Agentic Commerce:
The Next Frontier in AI-Driven Transactions

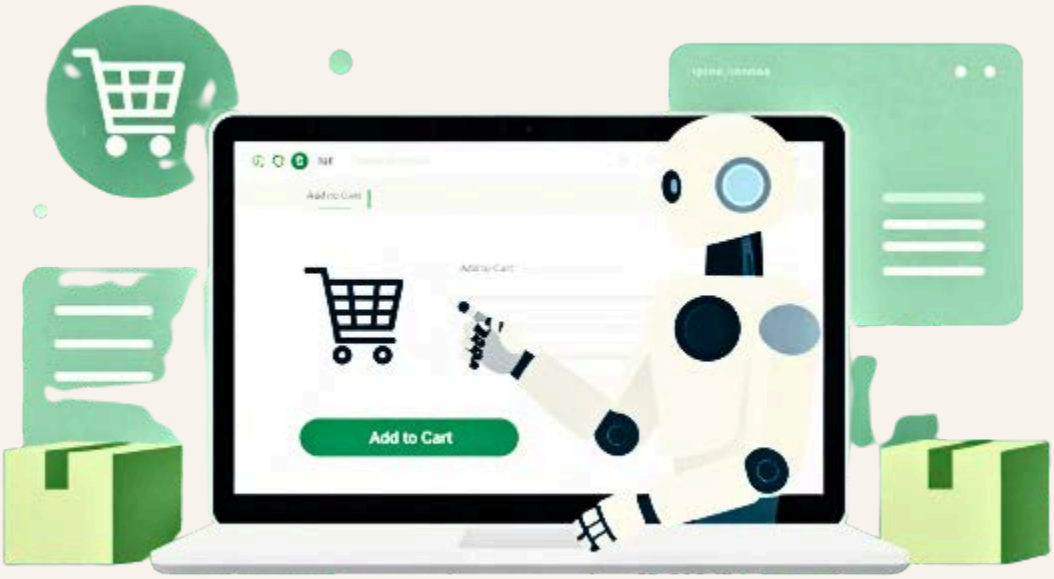

Perception, Implementation, Players, Protocols, Guardrails, and the Future of Autonomous Commerce

payhuddle® | Count on us

Meet the Contributors

Sabapathy Narayanan

Head of Marketing

Praveenkumar G

Senior Product Lead Engineer

Contents

Executive Summary

Agentic commerce, the ability of AI agents to autonomously discover, negotiate, and complete purchases on behalf of humans, has rapidly transitioned from a theoretical concept to a live, operational reality in the span of 18 months. What was once described as the future of shopping is now being tested in production environments spanning payments giants, technology platforms, consumer brands, and regulators across the globe.

This white paper synthesizes the current landscape of agentic commerce: how it is perceived by practitioners, technologists, and consumers; how leading implementations technically operate; who the key players are; what protocols are emerging to govern agent identity and transaction security; and what compliance frameworks and guardrails are being built to ensure that autonomous commerce remains trusted, transparent, and consumer-controlled.

Among the major themes:

- Global payment networks — Visa, Mastercard, and American Express have each launched dedicated agentic commerce programs in 2025–2026, anchored in tokenization, agent registration, and spending controls.
- Google's Agent Payments Protocol (AP2), developed with Coinbase and 60+ partners, and subsequently donated to the FIDO Alliance, is emerging as the leading open standard for agent-initiated payments.
- India has become a remarkable early leader, with Razorpay and the National Payments Corporation of India (NPCI) launching production-grade agentic payments on Anthropic's Claude platform, powered by UPI Reserve Pay.

- Consumer trust remains the primary adoption barrier: surveys show that while over 47% of U.S. shoppers use AI for at least one shopping task, only 47% are comfortable with fully autonomous agent-initiated purchases.
- Regulatory frameworks are still catching up: the CFPB, NIST, UK Competition and Markets Authority, and the EU AI Act are each grappling with how existing consumer protection laws apply to agent-initiated transactions.

This paper is organised as follows: (1) Introduction, (2) How Agentic Commerce is Perceived, (3) How Implementations Work, (4) Key Players and Initiatives, (5) Protocols and Standards, (6) Payment Authorization and Security, (7) Compliance and Regulation, (8) Guardrails and Safeguards, (9) The Road Ahead, (10) Conclusion, and (11) References.

01 Introduction: What is Agentic Commerce?

Commerce has always been shaped by the tools available to buyers and sellers. From barter to currency, from storefronts to catalogs, from e-commerce to mobile payments, each era has reorganized who does what in a transaction. Agentic commerce represents the next such reorganization: a world in which AI agents not only recommend products but also complete the purchase.

Praveenkumar G, an engineer with Payhuddle, offers a succinct illustration:

"Agentic Commerce — In the future, commerce could happen with a prompt. Like, book me a bus from Bengaluru to Attur on 30th April after 9 PM in a sleeper class and budget 1.5k along with the return ticket on May 3rd with the same constraints."

In this model, the consumer expresses intent in natural language. The AI agent, equipped with knowledge of the user's preferences, spending limits, and authenticated payment credentials, searches, negotiates, selects, and completes the transaction. The human controls the parameters; the agent executes within them.

Harshil Mathur, CEO and Co-Founder of Razorpay, captures the philosophical shift succinctly: "AI shouldn't stop at recommendations; it should finish the job." (BusinessToday, February 2026)

Luke Gebb, Executive Vice President and Global Head of Innovation at American Express, describes the moment more broadly: “Agentic commerce is a sea change… a massive moment in commerce similar to the advent of the web and to mobile.” (Digital Commerce 360, April 2026)

Jack Forestell, Visa’s Chief Product and Strategy Officer, frames it in terms of network trust: “Soon people will have AI agents browse, select, purchase, and manage on their behalf. These agents will need to be trusted with payments, not only by users, but by banks and sellers as well.” (Visa Press Release, April 2025)

