



The Autonomous Economy

How AI Agents, Stablecoins and
Payment Infrastructure Are Shaping the
Next Generation of Commerce

September 2026



Disclaimer

This white paper is a synthesis of perspectives shared during an industry roundtable titled “Stablecoin-linked Agentic Commerce and Payments”. The insights, views, and takeaways contained herein represent a synthesis of the dialogue among individual participants and do not necessarily reflect the formal corporate positions, official policies, or product roadmaps of the participating organisations.



Foreword

Tianwei Liu

CEO & Co-Founder, StraitsX
Chair, Payments Subcommittee,
Singapore FinTech Association

The conversation around AI is increasingly moving from what AI can create to what AI can do. As agents begin to act on behalf of people and businesses, one question becomes increasingly important: how will these agents participate in the economy and move value safely and efficiently?

We believe this question deserves to be explored now, while the infrastructure and standards for agentic commerce are still taking shape. Stablecoins, payment networks and financial institutions are each evolving, but the connections between them, and the controls needed to allow autonomous systems to transact, are still being developed.

This white paper brings together perspectives from across the ecosystem to examine these questions and contribute to the conversation on what needs to be built for agentic commerce to develop responsibly.

Holly Fang

President
Singapore FinTech Association

Technology rarely transforms an industry in isolation. The most meaningful changes happen when new technologies meet the infrastructure, institutions and regulatory frameworks that allow them to operate in the real economy. The emergence of AI agents presents one such moment for financial services.

As the fintech ecosystem explores this next phase, the national FinTech association for Singapore, the Singapore FinTech Association (SFA), has an important role to play in convening stakeholders across technology providers, financial institutions, regulators and startups. Our role is to help surface shared challenges and facilitate cross-sector dialogue.

This white paper is intended to contribute to that collective conversation, encouraging further dialogue and collaboration on how Singapore and the wider region can responsibly prepare for an increasingly agent-driven economy.

Adeline Kim

Group Country Manager
Regional Southeast Asia, Visa

The next evolution of commerce may be defined by who, or what, is making the decision to pay. As AI agents become increasingly capable of acting on behalf of consumers and businesses, the payments industry has an opportunity to shape how this new form of commerce develops from the outset.

There are important questions to address early: how will agents interact with merchants, how will they be trusted and authorised, and how can existing payment infrastructure support transactions initiated by machines? These questions sit at the intersection of technology, payments and financial services.

By examining the emerging infrastructure and governance considerations around agentic commerce, this paper aims to support a constructive industry dialogue about how the payments ecosystem can prepare for the next chapter.

Table of Contents

Executive Summary	06
Defining Agentic Payments	09
2.1 From Automation to Agency	10
2.2 The Agentic Commerce Maturity Spectrum	11
2.3 From Transaction Approval to Delegated Outcomes	12
2.4 The Emerging Protocol Layer	14
Infrastructure for Agentic Commerce	16
3.1 The Adoption Matrix: Mapping Complexity and Frequency	17
3.2 From Payment Rails to Orchestration	19
3.3 The Role of Stablecoins	21
3.4 The Economics of Autonomous Transactions	23
Trust, Governance, and Industry Priorities	24
4.1 Defining Authority and Liability	26
4.2 Identity, Access and Runtime Controls	27
4.3 Compliance by Design	29
4.4 From Pilots to Production	30
Conclusion & Strategic Priorities	31

Chapter 1

Executive Summary



Artificial intelligence is moving from assisting human decision-making towards executing tasks on behalf of users and organisations. As this transition extends into commerce, AI agents will increasingly need to interact with merchants, financial institutions and payment infrastructure to complete transactions autonomously.

This creates a new challenge for the payments ecosystem. Enabling an agent to make a payment is technically different from enabling it to make a payment **within clearly defined authority, through an appropriate payment mechanism, and with sufficient accountability.**

Agentic commerce is likely to develop progressively. Early adoption may concentrate on high-frequency, lower-complexity transactions, where agents can execute repetitive activities within clearly defined parameters. As trust, identity, authorisation and interoperability mechanisms develop, more complex commercial workflows may become possible.

