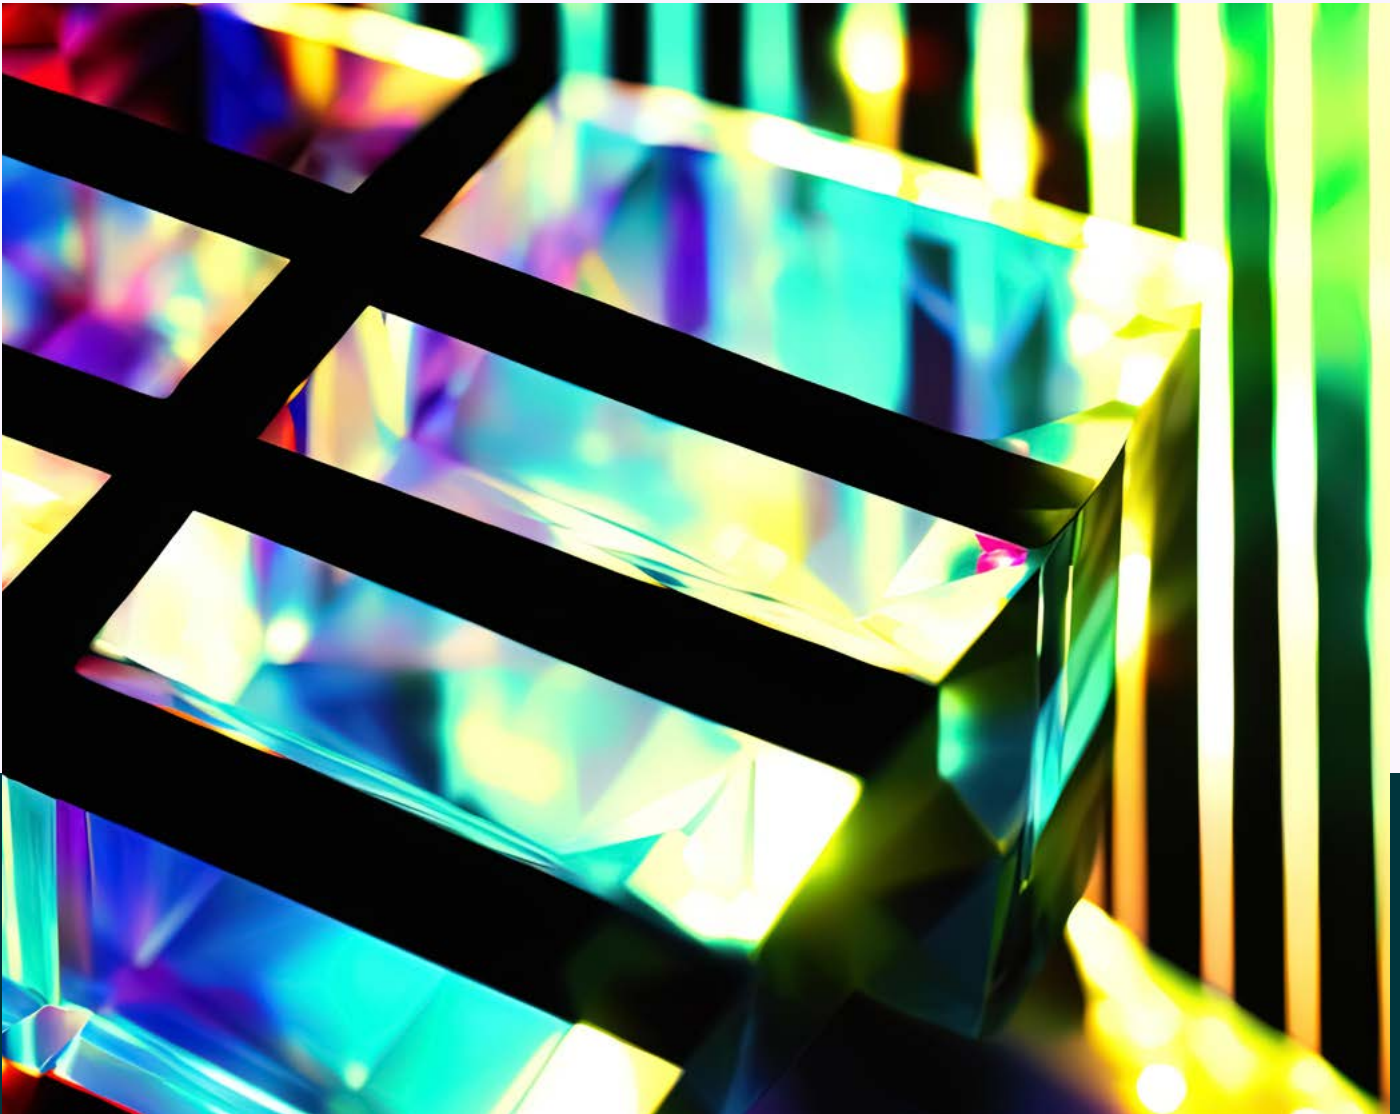

The Definitive Guide to Adopting Agentic Commerce in 2025

How merchants and ecommerce leaders can unlock growth with AI while building trust and safety



The Definitive Guide to Adopting Agentic Commerce in 2025

Table of Contents

1

Introduction 3

2

Building Readiness for Agentic Commerce 8

3

Building your Trust Stack 18

4

Where HUMAN Fits 21

5

Trust as the Accelerator of Growth 22



1.

Introduction

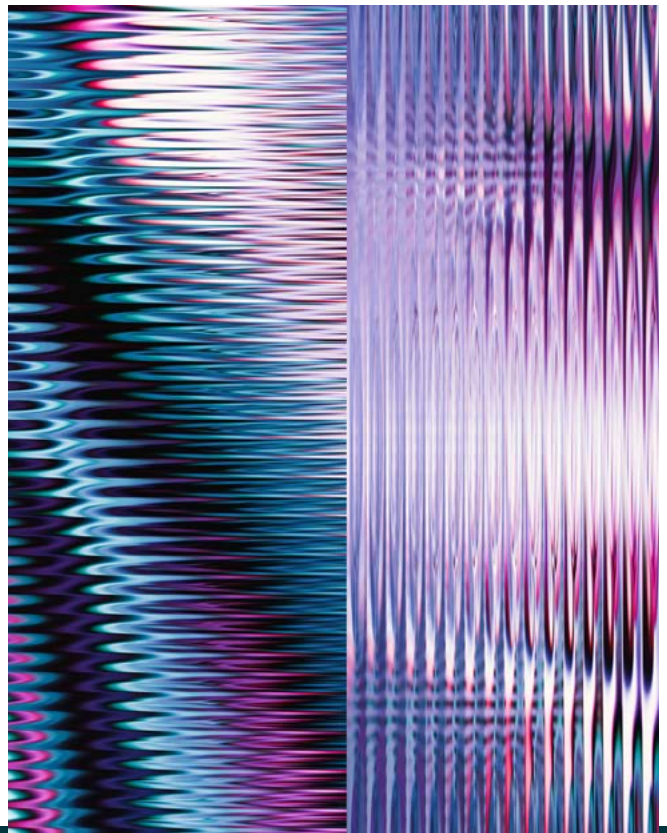
Every major digital shift has rewritten the rules of discovery and purchase. Search made information universally accessible. Social platforms reshaped how people discovered products through networks and influencers. Mobile put commerce directly into the customer's pocket. Each shift rewarded the businesses that adapted early—and left laggards behind.

