

Who Controls the Agent Controls the Customer

An executive briefing on agentic commerce for retail, consumer brand and private equity leaders

393%

AI traffic growth to US
retail sites, Q1 2026

10–25%

of US e-commerce
agent-led by 2030

1

recommendation replaces
the ranked shelf

AI agents are becoming a new layer between the consumer and checkout

Whoever runs that layer gains leverage over discovery, ranking and the customer relationship. The shift is early and uneven — but the rules are being set now, in courtrooms and in platform terms of service.

The shelf becomes an answer

An agent often returns one recommendation, or reorders without showing alternatives. Visual merchandising, packaging and impulse placement lose leverage; structured data and verifiable claims gain it.

Brand strength matters more — if pre-loaded

Brands consumers name explicitly win the “default” instruction. Undifferentiated mid-tier brands face commoditization.

Retailers split between open and closed

Amazon is blocking third-party agents; Walmart, Target and Shopify are experimenting with them. Neither posture is free.

Amazon’s blockade is strained long term

A federal appeals court has weakened its strongest legal tool. The likely endgame is licensed, identified agent access on Amazon’s terms.

Agent exposure is now a diligence item


Retail media revenue, marketplace take rates and mid-tier brand premiums are the most exposed value pools.

The window

Preparation is cheap; invisibility compounds. The next 12 months decide which brands are reachable, readable and named when an agent is asked to choose.

From shelf to answer


The “shelf” moves from a page a human scans to a ranked answer an agent produces. Three dynamics follow.



The more a purchase is routine and specifiable, the faster the agent moves from advising to executing.

How big, how fast

Agent commerce is small today but growing at triple-digit rates. Credible forecasts put it at roughly 10–25% of US e-commerce by 2030.



Where it stands now

AI-source traffic to US retail sites grew 393% year over year in Q1 2026, after rising 693% over the 2025 holiday season (Adobe). Absolute share is still undisclosed — the base remains small.

The quality flipped

In March 2026 AI-referred visits converted 42% better than non-AI traffic — versus 38% worse a year earlier — with revenue per visit 37% higher (TechCrunch).

The readiness gap

Retail homepages average 75% on Adobe’s AI visibility checker: about a quarter of their content can’t be read by the models sending this traffic.

The spread reflects definitions, not disagreement

Morgan Stanley and Bain midpoints both land near 15–17% of US e-commerce. The binding constraint is trust.

Source	2030 forecast	What it counts
Morgan Stanley	\$190B–\$385B · 10–20% of e-commerce	Stricter: the agent takes meaningful autonomous action
Bain	\$300B–\$500B · 15–25% of e-commerce	Initiated, influenced or completed by agents; excludes AI-assisted search alone
McKinsey / ICSC	Up to \$1T in B2C retail	Broad: all agent-orchestrated revenue

Why the numbers differ

Each house draws the boundary of “agentic” in a different place — autonomous purchase, agent-influenced purchase, or agent-orchestrated revenue. Compare definitions before comparing dollars.

The constraint is trust, not capability

Most shoppers still aren’t comfortable letting AI complete a purchase end to end (Bain). 58% would cap agent spending on electronics at \$100, and half of Americans worry about payment security (Global Payments).

No single agent controls the consumer yet

Control is splitting between a handful of consumer agents, the commerce protocols they back, and the payment networks underneath.

Player	Agent and approach	Checkout posture
Meta	Muse, launched Sep 8, 2026; topped Apple’s US App Store	Buys across the web — Stripe Link at 1M+ businesses, Shop Pay, PayPal, Expedia, Instacart. Blocked by Amazon
Amazon	Alexa for Shopping replaced Rufus in May 2026; Rufus reached 300M+ users in 2025	Closed to outside agents — yet its own Buy for Me feature purchases from non-Amazon retailers
Google	Search AI Mode, Gemini, and a cross-merchant Universal Cart	In-surface checkout via UCP; the retailer stays seller of record
OpenAI	ChatGPT shopping, now discovery-first	Retired in-chat Instant Checkout in March 2026; purchases route to merchant apps and sites
Perplexity	Comet browser agent	Shops Amazon via the user’s own session; litigation ongoing

Sources: Forkast, Axios, Prime Team, Efficiently Connected, Stellagent

The protocols and rails are layering, not competing

That lowers integration cost for merchants — Shopify merchants are already reachable by default.