02 How agentic commerce is perceived

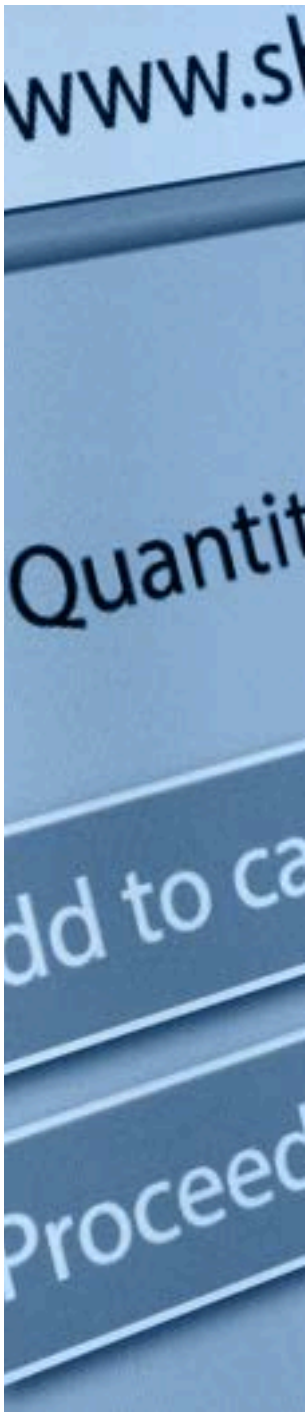

Industry optimism and investment

The investment momentum behind agentic commerce is significant. According to KPMG's Q1 2025 AI Quarterly Pulse Survey, 65% of organizations were piloting AI agents in Q1 2025, up from 37% the previous quarter, and 99% of executives surveyed said their organizations planned to deploy AI agents. (Digital Commerce 360, May 2025)

Salesforce's Connected Shoppers Report (March 2025) found that 66% of shoppers expressed interest in using AI agents to secure high-demand items before they sell out, and 65% wanted agents that could buy products once they reached a target price. Visa's own research found that nearly half of U.S. shoppers (47%) already use AI tools for at least one shopping task, from price comparisons to personalized recommendations. (Visa, December 2025)

Agentic AI is expected to handle up to 20% of e-commerce tasks in 2025, and Mastercard's CEO Michael Miebach declared at the company's October 2025 earnings call: "Agentic commerce is here, and we're at the center of it." (Digital Commerce 360, October 2025)

Mixed consumer sentiment

Consumer enthusiasm is real but calibrated. While 22% of shoppers now use AI tools like ChatGPT to research products before buying, only 14% trust AI recommendations alone to complete a purchase, according to Salsify's 2026 Consumer Research study of nearly 2,800 shoppers across the U.S., U.K., and Canada. (Savannah Brentnall, LinkedIn, 2026)

Trust barriers center on two issues. First, data security: consumers worry about the misuse of stored payment credentials and autonomous financial transactions. Second, the emotional dimension of shopping: as Ujjwal Keshri, an e-commerce AI specialist, argues:

> *"Humans don't just shop to acquire things. We shop to feel something: the thrill of discovering a product, the satisfaction of making the right choice, the joy of unboxing. Autonomous agents optimize for logic. But purchasing decisions are emotional."*

This perspective highlights an important nuance: agentic commerce is not a universal replacement for human-driven shopping. It is most compelling for high-frequency, low-emotional-investment transactions, such as groceries, recurring purchases, utility payments, and B2B procurement, where efficiency outweighs the pleasure of the shopping experience itself.

The B2B opportunity

While consumer-facing agentic commerce commands most headlines, practitioners consistently identify business-to-business procurement as the more immediately practical application.

Praveenkumar G points to inventory management as a high-value use case: "The agents can track the inventory and place an order. AI can bring in the advantage of predicting the future sales with the pattern in the past sales data and place the right order." (LinkedIn, 2026)

Ramp, the corporate spend management platform, already uses Visa Intelligent Commerce for B2B payments, streamlining corporate bill pay and allowing customers to capture card cashback on automated transactions. (Visa, December 2025)

India's distinct vantage point

India occupies a singular position in the global agentic commerce landscape. Khilan Haria, CPO of Razorpay, has argued: "The opportunity for agentic commerce in India is potentially as large as, if not larger than, UPI." (Economic Times, 2026)

The alignment of infrastructure (UPI's real-time, mandate-enabled rails), consumer behavior (high-frequency digital transactions in food and grocery), and AI platform penetration (Claude, ChatGPT, Gemini) creates rare conditions for agentic commerce to scale. Australia, too, is seeing early adoption: according to the National AI Center, over 45% of SME retailers have already implemented AI, outpacing finance and professional services. (Mehra Rajabova, AWS, LinkedIn, 2026)

03

How agentic commerce implementations work

Salesforce's Connected Shoppers Report (March 2025) found that 66% of shoppers expressed interest in using AI agents to secure high-demand items before they sell out, and 65% wanted agents that could buy products once they reached a target price. Visa's own research found that nearly half of U.S. shoppers (47%) already use AI tools for at least one shopping task, from price comparisons to personalized recommendations. (Visa, December 2025)

The agentic transaction lifecycle

A complete agentic transaction typically moves through the following stages:

Stage	Step	Description
1	**Intent Capture**	The consumer expresses a goal in natural language (e.g., 'order my usual groceries'). The AI agent interprets the request and infers parameters such as budget, preferences, and merchant.
2	**Agent Authentication**	The merchant or payment network verifies that the agent is registered, legitimate, and authorized to act on behalf of a specific consumer.
3	**Cart Construction**	The agent searches merchant catalogs, assembles a proposed cart, and presents options to the consumer for confirmation.
4	**Intent Verification**	The consumer's authenticated purchase intent is captured and transmitted to the payment network, creating a verifiable, tamper-evident record.
5	**Payment Credential Delegation**	Tokenized payment credentials are invoked within pre-authorized spending limits. No raw card data is transmitted.
6	**Transaction Settlement**	The payment is processed through existing card or UPI rails, with the agent's identity and transaction context recorded in the audit trail.
7	**Dispute Resolution**	Cart context and intent records are available to resolve disputes or chargebacks, enabling the payment network to adjudicate between consumer error, agent error, and merchant error.

The Razorpay-NPCI-Claud implementation (India)

The most fully documented consumer-facing implementation is the Razorpay-NPCI pilot on Anthropic's Claude platform, announced at the India AI Impact Summit in New Delhi on February 20, 2026.

