural to wonder how AI fits into that hard-won progress.

In reality, agents don't compete with the work retailers have done; they complement it. On-site discovery agents apply the same principles of personalization, convenience, and confidence that already define great ecommerce — turning static experiences into dynamic, guided ones. They help shoppers ask questions in their own words, find the right products faster, and feel more certain about their choices. Off-site or inbound agents, such as ChatGPT, Perplexity, or Gemini, are beginning to influence discovery, too — surfacing product information and comparisons before a shopper ever reaches your site.

Each wave of digital innovation — from search to mobile to social — once felt disruptive, but ultimately expanded how customers engaged and bought. Agentic commerce will likely follow that same path. For retailers, the opportunity isn't to defend against AI, but to use it to serve customers better, convert more confidently, and make every visit feel intuitive and supported.

The evolution of product discovery

How people shop has always evolved with the channels available. Each new era didn't erase the last — it layered on top, creating new behaviors and expectations while adding convenience and expanding options.

From flipping through catalogs at the kitchen table to scrolling social feeds and now delegating tasks to AI, every shift opened up new ways to discover and

purchase products outside of traditional, in-store shopping.

Agents (2024 🚀)

Discovery gets delegated. Everyone gets a personal shopping assistant. Tell AI your goal and send the agent on its mission to search, compare and even buy for you, through an LLM like ChatGPT or a retailer's own on-site agent.

Social Commerce (Late 2010s 📱)

Shopping found the feeds. Products appeared between photos and livestreams, blending entertainment with discovery, powered by creators and trends.

Retail Media (2010s 📺)

Shopping met advertising. Sponsored placements on retailer sites turned search results into ad networks, blending personalization with promotion.

Mobile Apps & Web (2010s 📱)

Shopping hit the home screen. Apps tied loyalty, push notifications, and in-store scans together, enabling anytime/anywhere discovery.

Ecommerce Websites (1995 🌐)

Shopping went digital. Online catalogs appeared, extending store assortments, bringing new "digital pureplay" options, and introducing a new discovery tool – the search bar.

Catalogs (1870s 📖)

Shopping came home. People received thick books of curated products, mailed in forms, and waited weeks. It was the first time discovery happened without visiting a store.

Every new channel begins as a nascent novelty, matures through adoption, and remains in the

ever-evolving retail ecosystem. Each wave expands consumer options, rather than replace them.

Adoption curves show patterns

While the advent of online shopping remains the most seismic shift in the way we shop, mass adoption didn't happen overnight. Both consumers and retailers embraced the digital channel in waves. In the mid-90s, not everyone had a computer or Internet connection, and only a small portion of connected consumers trusted the idea of typing a credit card number into a browser window. For retailers, ecommerce meant a hefty investment in web design, development and new enterprise systems. It took years for the industry to mature and for "omnichannel selling" to fully saturate.

Adoption curves tend to look like this: slow skepticism at the start, a tipping point of utility, then a rapid acceleration that quickly feels inevitable.

Mobile commerce and social commerce followed that script almost exactly. Early branded apps felt clunky and optional. But as smartphones became universal, retailers caught up with adaptive websites and eventually "mobile-first" responsive design. Likewise, early experiments in 'shop the post' were

awkward until social platforms developed shoppable tags, native checkout, and ad formats that made discovery feel natural inside the feed. Retail media is still in that steep climb today, with ad budgets flooding into retailer platforms and closed-loop measurement accelerating its growth.

Seen in this light, the agent era isn't without precedent. Just as catalogs brought shopping into the home and mobile apps put it into our pockets, agents introduce a new mode of discovery: delegation. Right now, agentic shopping is still experimental and may feel risky or unfamiliar for both shoppers and retailers. Whether it reaches mainstream adoption will depend on the same factors that shaped earlier waves — consumer trust, proven utility, and the readiness of supporting infrastructure.

History suggests that new discovery channels often begin with skepticism and gradually normalize as the value becomes clear. Agentic shopping may follow a similar curve, or take a slower, more selective path. What's certain is that discovery keeps evolving, and its shape will follow shopper preferences, not industry plans.



