cognizant

Agentic commerce: From click-to-pay to intent-to-fulfill

A strategic point of view for banks, card networks, acquirers, PSPs and fintechs navigating the next major evolution in digital commerce

This paper is intended for executives in banks, card networks, acquirers, PSPs, fintechs and merchant platforms who must make agentic commerce decisions before standards settle. It sets out how the commerce stack changes, what each ecosystem player stands to gain or lose, what trust and governance architecture is required, and what immediate readiness agenda looks like. It does not endorse a single protocol; rather, it argues for a protocol-agnostic posture and a deliberate choice of which layer of the agentic stack to own.

Executive summary

Agentic commerce marks the next major evolution of digital commerce—a shift from human-led "click-to-pay" journeys to autonomous "intent-to-fulfill" flows, through which AI agents discover, compare, negotiate and transact on behalf of consumers or enterprises.

The point of control moves away from the visible checkout page and toward machine-readable protocols, cryptographic mandates, agent identity, programmable payment credentials and automated evidence. For payments players, this is not merely a new user interface; it is a structural change in how demand is routed, payment choice is made, liability is proven and value is captured.

Market signals indicate that agentic commerce could become a multitrillion-dollar opportunity by 2030, with research from McKinsey estimating up to $1 trillion of orchestrated US B2C retail revenue and up to $5 trillion globally. Other market perspectives project that a substantial share of online commerce may be mediated by agents. The strategic implication is clear: Institutions must become visible, trusted and selectable not only by human customers, but by software agents as well.

Winners in this market will be defined not by a single protocol bet, but by the layer they choose to own: agent identity, programmable mandates, protocol abstraction, agent-ready data, trust infrastructure, dispute evidence or managed agentic operations. That choice is the first decision leadership must make, because it determines where investment, liability and monetization are focused. Moving from pilots to production, then, depends on a disciplined progression—advise, build, integrate, secure/govern, operate and scale—beginning with an agenda that proves a single, complete intent-to-evidence journey.

Table of contents

The agentic shift	4
Ecosystem architecture: Three interaction models	5
Agent-to-site	5
Brokered agentic commerce	5
Agent-to-agent	5
The standards war: Why protocol optionality matters	6
Market momentum and why the window is now	6
The new commerce stack	7
Strategic implications by ecosystem player	8
Card networks	8
Issuing banks	8
Fintechs and PSPs	9
Acquirers	9
Trust, security and governance architecture	10
Operating model implications	10
Building the delivery capability	11
Conclusion	12

The agentic shift

Digital commerce has already moved through two major eras. In the e-commerce era, the browser was the interface, and the human user navigated pages, forms and checkout screens. In the m-commerce era, apps, mobile wallets and biometrics reduced friction but still assumed that a human tapped, reviewed and approved the transaction.

Agentic commerce introduces a third era: The interface becomes an API. The action becomes an agent call. And the unit of authority becomes an intent rather than a click. This changes the economics of digital commerce. Agents optimize for cost, certainty, rewards, delivery speed and user-defined constraints. They do not carry legacy habits, brand attachment or interface loyalty in the same way humans do. A payment method that is not machine-readable, a reward that is not computable in real time or a credit offer that is not exposed at the moment of intent may simply be invisible in the agent's decision process.

Ecosystem architecture: Three interaction models

Model A: Agent-to-site

In the near term, agentic commerce volume is likely to be dominated by browser-based agents navigating existing merchant websites. But over time, a growing share of transaction volume will migrate toward merchant APIs, agent-to-agent protocols and brokered trust networks. This model offers broad reach—because every website is theoretically accessible—but it is fragile. User interface changes can break agent flows, and security tools such as bot managers and web application firewalls may incorrectly classify authorized agents as malicious traffic. The business risk is that funded demand is blocked before it even reaches checkout.

Model B: Brokered agentic commerce

In the brokered model, intermediaries provide the trust layer, standard protocol, discovery environment and payment orchestration between agents and merchants. This can accelerate adoption by reducing integration complexity, but it also introduces a protocol tax and gives the broker privileged access to customer intent data. For banks, networks and PSPs, the strategic question is whether they become the broker, integrate into the broker or are reduced to downstream funding and settlement infrastructure.