The architecture combines three layers: Razorpay's Agentic Payments stack, NPCI's UPI Reserve Pay infrastructure, and Claude's conversational intelligence. When a user says to Claude, "I am craving a mid-evening samosa and chai. Can you order this for my team and me from my favorite restaurant near my office?", the agent checks available options on Zomato, presents them, and with a single user confirmation, places the order through Razorpay's UPI-powered Agentic Payments. (Razorpay blog, March 2026)

The critical enabling technology is UPI Reserve Pay, built on NPCI's Single Block Multi Debit (SBMD) framework. This allows a user to set a one-time spending limit for a merchant. Within those pre-authorized limits, the AI agent can execute multiple transactions without requiring a PIN for each one, while the user retains real-time visibility and the ability to revoke consent instantly. (Razorpay blog, March 2026)

Sohini Rajola, Executive Director – Growth at NPCI, articulated the design philosophy: "UPI was built to make digital payments simple, secure, and universal, and Agentic Payments takes that vision into the next era. With UPI Reserve Pay, users can give consent once and allow intelligent systems to transact on their behalf in a controlled, transparent way." (BusinessToday, February 2026)

Anthropic's "Project Deal" experiment

Anthropic conducted an internal experiment called 'Project Deal,' a Craigslist-like marketplace in which both buyers and sellers were AI agents representing employee volunteers. Each agent was given a $100 budget, interviewed their human to understand needs, then negotiated with other agents to complete transactions.

Anthropic reported the project was successful and that employees expressed interest in a real-world version of the service. This experiment provides an early proof of concept for agent-to-agent commerce without human intermediation at the transaction level.

Visa intelligent commerce architecture

Visa's implementation, announced at the Visa Global Product Drop on April 30, 2025, organizes agentic commerce capabilities into five modular APIs that developers can adopt selectively:

- Authentication — confirms the AI agent is authorized to act for a specific consumer, extending identity verification to AI commerce.
- Tokenization — replaces card data with network tokens that work anywhere Visa is accepted, keeping sensitive details under wraps.
- Payment Instructions — lets the user preset dollar limits, merchant categories, or real-time approval prompts, loaded directly into VisaNet.
- Personalization — with the shopper's consent, shares basic spend patterns so the agent can rank offers by preferred airline, hotel price, or dining habits.
- Dispute Management — imports three decades of Visa's AI fraud models into agent-driven transactions, with Visa reporting its AI engines blocked roughly $40 billion in fraud in 2024. (PYMNTS, May 2025)

By April 2026, Visa expanded this framework through Intelligent Commerce Connect, a protocol- and token vault-agnostic 'on ramp' supporting Trusted Agent Protocol, Machine Payments Protocol (MPP), Agentic Commerce Protocol (ACP), and Universal Commerce Protocol (UCP). Pilot partners include Aldar, AWS, Diddo, Highnote, Mesh, Payabli, and Sumvin. (Visa, April 2026)

Mastercard agent pay architecture

Mastercard's Agent Pay program, announced April 29, 2025, introduces Agentic Tokens — dynamic digital credentials that empower AI agents to transact safely and transparently on the consumer's behalf, building on the same tokenization technology that powers mobile contactless payments. (PYMNTS, April 2025)

The Mastercard Agent Pay Acceptance Framework, launched in October 2025, is designed for merchants: once an agent is registered and verified, it uses a Dynamic Token Verification Code formatted for standard card payment fields, enabling "no-code" acceptance. "We've made it easy for merchants across the globe to benefit on day one," said CEO Michael Miebach. (Mastercard, October 2025)

By November 2025, all U.S. Mastercard cardholders were enabled for Agent Pay, with Citi and U.S. Bank cardholders among the first participants. OpenAI's launch of Instant Checkout in ChatGPT, powered by the Agentic Commerce Protocol (ACP), brought Mastercard Agent Pay onto OpenAI's platform. (Mastercard, 2025)

American Express ACE developer kit

American Express launched its Agentic Commerce Experiences (ACE) Developer Kit on April 14, 2026, offering five integrated services that are architecturally distinctive because of Amex's closed-loop network, where the company operates simultaneously as a card issuer, payment network, and merchant acquirer.

- Agent Registration — AI agents are registered with Amex and receive a verified identity before they can transact on the network.
- Account Enablement — Card Members register their cards for agentic transactions and link their accounts for personalized Membership experiences.
- Intent Intelligence — the consumer's authenticated purchase intent is captured and transmitted to American Express, creating the evidentiary basis for authentication and dispute resolution.
- Payment Credentials — verified agents complete payments using tokenized credentials, without ever handling raw card data.
- Cart Context — merchants can optionally share cart details with Amex pre- or post-transaction to smooth dispute investigation. (BusinessWire, April 2026)

The closed-loop model gives Amex a unique advantage: it can observe both merchant and consumer data for each transaction, enabling richer intent verification and more precise dispute adjudication than four-party networks can. (American Banker, April 2026)

04 Key players and the competitive landscape

Visa intelligent commerce architecture

The dominant card payment networks have each staked out positions in agentic commerce, with overlapping but differentiated strategies:

Network	Program	Launch	Key Differentiator
Visa	Intelligent Commerce	April 2025	Modular 5-API suite; Intelligent Commerce Connect as protocol-agnostic on-ramp; 100+ global partners.
Mastercard	Agent Pay	April 2025	Agentic Tokens; no-code merchant acceptance framework; OpenAI/ChatGPT integration via ACP.
American Express	ACE Developer Kit	April 2026	Closed-loop advantage; Amex Agent Purchase Protection (industry-first consumer coverage for agent errors).
UnionPay	Agentic Payment Open Protocol	April 2026	Open protocol for agent registration, user identity, and authorization; designed for global interoperability.

Technology platforms

- Anthropic (Claude): The Claude platform serves as the conversational interface for the Razorpay–NPCI pilot in India and the Visa Intelligent Commerce partner ecosystem.
- Google: Beyond its AI models, Google has developed the Agent Payments Protocol (AP2) and the broader Agent Commerce Kit (ACK), encompassing the Universal Commerce Protocol (UCP) and the Agent2Agent (A2A) communication protocol.
- OpenAI (ChatGPT): Launched Instant Checkout in ChatGPT, powered by the Agentic Commerce Protocol (ACP), enabling Mastercard Agent Pay transactions directly within ChatGPT conversations.

- Microsoft: Partnering with Mastercard to integrate Azure OpenAI Service and Microsoft Copilot Studio with Mastercard's payment solutions; contributing to the Trusted Agent Protocol.

