ACTIVANT RESEARCH

# Agentic Commerce

## Who Gets To Transact

July 2026

AUTHORS:


Steve Sarracino
Founder & Partner

Torben Wiesbach
Senior Research Analyst

Scott Watson
Vice President, Research


**About Activant**

Activant is a research-led global investment firm that partners with high-growth companies. Since 2015, we have invested in category-defining businesses during their most critical phases of growth, partnering with founders who have won their initial battles and are ready for the next challenge.

Our approach pairs deep, proprietary research with patient, flexible capital and hands-on operational partnership. We work alongside founders and leadership teams to refine strategy, strengthen operations, and accelerate sustainable growth.

Activant Research is dedicated to uncovering the most exciting emerging technologies, sectors, and companies we believe will shape the future. Our research-driven perspective informs everything we do, helping us invest at meaningful inflection points and support founders in building enduring, category-leading businesses.

You can find out more about Activant and our research at https://activantcapital.com/.

In the first piece of this series, we argued that agentic commerce is not a single wave but a stack of three layers and that the enablement layer is where near-term capital is flowing because nothing above it works until the plumbing is built. This piece goes inside that layer: who is building the bridge between capable agents and checkout pages that were never designed to talk to them, what they are building, and where the returns in this cycle are most likely to accrue.

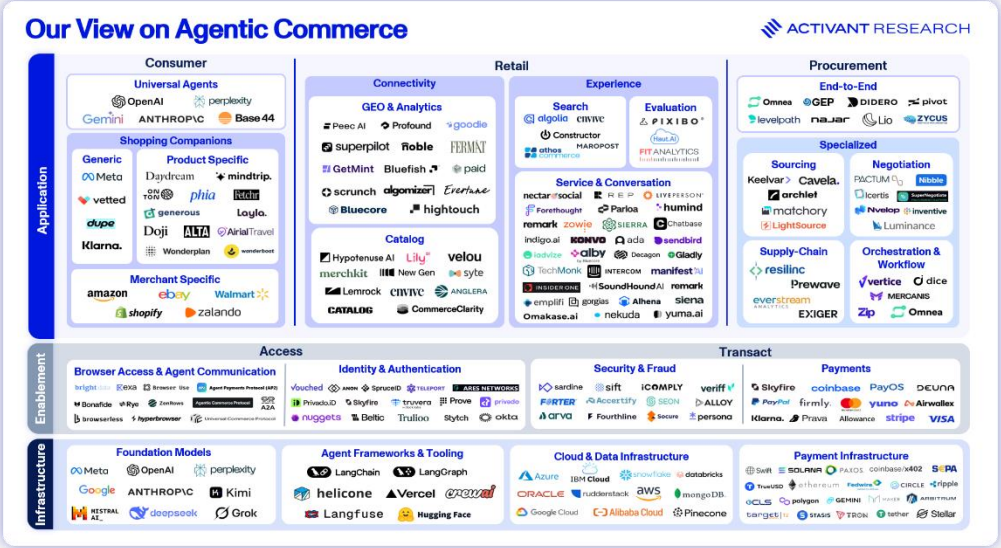


The short version: the enablement layer is not one market. It is four overlapping problems being solved by different categories of companies, with different return profiles and different timelines. ***Access*** (can the agent reach the checkout?) is the most tractable and closest to being solved. ***Identity*** (can the agent prove it is acting on behalf of a verified human with spending authority?) is harder because it sits between parties with conflicting incentives. ***Settlement*** (can the transaction complete, with the right money moving to the right place?) is being rebuilt from scratch for a signal environment that traditional payment rails were not designed for. And ***governance*** (who is liable when something goes wrong?) is not a technical problem at all, which means the companies that succeed in this cycle will be the ones who design around a legal regime that does not yet exist.