Model C: Agent-to-agent

In the target-state model, a consumer or enterprise agent negotiates directly with a merchant, bank, PSP or network agent through structured APIs. This enables dynamic pricing, loyalty optimization, contextual credit, negotiated terms and richer automation. It also creates unresolved questions around standards, liability, mandate interpretation and ownership of errors when two agents negotiate incorrectly.

These three models are not alternatives from which to choose; they are a sequence to plan for. Agent-to-site traffic is arriving now, and it must be handled through acceptance and bot-management policy. Brokered models are forming next, and they determine whether an institution is a participant in the trust layer or a supplier to someone else's. Agent-to-agent is the target state, and it should shape architecture decisions made today, because retrofitting identity, mandate and evidence primitives later is far more costly than designing for them now. The practical test for any institution is whether one intent, mandate and evidence contract can serve all three models without a separate build for each.

The standards war: Why protocol optionality matters

The agentic payments landscape is fragmenting quickly. Multiple protocols and commercial approaches are emerging across card networks, AI platforms, cloud providers, PSPs and crypto-native rails. (Among them are Visa's agentic commerce initiatives, Mastercard Agent Pay and Verifiable Intent, Stripe and OpenAI's Agentic Commerce Protocol, Google's Agent Payments Protocol and x402-style machine-to-machine settlement.) While approaches differ, they converge on common primitives: verified agent identity, scoped delegation, tokenized credentials, signed intent, mandate evidence and a way to resolve liability when the human is not present at checkout.

Since no single protocol has proven definitively superior, payments institutions should avoid locking their strategy to one rail. The most resilient posture is protocol-agnostic: normalize multiple external protocols into a common internal intent contract, expose consistent identity and mandate controls, then allow new rails to be added in weeks rather than through multi-quarter core programs.

Market momentum and why the window is now

The pace of activity across the ecosystem suggests that agentic commerce is moving from concept to infrastructure. McKinsey describes agentic commerce as a "seismic shift" in which AI agents can anticipate needs, navigate options, negotiate deals and execute transactions in alignment with user intent. Its research estimates that by 2030, AI agents could orchestrate up to $1 trillion of US B2C retail revenue and mediate up to $5 trillion of global consumer commerce—with the caveat that these figures reflect goods and exclude services and B2B commerce.

For payments institutions, the implication is clear: agentic commerce should not be treated as a distant channel experiment, but as an emerging operating environment that will influence traffic routing, payment choice, authorization models, fraud controls, identity and customer engagement well before full autonomy becomes mainstream.

The adoption curve is being accelerated by three forces:

1. **Consumers'** behavior is shifting upstream, as customers increasingly use AI to compare products, summarize reviews and shortlist options before reaching a merchant or bank-owned interface
2. **Infrastructure providers** are converging on the same trust primitives: verifiable identity, scoped spending authority, tokenized credentials and auditable consent
3. **Merchants and payments players** are realizing that agent traffic can either become a new demand source or a new disintermediation vector, depending on whether they expose machine-readable value propositions

For financial institutions, this creates a board-level urgency. Agentic commerce compresses the time between intent, comparison, financing, payment selection and fulfillment. Decisions that once happened across separate surfaces—search, banking app, wallet, checkout, dispute portal and service center—begin to collapse into a single agent-mediated flow. If the institution is not callable by that flow, its products may be excluded from consideration, even when they remain functionally competitive.

The new commerce stack

Agentic commerce requires a different architecture than traditional checkout. The stack begins with Intent capture, where the user delegates a bounded objective to an agent. It then moves to Discovery and comparison—where the agent interprets product, merchant, reward, credit, fulfillment and risk data. A third layer is negotiation and Decisioning, where the agent evaluates options against user constraints. A fourth layer is payment Authorization, in which credentials, mandates and signatures prove that the agent is allowed to act. The final layers are Evidence and Operations, where receipts, logs, manifests and telemetry support reconciliation, disputes and compliance.