Indian fintech ecosystem

India's fintech infrastructure players are among the most operationally advanced globally in agentic commerce:

- Razorpay: Launched Agentic Payments on Claude with NPCI; earlier partnership with OpenAI at Global Fintech Fest 2025.
- NPCI: Developed UPI Reserve Pay (SBMD framework) as the consent-based mandate infrastructure for agentic transactions.
- Cashfree Payments: Launched 'Cashfree Here', a payments extension for AI applications built on OpenAI's Apps SDK and Anthropic's MCP framework, using passkey-enabled biometric authentication supported by Mastercard. (YourStory, February 2026)
- Platforms (Zomato, Swiggy, Zepto, Flipkart, Firstcry, Bigbasket, Ajio): Building AI storefronts to ensure their products are discoverable and purchasable by LLMs. (Economic Times, 2026)

Merchant-side enablers

A new category of agentic commerce enablers bridges AI platforms and payment networks:

- Skyfire: Enabled Consumer Reports' product recommendation agent to demonstrate purchases via browser automation on Visa's network.
- Nekuda: Allows fashion platform Gensmo's app to move from AI-styled looks to purchase from Fabrique in a single tap via Rye's checkout API; also enabling Henry Labs' one-click checkout at Price.com and Honeylove.
- PayOS: Providing BeyondStyle with payment infrastructure for agent-driven checkout at online retailer Jomashop; also expanding Ramp's B2B agentic payments.
- Firmly.AI, Basis Theory: Early agentic commerce enablers for Mastercard's network. (Visa, December 2025; Mastercard, September 2025)

05 Protocols and open standards

The agentic commerce ecosystem faces a risk of fragmentation: multiple competing protocols could create incompatible silos, increasing integration costs and slowing adoption. A fierce standardization contest is underway, with early signs of convergence around open, industry-governed frameworks.

Google's Agent Payments Protocol (AP2)

Announced by Google Cloud and Coinbase in September 2025, AP2 is the most comprehensive open protocol for agentic commerce. It gives autonomous agents a wallet, a programmable settlement rail, and auditable proofs, enabling them to price, purchase, and receive payment without human-in-the-loop friction. (AP2 Lab, 2025)

AP2 builds on the Agent2Agent (A2A) protocol for messaging and capability discovery. Where A2A allows agents to talk to each other, AP2 allows them to transact. The protocol is payment-agnostic, supporting traditional cards, bank transfers, and stablecoins via its x402 extension (co-developed with Coinbase, Ethereum Foundation, and MetaMask). (Everest Group, January 2026)

AP2 builds trust through three core mandate types:

- Intent Mandates — user-signed proof of the consumer's original purchase intent.
- Cart Mandates — tamper-evident records of the specific items the agent has selected.
- Payment Mandates — authorization for the agent to execute the financial transaction within defined limits.

In April 2026, Google donated AP2 to the FIDO Alliance and released AP2 v0.2, introducing “Human Not Present” payments, enabling agents to securely execute autonomous purchases (such as securing limited-run tickets the moment they go on sale) based on pre-authorized user instructions. (Google Blog, April 2026)

Simultaneously, Google and Mastercard co-developed Verifiable Intent, an AP2-compatible standard also donated to FIDO that creates a tamper-proof log of user-authorized agent actions to ensure accountability. (Google Blog, April 2026)

At launch, AP2 had 60+ organizational partners, including Accenture, Adobe, Adyen, and a range of payment networks and technology providers. (Google Cloud Blog, September 2025)

Trusted Agent Protocol (TAP)

In October 2025, Visa and more than 10 partners introduced the Trusted Agent Protocol (TAP), an open framework built on existing web infrastructure that helps merchants distinguish between malicious bots and legitimate AI agents acting on behalf of consumers. Akamai later integrated TAP with its edge-based behavioral intelligence, bot protection, and user recognition capabilities. (Visa, December 2025)

Agentic Commerce Protocol (ACP) and Universal Commerce Protocol (UCP)

ACP, developed in the OpenAI–Stripe–Mastercard ecosystem, powers ChatGPT’s Instant Checkout. UCP is part of Google’s Agent Commerce Kit (ACK), designed to enable AI agents to communicate with financial systems securely, autonomously, and auditably. Visa’s Intelligent Commerce Connect supports all four major protocols: TAP, Machine Payments Protocol (MPP), ACP, and UCP, reflecting a deliberately protocol-agnostic stance.

FIDO Alliance Governance

The donation of AP2 and Verifiable Intent to the FIDO Alliance is significant: it signals that the major players are prioritizing interoperability over proprietary lock-in, and that a respected, neutral standards body will govern the evolution of these protocols. The FIDO Alliance has previously standardized passkeys and FIDO2, giving it credibility and global reach in authentication infrastructure.

EMVCo and Payment Security Standards

American Express is actively engaging with EMVCo, the joint venture of global card networks that maintains EMV payment standards, to develop agentic commerce specifications that integrate with existing PCI-DSS compliance frameworks. Amex is also contributing to Google's AP2 and working with Cloudflare's Web Bot Auth Protocol. (AmericanExpress.io, November 2025)

06 Payment authorization and security

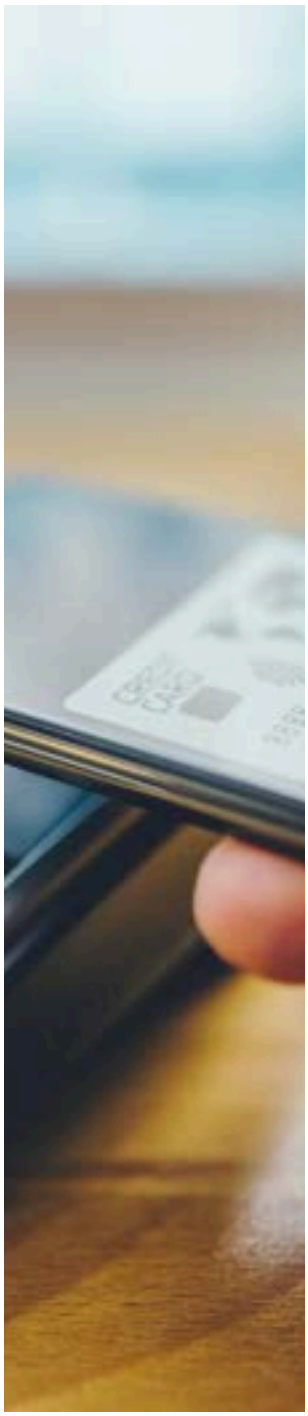

Payment authorization in agentic commerce must solve a fundamentally new problem: how does a payment network know that a non-human actor has legitimate authority to spend a human's money? Existing authorization models, such as PINs, passwords, and 3DS, are designed for humans who are present and actively consenting. Agentic commerce requires authorization models that work when the human is not actively present.

