

The Future of Shopping with Agentic AI

The retail landscape is undergoing its most significant structural shift since the advent of mobile commerce: the transition from Generative AI (content creators) to Agentic AI (autonomous actors). While 2023-2024 were defined by chatbots that could describe products, 2025-2026 marks the era of "Actionbots" that can buy them. Agentic AI represents a fundamental inversion of the e-commerce model. Instead of humans browsing catalogs curated for eyeballs, autonomous software agents now navigate digital aisles on behalf of consumers—negotiating prices, managing returns, and executing complex multi-step workflows.

This comprehensive research document analyzes the rise of "Machine Customers" (Custobots), the disruption of traditional SEO by "Agent Optimization," and the technical architecture enabling this autonomy. With agentic AI projected to redefine the trajectory of global e-commerce spending, retailers who view this merely as a "chatbot upgrade" risk obsolescence. The primary customer of the future is not a human, but an algorithm acting on their behalf.

Rick Spair | DX Today | January 2026

From "Search" to "Solved": The Commerce Revolution

The Old Model

For two decades, e-commerce has relied on a pull-based model where users search, filter, click, and check out. This human-in-the-loop friction has been the primary bottleneck in digital commerce.

- Manual product research
- Repetitive comparison shopping
- Time-consuming checkout processes
- Limited decision optimization

Unlike Large Language Models which are probabilistic text predictors, AI Agents are designed for agency. They possess perception to read dynamic web content, reasoning to break complex goals into sub-tasks, tool use to call payment and booking APIs, and memory to persist user preferences. We are moving from a web of information to a web of action, where the friction between desire and fulfillment approaches zero.

The Agent-Driven Future

Agentic AI removes bottlenecks by introducing Goal-Oriented Architectures that fundamentally transform how commerce operates in the digital age.

- Autonomous product discovery
- Intelligent price negotiation
- Seamless transaction execution
- Continuous optimization learning

The Four Pillars of Agentic AI



Perception

Advanced ability to read and interpret dynamic web content, inventory databases, and real-time pricing information across multiple platforms simultaneously.



Reasoning

Sophisticated planning capabilities to decompose complex goals like "Plan a ski trip for under \$2k" into actionable sub-tasks and decision trees.



Tool Use

Seamless integration with payment gateways, booking systems, inventory management, and logistics APIs to execute transactions autonomously.



Memory

Persistent storage of user preferences, past interactions, purchasing patterns, and contextual understanding that improves over time.

These four capabilities combine to create a fundamentally new type of digital entity: one that can perceive opportunities, reason about optimal solutions, execute complex workflows, and learn from every interaction. This represents a quantum leap beyond the passive recommendation engines and content generation tools that defined earlier waves of retail AI.

Three Waves of AI in Retail Evolution



Wave 1: Predictive AI (2010-2022)

Function: Recommendation engines, demand forecasting, and inventory optimization. These systems analyzed patterns to predict what customers might want.

Limitation: Entirely passive. Required human action to operationalize insights. Could suggest but never execute.

Wave 2: Generative AI (2022-2024)

Function: Content creation including automated product descriptions, personalized marketing copy, and conversational search interfaces.

Limitation: Hallucinations and lack of execution capability. Could write compelling emails but not send them; could suggest products but not purchase them.

Wave 3: Agentic AI (2025-Present)

Function: Autonomous execution and decision-making. Agents operate as digital employees or shopping concierges with full transaction authority.

The Shift: Integration of browser, wallet, and decision-maker into unified autonomous systems capable of end-to-end commerce workflows.

Market Dynamics: The Custobot Economy



Explosive Growth Trajectory

The economics of delegation are driving unprecedented adoption rates across both consumer and enterprise segments. Market research reveals a fundamental shift in how transactions will be initiated and completed.

Gartner projects that **40% of enterprise applications** will include task-specific AI agents by the end of 2026. This represents a massive acceleration from less than 5% penetration in early 2024.

Consumer willingness to delegate low-reward tasks is particularly striking. Early adopters report time savings of 60-80% on routine purchasing decisions, with satisfaction rates remaining consistently high.

40%

Enterprise Adoption

Applications with embedded AI agents by end of 2026

\$2.1T

Influenced Spending

Global e-commerce volume impacted by agentic AI

73%

Consumer Interest

Willingness to use AI agents for routine purchases

The Agentic Storefront Revolution

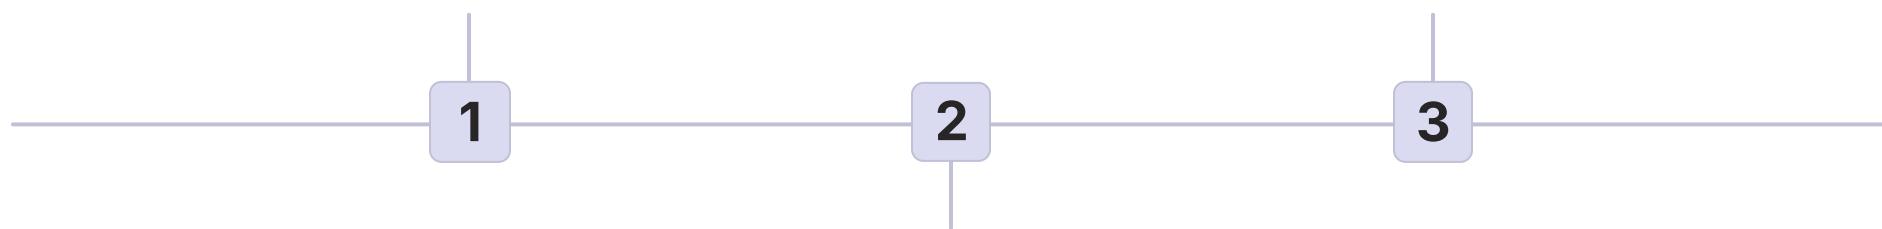
Traditional "Headless Commerce" architectures are rapidly evolving into "Agentic Commerce" platforms. This transition represents not merely a technical upgrade but a fundamental reimaging of how brands connect with customers—or more accurately, with the AI agents representing those customers.