Stack layer	Core capability	Why it matters
Intent	Human goal translated into machine-readable instruction	Creates the authority boundary for the agent
Discovery	Agent-readable catalogs, offers, pricing, eligibility and rewards	Determines whether products and payment options are visible to agents
Decisioning	Real-time optimization across cost, loyalty, credit, risk and delivery	Moves value selection upstream of checkout
Authorization	Signed mandates, tokenized credentials and cryptographic proof	Replaces human presence with verifiable authority
Evidence	Intent receipts, cart manifests, logs and telemetry	Supports disputes, audit, compliance and forensic reconstruction
Operations	Monitoring, reconciliation, drift control, liquidity guardrails and incident response	Turns pilots into reliable production operating models

The most important architectural shift is that payments become embedded into broader agentic decisioning. Card authorization alone is insufficient if the agent's mandate, cart content, identity and toolchain cannot be trusted. Similarly, a clean payment may still become a commercial failure if the agent chose the wrong product, misunderstood a return policy or accessed outdated pricing. Institutions need architectures that treat commerce, payment, identity, risk and evidence as a single, integrated control fabric.

Strategic implications by ecosystem player

Card networks

For card networks, the rail itself becomes more contestable. Agents can route to cards, pay-by-bank, account-to-account rails or stablecoin-based settlement depending on cost and certainty. The defensible moat then shifts from the card rail alone to the trust layer: agent registries, programmable tokens, revocation lists, evidence standards and scheme rules that define liability for agent-not-present transactions. Networks that bridge protocols and own the rulebook remain central; those that defend closed standards risk being bypassed.

Strategically, networks should treat Know Your Agent (KYA) as franchise infrastructure. A network-operated agent registry can create the same kind of ecosystem leverage that tokenization and scheme rules created in earlier payment eras. The registry should support agent onboarding, attestation, scoring, revocation, provenance and liability routing. Programmable network tokens should carry metadata such as agent ID, user mandate, spend ceiling, merchant or category binding, time-to-live and intent hash. This transforms the token from a static credential substitute into a dynamic authorization container.

Issuing banks

For issuers, agentic commerce reaches directly into the P&L. Interchange, FX, credit distribution, deposit stability, rewards and liability all become exposed to algorithmic decisioning. Credit must move upstream into the agent's negotiation. Rewards must become machine-readable and computable, and mandates must allow customers to scope, age out and revoke delegated authority in real time. The winning issuer will not merely fund someone else's agent; it will become the programmable decisioning layer.

The issuer opportunity is to become the customer's trusted agentic financial controller. This means exposing credit, rewards, liquidity and risk preferences as callable services. It also means redesigning customer consent around live delegation: "Allow this agent to book travel under a defined budget," "Use the best net-value payment method" or "Optimize savings unless balance falls below a floor." The issuer that owns these mandates becomes deeply embedded in daily decisioning, whereas the issuer that merely provides a credential becomes easier to substitute.

Fintechs and PSPs

For PSPs and fintechs, the checkout page is no longer the main product surface. The moat moves to orchestration, protocol abstraction, agent-grade risk, merchant enablement and automated dispute evidence. PSPs can monetize their agents' readiness by:

- Providing machine-readable catalogs, pricing and returns schemas
- Offering rail orchestration through a single merchant contract
- Creating new revenue pools in attestation, autonomous reconciliation, dispute automation and B2B micro-settlement.

For PSPs, the strongest commercial opportunity is to become the merchant's agentic operating system. This includes schema quality, catalog readiness, agent checkout endpoints, protocol routing, fraud scoring, dispute evidence and reconciliation. The PSP can also help merchants shift from campaign-led merchandising to rule-based offer execution, where an agent can evaluate discount logic, loyalty stacking, financing cost and fulfillment certainty in real time.

Acquirers

For acquirers, the first-order challenge is acceptance. Existing bot defenses may block the very agents that represent future demand. At the same time, accepting agent traffic blindly creates fraud and liability exposure. Acquirers must distinguish authorized agents from scrapers, enable merchants with agent-readable data and checkout endpoints, and automate representment through signed mandates and cart manifests.