We are now at the start of another transformation: agentic commerce. AI-powered agents are beginning to collapse the traditional funnel, moving directly from intent to purchase in machine time. Instead of navigating ads, search results, and checkout flows, customers increasingly rely on agents to find, evaluate, and transact on their behalf.

This creates a two-sided challenge. On the one hand, brands will be discovered and recommended only if their systems can be read, trusted, and transacted with by autonomous agents. On the other hand, they must also be able to trust the agents themselves—knowing which are legitimate partners and which are attempting to exploit systems at scale.

Winning in agentic commerce depends on more than just having the right data and APIs. Your catalog must be machine-readable to even be considered, your APIs must respond in milliseconds or risk being skipped, your infrastructure must scale for machine-speed queries, and your trust controls must block abuse without slowing legitimate transactions.

This guide provides a framework for building readiness on both fronts. We'll outline the practical steps merchants can take today, the emerging practices to prepare for, and the trust challenges that must be solved to unlock growth without exposing new vulnerabilities. The future of digital commerce is arriving quickly. Merchants that invest now in discoverability, reliability, and trustworthy governance will secure a place in the agent-driven economy; those that don't will see their visibility and sales shrink as AI agents favor better-prepared competitors.





Why Agentic Commerce Matters

The Collapsing Funnel

For decades, marketers have mapped the customer journey as a funnel:

1. **Awareness** – triggered by ads, social content, or word of mouth.
2. **Consideration** – shaped by reviews, blogs, and brand research.
3. **Conversion** – on a merchant site, marketplace, or app.
4. **Loyalty** – nurtured through service, promotions, and re-engagement.

Every stage required investment – ad spend, content marketing, retargeting – and each was a point where customers could drop off.

Agents compress this journey into a single interaction. A customer asks, “What’s the best running shoe for under \$150 that ships quickly?” The agent interprets intent, filters options, weighs reviews and policies, and presents one or two purchase-ready answers. Conversion can happen immediately, without the traditional multi-click path.

For merchants, this changes the dynamics entirely:

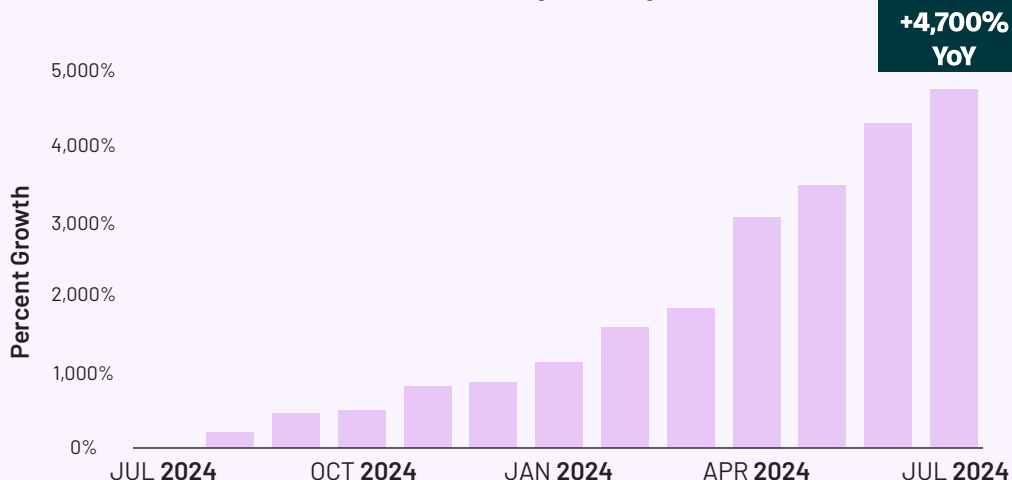
- **Fewer opportunities to intercept:** If you’re not in the first answer, you may not be seen at all.
- **More reliance on structured truth:** Agents will favor sources that are clear, factual, and authoritative.
- **Greater rewards for inclusion:** Once surfaced, friction is nearly eliminated, raising conversion rates.
- **Fewer chances for upsell:** Agents fulfill the exact request. If a shopper asks for shoes under \$150, the agent won’t suggest a t-shirt on sale unless such offers are structured and surfaced as clear value.

The Scale of Agentic Opportunity

Agentic commerce is no longer speculative. Major platforms already deploy shopping assistants that intermediate the journey—Amazon’s Rufus and Walmart’s Sparky among them—while general-purpose agents like ChatGPT Agent and agentic browsers like Perplexity Comet and Claude for Chrome influence what shoppers see first.

As this first round of AI agents become publicly available, consumer adoption is starting to move from curiosity

Growth in AI-Driven Visit Share (Retail)



Source: Adobe Digital Insights, July 2024–July 2025

to habit. Recent [research from Adobe](#) found that **38%** of consumers have used GenAI for shopping. Among those that have tried agentic shopping, **85%** said it improved the experience, and **73%** said it is now their primary source of research.

And when an AI assistant pre-qualifies a product and provides it as a direct answer, customers arrive with higher intent, and are therefore far more likely to convert. A [Semrush search study](#) found that users who arrive via AI-generated answers **convert 4.4x better** than those coming from traditional search results.

Analysts are also bullish: Gartner® predicts that “[by] 2028, AI agent machine customers will replace 20% of the interactions at human-readable digital storefronts, and that “by 2035, 80% of internet traffic could be driven by AI Agents.”¹

What Agentic Adoption Looks like Today

Our own telemetry confirms this acceleration. HUMAN verifies over 20 trillion digital interactions weekly, giving us a unique vantage point into traffic trends. In just eight months of 2025, **agentic traffic increased more than 1,300%**, driven largely by the releases of commercial agents like ChatGPT Agent and Perplexity Comet.

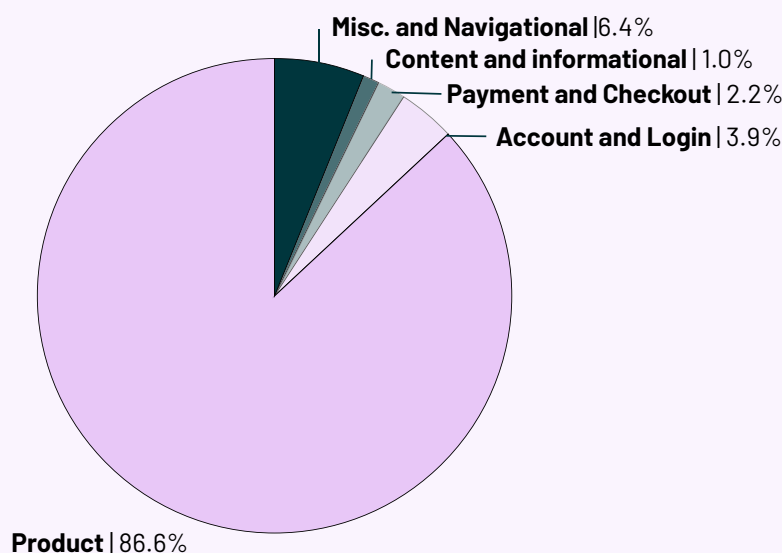
By examining the page paths and interactions of AI agents, we can observe that the vast majority of agentic interactions signal commercial intent. Approximately **87% of all pages browsed by agents were related to products.**

This means that blocking agent traffic turns away potential customers. Most agents are acting on behalf of real users; if they can’t access your product pages, they simply move on to another merchant that allows them to complete the task. In other words, every blocked request risks a lost sale.