Traditional Model

Brands optimize storefronts for human eyeballs through Google SEO, social media advertising, and visually appealing interfaces designed to capture attention and drive clicks.

Agentic Future

Brands optimize primarily for Agent Discovery. AI agents like OpenAI's Operator or Amazon's Rufus become the primary interface, with human experiences becoming secondary.



Transitional Phase

Hybrid approaches emerge where brands maintain human-facing experiences while building parallel agent-accessible APIs and structured data feeds.

The implications are profound. If an AI agent cannot parse a retailer's product catalog, negotiate terms, or execute transactions programmatically, that retailer becomes effectively invisible in the agentic economy. This creates an existential imperative for digital transformation that goes far beyond traditional e-commerce optimization.

Agent Optimization: The New SEO

Traditional SEO

- Keyword density and placement
- Backlink authority building
- Meta tags and descriptions
- Page load speed optimization
- Mobile responsiveness
- Content freshness signals

Optimized for human search behavior patterns and Google's ranking algorithms.

Agent Optimization

- Structured data schemas (JSON-LD)
- API accessibility and documentation
- Machine-readable product attributes
- Programmatic transaction workflows
- Real-time inventory signals
- Agent-friendly authentication

Optimized for algorithmic interpretation and autonomous execution capabilities.

The shift from Search Engine Optimization to Agent Optimization represents a complete paradigm change in digital discoverability. Where SEO focused on gaming Google's algorithms to appear in the top ten results for human eyeballs, Agent Optimization requires fundamental structural changes to how product information is exposed, priced, and transacted. Retailers must think less like marketers and more like API providers, ensuring their entire catalog is machine-readable, real-time accurate, and programmatically accessible.

Technical Architecture: Enabling Autonomy

01

Goal Interpretation

Natural language processing systems convert user intentions ("Find the best noise-canceling headphones under \$300") into structured computational goals with defined parameters and constraints.

03

Constraint Evaluation

Advanced filtering algorithms evaluate options against user preferences, budget limits, delivery timelines, and historical satisfaction data to narrow possibilities.

05

Transaction Execution

Once optimal choices are identified, agents autonomously execute purchases through secure API calls, managing authentication, payment, and confirmation workflows.

02

Multi-Source Discovery

Agents query multiple data sources simultaneously—retailer APIs, price aggregators, review databases, and inventory systems—to build comprehensive option sets.

04

Optimization Reasoning

Multi-objective optimization balances competing factors like price, quality, delivery speed, and brand reputation using weighted scoring models.

06

Continuous Learning

Outcome tracking and feedback loops enable agents to refine future decisions based on actual user satisfaction with completed transactions.

The Operator Platform: OpenAI's Entry

OpenAI's Operator platform represents one of the most significant commercial implementations of agentic AI for consumer commerce. Launched in early 2025, Operator combines advanced vision models with web automation capabilities to create a truly autonomous shopping assistant that can navigate complex retail environments.

Core Capabilities

- Visual understanding of web interfaces
- Form filling and checkout automation
- Price comparison across platforms
- Deal hunting and coupon application
- Cart abandonment recovery
- Multi-step purchase orchestration

Technical Innovation

Operator leverages Computer Use API (CUA) technology to interact with websites as a human would—clicking buttons, filling forms, and navigating menus—but at machine speed and with perfect consistency.

The system maintains persistent sessions, remembers user preferences across shopping trips, and can execute complex multi-vendor purchases in a single workflow.

What distinguishes Operator from earlier shopping assistants is its ability to handle edge cases and exceptions. When a desired product is out of stock, it can automatically identify substitutes. When checkout processes fail, it can retry with alternative payment methods. This resilience transforms it from a convenience feature into a genuinely autonomous purchasing agent.

Amazon Rufus: Retail Giant's Response

Conversational Commerce

Rufus serves as Amazon's entry into agentic shopping, initially focused on conversational product discovery and recommendations within the Amazon ecosystem.

While Operator focuses on cross-platform autonomy, Rufus represents a walled-garden approach that keeps users within Amazon's ecosystem. This strategic difference reflects fundamentally different visions for agentic commerce: open versus closed, cross-platform versus proprietary. Amazon's advantage lies in controlling the entire transaction stack from discovery through delivery, enabling seamless agent experiences that external platforms struggle to match. However, this creates dependency that some consumers and regulators view with concern.

Voice Integration

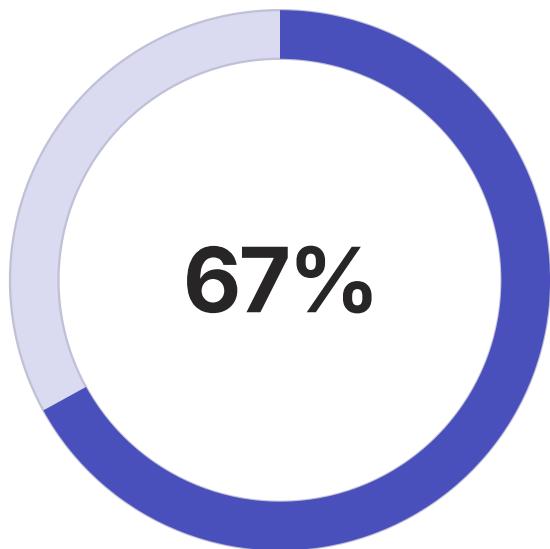
Deep integration with Alexa enables voice-first shopping experiences where users can delegate entire purchase workflows through natural conversation.