Acquirers should also rethink merchant onboarding and risk underwriting. Agent-mediated traffic may appear anomalous by traditional models: high-velocity, non-human timing, repetitive API calls and unusual basket patterns. Rather than treating these indicators as uniformly suspicious, acquirers need agent-aware risk classification, verified-agent allow-listing and evidence-led representment. The acquirer that can safely admit good agents while blocking malicious automation will turn acceptance policy into a revenue advantage.

Trust, security and governance architecture

Agentic commerce changes the threat model, because fraud begins upstream of the payment. The exposure is not only whether funds are available, but whether the agent was authentic, the mandate was valid, the toolchain was trustworthy and the transaction reflected the user's original intent. Threats include indirect prompt injection, tool and context poisoning, key compromise, broad-mandate abuse, collusion, herding and machine-rail laundering.

A resilient trust architecture requires three control planes:

1. **Identity:** A KYA registry, hardware-bound keys, attestations, provenance and revocation
2. **Authority:** Dynamic mandates that are narrow, time-bound, merchant-aware and revocable in real time
3. **Evidence:** Signed intent receipts, cart manifests and transaction telemetry that can reconstruct what the agent saw, selected, authorized and executed

Together, these replace human presence with cryptographic accountability.

Operating model implications

Agentic commerce is not only a technology transformation. It requires new accountabilities across product, risk, compliance, operations, engineering and customer service. Traditional handoffs among digital channels, payments operations, fraud teams and service desks are too slow for machine-speed commerce. Institutions need an operating model that treats agent identity, mandate governance, evidence management and continuous monitoring as standing capabilities.

New roles may include agent identity steward, agentic commerce product owner, machine-risk analyst, agentic dispute specialist, protocol product manager and agent operations lead. These roles are not necessarily new departments; they are new accountability lenses that span across existing teams. The key is to clarify who can approve an agent, who defines autonomy limits, who owns revocation policy, who monitors agent behavior, who handles disputes and who signs off on liability positions.

Performance management must also evolve. While traditional metrics such as authorization rate, fraud rate and dispute ratio remain important, they need agent-specific extensions: verified-agent acceptance rate, good-agent false positive rate, mandate revocation latency, agent checkout completion rate, signed-evidence capture rate, protocol share, autonomous reconciliation match rate and agent-driven cost-to-serve reduction.

Building the delivery capability

Executing on this agenda requires three interwoven capabilities:

1. **Payments domain depth:** An understanding of tokenization, authorization, scheme rules, dispute mechanics and settlement that is sufficient to redesign them for machine-initiated flows
2. **Agentic AI engineering:** The ability to design, build, test and lifecycle-manage agents, and to produce the agent-readable data and context products that discovery and decisioning depend on
3. **Enterprise-scale governance:** The trust, safety, accountability, compliance and auditability controls that allow autonomy to be introduced inside a regulated estate rather than alongside it

A structured transformation approach should connect payment domain depth, agentic AI engineering and enterprise-scale governance identifying value sources, redesigning processes, building and governing agents, reengineering IT capability, transforming the operating model and enabling adoption.

The practical objective is speed with control. That means quantifying value at stake, selecting no-regret investments, designing the trust architecture, integrating payment rails, creating agent-readable data products, building adversarial test harnesses, defining operating model roles and standing up managed operations.

This balance matters most for institutions that cannot afford to wait for standards to settle but equally cannot expose customer funds and regulated processes to unmanaged autonomy. Whether these capabilities are built internally, acquired or sourced from delivery partners is a separate decision, but the capability set itself is not optional.

Conclusion

Agentic commerce is not simply checkout automation. It is a replatforming of commerce around intent, autonomy, identity and evidence. The browser becomes less important than the protocol. The click becomes less important than the signed mandate. The visible wallet becomes less important than the programmable decisioning layer an agent can evaluate.