What defines the "agentic era" of discovery?

When we talk about agents in commerce, we're not just describing smarter recommendation engines or chatbots.

An agent is an intelligent system that can understand a shopper's goal, break it down into steps, and take actions using the tools it has available — whether that means querying a catalog, checking delivery windows, or drawing on external knowledge to combine product data with practical, real-world context. In other words, agents don't just surface information, they operate on it.

That distinction is what makes this moment feel like a new era. Discovery is no longer limited to people typing queries, clicking filters, or scrolling feeds. Shoppers can describe an outcome — "find me a black cocktail dress under \$200 that ships by Friday" and delegate the search, comparison, and transaction to an agent.

For retailers, that means success is no longer defined only by how engaging your storefront looks to human visitors. It also depends on how clearly your catalog, inventory, and policies can be read, reasoned over, and acted on by machines.

Agents in practice: two sides of strategy

Agent adoption isn't a single move — it plays out on two fronts. On the one hand, retailers need to make their products and services legible to external agents: ChatGPT, Perplexity, and other LLM-driven

shopping tools that consumers increasingly turn to. These systems behave like new search engines, but instead of sending traffic to ten blue links, they surface structured answers and sometimes transact on behalf of the shopper. Retailers who want to be found must expose clean product data, availability, and fulfillment options in a way that external agents can parse, trust, and act on (more on this later).


At the same time, there's an opportunity to bring agents inside your own walls. On-site agents (or "owned" agents) can act like digital concierges, fielding natural language questions about a product, surfacing comparison points, or helping shoppers troubleshoot sizing and compatibility in real time. At the product detail page level, they can function as a Q&A layer that sits between static attributes and live customer service, answering the "does this work for me?" questions that drive conversion.

These embedded agents serve as continuous companions across the site experience, providing the value of an AI assistant without the friction of bouncing between your site and a third-party LLM. This enriches your site experience, turning it from a static catalog to a dynamic discovery destination — a place shoppers trust as a helpful, confidence-building companion in their buying decisions.

Agents in action: the new customer journey

The traditional customer journey has always been drawn as a funnel — awareness at the top, purchase at the bottom. But in an agentic world, journeys look less like funnels and more like networks. A shopper may begin by describing a goal to an off-site agent, like ChatGPT or Perplexity, which in turn queries retailer catalogs, compares options, and narrows the field. That same shopper might then shift to an onsite agent embedded in a brand's website to validate details, ask personalized questions, or confirm compatibility before purchasing.

This flow between off-site discovery and onsite confirmation creates a hybrid journey. External agents act like scouts — mapping the landscape, comparing features, and surfacing candidates — while onsite agents serve as concierges, deepening engagement and handling the last-mile of reassurance and conversion. The new, agentic customer journey is not a linear path but a series of hand-offs, where the quality of data and the consistency of answers determine whether the shopper completes checkout (or just 'checks out' to a competitor).




How agents execute a discovery mission

When a shopper tells an AI agent what they want, the agent doesn't type into a search box the way a human would. Instead, it connects through structured interfaces — APIs, Model Context Protocol (MCP)*, or other emerging standards — to pull data directly from retailer systems. In practice, this means the agent isn't struggling with keywords or filters, but with whether the retailer exposes data that is rich enough, fresh enough, and structured enough to reason over.

The agent's discovery mission looks less like browsing and more like planning: breaking down a goal into steps, calling the right endpoints, comparing results, and filtering against constraints like budget, delivery window, or compatibility. If all the agent finds is a shallow, static catalog, it's stuck. But when it can access detailed attributes, real-time inventory, and context that explains when and why a product fits, it can confidently recommend and even transact.


* Model Context Protocol (MCP) is an open standard that lets AI agents connect to external data and tools via secure, two-way interfaces — think a "USB-C for AI" that standardizes how apps expose context and actions to models.



Are consumers ready for agentic shopping?