Personalization Engine

Leverages Amazon's vast purchase history data to provide hyper-personalized recommendations that improve with every interaction and transaction.

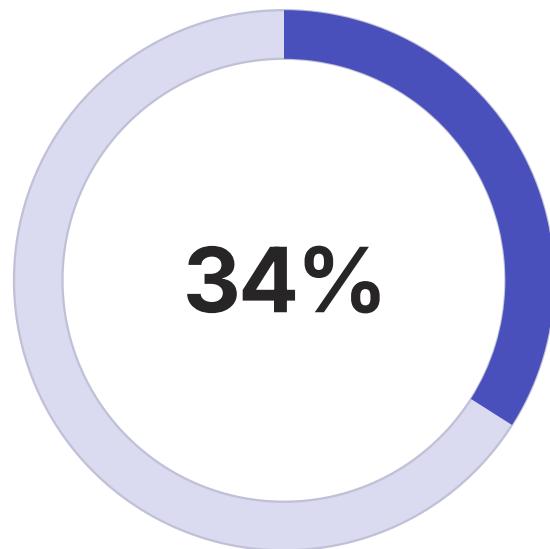
Consumer Trust and the Delegation Dilemma

The success of agentic AI in commerce hinges not on technical capability but on psychological acceptance. Consumers must overcome deeply ingrained habits of manual control and develop comfort with algorithmic decision-making on consequential purchases. This trust gap represents the primary barrier to widespread adoption.



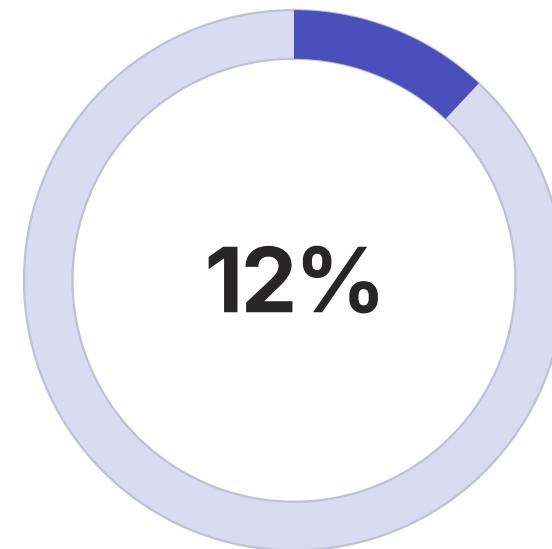
Comfortable with Basics

Consumers willing to delegate routine, low-stakes purchases like household staples and subscriptions to AI agents.



Open to Complex Tasks

Users ready to trust agents with multi-step workflows like vacation planning or electronics purchases exceeding \$500.



Full Delegation

Early adopters willing to grant broad autonomy across all purchase categories with minimal oversight requirements.

Trust building follows predictable patterns. Users typically start with low-stakes experiments—allowing an agent to reorder paper towels or find the cheapest gas station. Positive experiences create confidence that expands delegation boundaries. However, a single significant failure can trigger complete withdrawal, making reliability paramount. The industry must solve not just for average performance but for edge case handling and graceful failure modes that maintain user confidence even when things go wrong.

The Privacy Paradox

Data Requirements

Effective agentic AI requires unprecedented access to personal information to function optimally:

- Complete purchase history analysis
- Payment method credentials
- Delivery address preferences
- Household composition and needs
- Budget constraints and priorities
- Brand loyalties and aversions
- Quality versus price trade-off preferences

The privacy paradox emerges: consumers want intelligent, personalized agent behavior but resist the data collection that enables it. This creates a fundamental tension that the industry must navigate carefully. Solutions include federated learning architectures that keep personal data local, differential privacy techniques that obscure individual records, and transparent consent frameworks that give users granular control. However, the economic incentives favor centralized data collection, creating inherent conflicts between user interests and platform profitability.

Consumer Concerns

This data intensity triggers legitimate privacy anxieties:

- Surveillance capitalism fears
- Data breach vulnerability
- Manipulation through targeting
- Loss of purchase anonymity
- Corporate control over access
- Algorithmic bias concerns
- Difficulty auditing decisions

Regulatory Landscape and Compliance

1

Consumer Protection Mandates

Regulators are extending existing consumer protection laws to autonomous agents. Questions include: Who is liable when an agent makes an unauthorized purchase? How do cooling-off periods apply to automated transactions? What constitutes adequate disclosure of agent involvement?

2

Data Governance Requirements

GDPR, CCPA, and emerging frameworks impose strict requirements on how agent platforms collect, store, and utilize personal data. Right-to-explanation mandates force platforms to provide human-understandable rationales for agent decisions.

3

Anti-Manipulation Standards

Regulators are developing guardrails against manipulative agent design—prohibiting dark patterns that exploit behavioral biases or steer users toward higher-margin products against their stated preferences.

4

Interoperability Requirements

Some jurisdictions are considering data portability mandates that would allow users to move their agent training data between platforms, preventing lock-in and promoting competition among agent providers.

The regulatory environment remains highly fluid and fragmented across jurisdictions. Platforms operating globally must navigate a complex patchwork of requirements while anticipating future restrictions. The industry's self-regulatory efforts will determine whether governments adopt light-touch frameworks or impose heavy-handed controls that could stifle innovation.

Business Model Disruption

Agentic AI fundamentally destabilizes existing e-commerce business models by shifting power from merchants to consumers (via their agents). This rebalancing threatens profit margins, marketing effectiveness, and competitive positioning across the retail ecosystem.

Advertising Obsolescence

If agents make purchase decisions based on objective criteria rather than persuasive marketing, the \$600 billion digital advertising industry faces existential threat. Brands accustomed to buying attention may find that agent-mediated commerce rewards product quality and value over promotional spending.