DISCOVERY, CART, CHECKOUT

Google UCP

Covers the full path. Amazon, Meta, Microsoft, Salesforce and Stripe sit on its tech council. Shopify made it self-serve for any agent developer from June 17, 2026.

AGENT-TO-MERCHANT CHECKOUT

OpenAI + Stripe ACP

Handles the transaction handshake. The protocol outlived OpenAI's own checkout product — a signal that rails matter more than any one agent.

PAYMENT AUTHORIZATION

AP2

Proves a human authorized the agent's payment. Launched with 60+ partners including Mastercard, PayPal and American Express.

CARD NETWORKS

Visa Intelligent Commerce · Mastercard Agent Pay

Scoped tokens mean agents never hold card numbers. Visa's single integration accepts agent payments across four protocols.

The working model today: discover in AI, buy on site.

Walmart found ChatGPT checkout converted roughly 3x worse than its own site. Reachability comes first; in-surface checkout follows the trust curve.

For consumer brands: named demand is the moat

Brands consumers name explicitly gain power; brands that relied on shelf presence and persuasion lose it.

Brand equity is worth more — if it is pre-loaded

Emotional, identity-driven and category-defining brands survive disintermediation. “Good enough, well-merchandised” mid-tier brands are the most exposed.

Product data becomes marketing

Agents read attributes, specs, ingredients, sizing, compatibility, availability, return terms and reviews. Incomplete or inconsistent data is the new out-of-stock.

Proof beats adjectives

“Premium quality” means nothing to an agent. Certifications, test results, warranty terms and review sentiment do.

Open-web reputation is a ranking input

Agents synthesize reviews, forums, press and expert content. Brands can influence that corpus — but they cannot buy it.

Retail media gets repriced

Much of today’s spend buys placement in front of human eyes. Expect agent platforms to sell sponsored ranking — and expect disclosure fights over it.

The test for every brand team

If a shopper asked an agent for your category today, would it name you — and could it defend the choice with data you published?

For retailers: compete on what an agent can evaluate

When an agent does the browsing and comparing, retailers compete on price, availability, speed and returns — plus experiences agents cannot replicate.

Marketplaces and aggregators face the most direct threat

Much of their value is discovery and comparison. A cross-retailer agent turns them into one fulfillment option among many.

Operational excellence becomes the differentiator

Inventory accuracy, delivery promise, return friction and price are what an agent can evaluate. Curation and experience must become machine-legible — or pull customers into owned channels.

Every retailer must choose a stance

Open APIs capture demand wherever it originates; restriction protects the relationship and the ad business. Most retailers cannot afford to be invisible — only the largest can afford to be closed.


Physical and experiential retail gains relative value

Stores become where discovery, trial and experience happen. Destination, curated and place-based retail is least substitutable, because the purchase is tied to a moment or a story.

Only the largest can afford to be closed. Everyone else has to be findable.

Routine purchases shift first; taste and identity shift last

Consumer willingness to delegate a purchase to an AI agent, by category.



Sources: Checkout.com (delegation willingness); Martechvibe / Zeta Global (electronics research, named by 33% of AI shoppers). Groceries and CPG already lead AI-driven purchases (Morgan Stanley).

Category	Pace of shift	Agent role
Grocery, household, CPG replenishment	First (now–2027)	Executes: auto-reorder and substitution
Consumer electronics, appliances	Early (2026–2028)	Researches and compares; human approves
Travel, tickets, bookings	Early (2026–2028)	Books within set parameters
Beauty, personal care	Middle (2027–2029)	Recommends; replenishes repeat SKUs
Apparel, footwear	Middle to late	Discovers and filters; fit remains a barrier
Furniture, home decor	Late	Inspiration and shortlist
Luxury, gifting, experiential retail	Last	Discovery aid at most

What each pace band means for strategy

The more a purchase is routine and specifiable, the faster the agent takes over execution.

Shifting first

Grocery · Household · CPG · Travel · Electronics

Win the default instruction — brand is re-competed on price and spec every cycle. Specs, reviews and compatibility data decide the shortlist; loyalty perks must be agent-readable.

Shifting in the middle

Beauty and personal care · Apparel and footwear

Ingredient and claim data plus review sentiment drive recommendation; replenishment is the lock-in point. Sizing data and return terms become ranking inputs.

Shifting last


Furniture and home decor · Luxury, gifting, experiential and destination retail

Visual and experiential selling remains decisive. Brand and place carry the purchase — these are the least substitutable categories.