Agentic commerce is not a single state. It exists along a spectrum of autonomous commerce, and where we sit on that spectrum today is largely determined by what the enablement layer already supports and what it does not. In the current state, agents assist. You ask a chatbot for a recommendation on noise-cancelling headphones that cost less than $300. The model either draws on its training data or performs a web search to compare specs, and surfaces three or four options. You read the output, click a link, and buy it yourself on the retailer's site. The agent managed the research. You managed the shopping. This is already mainstream. Roughly 49% of AI-assisted shoppers said they would consider a different brand

or product if an AI assistant recommended it.[1] But calling this agentic commerce is generous. The agent is a better search engine. The transaction still runs on human rails.

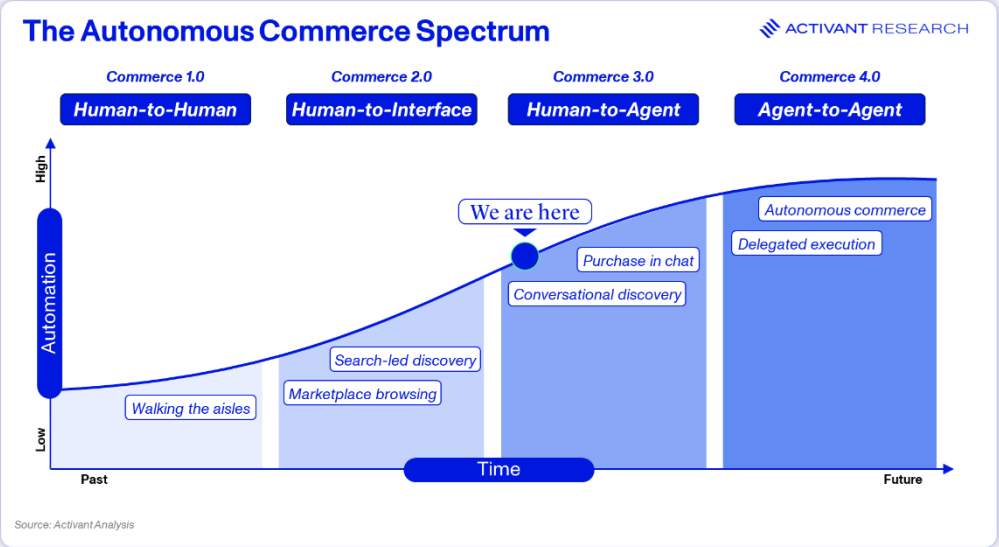


Further down the road is where things will get genuinely interesting. Agents fill carts, apply discount codes, navigate checkout flows, enter shipping details, and complete payments.

Market makers and high-frequency trading firms are one of the oldest examples of agent-to-agent commerce. Firms like Citadel Securities, and Two Sigma deploy algorithms that quote, lift, and hit prices on exchanges in microseconds. The interaction between a market-making algorithm and an order-routing algorithm from a broker is genuinely agent-to-agent: no human approves individual trades.

Programmatic advertising is another agent-to-agent negotiation happening today. When you load a webpage, a buyer-side agent bids against other agents in an auction run by a supply-side platform. Major players are involved: The Trade Desk, as well as Google's DV360 and Microsoft's Xandr. The negotiation is price-only, which simplifies things considerably. Agents bid on a fixed good (an impression) with no ability to negotiate terms, bundling, or delivery conditions. Effectively it's an auction, not a true negotiation.

# Removing the Human in the Loop

Software negotiates with software, payment settles, delivery is initiated, and no human is involved. Those markets were built with machines as the primary participants from the start. Consumer commerce was not. Every technical decision made in online retail over the past 25 years assumes a human is on the other end of the transaction: CAPTCHAs that require visual pattern recognition, login walls that trigger on unusual session behavior, dynamic pricing flows that detect non-human traffic, and payment authentication steps that demand biometric confirmation. None of it is malicious. It was built to prevent fraud. It also makes autonomous agent access unreliable, which means the near-term commercial opportunity in agentic commerce is currently stuck between a capable model and a checkout page that does not know how to talk to it. The bridge gets built in four places.