However, our data indicates that the majority of these interactions are related to product research and recommendations, not full transactions. Just **2.2% of Agents** interacted with shopping carts, checkout, and payment funnels.

This aligns with current adoption trends in agentic commerce, where users primarily utilize agents for product research and comparison. Most commercially available agents, however, require human interaction to complete checkout flows. But adoption is accelerating quickly, and the next two phases are already on the horizon: autonomous purchasing and agent-to-agent negotiation. Understanding this trajectory is critical for merchants planning their readiness strategies.

AI Agent Website Interactions by Page Type



1

Three Phases of Agentic Adoption

The acceleration of agentic commerce is not uniform. Different use cases will emerge at different speeds, and several models will coexist for years. What matters is recognizing the trajectory: from discovery, to autonomous purchasing, to full agent-to-agent negotiation.

These phases are not clean handoffs. Adoption of Phase 1 is already in full swing, while pilots of Phase 2 are beginning to surface, and experimentation around Phase 3 is underway in enterprise and procurement contexts. Merchants should expect overlap: human checkout flows will remain dominant for some time, even as autonomous agents start to handle routine purchases, and specialized agents begin testing direct negotiations.

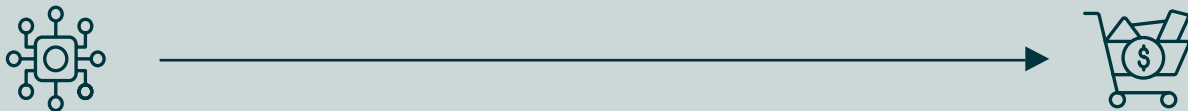
Phase 1: AI-Assisted Discovery - Current State



AI handles product research and recommendations, humans complete purchases through traditional checkout flows.

Example: You ask ChatGPT Agent to find a stand mixer on Amazon, it compares several options before recommending a specific appliance. You complete the checkout process manually.

Phase 2: Autonomous Purchasing - 2-3 Years



AI agents independently complete transactions using stored payment methods and user preferences, navigating existing e-commerce sites.

Example: Your Perplexity browser books flights and hotels for an upcoming vacation using your stored credentials.

Phase 3: Agent-to-Agent Negotiation & Purchasing - 5+ Years



Buyer-side agents negotiate directly with seller-side agents through APIs, bypassing human-facing interfaces entirely.

Example: Your procurement agent negotiates bulk pricing and shipping details with a supplier's agent.

New Frontiers for Commerce

While the collapsing funnel changes how customers arrive at your digital storefront, the next evolution of agentic commerce will shift where transactions happen altogether. Commerce is moving away from being a destination and becoming an embedded experience within the tools customers and businesses already use.

In-Situ Consumer Commerce:

Customers are beginning to discover and purchase products directly inside AI tools. For example, a user can ask Perplexity to find and book a hotel, completing the entire transaction within the app without ever being linked out to a merchant's website. This creates a new, high-conversion channel where your product can be bought the moment intent is expressed.

Agentic Browsers

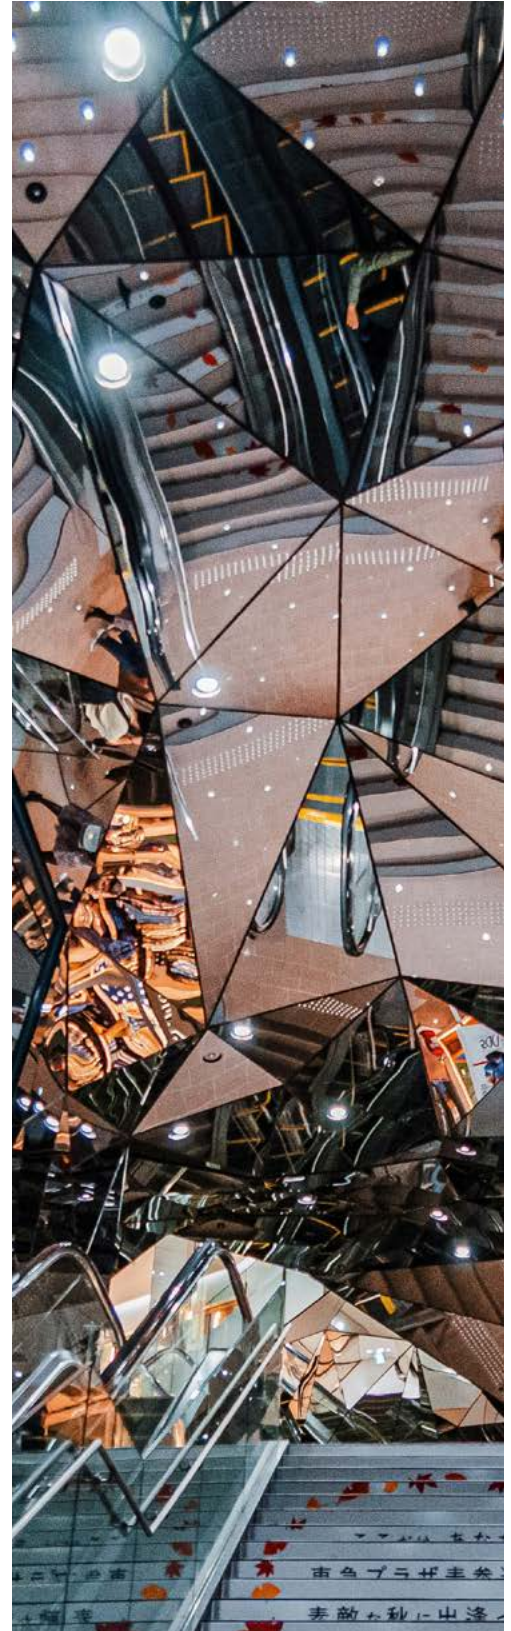
In parallel, AI-driven browsers, such as Perplexity Comet, are emerging that can autonomously navigate merchant websites on behalf of users: searching, selecting, and even purchasing without human clicks. Unlike in-situ commerce, these agents still rely on your existing storefront but consume it in a machine-first way. This creates both opportunity and pressure: opportunity to capture demand without requiring new marketplace integrations, and pressure to make your site agent-compatible by minimizing brittle front-end complexity, simplifying checkout flows, and exposing structured data or machine-facing endpoints to reduce errors.

Developer B2B Commerce

The opportunity extends beyond consumer products. Developers can now purchase services directly from their code editor, creating a new sales channel directly within the programming workflow.