Margin Compression

Agents with perfect price transparency and comparison capabilities will drive fierce competition on price, compressing retail margins toward commodity levels. Differentiation through branding becomes harder when algorithms prioritize functional specifications.

Loyalty Program Disruption

Traditional loyalty schemes that lock customers into ecosystems lose effectiveness when agents optimize across all options. Points and rewards that create switching costs are easily arbitrated by algorithms that calculate true value.

Disintermediation Risk

Retailers who merely aggregate products face displacement by agents that source directly from manufacturers. The value of curation diminishes when algorithms can evaluate millions of options instantly.

Strategic Responses for Retailers



API-First Architecture

Rebuild commerce platforms with agent access as the primary interface. Expose complete product catalogs, real-time inventory, dynamic pricing, and transaction capabilities through well-documented, performant APIs that agents can easily consume.



Agent Partnership Programs

Develop strategic relationships with leading agent platforms. Offer preferential terms, early access to inventory, or exclusive deals to agents that drive significant transaction volume through verified performance.



Frictionless Fulfillment

Optimize post-purchase experiences including shipping speed, return ease, and customer service responsiveness. Agent algorithms incorporate fulfillment quality into decision models, rewarding operational excellence.



Value Over Marketing

Shift investments from advertising spend toward product quality, competitive pricing, and reliable service. When agents make decisions based on objective metrics rather than persuasive messaging, fundamentals matter more than promotions.



Structured Data Excellence

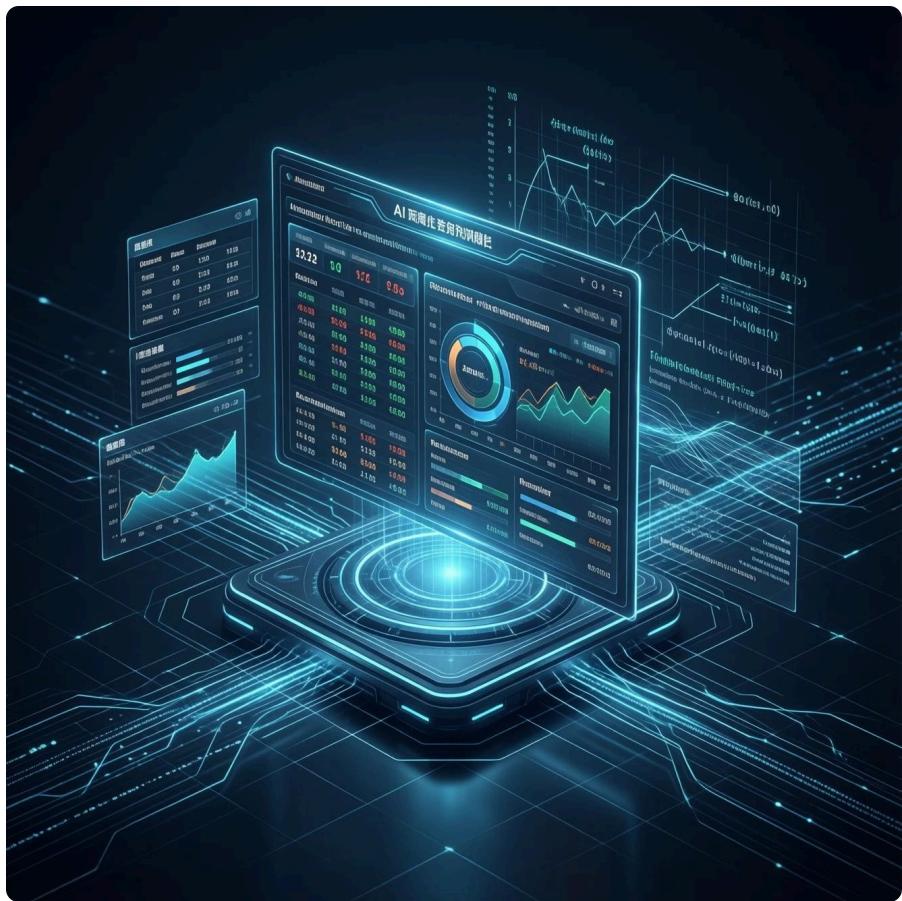
Implement comprehensive schema markup, rich product attributes, and machine-readable specifications that enable agents to accurately evaluate offerings. Poor data quality creates invisible competitive disadvantage.



Proprietary Agent Development

Build owned agents that guide customers toward products while maintaining brand relationships. First-party agents can balance optimization with strategic business objectives in ways third-party agents cannot.

The Negotiation Layer



Dynamic Pricing Meets Algorithmic Bargaining

One of the most disruptive capabilities of agentic AI is automated negotiation. Where human shoppers typically accept listed prices, agents can engage in rapid, automated bargaining that tests retailers' pricing flexibility and exposes hidden margins.

This creates a new strategic game where retailer pricing algorithms interact with buyer agent algorithms in microsecond negotiations.

Price Discovery

Agent queries multiple retailers to establish baseline pricing and identify outliers above or below market rates.

Offer Submission

Agent submits optimized counter-offer slightly below discovered price floor, testing retailer's willingness to negotiate.

1

2

3

4

Leverage Assessment

Algorithm evaluates bargaining power based on inventory levels, competitor pricing, historical price flexibility, and purchase volume potential.