## Access: The Browser Versus API Bet

In Part I, we introduced the fork between browser-based access and API-native protocols. Both are live, and neither wins outright. The question worth going deeper on is not which approach wins. It is whether merchants want to participate, and what happens when they decide they do not.

On one side, browser-based access tools let an agent navigate the existing web the way a human would, clicking through pages and filling forms programmatically. Browser Use and Browserless are the two most visible builders in this category. The virtue of the browser approach is reach. Any site an agent can load, it can attempt to transact on without the merchant doing anything. No integration, no cooperation, no waiting for retailers to ship API endpoints. For an agent trying to serve the long tail of merchants who will never prioritize agent integration, browser-based access is the only option.

On the other side, API-native checkout requires the merchant to participate. Merchants expose structured endpoints that agents can call directly, which is cleaner, faster, and more reliable than any browser workflow.

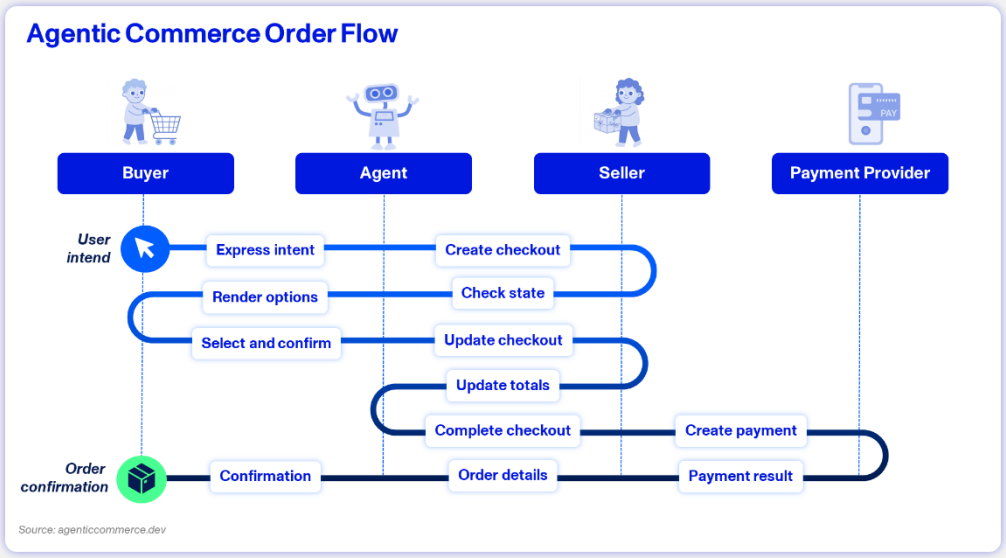


Before either protocol existed, every AI agent needed a custom integration with every merchant. That's an N×N problem that doesn't scale. The Agentic Commerce Protocol (ACP) and Universal Commerce Protocol (UCP) both solve it, but through different architectures that reflect the ecosystems backing them. The ACP developed by OpenAI and Stripe is centralized. It is tied to ChatGPT and Stripe's payment rails, and its scope is deliberately narrow: standardizing checkout session management and payment tokenization within an AI chat interface. That constraint is also its advantage. Merchants already on Stripe can integrate quickly, sometimes in a single configuration change.[2] The tradeoff is coverage. The UCP from Google, co-developed with Shopify, Walmart, and Target is the alternative. It is decentralized and allows merchants to publish a manifest at a standardized endpoint detailing their supported services. Any compliant AI agent can read and negotiate against it. The protocol covers the full commerce journey: discovery, cart management, checkout, loyalty, and post-purchase. More surface area means more implementation work, but also more reach.[3]