The payment infrastructure supporting this activity will need to extend beyond any single payment rail. Agents may need to interact with cards, bank transfers, local payment networks and stablecoin-based settlement depending on the requirements of each transaction. This creates a role for orchestration: connecting an agent's intended outcome with the appropriate payment and settlement mechanism.

Stablecoins represent one component of this emerging infrastructure. Their programmability, rapid settlement and suitability for frequent or small-value transactions may be particularly relevant to machine-to-machine activity. However, stablecoin settlement alone does not provide access to the wider economy. Merchant acceptance and interoperability with existing payment networks remain important considerations.

Trust and governance are equally fundamental. As financial execution becomes increasingly autonomous, the ecosystem needs mechanisms to establish **who an agent is acting for, what it is authorised to do, whether it remains within those parameters, and who is accountable when something goes wrong**. Identity, delegated authority, runtime controls, compliance and liability therefore become core infrastructure considerations rather than downstream controls.

Four priorities emerge for the development of agentic payment infrastructure:

1. **Enable interoperable payment infrastructure:** connect agents to multiple payment and settlement mechanisms rather than designing around a single rail.
2. **Establish clear authority and accountability:** define how agents are authorised to act and how responsibility is determined when they operate outside those parameters.
3. **Build controls into autonomous execution:** incorporate identity, access, transaction limits, monitoring and runtime safeguards into agentic payment workflows.
4. **Evolve standards and regulatory frameworks:** adapt existing frameworks and develop greater interoperability around agent identity, delegated authority, security and compliance.

The transition to agentic commerce is therefore not simply about making AI capable of making payments. It is about building financial infrastructure that can connect **intent, authority, payment execution and accountability** across an increasingly autonomous and interconnected economy.

Chapter 2

Defining Agentic Payments



2.1 From Automation to Agency

Agentic payments represent a shift from **executing predefined instructions to acting on delegated intent**.

Traditional payment automation executes known instructions under predetermined conditions. Recurring subscription payments, scheduled transfers and automated balance management can perform tasks without manual intervention, but the underlying logic remains defined in advance. The system executes an instruction; it does not determine how to achieve a broader objective.

An AI agent operates differently. It can interpret a high-level objective, determine the steps required to achieve it, interact with external tools and services, evaluate information as it becomes available, and take action on the user's behalf within defined parameters.

In a payment context, the distinction can be understood through five stages:

- **Intent:** A user or organisation establishes an objective rather than specifying every transaction step.
- **Planning:** The agent determines the sequence of actions required to achieve that objective.
- **Tool use:** The agent interacts with external platforms, APIs, data sources or merchant systems.
- **Adaptation:** The agent evaluates information received during execution and adjusts its actions where required.
- **Payment execution:** The agent initiates a payment as part of completing the broader objective.

“An agent is essentially a system that combines the intelligence of advanced AI models with access to tools to take actions on a user's behalf and under their control.”

— Joe Hwang, Regional Lead, APAC Web3 & AI, Google Cloud

Agentic payment therefore occurs when **financial execution becomes part of an agent's autonomous decision-making process**, rather than simply the execution of a predetermined payment instruction.

2.2 The Agentic Commerce Maturity Spectrum

Agentic commerce is not a binary transition from human-controlled to fully autonomous payments. It is better understood as a spectrum of increasing delegation and autonomy.

Maturity Stage	Role of the Agent	Human Role	Illustrative Use Case
1. AI-Assisted Commerce	Researches, compares and recommends options	Makes the final decision and payment	Researching products and preparing a shopping cart
2. Human-Delegated Commerce	Executes transactions within defined rules	Establishes intent, limits and conditions	Purchasing an item when predefined conditions are met
3. Machine-to-Machine (M2M)	Acts autonomously based on operational requirements	Sets the broader parameters	Purchasing API access, data or computing resources
4. Fully Autonomous Commerce	Multiple agents interact, negotiate and transact	Establishes overarching authority and controls	Agents coordinating workflows and settling value with one another

The first stage extends existing digital commerce by using AI to assist with discovery, comparison and decision-making while the human remains responsible for the final transaction.