Tokenization as the foundation

Every major implementation relies on tokenization as the security bedrock. Raw card numbers are never transmitted to agents or merchants. Instead, AI-ready tokenized credentials replace card data, ensuring that even if an agent is compromised, the underlying payment instrument is protected.

Mastercard's Agentic Tokens add a layer of programmability: they can encode spending constraints, merchant category restrictions, and time limits directly into the token itself, so enforcement happens at the credential level without requiring real-time human approval for every transaction. (Mastercard, 2025)

Agent identity and registration

A core requirement across all implementations is that agents must be registered and verified before they can transact. Unregistered agents cannot initiate transactions on Visa, Mastercard, or Amex networks. This registry function is architecturally analogous to merchant acquiring: just as a merchant must be onboarded to accept cards, an AI agent must be onboarded to initiate payments.

American Express's Luke Gebb describes the process: "American Express gives them an ID so we know what agent we're dealing with." (Digital Commerce 360, April 2026) Mastercard's program requires every agent's identity and authorization data to be passed through the transaction for traceability. (Digital Commerce 360, October 2025)

Spending controls and limits

All implementations require user-defined spending controls:

- Visa allows users to preset dollar limits, merchant categories, or real-time approval prompts directly into VisaNet.
- Mastercard's Agentic Tokens encode constraints at the credential level.
- American Express requires that no purchase can occur without a directive from the Card Member: "If there's no directive, there is no authorization to purchase." (Digital Commerce 360, April 2026)
- UPI Reserve Pay allows users to set a single spending limit per merchant, within which multiple transactions can settle without a PIN.

The 3DS question: Towards 4DS?

Praveenkumar G raises an important forward-looking question about the evolution of 3-D Secure (3DS) authentication, the protocol that currently secures card transactions by authenticating the cardholder across three domains (acquiring bank, issuing bank, and interoperability domain):

"While 3DS secures 3 domains (Acquiring, Issuing, and Interoperable), with agents doing the payments, it has to secure Cardholders Money, and it can become 4DS."

The implication is that agentic commerce may require a fourth authentication domain, the AI agent itself, to be incorporated into the authentication protocol. EMVCo's engagement with American Express on agentic commerce standards suggests this evolution is actively being planned.

Dispute resolution and agent error protection

American Express' Amex Agent Purchase Protection is the industry's first explicit consumer protection for agent-initiated transactions. If a Card Member authorizes a registered AI agent with authenticated purchase intent, and the agent makes an error (e.g., purchasing the wrong item), American Express will cover eligible charges. Eligible agent errors are those where the purchase deviates from the Card Member's authenticated purchase intent. Claims are subject to review, and cases where intent is subjective ("best" or "really nice") are explicitly excluded. (BusinessWire, April 2026)

This protection model creates powerful incentives for standardized intent capture: the more precisely intent is recorded, the more clearly disputes can be adjudicated. Cart Context, the fifth component of the Amex ACE Developer Kit, exists specifically to document the agent's understanding of the user's request as a reference point for dispute resolution.

07

Compliance and Regulatory Framework

Agentic commerce presents genuine regulatory novelty. Unlike previous fintech innovations built on existing payment rails, autonomous AI-initiated transactions create new legal questions around authorization, consumer protection, data rights, and liability. Regulatory frameworks designed for human-initiated commerce are struggling to keep pace.

United States: CFPB, Regulation E, and NIST

The Consumer Financial Protection Bureau (CFPB) enforces Regulation E, which governs electronic fund transfers and provides consumers with error resolution rights. In August 2025, the CFPB sought comment on an advance notice of proposed rulemaking concerning personal financial data rights, specifically who can serve as a ‘representative’ acting on a consumer’s behalf. Agentic commerce’s availability and adoption will depend on the CFPB's answer to whether consumer-authorized AI agents waive error-resolution rights, a question with profound implications for consumer protection. (Center for Data Innovation, March 2026)

The architecture combines three layers: Razorpay's Agentic Payments stack, NPCI's UPI Reserve Pay infrastructure, and Claude's conversational intelligence. When a user says to Claude, "I am craving a mid-evening samosa and chai. Can you order this for my team and me from my favorite restaurant near my office?", the agent checks available options on Zomato, presents them, and with a single user confirmation, places the order through Razorpay's UPI-powered Agentic Payments. (Razorpay blog, March 2026)

The National Institute of Standards and Technology (NIST) announced plans to host a public-private conversation in April 2026 regarding standards for AI agents and barriers to adoption. (Center for Data Innovation, March 2026)

The FTC's "Operation AI Comply" has launched coordinated enforcement targeting deceptive AI marketing, signaling that autonomous AI systems engaging in commerce will face scrutiny under existing consumer protection law even before bespoke agentic regulations are enacted. (Authority Partners, March 2026)

The EU is simultaneously exploring GDPR simplification through its Digital Omnibus package, recognizing that compliance burdens may be inhibiting European innovation relative to U.S. and Asian competitors. (Privacy Law analysis, January 2026)

United Kingdom: CMA Guidance on Agentic AI

On March 9, 2026, the UK Competition and Markets Authority (CMA) published guidance on 'Complying with Consumer Law When Using AI Agents.' The guidance applies consumer law principles to agentic AI and sets out expectations, including: being transparent about agentic AI use through clear labeling;

conducting ongoing monitoring of real-world agent performance, including errors, bias, and unintended outcomes; maintaining human oversight; and promptly remediating identified issues. (Morrison Foerster, May 2026)

The CMA has signaled that it will assess compliance through established consumer law principles, and that failure to comply may result in fines of up to 10% of worldwide turnover or orders to compensate affected consumers. (Cooley, March 2026)

India: RBI, NPCI, and the UPI Framework

India's regulatory environment is arguably the most agentic-commerce-ready in the world. UPI's architecture incorporates a formal consent-and-mandate framework that maps naturally onto agentic transaction models. The Reserve Bank of India's regulatory sandbox and NPCI's open innovation program have enabled the Razorpay–NPCI pilot to operate within existing regulatory permissions, rather than requiring new legislation.

The UPI Reserve Pay feature operates within NPCI's Single Block Multi Debit (SBMD) framework, which already has established rules for consent, merchant authorization, and transaction limits, creating a compliant foundation for agentic payments.