The two protocols are complementary because they serve different ecosystems. ACP routes through OpenAI's platforms. UCP routes through Google's. A merchant running only one is invisible to agents operating on the other. Dual-stack configurations are increasingly becoming the default for large retailers like Walmart, which now attributes 20% of its referral traffic to ChatGPT.[4] The gap in agent-driven traffic between single-protocol and dual-protocol merchants appears to be widening. R[AI]sing Sun for example claims a 40% traffic uplift for dual-protocol vs single-protocol merchants.[5]

The actual question now is whether merchants want to participate at all. Merchants have weak incentives to open clean APIs. When they do, they lose the behavioral data that feeds retail media targeting and the incidental revenue that

comes from browsing. A consumer shopping for a blender might leave with a blender and a set of knives. An agent sent to buy a blender buys a blender, with no cross-sell, no impulse, no ad inventory impression. This is a real revenue hit, and it is why the first wave of merchant reactions to agentic commerce has been cautious.

The tension is eventually unstable. As the Amazon-Perplexity dispute demonstrated, agents will attempt to transact regardless of whether a merchant has built for it.[6] If merchants refuse to expose agent-friendly access, they do not prevent agent traffic. They just lose control over it. The choice is not whether to participate. It is whether to participate on your own terms or someone else's.

*The future will most likely be API-based. Most people building in this space agree on that. But merchants have not built for agents yet, so the browser approach wins today by default. In the end the consumer does not care whether the agent used a browser or called a clean API. They care whether it worked and whether it was fast. The protocol is invisible to them."*

Dasmer Singh
Founder & CEO, Allowance (YC P26)

Our read is that the browser-versus-API fight resolves as both. The largest merchants and the highest-volume categories move to API-native integration because the economics justify it. According to Salesforce, traffic referred by AI agents converted eight times higher than traffic from social media platforms.[7] This conversion premium pays for integration work quickly. But one should keep in mind that there are over 30 million e-commerce domains worldwide with 46% of stores offering less than 10 products.[8] It is unrealistic to believe that in the near term all of them will implement agent-ready APIs. Browser access tools are not a bridge to the API world. They are the permanent infrastructure for the long tail.

Access gets an agent to the checkout page. It does not get the transaction completed. That is a harder problem.

## Identity and Fraud: When the Signals Disappear

Current authentication systems assume the entity logging in is the same as the entity making decisions. Agents break that assumption, and the breakage cascades through every layer of the payment stack.

The first casualty is fraud detection. Device signals, behavioral biometrics, IP addresses, and clickstream patterns either disappear or reflect the agent rather than the user. Every fraud model built on those signals degrades the moment an agent is in the flow. This is not a gradual problem. Fraud models go from useful to actively misleading because an agent's behavior looks, by every metric these models were trained on, like automated fraud.

The second casualty is the verification chain. No single party sees the full picture of an agent-mediated transaction. The AI provider verifies the user at onboarding. The merchant sees only the agent. The card issuer works from signals that no longer reflect human intent. Spoofed agents, upstream account takeovers, and 2FA timing failures all exploit the gaps between actors rather than within any one system.

The third casualty is liability. AI providers have no regulatory obligation to share risk-relevant data with merchants. Agent mistakes and friendly fraud create chargebacks merchants cannot contest. Agentic behavior quietly degrades users' risk profiles over time, generating false positives on future legitimate transactions. The exposure is real, and it is largely invisible until the damage is done.

The open question is whether models built on years of human behavioral data can be extended far enough to handle traffic that looks nothing like what they were trained on. Agent-to-agent payments and agentic commerce are about to reshape how money moves — the biggest step-function shift since peer-to-peer payments. And like every such shift, it opens gaps: today, agentic browsers all look like bots, so a legitimate, fully authorized purchase gets blocked as a false positive. The answer is to move fraud detection from the point of checkout to the point of intent, and to give both merchants and issuers smarter controls.