The short answer: yes, but with caveats. Familiarity with GenAI is already widespread — 64% of consumers say they've tried ChatGPT, Bing Chat, or Google's Gemini (up from 51% last year and

only 29% two years ago). As more native shopping capabilities surface within LLM applications, like ChatGPT's existing in-chat product links and rumored Orders tab, delegating part or all of the shopping journey to an AI agent will feel more normalized.



64% have tried GenAI (2025) vs 51% (2024) → +13 pp YoY


On-site use cases in particular resonate with today's consumer. When asked what would improve their shopping experience, 43% said they want an AI-powered search bar or chatbot to explain what they need in their own words. Roughly a third also pointed to FAQ-style product page support and AI help comparing items as valuable. And when asked directly about willingness to engage with a true "AI agent" on a retail site, the majority leaned positive: 23% would definitely try it, 37% probably would. Only 13% said they'd likely refuse.

This interest is already translating into behavior. Nearly 38% of shoppers report using agents like Amazon's Rufus or Walmart's Sparky, meaning the concept isn't speculative — it's being tested at scale by market leaders.

Retailers Already Seeing Real Result

Retailers such as Belk have already unlocked measurable value from AI Shopping Agent. On its mobile app — which drives nearly 40% of total sales — shoppers who engage with the agent are converting at 2x the rate of standard search users.

As more retailers adopt on-site agentic features, online shoppers will in turn become more familiar and comfortable using them, and we can expect a similar growth trajectory to what we've seen with previous discovery eras: mobile, social and retail media.



58% are comfortable with AI shopping help (2025) vs 52% (2024) → +6 pp YoY

Preparing for off-site agent discovery

Off-site agents access data differently. Rather than navigating pages in real time, they rely on internal retrieval indices built from prior crawls, product feeds, and public datasets. In practice, this means they often “see” your catalog through cached snapshots of structured data rather than live interaction with your site. In some cases, they may deploy a headless browser or structured data fetcher to validate or supplement what they already know — especially for freshness checks or high-confidence shopping queries — but those sessions are the exception, not the rule. For most prompts, an agent’s view of your brand depends entirely on what’s already been indexed or exposed through machine-readable feeds and metadata.

Visibility now depends on structure. If your catalog data is incomplete, inconsistent, or difficult to parse, your products may never appear in an agent’s results — even if your site ranks well in traditional search. The challenge for retailers is no longer

Owned, on-site agents represent one side of the agentic equation; off-site or inbound agents — such as ChatGPT, Perplexity, and Gemini — represent the other. These external systems now influence how products are discovered, but they don’t behave like shoppers.

limited to optimizing the on-site journey; it’s ensuring that product data is readable, trustworthy, and accessible to discovery engines operating beyond the storefront.

In this environment, discovery doesn’t begin with a human search — it begins with a prompt. An LLM agent interprets that intent by scanning structured signals it can consume: product attributes, availability, pricing, and rules of sale. That makes your data layer as critical as your design layer. The real work now lies in the foundation that allows agents to read, reason, and act on your catalog.

Exposing accurate, machine-readable data isn’t a technical detail — it’s what makes your products visible in agentic journeys. Without it, you won’t even appear in an agent’s decision set. Investing in this “plumbing” today gives your brand the best chance to capture revenue as this new discovery channel takes shape.



Removing the discovery bottleneck

Without the right data pipes, agents will get stuck in the discovery process. Agents don't type queries into a search bar — they connect through APIs and emerging protocols like MCP and A2A to retrieve data directly.

- **MCP (Model Context Protocol):** An emerging open standard created by Anthropic that lets LLMs securely access external data and tools. Instead of scraping or guessing, an agent can call MCP endpoints to pull structured, real-time product information.
- **A2A (Agent-to-Agent communication):** The ability for different agents (such as a shopper's assistant and a retailer's discovery agent) to talk directly, exchange information, and even negotiate.
- **APIs (Application Programming Interfaces):** The existing "plumbing" of digital commerce. APIs are how systems share structured data, like inventory status, prices, or shipping options, in formats machines can reliably use.

This same infrastructure underpins new experiences like "Buy it in ChatGPT," where agents pull live product data directly from retailer systems rather than scraping static pages. It's an early glimpse of how agentic commerce is already emerging through real-time, machine-to-machine connections.