The second stage introduces **delegated authority**. A user establishes an objective and the rules governing execution, allowing the agent to complete transactions without requiring approval for every individual action.

At the machine-to-machine stage, software or hardware agents can initiate transactions in response to operational requirements. These may include purchasing data, API access, computing resources or other digital services.

At the furthest end of the spectrum, multiple agents may communicate, negotiate, validate and settle transactions with one another. The agent is no longer simply assisting a human participant in commerce; it becomes an active participant in the transaction itself.

The infrastructure requirements increase with each stage. As authority moves from the human towards the agent, systems require stronger mechanisms to **define authority, control transactions, verify participants and maintain accountability**.

2.3 From Transaction Approval to Delegated Outcomes

The progression towards agentic commerce changes the fundamental interaction between users and payment systems.

In conventional commerce, the user typically specifies and approves the transaction: selecting a product, choosing a payment method and authorising the purchase. In agentic commerce, the user increasingly specifies an **outcome**, while the agent determines the steps required to achieve it.

For example, an instruction such as “purchase the best available option within this budget” requires the agent to evaluate available options, determine which meets the user's criteria and execute the resulting transaction. Similarly, an enterprise agent may acquire data, computing resources or other digital services as part of a broader operational workflow.

This creates a new requirement for payment infrastructure: the ability to distinguish between **what the user intends, what the agent is authorised to do, and what the agent ultimately executes**.

As autonomy increases, transaction-by-transaction approval becomes less practical. Control therefore needs to move towards mechanisms that allow users and organisations to establish the boundaries within which an agent can operate.

These boundaries may include:

- **Who** the agent is acting for;
- **What** the agent is permitted to purchase or transact;
- **Where** it can transact;
- **How much** it can spend; and
- **Under what conditions** the authority remains valid.

The central challenge is therefore not simply enabling an AI agent to make a payment. It is enabling autonomous execution **without losing the connection between human or institutional intent and the resulting financial transaction.**

This creates the foundation for the infrastructure questions that follow: how agents are identified and authorised, how transactions are controlled, how payment routes are selected, how value is settled, and how actions remain traceable and accountable.



“The question that consumers will ask is, is the agent making the payment correctly on my behalf?”

— Adeline Kim, Group
Country Manager,
Regional Southeast Asia,
Visa

2.4 The Emerging Protocol Layer

As agentic commerce develops, a range of protocols and standards is emerging to address different parts of the interaction between AI agents, applications, merchants and payment infrastructure. These initiatives are at different stages of development and adoption and address distinct technical requirements.

Rather than representing a single new payment system, they illustrate the **multiple layers of interoperability required for agentic commerce**.

- **Universal Commerce Protocol (UCP):** Defines a common way for AI agents and commerce systems to interact across activities such as product discovery, checkout and order fulfilment.
- **Model Context Protocol (MCP):** Provides a standardised way for AI applications to connect with external tools, data sources and systems, allowing agents to access information and perform actions beyond the model itself.
- **Agent Payments Protocol (A2P) / Machine Payments Protocol (MPP):** Explore mechanisms for communicating payment intent, authorisation and transaction information between agents and payment infrastructure.
- **HTTP 402 / x402:** Uses the HTTP 402 Payment Required mechanism to enable web services and APIs to request payment from an agent before providing access to a resource or service, including potential machine-to-machine and micropayment use cases.
- **ERC-8004:** Proposes an on-chain framework for agent identity, verification and reputation, addressing how autonomous agents may be identified and evaluated when interacting with other systems.

Together, these developments point to an emerging **protocol layer around agentic commerce**, spanning agent identity and tool access through to commerce interaction and payment execution. Their respective approaches, adoption and interoperability remain areas of ongoing development.