Another critical concerning point for issuers is how to prevent false chargebacks from customers who might come back and claim that they never authorized an agent to shop on their behalf. To answer this, Sardine (an Activant portfolio company) provides a novel product, an Access Control Server (ACS) that is made for the agentic economy.

*Traditional ACS solutions simply execute the 3DS protocol. Sardine's ACS goes further: it combines device intelligence from the 3DS Method URL with our fraud graph and behavioral biometrics to recognize when a trusted agent is acting on a real customer's instruction — approving more legitimate transactions, reducing unnecessary challenges, and stopping advanced agent-driven fraud, all while maximizing liability shift."*

Soups Ranjan
Co-Founder & CEO, Sardine

This means that when a customer has directed an agent to shop on their behalf, the issuer can confidently let it through; when something looks off, it can step up with strong customer authentication.

Getting fraud infrastructure agent ready means extending existing human-behavior models to cover agentic traffic, running separate models trained on agentic data, layered on top of existing know-your-customer (KYC) and risk-based models.[9] Essentially running agentic commerce and payment fraud models in parallel, routing transactions based on the nature of customer.

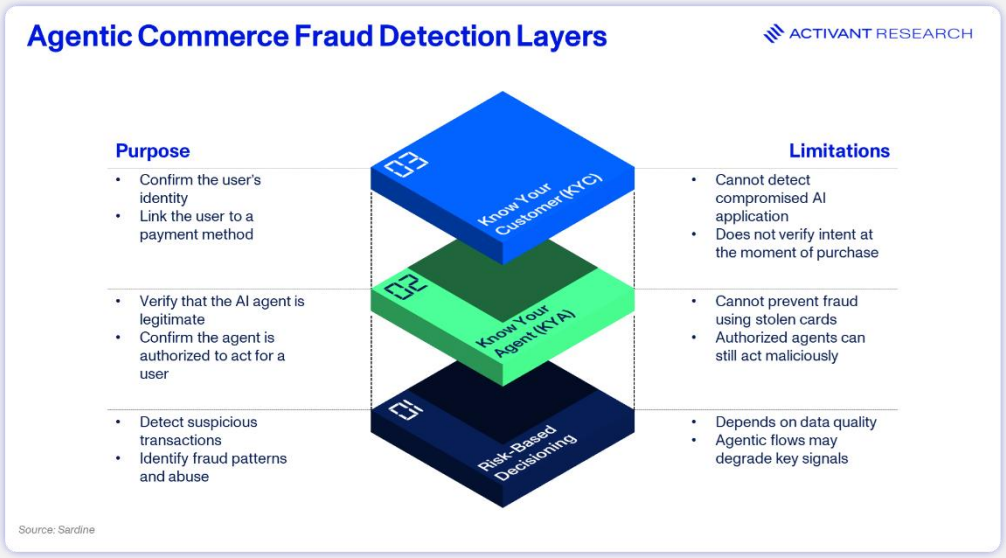


Other companies like Forter (an exited Activant portfolio company) and Sift have also started to extend models to handle agentic traffic. This is not a cosmetic difference. It is a different bet on a new fraud problem. The bet matters because retrofitting would assume agentic behavior is a variant of human behavior with some signal degradation, which it is not. An agent moves through a transaction flow in patterns that share almost nothing with a human shopper. The mouse

movements, the timing, the sequence of page views, the decision latency, are all categorically different. A model trained on human signals will either flag every agent transaction as fraud (false positives), or it will be quietly tuned to ignore the agent signal entirely, at which point it is not really detecting fraud anymore.[10] Segmenting the traffic and building purpose-built models is the more rigorous answer, and our view is that the companies who made this bet early will have a structural advantage as the market grows. It is part of why Sardine sits in the portfolio, and why we are watching carefully for other companies taking similar architectural bets.