Seller-Side Agents for Business Administration

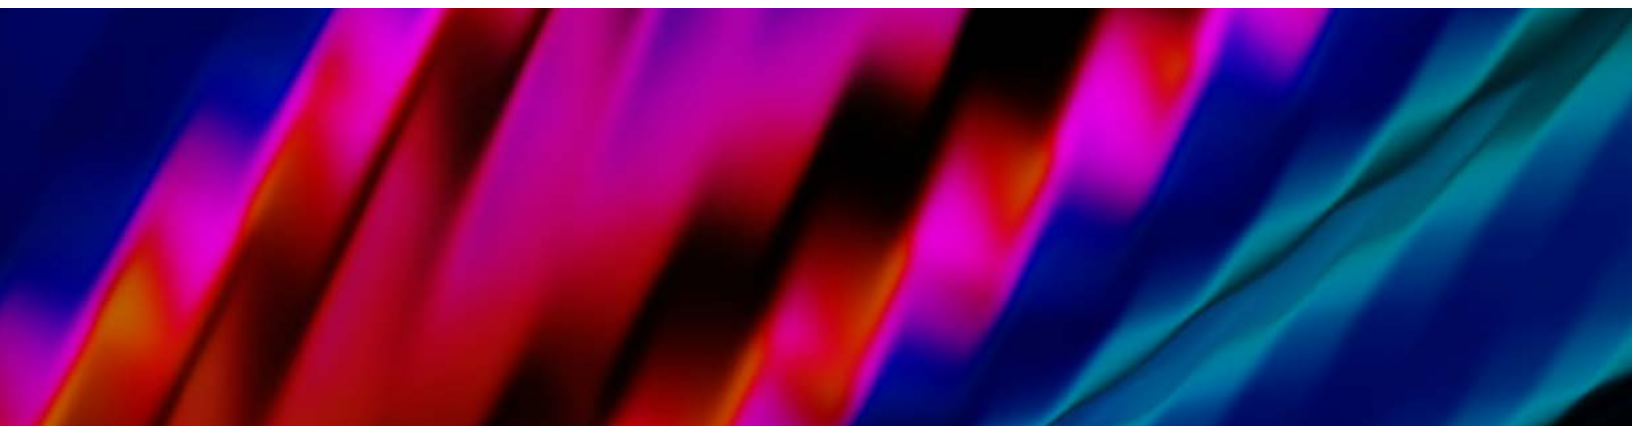
Agents are also handling commerce-related administrative tasks. For instance, customer support agents may integrate with payment processors via protocols like MCP to handle refunds, find invoicing information, or manage subscription cancellations without human intervention.



2.

Building Readiness for Agentic Commerce: A Nine- Step Guide

Success in this new era depends on whether your digital properties can be read, trusted, and transacted with by autonomous systems. The following framework brings together the essential disciplines for preparing your business to thrive in a world where agents drive discovery and purchase.



01

Making Your Catalog Machine-Readable

Agents can't recommend what they can't parse. Traditional web pages were built for people; agentic commerce requires product and policy data expressed in formats machines can reliably consume.

Structured Markup

At a minimum, every product page should emit JSON-LD using the schema.org vocabulary. This allows agents to extract product name, description, price, availability, identifiers, and reviews without guessing.

Identifiers and consistency

Each SKU must carry stable identifiers (SKU, GTIN, or MPN) across pages, feeds, and APIs. Inconsistencies, such as a GTIN missing in one place or mismatched SKUs between a feed and the site, create "ghost products" that agents can't reconcile and may drop altogether.

Policy coverage

Key terms like return windows, shipping details, and warranties should be explicitly marked up using the relevant schema properties. For example, return policies can be expressed in MerchantReturnPolicy objects linked from each offer. If your page copy says 30 days but your markup says 45, agents won't know which is correct and will likely exclude you.

Discovery surfaces

Sitemaps and product feeds should reflect the full live catalog and stay synchronized with your markup. Outdated or incomplete feeds are one of the fastest ways to vanish from agent answers. Robots.txt should allow discovery of product and policy pages while protecting sensitive endpoints like checkout and admin.

02

Build Visibility Into AI Traffic Signals

AI systems are already discovering, evaluating, and recommending your products, but most merchants have no idea when it's happening. When ChatGPT answers a shopping query or Perplexity provides product recommendations, they're pulling information through **RAG (Retrieval-Augmented Generation) bots** that crawl and fetch content in real-time. This traffic represents actual user intent, but remains invisible in traditional web analytics.

Unlike traditional search, where you see impressions and clicks, this happens in what we call the "zero-click discovery" layer. The only signal traditional web analytics might pick up is if the human ultimately clicks through, but LLM clickthrough rates are vanishingly small—0.67%, according to [recent research from Tollbit](#). LLM clickthrough rates are vanishingly small—0.67%, according to [recent research from Tollbit](#).

AI traffic as intent signal

Every AI visit represents potential customer intent. When "ChatGPT-User" or "Claude-Web" hits your product pages, a real person has likely asked a question that surfaced your brand. This is your visibility into the new top of the funnel.

These signals provide:

What this traffic tells you:

Real-time intent: When ChatGPT-User hits your product page, someone just asked about your category.

Discoverability confirmation: Regular RAG retrieval means you're in the consideration set.

Content quality signals: Patterns in what gets retrieved (and what doesn't) reveal how AI systems parse your content.

Category gaps: Products never accessed by RAG systems won't appear in AI recommendations — check structured data and crawlability.

Building measurement infrastructure

Create dedicated dashboards that separate AI traffic from both human visitors and traditional bots. AI scraper traffic isn't noise to filter out, it's a signal to amplify. Track:

- **User-agent identification:** Flag and segment traffic from known RAG bots (ChatGPT-User, Claude-User, Perplexity-User, etc.)
- **Volume and velocity of AI scraper visits**
- **Correlation between RAG access, subsequent referrals, and ultimately human conversion.**
- **Products with high AI traffic but low human follow-through (optimization opportunities)**
- **Categories never touched by AI systems (potential structured data gaps)**

03

Build Clear Interfaces for Agents

Modern agentic commerce requires merchants to support two distinct modes of interaction: browsing (product discovery, availability, and price checks) and purchase (secure initiation of checkout and payment). Treating these as separate phases allows merchants to give agents reliable, verifiable access without resorting to brittle scraping or UI automation.

While APIs and feeds are the long-term foundation for agent interactions, many agents today rely on browser automation to interact with existing websites. Merchants can improve outcomes for these agents by streamlining product detail pages, minimizing unpredictable DOM changes, avoiding interstitials and modal blockers such as pop-ups, and ensuring checkout flows are simple and linear.

Browsing Interfaces

Agents often begin by gathering product details before a purchase decision is made. Merchants should provide machine-friendly ways to retrieve this information, with minimal ambiguity and strong integrity guarantees.

Emerging best practices include:

- **Structured Catalog and Availability Feeds:** Product, inventory, and pricing data should be available through machine-readable feeds (JSON, XML, GraphQL).
- **Lightweight Read APIs:** Endpoints exposing prices, availability, shipping terms, and return policies should be simple to query and well-documented.
- **Signed Requests for Legitimacy:** Short-lived, signed request headers can help merchants and CDNs distinguish approved agents from bot traffic and control which operations are allowed on product detail pages.