Key compliance requirements for organizations

Organizations deploying agentic commerce solutions should address the following compliance dimensions:

- Agent Disclosure: Consumers must know when they are interacting with an AI agent, and when that agent is authorized to make purchases on their behalf.
- Consent Architecture: Every agentic transaction must trace back to explicit, authenticated consumer intent. Implied or ambient consent is insufficient.
- Data Protection: Purchase intent data, spend patterns, and behavioral data shared with agents for personalization must comply with applicable privacy laws and be subject to meaningful consent.
- Anti-Money Laundering (AML) and Know Your Customer (KYC): Agents transacting on behalf of consumers must operate within the same AML/KYC frameworks as the consumers themselves. Agent registration processes must not create new vectors for identity fraud.
- Strong Customer Authentication (SCA): EU regulations requiring SCA for electronic payments will require reconsideration as agent-initiated transactions scale. The Razorpay–Mastercard passkey biometric model, which replaces OTPs with biometric confirmation, is one compliant approach.
- Liability Allocation: Clear contractual frameworks between consumers, agents, developers, payment networks, and merchants are required to allocate liability for agent errors, fraud, and unauthorized transactions.

08 Guardrails and Safeguards

Guardrails in agentic commerce operate at multiple layers: technical controls embedded in the agent architecture, contractual requirements imposed by payment networks, and organizational policies maintained by deploying businesses. Effective guardrails are not optional — they are the precondition for consumer trust and regulatory compliance.

Spending Controls (Consumer Level)

Every major platform has implemented user-configurable spending controls as the primary guardrail against overspending and unauthorized transactions. These typically include:

- Maximum transaction amount (per transaction and/or per period)
- Merchant category restrictions (e.g., food delivery only, no travel bookings)
- Allowlist/blocklist of specific merchants
- Real-time approval prompts for transactions above a threshold
- Time-limited authorizations that expire automatically

Agent Authentication and Verification (Network Level)

Payment networks serve as the trust anchor for the agent ecosystem. By requiring agent registration and verification before any transactions can occur, they create a closed ecosystem in which only vetted agents can transact. This is functionally analogous to the role that PCI-DSS compliance plays in securing merchant-side card acceptance.

The Trusted Agent Protocol (TAP) and Cloudflare's Web Bot Auth Protocol add a further layer: distinguishing legitimate AI agents from malicious bots at the network and edge-computing level, using behavioral intelligence and cryptographic attestation. (Visa, December 2025; AmericanExpress.io, November 2025)

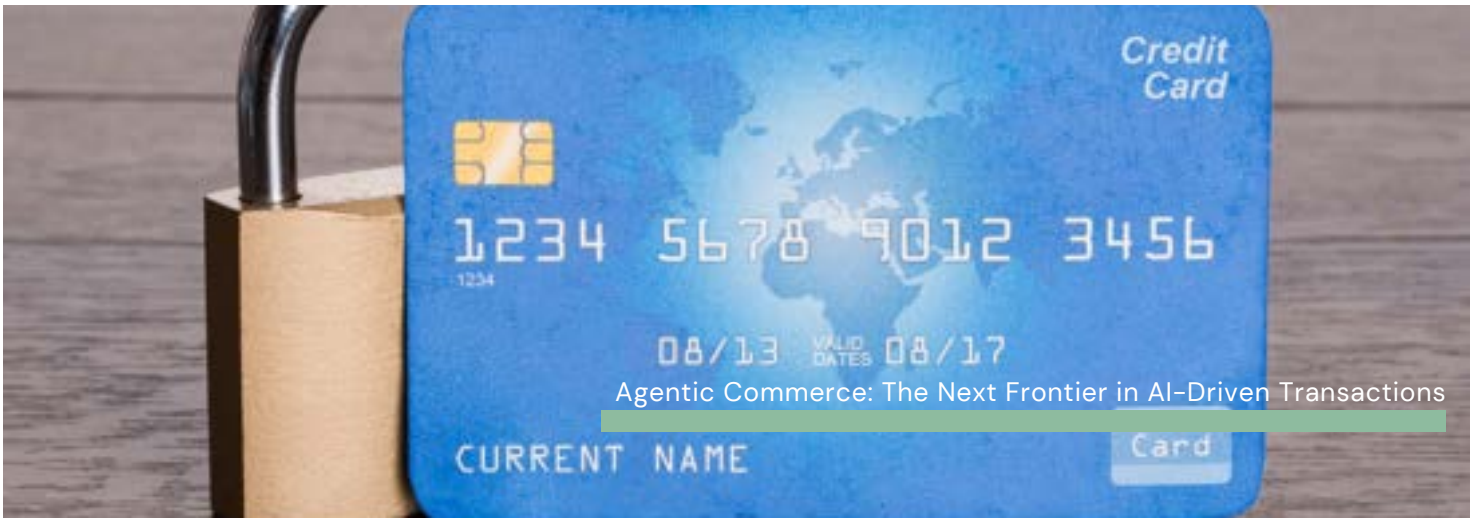

Intent Verification (Transaction Level)

The tamper-evident capture and transmission of consumer purchase intent through AP2's Intent Mandates, Amex's Intent Intelligence, or equivalent mechanisms creates a non-repudiable audit trail that links every transaction to an authenticated human decision. This serves three guardrail functions:

- Fraud prevention: agents cannot initiate transactions that deviate from expressed intent without creating a detectable inconsistency.
- Dispute resolution: the intent record serves as the reference against which agent behavior is evaluated.
- Regulatory compliance: the audit trail demonstrates that consumer-authorized agents, not unauthorized actors, initiated transactions.

Cart Verification and Human Confirmation

Most implementations include a human confirmation step before the final purchase, particularly for new or large transactions. In the Razorpay–Claude model, the agent presents a proposed cart for user approval before processing the payment. In the Amex model, the mobile app serves as a hub through which the Card Member confirms the agent's proposed purchase. (eMarketer, April 2026)

As trust accumulates and spending patterns are established, confirmation requirements may be progressively relaxed. The UPI Reserve Pay model of 'confirm once, transact many times within limits' represents one such graduated trust model.

Prompt Injection Defense

A significant technical risk in agentic commerce is prompt injection: malicious instructions embedded in product listings, web pages, or merchant content that attempt to manipulate an AI agent into making unauthorized purchases or exfiltrating data. Guardrails against prompt injection include:

- Treating all external content (web pages, product descriptions, merchant catalogs) as untrusted data that must be verified against the user's authenticated intent before any transaction is initiated.