Where the agent stays a discovery tool, the brand keeps the relationship. Where it executes, the standing instruction is the asset.

Amazon's agent blockade

Sustainable for now — but its legal footing has weakened and the economics shift as agent checkout grows.



Why the block works today

Consumers mostly use agents to research, not buy, so Amazon loses little near-term business. Scale and Prime mean agents need Amazon more than Amazon needs any one agent — and it has a credible security rationale after alleging Muse captured credentials.

What erodes it

The legal route narrowed when the Ninth Circuit found an injunction would impair consumer choice; Amazon has fallen back on contract terms. Competitors are opening up, enforcement is an arms race, and a dominant retailer blocking tools consumers choose invites antitrust scrutiny.

Likely endgame: not “no agents,” but “identified agents, through our channels, on our terms.” Treat the block as a negotiating posture, not permanent policy.

Agent exposure belongs in diligence now

Agentic commerce is a hold-period risk for some consumer assets and a value-creation lever for others.

Asset profile	Exposure	What to test in diligence
Mid-tier brands in replenishable categories	High — price and spec re-competed every cycle	Share of revenue from explicit brand search vs. generic; the real repeat-purchase driver
Marketplaces, aggregators, comparison sites	High — discovery value intermediated	Share of traffic from search and discovery; direct and app traffic trend
Retail media-dependent retailers	Medium-high — inventory assumes human browsing	Retail media share of EBITDA; sensitivity if 20% of sessions become agent-led
Operationally strong retailers and 3PLs	Positive — wins on speed, price, returns	Inventory accuracy, delivery promise, API readiness
Iconic, experiential and destination brands	Low to positive	Strength of named demand; owned community and membership

Four value-creation plays

Low-cost moves that protect the exit narrative as buyers start asking how an asset fares in agent-led discovery.

01

Product data as a 100-day priority

Cleaning and structuring catalog data is cheap relative to its impact on agent visibility. Name an executive owner.

02

Agent-channel readiness

Support emerging agent protocols and payment rails so portfolio companies are reachable wherever agents transact.

03

Owned relationships

Subscriptions, memberships, apps and services create demand agents cannot intermediate — and support exit multiples.

04

The exit narrative

Buyers will ask how the asset performs in agent-led discovery. Having an answer, with evidence, protects valuation.

The through-line

Every play buys the same thing: the ability to be found, read and named by a machine that is making the shortlist on the consumer's behalf. None of it requires betting on which agent wins.

Where the thesis may be overstated

The shift is real but uneven. Plan for gradual change in commoditized categories and slow change elsewhere.

Delegation varies sharply by category

Consumers hand off paper towels and batteries far more readily than fashion, gifts or furniture, where taste and identity drive the purchase.

Trust is still early

Agents make mistakes, and consumers are mostly using them to research rather than to buy. Adoption tracks confidence, not capability.

The rules are unsettled

Disclosure of paid ranking, liability for bad purchases and data access are open questions for courts and regulators.

Gatekeepers may fragment

Several competing agents — Meta, OpenAI, Google, Amazon, banks, retailers — would dilute any single platform's control.

Why it still warrants action now

None of these caveats reverses the direction of travel — they change its speed. The preparation is cheap, reversible and useful even if agents plateau; invisibility is not.

Six moves for the next 12 months

The cost of preparing is low and the cost of being invisible compounds.

1

Treat product data as a core asset

Complete, structured and consistent across every feed — owned by a named executive.

2

Get onto the rails

Support emerging agent-commerce protocols and payment standards rather than betting they won't matter.

3

Build relationships agents can't intermediate

Communities, memberships, owned apps, subscriptions, services and experiences.

4

Win the default instruction

First purchase, onboarding and replenishment matter most — agents repeat what worked.

5

Audit how agents describe you

Query the major assistants about your category and brand every quarter, as you once audited search rankings.

6

Set your agent stance deliberately


Open, partial or closed — with a clear view of what data and margin you trade for reach. Brands dependent on Amazon should keep strong listings on open channels.

Move on all six in parallel: none depends on knowing which agent wins.

THE BOTTOM LINE

The contest is over which instruction gets written — and who gets to write it.

Agent adoption will be gradual in commoditized categories and slower elsewhere, but the rules, rails and default instructions are being set in the next 12 months. Being reachable, readable and named costs little today; becoming invisible compounds.



THE GRAYSON COMPANY

Retail and consumer industry advisors · New York · Los Angeles · thegraysoncompany.com

Full source list available
on request