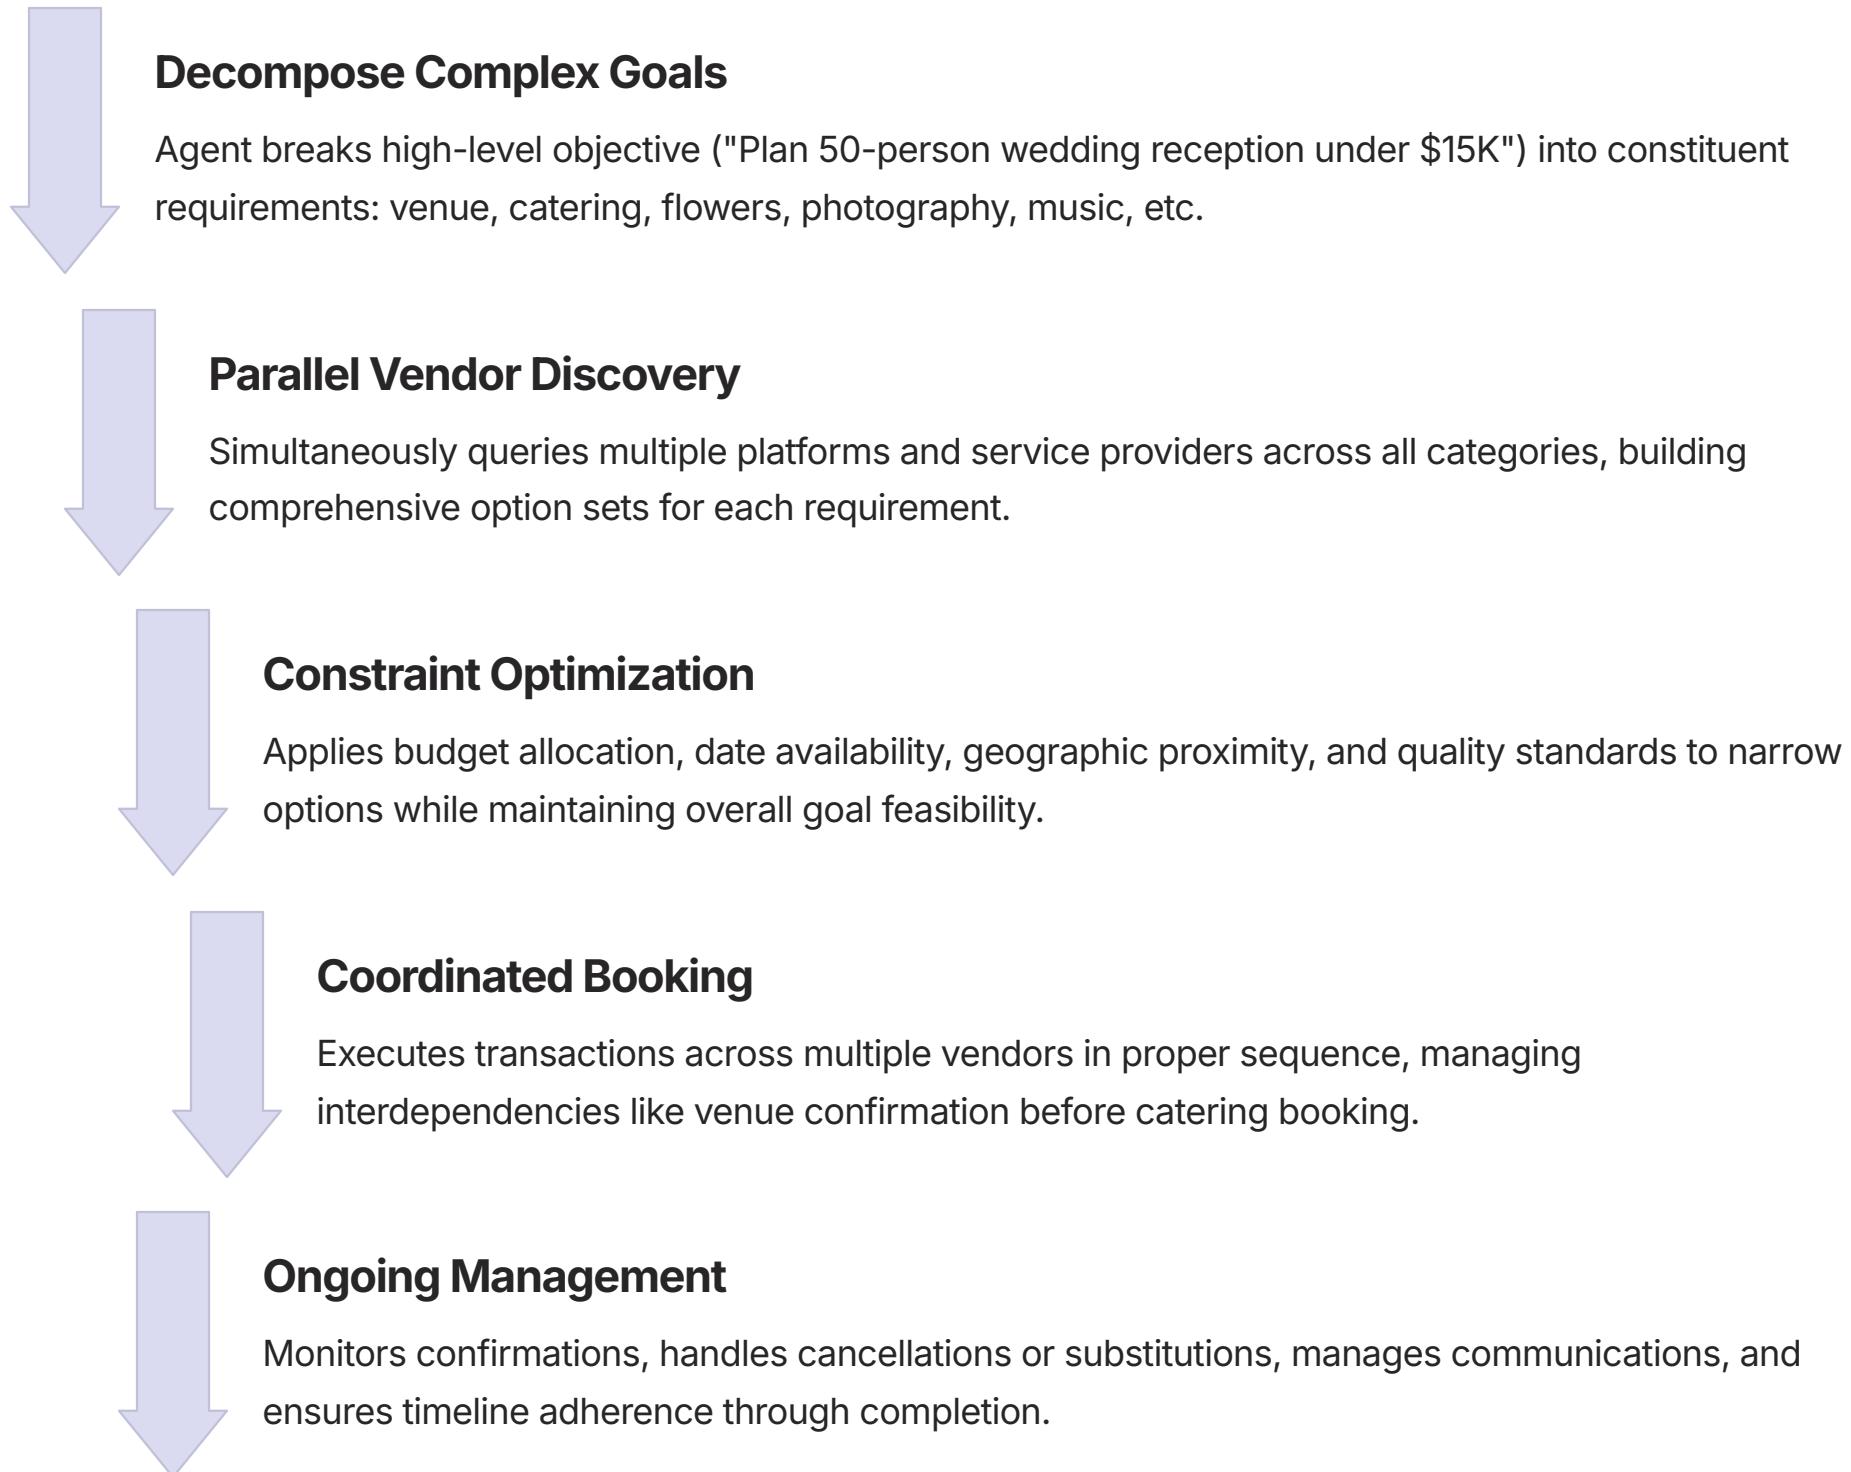
Iterative Refinement

If rejected, agent refines offer based on response patterns. If accepted, updates models with successful negotiation parameters.

This negotiation layer introduces volatility and complexity into retail pricing. Fixed pricing gives way to personalized, context-dependent offers determined by algorithmic bargaining. Retailers must develop sophisticated dynamic pricing systems that can respond to agent negotiations without eroding margins unsustainably.

Cross-Platform Orchestration

The most powerful agentic AI applications orchestrate complex, multi-vendor workflows that would be impractical for human shoppers. Consider planning a wedding, remodeling a kitchen, or organizing a corporate event—scenarios requiring dozens of coordinated purchases from disparate vendors with timing dependencies and budget constraints.



This cross-platform orchestration creates value that far exceeds simple price comparison. Agents serve as project managers, coordinators, and troubleshooters—roles that previously required human event planners or consultants charging substantial fees.

The Subscription Management Revolution

The Subscription Sprawl Problem