- Sandboxed execution environments that isolate agent transaction logic from general content browsing.
- Network-level blocking of anomalous transaction patterns through the same fraud detection models that Visa has used for decades. (PYMNTS, May 2025)

Revocability and Transparency

Every implementation enables consumers to revoke agent authorization in real time. The Razorpay–UPI Reserve Pay model allows instant revocation of consent at the merchant level. American Express envisions a mobile app interface showing users "all the places in which your account is stored and all the places in which you have an outstanding intent," creating a unified dashboard for managing agent authorizations. (Digital Commerce 360, April 2026)

Full transaction visibility, such as real-time notifications, transaction histories, and clear identification of agent-initiated vs. human-initiated purchases, is a baseline requirement across all responsible implementations.

Fraud Detection at Scale

Visa's AI fraud detection infrastructure, which blocked approximately $40 billion in fraud in 2024, is being extended to cover agent-initiated transactions through Visa Intelligent Commerce. Mastercard's collaboration with OpenAI, Google, and Cloudflare is developing new safety and authentication standards for autonomous transactions. The challenge is that traditional fraud detection models optimized for human behavioral signals will need recalibration for AI behavioral patterns. (Digital Commerce 360, October 2025)

09 The Road Ahead

From Pilot to Mainstream

The trajectory from pilot to mainstream in agentic commerce is already mapped. Visa declared that “2025 will be the final year consumers shop and checkout alone,” and by December 2025, hundreds of controlled, real-world agent-initiated transactions had been completed across Visa’s network. Visa predicts that millions of consumers will use AI agents to complete purchases by the 2026 holiday season. (Visa, December 2025)

Agent-to-Agent Commerce

The Anthropic ‘Project Deal’ experiment and the emergence of protocols like AP2 point toward a future of agent-to-agent (A2A) commerce in which no human is present on either side of a transaction. In such a model, a buyer agent representing a consumer negotiates with a seller agent representing a business, with both operating within pre-authorized constraints set by their respective humans.

Visa’s Rubail Birwadker envisages this future: “In 2026, AI agents won’t just assist your shopping; they will complete your purchases.” (Visa, December 2025)

Convergence with Voice and the Metaverse

India already conducts over 1 billion voice searches every month, and Razorpay is enabling voice-driven commerce in which a spoken confirmation can trigger a Reserve Pay-backed transaction. (Razorpay blog, March 2026)

Looking further ahead, agentic commerce is likely to converge with immersive commerce in the metaverse, where virtual avatars already purchase items from brands like Nike, Adidas, Samsung, and Walmart on behalf of their human counterparts. The infrastructure for agentic payments, such as tokenized credentials, spending controls, and agent authentication, provides the same building blocks required for avatar-mediated commerce.

The Critical Role of Data Quality

A significant headwind noted by American Banker in its coverage of the Amex ACE launch is data quality: “In 2025, agentic commerce rode the tailwinds of large language model breakthroughs. Today, it faces the headwinds of messy merchant data and building consumer trust.” (American Banker, April 2026)

For AI agents to accurately represent products and make appropriate purchase decisions, merchant data, such as product descriptions, specifications, pricing, and availability, must be machine-readable, accurate, and consistent across platforms. Salsify’s research found that 45% of shoppers returned online purchases because product information was incorrect, a problem that becomes catastrophic when agents are making purchase decisions autonomously. (Salsify, 2026)

The Regulatory Race

The regulatory environment for agentic commerce will be shaped significantly by decisions made in 2026–2027. The CFPB’s rulemaking on representative authority, EMVCo’s work on agentic authentication standards, NIST’s public-private convening, and the CMA’s consumer law guidance will collectively define the compliance baseline. Organizations that engage proactively with these processes, as American Express and Google are doing, will have a significant advantage in shaping standards that work for their architectures.

10 Conclusion

Agentic commerce is not a distant prospect. It is operating today, in production, at scale, across the world's most important payment networks. The question is no longer whether AI agents will transact on behalf of humans, but how quickly the infrastructure of trust, standards, and regulation can be built to support that reality at a global scale.

The early evidence is encouraging. Payment networks have moved with remarkable speed to build agent authentication, tokenized credential delegation, and consumer protection frameworks. Open protocols, particularly Google's AP2, now under FIDO Alliance governance, are creating the interoperability layer that could prevent fragmentation. India's UPI infrastructure has proven that consent-based agentic payments can be deployed at scale within existing regulatory frameworks.

The challenges ahead are real but tractable. Consumer trust requires consistent positive experiences, clear disclosure, and meaningful control mechanisms. Regulatory frameworks require updating to address agent-initiated transactions without inadvertently stripping consumers of error resolution rights. Data quality must improve for agents to reliably represent products and make sound purchase decisions. And the technical guardrails against prompt injection, unauthorized agent behavior, and fraud must evolve as adversarial actors probe the new attack surface.

What is clear is that agentic commerce represents a genuine paradigm shift in how commercial transactions are initiated, authorized, and executed. Organizations across the commerce value chain, such as retailers, payment processors, AI developers, financial institutions, and regulators, must proactively engage with this shift. Those who do will help shape a trusted, interoperable, and consumer-empowering agentic commerce ecosystem. Those who do not risk being bypassed by it.