Documentation and developer support

Agents increasingly come from third-party platforms—marketplaces, voice assistants, AI copilots. Clear documentation of your endpoints helps ensure your products are represented correctly. If integration is unclear, agents will default to competitors who make it easier.

Purchase Interfaces

When an agent moves from research to transaction, the interaction should shift to secure, state-changing endpoints that handle checkout flows safely.

Emerging best practices include:

- **Cart and Checkout APIs:** Allow agents to add items, apply promotions, and confirm orders programmatically, rather than forcing front-end automation.
- **Trusted Payment Initiation:** Provide a way to submit payment credentials securely, with validation to prevent replay or misuse.
- **Controlled Interact Endpoints:** Documented, authenticated APIs for actions such as placing orders, canceling, or initiating returns, so that agents have a reliable path to complete the transaction.

Emerging practice: agent landing zones

Some merchants are experimenting with consolidating browsing and purchasing interfaces into dedicated subdomains—what we call agent landing zones. While not yet industry standard, they illustrate a natural next step: creating intentional, agent-specific entry points instead of leaving systems to be reverse-engineered.

04

Optimize for Performance and Machine-Time Latency

Human customers are relatively forgiving of slow pages. Research shows bounce rates rise sharply after three seconds, but people will sometimes wait if they're invested. Agents operate differently. They are designed to move through large volumes of information quickly and select the fastest, most reliable source. If your response times lag by even a few hundred milliseconds, agents may skip your content entirely.

The Shift to "Machine Time"

In traditional e-commerce, performance has meant shaving seconds off load times. In agentic commerce, the standard tightens to milliseconds. Agents don't wait around—they will rank, summarize, and transact using whichever endpoint responds the fastest. Merchants must prepare infrastructure for this new latency budget.

What to prioritize

API optimization: Design endpoints that return exactly what agents need, without unnecessary payloads or redirects.

Global distribution: Use CDNs and edge networks to ensure low-latency responses regardless of where an agent request originates.

Caching strategies: Implement headers like ETag and Last-Modified to reduce redundant calls and accelerate repeat queries.

Load resilience: Prepare systems to degrade gracefully under traffic spikes. A partial response delivered quickly is more valuable than a complete response delivered too late.

05

Maintain Integrity and Consistency

Agents don't evaluate your brand in the way a human would. They look for patterns, reconcile signals across multiple sources, and favor whichever information appears most reliable. If your product catalog, policy pages, and feeds present conflicting facts, agents will down-rank you or ignore you altogether. Consistency isn't just a brand virtue in agentic commerce — it's a prerequisite for being surfaced at all.

Why Inconsistencies Create Risk

Contradictory policies: A return window listed as 30 days on one page and 45 days in a product feed creates uncertainty.

Price mismatches: Different values across markup, APIs, and human-visible content can be flagged as errors or fraud-like behavior.

Drifting product details: Attributes such as size, material, or warranty often get updated in one system but not another.

How to Protect Integrity

Single source of truth: Product and policy data should flow from one authoritative system into every channel (markup, feeds, APIs, and on-page content).

Governance processes: Establish checks that flag discrepancies between structured data and what's displayed on the site.

Policy harmonization: Standardize return, shipping, and warranty language so every system reflects identical facts.

Authentic reviews and ratings: Encourage verified, structured reviews. Suppressing or inflating reviews introduces noise that agents are increasingly adept at detecting.

06

Govern and Secure Transactions

Discovery is only half of agentic commerce. For agents to become truly useful, they must be able to complete transactions — adding items to a cart, applying promotions, processing returns, or checking loyalty balances. But opening these flows without governance invites exploitation. The challenge is to make transactions accessible enough for trusted agents while resilient against malicious automation.

Why Governance is Necessary

- **Exposed transaction flows:** Whether through APIs, automated browsers, or embedded checkout surfaces, any machine-facing interface can be probed by attackers as easily as by trusted agents.
- **Promotion and loyalty abuse:** Unchecked endpoints for coupons or rewards can be scripted at scale, draining value.
- **Refund fraud:** Automated requests can overwhelm systems if not monitored and controlled.

Key controls

Authentication and tokens: Where possible, establish agent identity through cryptographic HTTP signatures, OAuth 2.0, or tokenized identity. Emerging approaches include issuing short-lived tokens for search, payment, or intent declaration, giving agents scoped, auditable access to exactly the operations they need.

Rate limiting and replay protection: Prevent brute-force coupon testing, scripted returns, or repeated transactions.

Segmentation of access: Separate read-only endpoints (catalog, availability) from transactional endpoints (checkout, returns) and from administrative functions.

Monitoring and anomaly detection: Transactional flows should adapt based on how agents behave, not just who they claim to be. Track velocity and behavioral patterns in transactional traffic. Spikes, repeated failures, or agents attempting unpermissioned purchases should trigger alerts.

Note: Reconsidering Chargeback Protections:

Most fraud and chargeback tools were built for human commerce and rely on device fingerprints, IP reputation, household data, and cardholder history. These signals are often weak or missing when traffic comes through agents. Merchants should assess whether their systems can tell trusted agents from bots and adapt models to use agent identity, cryptographic signatures, and machine-specific behavioral baselines.

07

Monitor and Manage Agent Activity

Once agents are active on your systems, the work doesn't stop at setup. Visibility into their behavior is essential to capture value while protecting against misuse. Treating agent traffic as background noise is risky—without observability, you can't distinguish beneficial agents from malicious automation.

Why Visibility is Critical

- **Opaque traffic:** Without classification, agent requests look similar to bots or even human users, making it impossible to govern accurately.
- **Escalating abuse:** Malicious actors can mask themselves as legitimate agents, using access to scrape, drain promotions, or amplify fraud.
- **Optimization opportunity:** Beneficial agents can drive growth, but only if you can see which interactions lead to conversions and which stall.

What Good Monitoring Looks Like

Actor classification: Identify whether incoming traffic is human, bot, or AI agent, and segment reporting accordingly.

Agent trust scoring: Establish agent trust scores based on agent identity and behavior. Assign permissions based on trust and scoped intent.

Dashboards and analytics: Provide a consolidated view of agent activity—volumes, endpoints accessed, success/failure rates, and transaction outcomes.

Anomaly detection: Flag unusual spikes, abnormal session behavior, or repeated policy queries that may indicate scraping or abuse.

Synthetic probes: Deploy test agents that mimic real interactions to continuously validate the reliability and resilience of your websites, applications, and APIs.

08

Secure and Enrich First-Party Data

As agents mediate more of the customer relationship, first-party data is a competitive advantage. If the accounts, profiles, and loyalty records you hold are polluted with fake or low-quality entries, every downstream process—personalization, retention analysis, and agent optimization—suffers. Protecting this data while enriching it with genuine customer insights is therefore critical.

Why this Matters Now

Fraudulent accounts: Fake sign-ups inflate acquisition metrics and drain promotional budgets.

Credential stuffing: Automated login attempts can compromise accounts at scale, distorting both trust and analytics.