Average households now maintain 12-15 active subscriptions across streaming, software, groceries, and services. Managing renewals, comparing plans, and optimizing spending requires constant vigilance that few consumers maintain.

This creates subscription fatigue where unused services persist indefinitely, costing households hundreds monthly in forgotten charges.

Subscription management represents an ideal agentic AI use case: highly repetitive, rules-based, with clear optimization metrics. Agents can deliver immediate, quantifiable value through cost reduction while requiring minimal user oversight. This "killer app" drives initial adoption and builds trust for more complex delegation scenarios. Several startups have emerged specifically targeting subscription optimization, demonstrating dedicated market demand for this capability.

Agent-Managed Solutions

Agentic AI excels at subscription optimization through:

- Usage monitoring and utilization analysis
- Automatic cancellation of unused services
- Plan comparison and downgrade recommendations
- Bundle optimization across providers
- Negotiation of renewal terms and pricing
- Trial management to avoid unwanted conversions

Personalization at Machine Scale

Individual Preferences

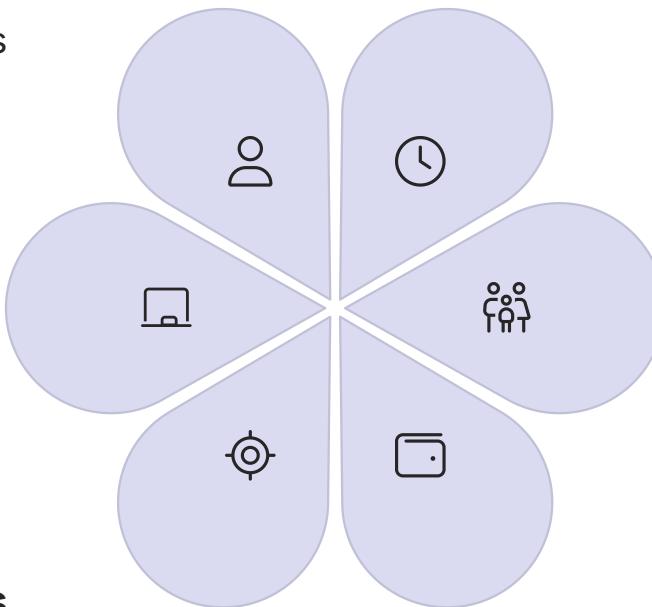
Detailed taste profiles including brand preferences, quality standards, aesthetic sensibilities, and dietary restrictions