An agent's "crawl" isn't a human-style browse, it's a structured request for attributes, availability, pricing, or compatibility. It's machine-to-machine discovery.

What creates the bottleneck: technical challenges

Most retailers still publish static catalogs that were never designed for agentic discovery — nightly batch feeds or crawlers that describe SKUs at a single point in time. That model works when a human shopper is searching and filtering. But agents need to reason over live conditions, not stale snapshots.

A product record frozen in a spreadsheet or outdated API call can't answer whether an item is back in stock, whether a flash promotion is live, or whether a new release just dropped. Static catalogs also struggle with edge cases — bundles, substitutions, or configurable products — that agents must parse in order to satisfy real shopper goals.

Freshness is the new currency. Inventory, price changes, and promotional availability need to be broadcast in real time if agents are to recommend with confidence. Without these live signals, agents risk showing items that are unavailable, or worse, misrepresenting what a retailer can deliver. That erodes trust for both the consumer and the agent ecosystem.

Without addressing these technical challenges, retailers risk invisibility. Agents will default to sources that can provide clean attributes, context, personalization signals, and accurate inventory data in real time. In other words, the "agent-ready" retailer will win the recommendation battle by being not just discoverable, but dependable.

What the discovery data layer must deliver

To be “agent-ready” with clear plumbing, the data layer must do these things well:

1. Make products machine-readable

Agents want structured attributes that describe products in ways a model can reason over: size, materials, compatibility, use cases, even substitution rules. Ensure Schema markup (e.g. Product, Offer, and AggregateRating schemas) is implemented and validated to give agents a consistent, machine-readable view.

2. Carry context, not just specs

A shirt isn’t just “cotton, blue, size M” — it might be “casual summer wear for travel” or “office-appropriate under a blazer.” Contextual metadata, semantic tags, and even natural-language copy in titles and descriptions help agents map products to shopper goals and scenarios. Including FAQs, comparison phrasing, and “Best X for Y” boosts semantic context.

3. Keep data live

LLMs often pull first from trusted 3P channels like Google Shopping, Klarna, Meta, or Shop.app, so your feeds need to be clean and up-to-date there before agents can find you directly. Direct feed submissions, where supported, give agents the most reliable source of truth and widen your discovery net.

4. Capture human signals

Reviews, FAQs, and user-generated content are not fluff — they’re training material for agents. This is how models learn which products are durable, whether sizing runs small, or what “best budget option” really means. Feeding these natural-language signals into your catalog enriches what agents can say about your products.

5. Encode business rules

Margin floors, bundling constraints, geo restrictions, shipping cutoffs — these rules shape what a retailer can and cannot offer. If they’re only buried in human policy docs, agents will miss them. Exposing business logic in machine-readable form ensures agents don’t make promises you can’t keep.

6. Optimize your media

Agents don’t browse your front end like humans do, they extract and re-display images in their own UI containers (carousels, product cards, answer boxes). If media doesn’t fit their standards, it may be cropped, distorted, or dropped entirely. Prioritize square, high-res images (1024×1024+), keep aspect ratios consistent, compress file sizes under ~300KB, and use semantic filenames.

7. Make sure LLM bots can crawl your site

Check your robots.txt to confirm that discovery crawlers have access. At minimum, allow OpenAI’s OAI-SearchBot, but also make sure Googlebot, Bingbot, and CCBot (Common Crawl) are unblocked. These are the sources most LLMs draw from today, and blocking them could make your catalog invisible.

The personalization divide

Large-language-model agents like ChatGPT, Perplexity, or Gemini can surface products, compare options, and even summarize reviews — but they can't truly personalize. They don't know who the shopper is, what they've bought before, or how they interact with your brand. Their answers are generalized, drawn from aggregated data rather than individual context.

That's where on-site agents stand apart. Tools like AI Shopping Agent (ASA) and Product Insights Agent (PIA) can connect to first-party data — loyalty history, browsing behavior, preferences, and affinities — to deliver experiences that feel tailored and human. They can recognize returning shoppers, remember context across sessions, and refine results based on intent and behavior in real time.