Agent mediation: If an AI assistant queries a polluted dataset, its recommendations will reflect inaccuracies.

Protecting the Foundation

- **Account integrity:** Harden sign-up and login flows against fake accounts and credential stuffing with layered authentication and anomaly detection.
- **Policy alignment:** Ensure that privacy, consent, and retention practices comply with regulations and are transparent to customers.
- **Data hygiene:** Continuously monitor for duplicates, stale entries, or suspicious clusters that skew customer insights.

Enriching the dataset

Zero-party data collection: Encourage customers to share preferences directly, in clear and value-exchanged ways (e.g., loyalty rewards for sharing dietary preferences or sizing information).

Loyalty programs: Incentivize authenticated, direct engagement rather than anonymous interactions.

Trustworthy analytics: With clean, verified records, merchants can power more accurate personalization and improve their chances of being included in agent recommendations.



09

Align on Trust as Infrastructure

Agentic commerce is not a siloed initiative. Success requires marketing, product, and security teams to work in concert. Without alignment, the risk is fragmentation: marketing optimizes for visibility, product builds new APIs, and security clamps down on threats—each in isolation. The result is friction, missed opportunity, or exposure to abuse.

Why Alignment is Critical

- **Marketing needs trust:** Inclusion in answer surfaces depends on structured truth and consistent policies, but it also depends on experience quality. Marketing teams need trustworthy data and signals to optimize campaigns, craft agent-readable offers, and design experiences that convert when surfaced. Without security's support, visibility can be undermined by fraud or inconsistencies.
- **Product needs trust:** All machine-facing entry points — from APIs to checkout pages to data feeds — must be carefully governed. If opened without oversight, they can become vectors for exploitation.
- **Security needs trust:** Security functions alone can't drive growth; they need to integrate their controls with business objectives to ensure protection fuels opportunity, not stifles it.

How to Embed Trust as Infrastructure

Cross-functional working groups: Establish teams that cut across marketing, product, and security to oversee agentic readiness.

Shared metrics: Track growth and safety together, including answer inclusion rates, API or agent adoption, and fraud prevention success, in a single source of truth that both product and risk teams use to guide decisions.

Leadership framing: Position trust initiatives as growth enablers, not compliance costs. This shifts the mindset from "blocking risk" to "unlocking opportunity safely."

3.

Building Your Trust Stack

As we have seen, agentic commerce presents revolutionary growth opportunities, but growth at scale requires trust. Without a trust foundation, brands and merchants risk invisibility in agent-driven commerce, polluted funnels and broken attribution, fraud and abuse at scale, and ultimately, competitive disadvantage.

The solution is to build a **trust stack**: a layered set of capabilities that ensures every interaction—whether human, bot, or agent—is understood, governed, and safe. This gives merchants the confidence to let beneficial agents in, while shutting down malicious automation before it causes harm. This also makes it easier for agents to trust the merchant's brand as a reliable source—increasing the likelihood that their products are discovered, recommended, and purchased.

Why a Trust Stack Matters

Digital commerce has always been built on trust: payments must be secure, reviews must be authentic, and product information must be accurate. In the agentic economy, that trust must also apply to AI systems that search, recommend, and initiate transactions on behalf of customers.

The trust stack answers three fundamental questions for brands and merchants:

1. **Who or what is interacting with me?** Can you identify every actor across sessions and accounts? Can you classify humans, bots, and AI agents, and attribute them to known providers or platforms?
2. **Is this interaction legitimate?** Can you determine whether behavior matches real customer intent, complies with business rules, and avoids fraud or exploitation?
3. **How do I enable the beneficial while stopping the harmful?** Can you apply policies that permit, block, or throttle actions in real time? Can you let beneficial agents transact while preventing abuse and protecting revenue?

When these answers are built into the system, merchants can safely open their sites and APIs for growth. Good agents can transact freely, bad actors are blocked, and every interaction drives measurable value.

The Four Layers of the Trust Stack

A modern trust stack works as a system: each layer builds on the last. Together, they give merchants the ability to open their data and transaction surfaces to agents without losing control or exposing themselves to abuse.

1. Identity and Data Trust

Trust begins with knowing exactly who or what is interacting with your systems and ensuring the information they act on is reliable. This layer establishes a clean, consistent foundation for both human and non-human participants.

What this requires:

Verified consent and privacy:

Make privacy and consent signals transparent, machine-readable, and aligned with regulations so both agents and humans can trust the data they receive.

Structured data quality:

Keep product information, policies, and pricing consistent across pages, feeds, and APIs. Inconsistent or incomplete data creates uncertainty that can lead agents to drop you from recommendations entirely.

Actor integrity:

Accurately distinguish between humans, bots, and AI agents at the point of interaction. Where possible, attribute agents to known providers or platforms to understand who is driving requests.

Account integrity:

Protect signup and login flows against fake accounts, credential stuffing, and synthetic IDs that distort metrics and drain promotional budgets.

2. Commerce and API Trust

As agents shift from discovery toward transacting, your APIs and structured feeds become the new storefront. They need to be accurate, secure, and resilient enough to serve both high-volume queries and sensitive transactions.

What this requires:

API security and governance:

Authenticate requests with scoped tokens or signatures, enforce rate limits, and detect anomalies before they cascade into exploitation.

Structured and controlled data access:

Provide complete, machine-readable product and policy feeds (e.g., JSON-LD, catalog APIs) while controlling who can query them and at what frequency.

Transactional governance:

Protect loyalty points, discounts, and refund systems from automated abuse. Segment read-only from transactional endpoints and enforce permissions so that only trusted agents can add to carts, redeem offers, or initiate returns. HUMAN ensures that beneficial agents can query it, while scrapers and fraud farms cannot drain the system.

3. Monitoring and AgentOps

Once agents are active in your systems, visibility and control become continuous requirements. You need to know which agents are present, what they are doing, and whether their behavior aligns with your policies. This layer turns agentic commerce from a black box into something you can observe and govern in real time.

What this requires:

Behavioral baselining:

Learn what “normal” agent activity looks like across search, add-to-cart, and checkout flows, so you can quickly flag anomalies like scraping, mass carting, or velocity abuse.

Real-time governance:

Apply policy controls dynamically — allow trusted agents to proceed, throttle excessive requests, and block harmful behaviors before they impact users or revenue.

Operational visibility:

Give product, security, and marketing teams a shared dashboard of agent interactions, including volumes, success rates, and policy enforcement outcomes.

Synthetic testing:

Run test agents that mimic real usage to validate that your APIs and controls are resilient and that trusted agents can still complete their tasks.

4. Fraud and Abuse Defense

Even with clean data, secure APIs, and active monitoring, determined attackers will keep probing for gaps. This layer is your persistent shield against the most sophisticated, adaptive threats—the ones designed to exploit open systems at scale.

What this requires:

Advanced bot management:

Detect and stop automated traffic attempting credential stuffing, account takeovers, scraping, or inventory hoarding.

Promotion and loyalty protection:

Block large-scale coupon drains, loyalty exploitation, and fake conversions so marketing spend and incentives drive real outcomes.

Transaction verification:

Validate that purchases, clicks, and redemptions are tied to legitimate customers, ensuring that marketing spend drives real value.