Identity verification is the layer below fraud, and it is where the real long-term answer sits. Even a good fraud model can only infer. Verified agent identity, where the agent itself carries cryptographically signed credentials proving it was issued by a known provider and authorized to act on behalf of a specific user, closes the inference problem entirely. Allowance is building a platform that provides agents with guardrails for spending, limiting amounts and merchants with time bound payments without exposing card details. Visa and Mastercard have both unveiled agent credentialing frameworks anchored to Web Bot Auth, a cryptographic standard from the IETF.[11,12] The earliest version was expected in April 2026, with meaningful adoption across merchants and issuers unlikely before year end. Until then, the gap stays open, and every merchant transacting with an agent is absorbing risk they cannot fully price. American Express is not standing still either and is using its venture arm to back various companies in the agentic commerce space, betting heavily on agents becoming a new type of customer.[13] However, at Activant we believe to really solve the underlying root cause, not the symptoms, it needs fundamentally new technology that provides agents with their own credentials to prove their identity. This is why Skyfire is building a payment and identity layer specifically for AI agents, letting agents carry verifiable spending authority.

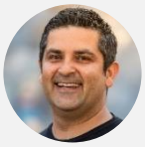

*The internet was built for humans, not autonomous agents. Today's internet security models assume a human is sitting behind the keyboard. As agents begin making purchases, accessing services, and interacting with businesses on our behalf, the fundamental challenge becomes trust. Otherwise, agents don't have the access needed to reliably act on behalf of humans. The long-term answer isn't better bot evasion, it's verified identity. If businesses can reliably know who an agent is, who is operating it, and what it's authorized to do, agents can participate safely in the digital economy at scale."*

Amir Sarhangi
Founder & CEO, Skyfire

Solving identity makes the transaction technically possible. It does not make it commercially clean. Settlement is where commerce finishes, and that layer is being rebuilt from scratch.

## Settlement: Agent-Native Payments

Traditional payment rails were built for predictable, periodic, human-initiated transactions. Agent-mediated commerce has a different signal profile. Transactions can be high-frequency, low-value, programmatic, and chained across multiple merchants in a single user instruction. The rails that handle these transactions need to clear faster, authorize differently, and settle with less human intervention than the cards networks were designed to support. Two categories of companies are working on this, and they are pulling in different directions.

The first category is payment orchestration built for agent-native flows. Deuna (an Activant portfolio company) sits in this space. Deuna is a payment orchestration platform, and its Athia layer was built to bring agentic intelligence directly into payment operations. The distinction is that agentic transactions do not look like card transactions at checkout. They look more like API calls that need to be interpreted, authorized, and settled with context the traditional rails do not carry. Orchestration layers built around that reality, rather than retrofitted onto card-centric flows, have a structural advantage as the volume shifts.

The second category is global commerce. E-commerce allows users to buy from digital storefronts across the globe and agentic commerce will likely extend this trend. Companies like Airwallex (another Activant portfolio company) are preparing to become the payment layer and finance stack that helps merchants be part of agentic commerce by providing agent-ready infrastructure at scale.

Airwallex does so with its multi-currency wallet infrastructure that allows AI agents to autonomously transact and settle funds globally. In practice, an agent operating on behalf of a user in London, buying from a merchant in Seoul, settling in stablecoin, might not be the first use case we will see, but it is not far down the road. Companies building cross-border rails that could support that flow when demand arrives are positioned for a layer of the market that does not yet have a natural incumbent.

The third category is stablecoin-native settlement. As we highlighted in our stablecoin research "Stablecoins: Enabling Direct Flights", Stripe, PayPal, Coinbase, and protocols like x402 are all building settlement layers where the primary medium is a stablecoin rather than a card, with final settlement happening on public blockchains. This approach is not yet mainstream. It may become so in agentic commerce faster than elsewhere, because agents benefit asymmetrically from settlement that is low-cost and programmatic, and therefore does not depend on human-in-the-loop authorization.