Off-site agents may drive awareness and discovery, but on-site agents convert it into relevance and action. The future of retail personalization isn't about hoping external AI platforms get smarter — it's about ensuring your owned experiences already are.

What separates good agents from great ones

The best agents don't just compute for relevance — they learn from every interaction. Great agents continuously refine results using first-party behavioral signals such as click-throughs, dwell time, and repeat purchases. Over time, this feedback loop makes each session smarter than the last, turning static recommendations into adaptive personalization that reflects how real shoppers evolve.

Delivering that kind of adaptive intelligence requires a discovery platform built to unify, enrich, and act on data signals in real time.



🔍 how stain-resistant is this chair?



This chair is highly stain-resistant thanks to its performance fabric. The upholstery is made from a tightly woven polyester blend that naturally repels liquids, so spills don't soak in easily. Most everyday messes, like coffee, wine, or sauces, can be wiped clean with a damp cloth before they set. It's a great option for families or anyone who wants a low-maintenance dining setup.

Risk vs reward: is it time to act?

It's critical to evaluate agentic commerce not as a speculative bet, but as a series of incremental decisions: which foundations strengthen your business regardless of adoption, and which initiatives can be tested in low-risk ways while the channel matures.

Because standards like MCP and A2A are evolving, consumer adoption is uneven, and agentic tools themselves are changing quickly, retailers face a dilemma: move too fast and risk investing in a channel that stalls, or move too slow and risk being invisible if adoption tips. The best path forward is to balance preparation with pragmatism, guided by realistic expectations.

Setting expectations around agentic commerce

Realistically, agents won't replace human UX in the next year. Adoption will be gradual, and more relevant to some categories than others. Early use cases are narrowly focused — product detail page Q&A, guided discovery, replenishment, and simple comparison tasks. We're nowhere near full end-to-end autonomy yet.

Concerns around investing in agentic commerce today are valid. Retail has lived through more than one hype cycle that didn't deliver on its promises – and in today's economy, there's less room for greenfield experiments.

What if agents don't take off?

It is possible agentic commerce will remain niche, or even fade as a consumer behavior. If that happens, the risks for early movers are real:

- **Misallocated resources.** Overinvesting too early may tie up capital and teams in projects that don't return immediate value.
- **Opportunity costs.** Resources spent on agent readiness could detract from proven growth levers like core UX improvements, personalization and loyalty.
- **Protocol and vendor uncertainty.** Betting on a specific standard or integration too soon risks sunk costs if protocols change or adoption flattens.
- **Operational overhead.** Preparing and maintaining personalization, feeds, and permissions for agents creates complexity that may not deliver incremental traffic or conversion if consumers don't embrace agentic shopping.

The important nuance with AI readiness is that most of the prep isn't wasted even if agents stall. Structured product data, fresher inventory signals, and richer personalization pipelines strengthen core ecommerce, improve search and discovery, and power existing channels like marketplaces and retail media. The greatest risk may not be wasted effort, but the perception of wasted effort if agentic commerce doesn't scale as expected.

The risks of sidelining AI readiness

Underpreparation comes with its own set of opportunity costs. It rarely shows up as one big failure, but rather as a series of smaller choices — how long to wait, what to ignore, and where to over- or under-invest. These are a few of the most common missteps retailers can make in this time of transition.

Waiting for standards to finalize

It's tempting to hold off until agent protocols are finalized. But while MCP and A2A are still emerging, agentic discovery doesn't pause. External agents are already scraping, parsing, and experimenting with data. Many consumers are already using agents to shop. Waiting means failing to meet the needs of early adopters and ceding visibility to competitors. And the longer data and feeds remain unstructured, the harder it is to retrofit later.



Treating agents as external threats rather than channels

Some will view agents as competitors — gatekeepers that take traffic away from their storefront. But history suggests that resisting new intermediaries rarely works. Retailers who refused to list on Amazon lost reach; those who saw mobile as a distraction fell behind. Agents aren't adversaries, they're channels. The sooner retailers treat them as distribution, the sooner they can shape how their products are represented.