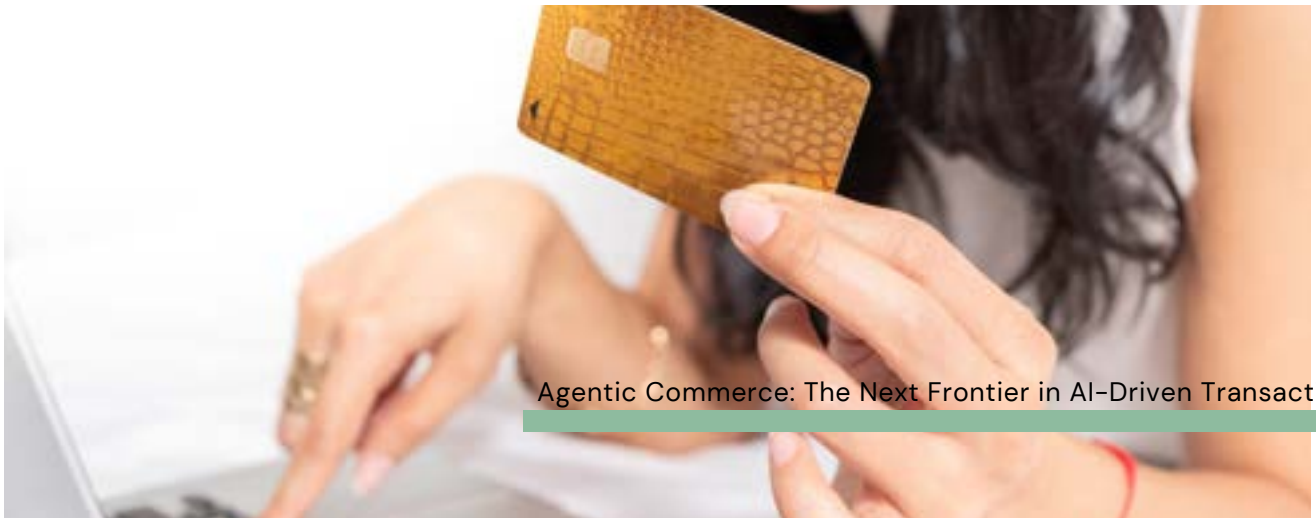

References and Citations

Payment Networks

- Visa (April 30, 2025). Find and Buy with AI: Visa Unveils New Era of Commerce. Visa Global Product Drop Press Release. corporate.visa.com
- Visa (December 18, 2025). Visa and Partners Complete Secure AI Transactions, Setting the Stage for Mainstream Adoption in 2026. Visa Newsroom. usa.visa.com
- Visa (April 2026). Visa Opens the Door to AI-Driven Shopping for Businesses Worldwide. Visa Investor Relations. investor.visa.com
- PYMNTS (May 14, 2025). Visa Gives AI Shopping Agents 'Intelligent Commerce' Superpowers. pymnts.com
- Mastercard (April 29, 2025). Mastercard Unveils Agent Pay, Pioneering Agentic Payments Technology to Power Commerce in the Age of AI. mastercard.com
- Mastercard (October 2025). Mastercard Agent Pay Acceptance Framework. mastercard.com
- Mastercard (December 2, 2025). Mastercard Unveils Agent Pay in Latin America and the Caribbean. mastercard.com
- Digital Commerce 360 (October 31, 2025). Mastercard: 'Agentic commerce is here.' digitalcommerce360.com
- PayPal Newsroom (October 27, 2025). Mastercard and PayPal Join Forces to Accelerate Secure Global Agentic Commerce. newsroom.paypal-corp.com
- American Express (April 14, 2026). Agentic Commerce Experiences | Developer Kit. americanexpress.com
- BusinessWire (April 14, 2026). American Express Debuts Agentic Commerce Experiences (ACE) Developer Kit and Announces Industry-First Protection for Registered Agent Purchases. businesswire.com
- Digital Commerce 360 (April 14, 2026). American Express launches developer kit, purchase protection for agentic commerce. digitalcommerce360.com
- American Banker (April 2026). Amex introduces agentic commerce developers kit. americanbanker.com
- The Paypers (April 2026). American Express rolls out agentic commerce kit and agent protection. thepaypers.com
- AmericanExpress.io (November 20, 2025). Shaping the Future of Agentic Commerce: How American Express is Helping Define the Standards of a New Era. americanexpress.io
- Digital Commerce 360 (May 6, 2025). Visa, Mastercard offer support for AI agents. digitalcommerce360.com

Protocols and Standards

- Google Cloud Blog (September 16, 2025). Announcing Agent Payments Protocol (AP2). cloud.google.com
- Google Blog (April 2026). Donating the Agent Payments Protocol to the FIDO Alliance. blog.google
- AP2 Protocol Documentation. Agent Payments Protocol (AP2). ap2-protocol.org
- Everest Group (January 12, 2026). Google's Agent Payments Protocol (AP2): A New Chapter in Agentic Commerce. everestgrp.com
- AP2 Lab (2025). Introducing the Agent Payments Protocol. ap2lab.com

Protocols and Standards

- Google Cloud Blog (September 16, 2025). Announcing Agent Payments Protocol (AP2). cloud.google.com
- Google Blog (April 2026). Donating the Agent Payments Protocol to the FIDO Alliance. blog.google
- AP2 Protocol Documentation. Agent Payments Protocol (AP2). ap2-protocol.org
- Everest Group (January 12, 2026). Google's Agent Payments Protocol (AP2): A New Chapter in Agentic Commerce. everestgrp.com
- AP2 Lab (2025). Introducing the Agent Payments Protocol. ap2lab.com

India Ecosystem

- Razorpay Blog (March 12, 2026). Razorpay & NPCI: Agentic Payments for UPI on Claude. razorpay.com
- Razorpay Blog (March 12, 2026). UPI Reserve Pay: Checkout Reimagined. razorpay.com
- BusinessToday (February 20, 2026). Razorpay, NPCI launch Agentic Payments on Anthropic's Claude; Zomato, Swiggy, Zepto Go Live. businesstoday.in
- YourStory (February 20, 2026). Razorpay's Harshil Mathur bets on AI to make commerce conversational again. yourstory.com
- BW Disrupt (February 20, 2026). Razorpay And NPCI Enable UPI Agentic Payments On Claude At India AI Impact Summit. bwdisrupt.com
- The Paypers (February 22, 2026). Razorpay and NPCI launch agentic payments on Claude. thepaypers.com
- Economic Times (2026). Platforms, brands accelerate agentic commerce push as fintechs plug payment gaps. economictimes.indiatimes.com

Compliance and Regulation

- Center for Data Innovation (March 23, 2026). Agentic Commerce is Coming, but Regulation Meant for Humans Will Slow It Down. datainnovation.org
- Cooley (March 26, 2026). AI Agents and Consumer Law: What Businesses Need to Know. cooley.com
- Morrison Foerster (May 2026). European Digital Compliance: Key Digital Regulation & Compliance Developments (May 2026). mofo.com
- Authority Partners (March 13, 2026). AI Agent Guardrails: Production Guide for 2026. authoritypartners.com
- eMarketer (April 2026). Amex debuts agentic shopper toolkit.emarketer.com

Industry Perspectives

- Praveenkumar G (2026). On Agentic Commerce and 4DS. LinkedIn.
- Ujjwal Keshri (2026). The Emotional Dimension of Shopping vs. Agent Efficiency. LinkedIn.
- Savannah Brentnall (2026). Citing Salsify 2026 Consumer Research. LinkedIn.
- Mehra Rajabova (2026). AI Adoption in Australian Retail. LinkedIn.
- Miraj Mor (2026). How AI Improves E-commerce Websites. LinkedIn.
- Salsify (January 21, 2026). 2026 Consumer Research — Daily Online Shopping Plunges, Shoppers Return to Certainty of Stores. salsify.com