Continuous Adaptation

Ongoing refinement through outcome tracking, feedback incorporation, and preference drift detection over time

Context Awareness

Integration of location, weather, events, and situational factors that influence needs and purchase timing



Temporal Patterns

Understanding of seasonal needs, life stage transitions, and cyclical purchase behaviors that repeat predictably

Household Dynamics

Recognition of family composition, individual member needs, and complex multi-stakeholder decision-making processes

Budget Intelligence

Sophisticated understanding of spending capacity, category allocations, and value-quality trade-off thresholds

This multi-dimensional personalization operates at a scale and sophistication impossible for human-staffed concierge services. Agents synthesize thousands of data points to make each recommendation, learning from both explicit feedback and implicit signals like purchase completion rates. The result is personalization that improves continuously, becoming more valuable the longer the user-agent relationship persists.

Ethical Considerations and Bias

Agentic AI systems inherit and potentially amplify biases from training data, design choices, and deployment contexts. When these biases affect purchasing decisions, they perpetuate discriminatory patterns with economic consequences. Addressing algorithmic bias in commerce agents requires multi-layered interventions.

Price Discrimination Risks

Agents might learn to accept higher prices for users in wealthy zip codes or with premium device types, perpetuating economic discrimination. Regulators increasingly scrutinize algorithmic pricing for disparate impact.

Product Access Inequity

Training data skewed toward majority populations may cause agents to overlook products serving niche communities, from religious dietary requirements to accessibility adaptations, effectively making these products invisible.

Vendor Favoritism

Agents trained on platforms with financial relationships to certain vendors may systematically favor those partners over objectively superior alternatives, converting optimization promises into disguised marketing.

Transparency Deficits

Complex neural network decision-making resists human interpretation, making it difficult for users to understand why agents made particular recommendations or to identify when bias influenced outcomes.

Mitigating these risks requires diverse training data, fairness constraints in optimization algorithms, regular bias audits, and transparent decision explanation interfaces. However, the economic incentives of agent platforms may not naturally align with fairness goals, requiring regulatory intervention or competitive pressure from ethics-focused alternatives.

The Merchant Perspective: Adaptation Strategies

Threatened Business Models

Traditional merchants face disruption across multiple dimensions:

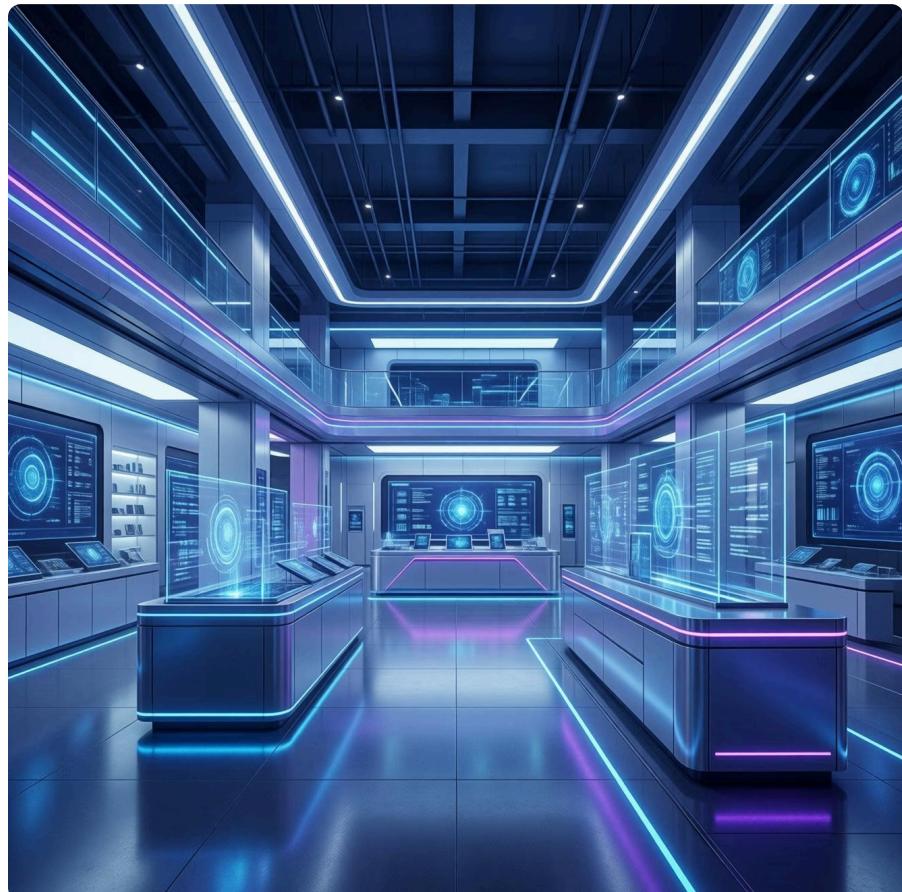
- Display advertising loses effectiveness
- Brand building through aesthetics diminishes
- Customer acquisition costs rise
- Direct relationships weaken
- Margins compress from price transparency
- Loyalty programs lose retention power



Survival Tactics

Forward-thinking retailers are pivoting strategies:

- Competing on fundamental value versus marketing
- Building API-first commerce architectures
- Developing proprietary agent capabilities
- Partnering with major agent platforms
- Optimizing for agent discovery algorithms
- Differentiating through service excellence



The most sophisticated retailers recognize that resisting agentic commerce is futile. Instead, they're racing to position themselves advantageously within the new paradigm—ensuring their products surface in agent searches, their checkout processes work seamlessly with automation, and their value propositions remain compelling when evaluated by algorithms rather than human emotion.