Over-focusing on storefront UX while neglecting data readiness

Agents don't evaluate design — they parse structured data. While LLM crawlers can browse storefronts and unpack JavaScript, they prefer to consume your site through lean bot templates: this means server-rendered HTML with key info (price, availability, reviews, returns) delivered as structured data.

Beautiful PDPs and clever navigation still matter, but in an agentic world, your "user" may not be a human clicking through menus. The misstep isn't investing in UX; it's over-indexing on UX at the expense of the underlying plumbing.

Assuming agent optimization is just an SEO extension

Agents don't return search results in the traditional sense. They reason over structured inputs, filter by constraints, and sometimes transact on behalf of the shopper. Treating agent optimization as a variant of SEO risks under-preparing for a deeper shift.

The practical risk is invisibility: if your data isn't legible to an agent, your products won't appear in its shortlist, no matter how strong your SEO rankings are today. The lesson isn't to abandon SEO, but to recognize its limits as agents change the discovery pipeline.

Building buy-in for agent-readiness

The challenge for most leaders isn't recognizing the potential, but building internal alignment around what to do next. Teams across merchandising, IT, finance, and compliance may see agent-readiness as either too speculative or too resource-intensive. Overcoming that requires reframing the conversation.

The most effective case for agent readiness isn't about chasing a hype cycle — it's about strengthening foundations that pay off regardless. Cleaner attribution improves search and personalization today. Real-time feeds reduce cart abandonment now. Structured data boosts SEO and marketplace performance as much as it helps agents. In other words, the ROI story is not "bet on agents" but "reduce technical debt and unlock flexibility."

Buy-in also means setting realistic expectations. Leaders should present agent-readiness as a phased roadmap, with low-risk pilots and incremental wins rather than sweeping transformations. Pilot a PDP Q&A agent and measure conversion lift. Enable structured feeds and validate they improve both SEO and agent visibility. Each experiment builds confidence and keeps momentum without overcommitting.

Framed this way, agent readiness isn't a gamble — it's disciplined preparation for multiple possible futures. And that's the bridge to the roadmap: how to sequence investments, stage pilots, and align teams to move forward with confidence.



From Readiness to Results: Taking the First Steps with On-Site Agents

On-site agents such as AI Shopping Agents (ASA) and Product Insights Agents (PIA) offer an immediate way to learn how agentic interactions behave in the real world. They help teams understand how customers express intent, what questions surface friction, and how structured data influences results.

Concerns around investing in agentic commerce today are valid. Retail has lived through more than one hype cycle that didn't deliver on its promises – and in today's economy, there's less room for unproven experiments. But because on-site discovery agents operate within a retailer's owned environment, they're low-risk, measurable, and entirely under brand control.

Build vs. Buy: Finding the Right Balance

For most retailers, there's no benefit to building an LLM-based discovery system in-house, from scratch. Developing proprietary models requires large-scale data, constant retraining, and dedicated engineering talent — a costly and slow path with limited competitive payoff. What matters isn't

Retailers don't need to wait for new protocols or external standards to begin experimenting with agents. The most practical place to start is on your own site — where you already control the data, design, and experience.

owning the model; it's orchestrating intelligence around your data in ways that are secure, adaptable, and measurable.

Partnering with a proven agentic platform delivers that advantage out of the box. These solutions are maintained by discovery experts, continuously improved to reflect new AI capabilities and shopper behavior, and come with built-in safeguards that minimize operational and economic risk.

Retailers can further enrich the agent with their own product data, documentation, and brand knowledge, creating a tailored experience that feels native while benefiting from a system that's already trained, maintained, and production-ready.

A "buy-and-extend" approach typically delivers the best balance: buy a turnkey agentic commerce foundation for core capabilities like reasoning, retrieval, and conversation — then extend it with your unique product context, tone, and business rules. You gain the benefit of ongoing product investment, a faster time to market, and ultimately, a faster path to value — all without the overhead of building or maintaining core AI infrastructure.