Adaptive defenses:

Continuously evolve detection logic to respond to new fraud patterns and attacker tactics.

When all four layers work together, they form a self-reinforcing loop: visibility allows you to open systems with confidence, which drives real growth, which strengthens trust signals that attract better customers and agents. This cycle accelerates over time.

The Roadmap to Building a Trust Stack

You don't need to deploy everything at once. But waiting too long at the early phases risks losing discoverability and opening the door to fraud. The roadmap below prioritizes the steps that protect visibility first, then builds toward full agent governance.

Phase 1 – Foundation

Stand up your structured product data, policy markup, and account integrity controls. This ensures agents can read your catalog and trust your signals, while keeping fake signups and credential abuse from polluting funnels.

Phase 2 – Enablement

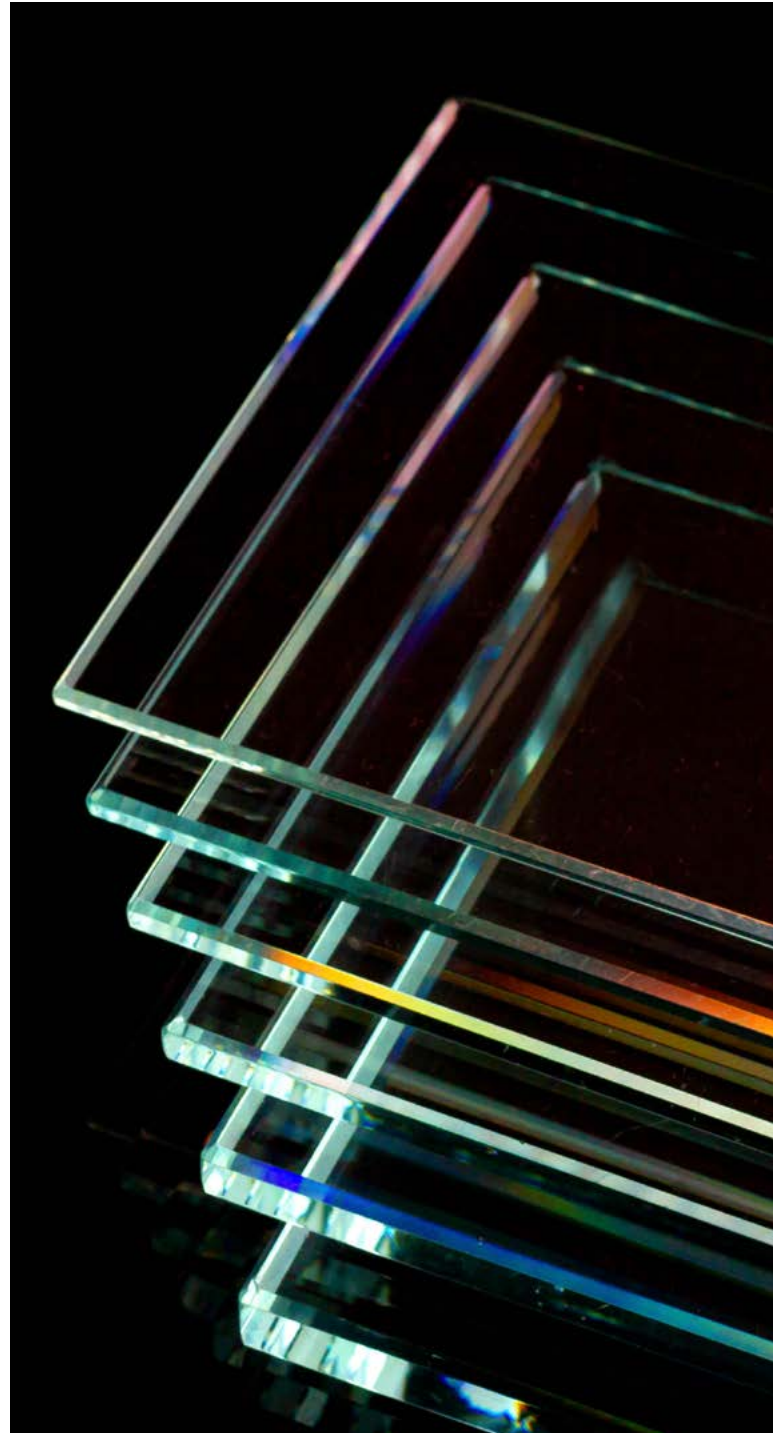
Open APIs and product feeds for machine access, with basic authentication and rate limiting in place. Begin measuring agent traffic separately to see where you're appearing in the new discovery layer.

Phase 3 – Governance

Deploy monitoring dashboards, classify agent traffic, and start applying real-time policy controls. This lets you safely open transactional flows — like carting and returns — without fear of abuse.

Phase 4 – Optimization

Tune policies based on outcomes. Align security and marketing around shared metrics like inclusion rates and conversion from agent-driven visits. Expand defenses to cover loyalty, promotions, and refund fraud as volume scales.



4.

Trust as the Accelerator of Growth

Agentic commerce is not a future concept. It is happening now. Customers are already asking AI assistants to make recommendations, copilots are surfacing products directly inside shopping experiences, and enterprises are beginning to delegate complex workflows to autonomous agents. This is the new interface of commerce: conversational, contextual, and powered by AI.

For brands and merchants, this shift is more than just another channel to optimize. It is a transformation of the entire digital journey. Discovery no longer happens only in search engines or marketplaces. Consideration no longer requires five clicks and three comparison sites. Conversion may not even take place on your website. Instead, purchase decisions are being shaped inside AI-driven conversations, where visibility and trust are the deciding factors.

The businesses that thrive in this new era will be those that see agentic commerce as a positive driver of growth. They will:

- Structure their catalogs and policies so that agents can easily surface them.
- Treat marketing as a discipline of agent optimization, securing presence in AI answers.
- Align security and marketing around a shared trust layer, ensuring that visibility leads to real, sustainable revenue.
- Invest in AgentOps and governance, so that beneficial agents are supported while malicious actors are controlled.

The opportunity is enormous. The next decade of digital commerce will not be won by those who wait and see. It will be won by those who act now.

Agentic commerce can expand reach, compress the funnel, and deliver personalized, frictionless experiences at scale. But the key enabler is trust. Customers must trust that the answers they receive are accurate. Merchants must trust that the traffic they see is genuine. Agents themselves must trust the data sources they rely on.

Trust is not a brake in innovation. **Trust is the accelerator.** With trust in place, merchants can confidently *embrace agentic commerce, unlock new growth, and build lasting loyalty* in a world where humans and AI agents shop side by side.

5.

Enabling Agentic Commerce with HUMAN

Agentic commerce is already shaping how customers find and decide. The practical question is whether agents will see you, trust you, and complete the actions that matter. That requires two things at once: clear visibility into non-human actors and the ability to govern what they can do in real time.

HUMAN provides the trust layer that operationalizes the framework in this guide. It makes every interaction visible, governed, and aligned with business outcomes so you can open systems to agents without losing control.