Our view on settlement: the first category (agent-native orchestration) is the near-term investable opportunity because it solves a problem that exists today in a form the existing rails cannot easily close. The second category (cross-border agentic rails) is a medium-term bet on where the market is going. The third category (stablecoin-native settlement) is the most speculative and the most transformational. We are closely watching whether agent traffic becomes the specific use case that moves stablecoin settlement from theoretical to operationally preferred for agentic commerce. If it does, the settlement layer gets rebuilt faster than anyone is currently pricing.

Access, identity, and settlement together resolve the technical problem. They do not resolve the legal one.

## Governance: The Liability Question

The hardest problem in agentic commerce is not technical. It is who pays when something goes wrong.

It is a question no regulator has yet fully answered. The consumer may argue they authorized general instruction, not a specific transaction. The merchant may argue the order was valid on its face. The platform may argue it executed within the parameters it was given. The AI provider may argue it only provided a model that acted based on the user's input. All these positions are plausible. None has been thoroughly tested in court.

Even the EU's AI Act, which is the most advanced AI regulation to date, does not provide an answer as it predates the agentic commerce use case. Additionally, regulatory frameworks still in the making, like PSD3, are currently not expected to fully close the gaps.[14]

The Amazon lawsuit against Perplexity offered the first real preview of how complex the legal aspects of agentic commerce can get. The harder cases will turn on authorization. A parrot triggering an Alexa order in 2018 was an early, trivial version of the same question.[15] Millions of agents acting simultaneously on behalf of millions of customers is not trivial. Even agents that do exactly what they are told create disputes. A user who instructs an agent to find the cheapest chainsaw may still contest the order: the delivery timeline was unacceptable, the seller's terms included conditions they would never have agreed to manually, or the price they expected and the price they got diverged. Intent and authorization are not the same thing, and existing consumer protection law was not written to distinguish them. Those cases have not yet been filed at scale. When they are, the liability regime they produce will determine which companies in the enablement stack capture value and which absorb it.

The regulatory framework most likely to emerge shifts the consent question from "did you approve this purchase?" to "did this purchase fall within the parameters you set?" You authorize an agent once, encoding specific rules: a price ceiling, approved merchants, product categories, delivery windows. The agent acts autonomously within those rails. Visa's and Mastercard's solutions already allow for agent-specific payment tokens with programmable spend controls baked into the credential itself.[16] The closest analogy is a corporate expense card. Companies do not approve every employee purchase in advance. They set a spend policy, issue a card with limits, and review the ledger after the fact. Agent tokens work the same way, except the policy is user-configured rather than company-wide. Routine, in-scope purchases clear without friction. Edge cases escalate. The credential enforces the rules at the point of payment.

Until final regulatory clarity arrives, agents will be designed conservatively. Confirmation steps will persist longer than the technology requires. The infrastructure will be ready before the governance is, and the gap between the two is where adoption runs slower than the projections suggest.

# What We Are Watching

Three observations as we close:

First, the enablement layer is moving faster than almost anyone realizes, and slower than most insiders would like. The technical problems are solvable. The governance problems are not solvable by companies alone. The dominant constraint is not engineering. It is the speed at which courts, regulators, and networks reach working equilibria, and that speed is measured in years, not quarters.

Second, the companies we find most compelling are the ones making architectural bets that assume the agent world is structurally different rather than incrementally different. These include Sardine's decision to segment agentic traffic entirely, Deuna's decision to build orchestration around agentic flow rather than retrofitting card-centric flow, Skyfire's decision to build identity for agents rather than extending human identity to cover them. These are the bets that compound if our thesis is right.

Third, the biggest risk to the thesis we suggest is that one of the major AI platforms integrates the entire enablement stack vertically and makes the independent layer unnecessary. It is a non-trivial risk. Amazon already has fraud, identity, and payments at scale, and if ChatGPT or Gemini decides to build the stack internally rather than use third-party providers, the independent enablement companies lose their best customers. Our view is that this is unlikely for reasons of regulatory neutrality and platform incentives. Unlikely, however, does not mean impossible and the enablement companies that survive long-term will be the ones that remain useful even in a world where major players try to absorb their function.