Future Scenarios: Three Possible Paths

Scenario 1: Fragmented Competition

Multiple agent platforms compete on features, privacy protections, and user control. Interoperability standards emerge, allowing users to switch agents and port data easily. Merchants adapt to agent-driven commerce, and the market self-regulates through competitive pressure for user-friendly, ethical behavior. This scenario maximizes innovation and user choice but may sacrifice some efficiency through fragmentation.

Scenario 2: Platform Dominance

Two or three major technology companies capture the majority of agentic commerce, leveraging network effects, data advantages, and ecosystem lock-in. These platforms wield enormous power over both consumers and merchants, extracting rents through gatekeeping. Regulatory intervention proves insufficient to prevent consolidation. Users gain convenience but lose privacy and choice, while merchants face monopolistic pricing pressures.

Scenario 3: Regulatory Equilibrium

Governments implement comprehensive frameworks mandating transparency, data portability, and algorithmic accountability. Agent platforms must adhere to strict fairness standards and allow third-party auditing. This creates a regulated middle ground where innovation continues within guardrails protecting consumer rights and preventing abuse. Market efficiency is somewhat reduced by compliance costs, but systemic risks are contained.

The path forward depends on regulatory choices made in the next 24-36 months while the technology remains nascent. Early regulatory frameworks will establish precedents difficult to reverse later.

Implementation Roadmap for Enterprises

01

Audit Current Architecture

Evaluate existing e-commerce infrastructure for agent accessibility. Identify gaps in API coverage, structured data implementation, and programmatic transaction support. Document technical debt that impedes agent integration.

02

Develop Agent Strategy

Choose strategic approach: build proprietary agents, partner with platforms, or optimize for third-party discovery. Consider competitive positioning, customer relationships, and resource constraints in decision-making.

03

Implement Technical Foundations

Expose comprehensive APIs for product discovery, inventory checking, pricing, and checkout. Implement rich schema markup and structured data across all product pages. Ensure real-time accuracy of all exposed information.

04

Pilot with Selected Partners

Launch limited agent integrations with leading platforms to validate technical implementation and identify friction points. Collect performance data on agent-driven transactions versus traditional e-commerce.

05

Optimize Based on Learning

Refine agent experiences based on pilot results. Adjust product data quality, API performance, and checkout workflows. Develop metrics for agent discovery ranking and conversion optimization.

06

Scale and Monitor

Expand agent commerce capabilities across entire catalog and all major platforms. Establish ongoing monitoring of agent-driven transactions, price negotiation outcomes, and customer satisfaction metrics.

Key Takeaways and Strategic Imperatives

Agentic AI represents a fundamental shift, not an incremental improvement

This is not simply better search or smarter recommendations. Autonomous agents that can execute transactions end-to-end invert the power dynamic between merchants and consumers, requiring complete strategic rethinking.

The primary customer of the future is algorithmic, not human

Retailers must optimize for agent discovery and decision-making rather than human attention. Marketing effectiveness, brand building, and loyalty programs all require reimagining for an algorithmic audience.

API-first architecture is now business-critical infrastructure

Merchants without comprehensive, performant, well-documented APIs become invisible in agent-mediated commerce. Technical infrastructure previously viewed as optional becomes existential requirement.

Trust and privacy concerns will pace adoption more than technical capability

The technology already works; consumer comfort and regulatory frameworks lag. Building trustworthy agent systems with strong privacy protections determines market success more than algorithmic sophistication.

Early movers gain compounding advantages through data and relationships

Agent platforms and retailers that establish early dominance in agentic commerce accumulate data and network effects that create increasingly insurmountable barriers to entry for late followers.

Conclusion: Preparing for the Agentic Future

The transition from human-mediated to agent-mediated commerce represents the most significant restructuring of retail since the emergence of e-commerce itself. Unlike previous shifts that enhanced existing models—mobile commerce remained fundamentally human-driven shopping—agentic AI introduces truly autonomous actors that bypass traditional consumer journeys entirely.

This transformation creates winners and losers with stark clarity. Retailers who embrace API-first architecture, optimize for algorithmic discovery, and compete on fundamental value rather than marketing prowess will thrive. Those who resist adaptation or underestimate the speed of change face marginalization as agent platforms route transactions around them to more accessible competitors.

Immediate Actions

- Assess agent-readiness of current systems
- Develop comprehensive API strategies
- Implement rich structured data markup
- Pilot integrations with leading platforms
- Retrain teams on agent optimization
- Establish agent transaction monitoring

Strategic Positioning

- Decide build, partner, or buy strategy
- Evaluate platform dependencies and risks
- Consider privacy and ethical frameworks
- Prepare for regulatory evolution
- Identify agent-resistant value propositions
- Plan for margin compression scenarios

The agentic future arrives not as distant speculation but as present reality. OpenAI's Operator, Amazon's Rufus, and numerous startup alternatives already process millions of transactions. The question is not whether agentic AI will reshape commerce, but how quickly—and whether your organization will lead, follow, or become obsolete in the transformation.

The time for strategic response is now. The infrastructure built, partnerships formed, and capabilities developed in 2026 will determine competitive positioning for the next decade of retail evolution. Those who move decisively today position themselves to thrive in the agentic economy. Those who wait risk finding themselves locked out of the future of shopping.

 **About this report:** This comprehensive analysis was produced by the DX Today research team to help business leaders understand and prepare for the transformative impact of agentic AI on retail commerce. For questions or deeper consultation on implementation strategies, contact the Senior Chief Editor at DX Today.