What you Gain with HUMAN

- **Visibility:** Session-level insight into how agents browse, call APIs, and attempt transactions, from discovery to checkout.
- **Governance:** Granular, real-time controls that let you allow trusted agents to act, throttle excess, and stop abuse before it affects customers or revenue.
- **Resilient growth:** Confidence to expose machine-readable catalogs and open APIs, protect promotions and loyalty value, and turn zero-click discovery into measurable outcomes.

How HUMAN Delivers

HUMAN brings the trust stack to life. Our platform combines proven cyberfraud defense with purpose-built agent governance, so merchants can stay open to opportunity without losing control.

Cyberfraud Defense

HUMAN Sightline Cyberfraud Defense protects accounts, funnels, and promotional value at internet scale. It blocks credential abuse, synthetic signups, scraping, and incentive drain before they distort metrics or deplete budgets – ensuring that every dollar spent and every customer acquired is real.

AgenticTrust

AgenticTrust extends this foundation to AI agents. It classifies and attributes agent traffic at the session level, gives your teams a live view of what agents are doing, and applies granular permissioning and policy controls. The result is an environment where trusted agents can browse, compare, and complete transactions while malicious automation is stopped in its tracks.

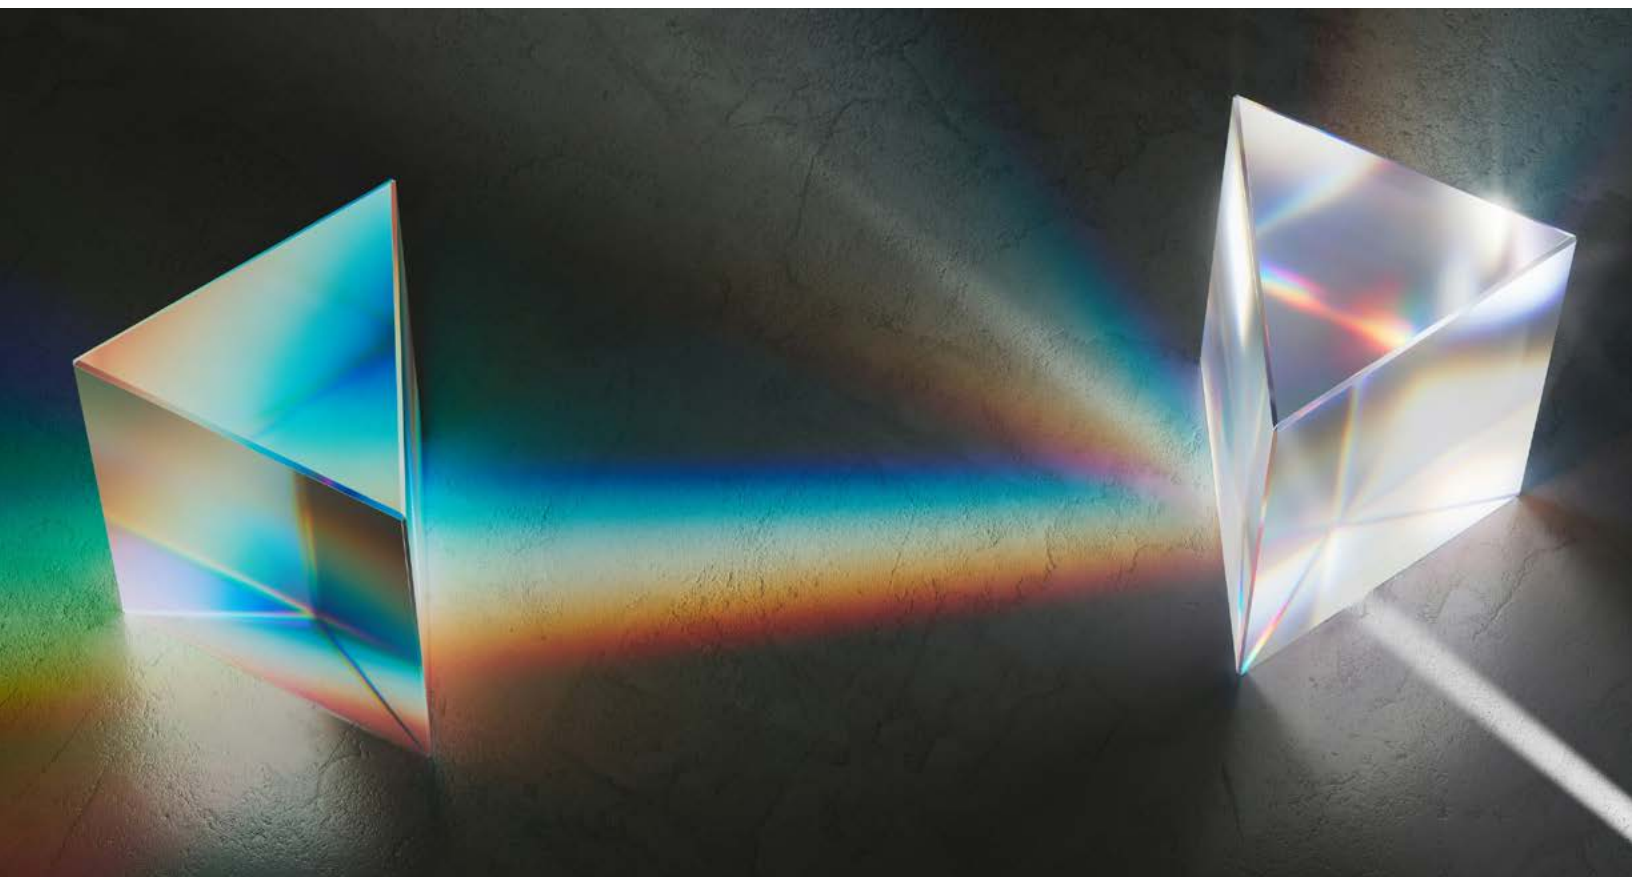
From Risk to Advantage: Embracing the Agentic Opportunity

Most merchants today respond to AI agents with either overexposure or blunt blocking. Both approaches leave money on the table, either by allowing exploitation or by shutting out legitimate demand.

HUMAN makes a third path possible: governed participation. You can see every agent session, apply policies at the point of action, and open APIs with confidence. Your product catalog becomes machine-readable without being scraped. Your promotions reach real customers instead of automated farms. And third-party agents don't get to set the rules for how your data and systems are used, you do.

Agentic commerce is still in its early stages, which means the window to shape it is open. Merchants who prepare now will become the preferred sources for trusted agents. With a trust stack in place, you can lead this shift, opening systems safely, capturing early share, and building durable advantage while others hesitate.

With HUMAN, trust stops being a cost center and becomes your competitive edge. Agents can find you, act within your rules, and drive measurable revenue — not risk.



About HUMAN



HUMAN is a leading cybersecurity company committed to protecting the integrity of the digital world. We ensure that every digital interaction, transaction, and connection is authentic, secure, and human. HUMAN verifies 20 trillion digital interactions, providing unparalleled telemetry data to enable rapid, effective responses to the most sophisticated threats. Recognized by our customers as a G2 Leader, HUMAN continues to set the standard in cybersecurity. For more information, visit www.humansecurity.com.