* * *

In part III of this series, we will expand our view outwards to the application layer to answer the question that brands and retailers should be asking: how do you compete for a sale when the buyer is software? We believe it depends on what you are selling. But the good news is that players like Peec AI, Catalog, Gorgias and many more are already building to help retailers keep up with technological change. How that works is where the third piece begins. If you are building or investing in this space, we want to hear from you. How do you think retailers need to adapt? Feel free to share any comments or thoughts by reaching out

[1] eMarketer, AI is rewriting the start of the shopping journey, April 2026

[2] Stellagent, UCP vs ACP – A Comprehensive Comparison of the Two Agentic Commerce Protocols, April 2026

[3] ibid

[4] Digiday, ChatGPT is now 20% of Walmart's referral traffic, September 25 2025

[5] R[AI]sing Sun, The 2026 Agentic Commerce Stack: What to Build Before the Window Closes, March 2026

[6] Reuters, Court temporarily allows Perplexity shopping 'agents' on Amazon, March 2026

[7] CMS Wire, Cyber Week 2025, December 2025

[8] Red Stag, How many ecommerce stores are there?, Accessed May 2026

[9] Sardine, Rethinking Fraud Prevention for Agentic Commerce, Accessed May 2026

[10] ibid

[11] Visa, Visa Introduces Trusted Agent Protocol, October 2025

[12] Mastercard, Powering the next frontier of commerce, Accessed May 2026

[13] Crunchbase News, From Credit Cards To An AI Concierge, May 2026

[14] European Business Magazine, Agentic Commerce: When AI Buys on Your Behalf, February 2026

[15] CNBC, This naughty parrot was caught ordering items off Amazon's Alexa, December 2018

[16] Eco, Visa Mastercard Agentic AI Commerce, Accessed May 2026

**Disclaimer:** The information contained herein is provided for informational purposes only and should not be construed as investment advice. The opinions, views, forecasts, performance, estimates, etc. expressed herein are subject to change without notice. Certain statements contained herein reflect the subjective views and opinions of Activant. Past performance is not indicative of future results. No representation is made that any investment will or is likely to achieve its objectives. All investments involve risk and may result in loss. This newsletter does not constitute an offer to sell or a solicitation of an offer to buy any security. Activant does not provide tax or legal advice and you are encouraged to seek the advice of a tax or legal professional regarding your individual circumstances.

This content may not under any circumstances be relied upon when making a decision to invest in any fund or investment, including those managed by Activant. Certain information contained in here has been obtained from third-party sources, including from portfolio companies of funds managed by Activant. While taken from sources believed to be reliable, Activant has not independently verified such information and makes no representations about the current or enduring accuracy of the information or its appropriateness for a given situation.

Activant does not solicit or make its services available to the public. The content provided herein may include information regarding past and/or present portfolio companies or investments managed by Activant, its affiliates and/or personnel. References to specific companies are for illustrative purposes only and do not necessarily reflect Activant investments. It should not be assumed that investments made in the future will have similar characteristics. Please see "full list of investments" at activantcapital.com/companies/ for a full list of investments. Any portfolio companies discussed herein should not be assumed to have been profitable. Certain information herein constitutes "forward-looking statements." All forward-looking statements represent only the intent and belief of Activant as of the date such statements were made. None of Activant or any of its affiliates (i) assumes any responsibility for the accuracy and completeness of any forward-looking statements or (ii) undertakes any obligation to disseminate any updates or revisions to any forward-looking statement contained herein to reflect any change in their expectation with regard thereto or any change in events, conditions or circumstances on which any such statement is based. Due to various risks and uncertainties, actual events or results may differ materially from those reflected or contemplated in such forward-looking statements.