For banks, networks, acquirers, PSPs and fintechs, the mandate is urgent: become agent-readable, agent-trusted and agent-selectable before autonomous demand arrives at scale. The institutions that move first will not be those that place the boldest protocol bet. Rather, it will be those that decide deliberately which layer of the agentic stack they intend to own, then build the identity, mandate, evidence and governance foundations to hold it.

Those foundations pay back whether autonomous volume arrives in 18 months or five years. The task now is to monetize the new value pools, secure the new threat surface and govern autonomy responsibly at enterprise scale—and to begin with a single, complete intent-to-evidence journey that proves the model before it is scaled.

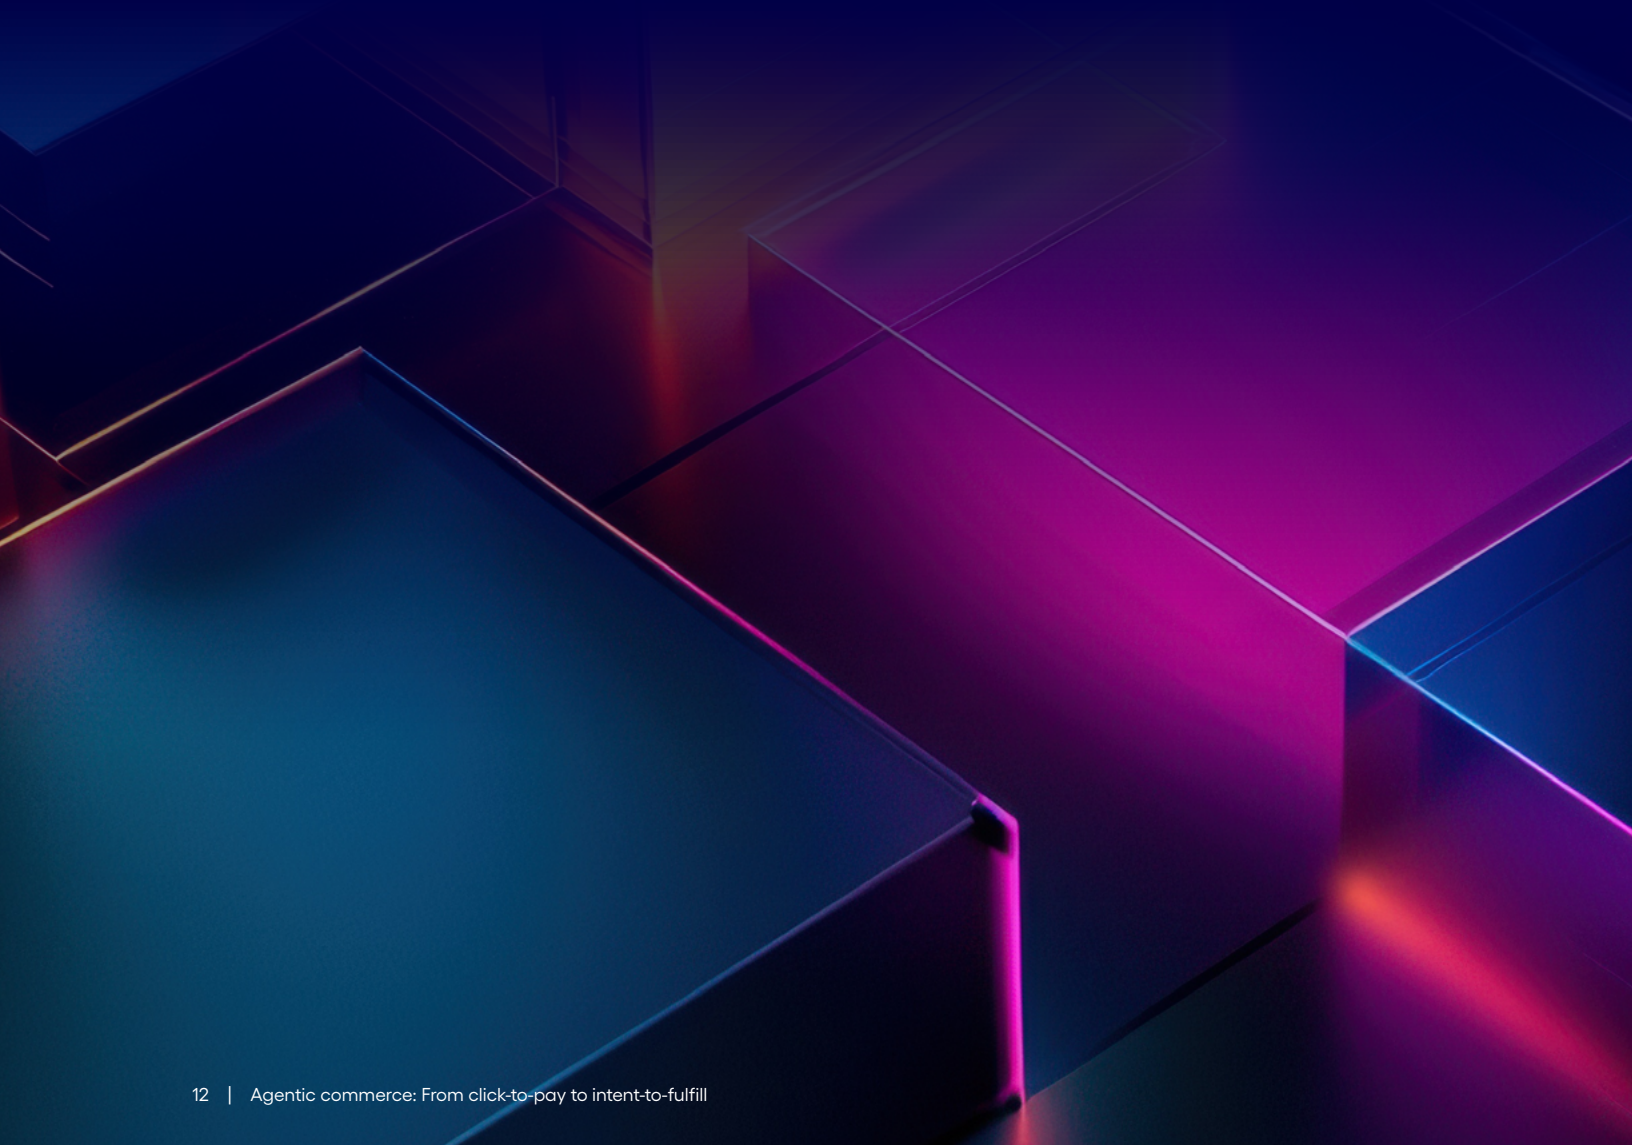

Author

Bala leads Global India Payments Consulting for Cognizant, driving strategic consulting engagements in payment and cards. He serves customers across all geographies based out of India. He has over 25 years of solutioning and global consulting experience advising global banks and payment networks and non-banking financial institutions in digital business and technology strategy and transformation. Bala has led consulting engagements across NA, US, UK, Continental Europe and APAC.

Bala Gurunathan

Associate Partner and Payments Lead,
Global India Consulting, Cognizant

cognizant

Cognizant (Nasdaq: CTSH) is an AI Builder and technology services provider, bridging the gap between AI investment and enterprise value by building full-stack AI solutions for our clients. Our deep industry, process and engineering expertise enables us to build an organization's unique context into technology systems that amplify human potential, drive tangible outcomes and keep global enterprises ahead in a fast-changing world. See how at cognizant.ai or @cognizant.

World Headquarters

300 Frank W Burr Blvd
Suite 36, 6th Floor
Teaneck, NJ 07666, USA
Tel: (201) 801-233

European Headquarters

280 Bishopsgate
London
EC2M 4AG
England
Tel: +44 (01) 020 7297 7600

India Corporate Office

Siruseri-Software Technology Park of India (STPI)
SDB Block – Ground floor north wing
Plot No H4, SIPCOT IT Park
Chengalpattu District
Chennai 603103, Tamil Nadu
Tel: 1800 208 6999

APAC Headquarters

1 Fusionopolis Link, Level 5
NEXUS@One-North, North Tower,
Singapore 138542
Tel: + 65 6812 4000

WF 5222450