How to Prove It: Run a Measured Pilot

Pilot with purpose. Structure an A/B test comparing an agentic PDP or guided-discovery flow against a control group. Focus on a handful of key metrics:

- Conversion rate or add-to-cart lift
- Time on page and session engagement
- Reduction in service tickets or live-chat volume
- Customer satisfaction or NPS changes

Within weeks, you'll have quantitative evidence and qualitative learning about where agents add value — data you can use to refine design, train content, and justify broader investment.

Looking Ahead

While on-site agents are the immediate frontier, off-site discovery is beginning to take shape through standards such as ChatGPT's Model Context Protocol (MCP) and other open interfaces. Forward-thinking retailers are already evaluating how their structured data and APIs could connect to these emerging ecosystems once consumer adoption scales. The best strategy now is to master on-site agents first — building fluency, trust, and measurable outcomes — so when off-site channels mature, you're ready to extend that same intelligence beyond your own walls.

But the impact of agentic AI won't stop at shopper assistance. As these systems evolve, agents themselves may begin to act as active participants in commerce — comparing options, transacting, and influencing demand alongside human shoppers. That broader shift will ripple across the entire retail value chain.

As David Dorf, Head of Retail Industry Solutions at AWS recently put it:

“Agentic AI in retail isn't just about shoppers delegating their purchases - it's going to reshape everything from merchandising to supply chain operations. If we reach a future where both humans and agents are 'shopping' side by side, retailers will need to rethink how ads, promotions, and discovery experiences function. This isn't the death of the website, but another major evolution in ecommerce. Just as we adapted from catalogs to stores, from web to mobile, and then to social commerce, we'll need to adapt again, ensuring our data, infrastructure, and strategies are ready for agents as the new gatekeepers.”



His point underscores the central takeaway: the future of retail isn't about replacing human shoppers — it's about preparing your business to engage both humans and agents with equal fluency.

Conclusion

At its core, agentic commerce isn't about technology — it's about serving customers better. Agents emerge because consumers want easier ways to research products, compare options, and make decisions. They want to delegate the tedious parts of shopping while maintaining control over the choices that matter to them.

When a customer-centric perspective guides how you approach agent readiness, the question isn't whether agents threaten your business model, but whether they help your customers accomplish their goals more effectively. If an on-site agent can answer product questions faster than browsing FAQ pages, that's better customer service. If it can surface the right product more quickly than navigating category pages, that's better discovery. If it can handle routine replenishment while alerting customers to new options, that's better convenience.

The retailers who succeed with agents will be those who see them as customer enablement tools rather than competitive threats. They'll focus on making their products and services more accessible to and helpful through agent interactions rather than trying to prevent agent adoption. They'll use agents to reduce friction, provide better answers, and create more personalized experiences.

This approach transforms agent readiness from a defensive strategy to an offensive capability. Instead of asking "How do we protect our traffic from agents?" the question becomes "How do we use agents to serve customers better than competitors can?" Instead of viewing delegation as a loss of control, it becomes an opportunity to be more helpful at the moment customers need assistance.

The future of commerce belongs to retailers who put customer needs first, regardless of which technologies enable those experiences. Agents are simply the latest tool for delivering value. The retailers who embrace this perspective will find ways to thrive in the agentic era and beyond.

If you're ready to embark on agentic transformation

Constructor can help. AI Shopping Agent (ASA) and Product Insights Agent (PIA) are turnkey solutions that are ready to integrate with your storefront. Together, they provide seamless agentic assistance across the shopping journey, from search and browse to product detail pages — creating a continuous thread of understanding from the first search to the final decision.

Unlike other conversational AI plugins on the market, Constructor's agentic solutions leverage behavioral data for truly personalized and optimized reasoning and recommendations. By grounding every interaction in real shopper behavior, Constructor's agentic AI turns intelligence into impact — driving higher engagement, stronger sales, and lasting loyalty.

Make us prove it

You can validate the real performance lift of AI Shopping Agent and Product Insights Agent. [Get a demo](#) to see how agents can work on your website.