For payment infrastructure, the significance is less about any individual protocol than the underlying requirement: **agents will need standardised ways to identify themselves, communicate intent, interact with commerce systems and initiate transactions across different environments.**

Chapter 3

Infrastructure for Agentic Commerce



The emergence of agentic commerce does not require a single new payment rail. It requires infrastructure capable of connecting autonomous agents to the payment mechanisms, financial institutions and merchant environments in which economic activity already takes place.

The infrastructure challenge therefore extends beyond enabling an agent to initiate a payment. It encompasses where agentic payments can be deployed, how the appropriate payment mechanism is selected, how value is settled, and how different payment environments can work together.

3.1 The Adoption Matrix: Mapping Complexity and Frequency

The adoption of agentic payments is likely to vary according to two characteristics: transaction frequency and complexity.

High-frequency, lower-complexity transactions represent a natural starting point because they involve repetitive decisions and clearly defined parameters. Examples discussed include consumer purchases and other recurring transactions where an agent can operate within relatively straightforward rules.

As the complexity of the transaction increases, the requirements for trust, authorisation and oversight also become more significant. Higher-value or more consequential transactions require greater confidence that the agent is acting within the authority delegated to it and that the resulting transaction can be traced and accounted for.

This suggests a progression in which agentic payments can extend from relatively simple, repetitive transactions towards more complex commercial workflows as the supporting infrastructure develops.

Frequency / Complexity	Low Complexity	High Complexity
High Frequency	Consumer purchases and repetitive digital transactions	Operational and machine-to-machine transactions
Low Frequency	Specialised or higher-value digital services	Complex enterprise and commercial transactions

The significance of this progression is that agentic payment infrastructure cannot be designed around a single use case. Different transactions place different requirements on payment acceptance, authorisation, settlement and risk management.

"Instead of looking at industries, you should look at a two-by-two spectrum: the complexity of the payment and the frequency. As a consumer or business, early adoption will concentrate in the quadrant defined by high frequency and low risk."

— **Adeline Kim**, Group Country Manager, Regional Southeast Asia, Visa

"Current agentic payment activity spans consumer-delegated purchases and machine-to-machine payments initiated on behalf of consumers. Wider adoption will require alignment across merchants, consumers and payment rails."

— **Karthikeyan Janakiraman**, Head of Product, Payments, Grab

3.2 From Payment Rails to Orchestration

An agent may ultimately need to access multiple funding sources and payment mechanisms depending on the transaction.

A consumer purchase, for example, may need to reach a merchant through an existing card or payment network. A machine-to-machine transaction may have different requirements for transaction size, frequency and settlement speed. A cross-border transaction may introduce additional considerations around currencies, liquidity and payment corridors.

This creates a role for an orchestration layer between the agent's objective and the underlying payment infrastructure.

A simplified framework can be considered across three functions:

- **Authorisation**
Establishes whether the agent is permitted to initiate the transaction and whether the transaction falls within the applicable rules or limits.
- **Orchestration**
Determines the appropriate payment or funding mechanism for the transaction, taking into account factors such as the use case, cost, speed and availability.
- **Settlement**
Transfers value to the receiving party through the selected payment or settlement mechanism.

.

Under this model, the agent does not necessarily need to understand or directly manage every underlying payment rail. It can express the intended outcome, while the infrastructure determines how that outcome can be achieved.

This is particularly relevant as agentic commerce develops across **cards, local payment networks, bank transfers and stablecoin-based settlement.**

Existing payment networks continue to provide important acceptance infrastructure, while emerging digital rails can provide alternative mechanisms for particular transaction requirements.

The objective is therefore not to replace one payment rail with another, but to make different rails **accessible and interoperable through a common infrastructure layer.**

“Agentic commerce can be viewed across three layers: authorisation, settlement and liquidity. Agents need defined mandates, the ability to choose the most efficient payment method, and access to liquidity when required.”

— **John Teo, Ecosystem Development, APAC, Aave**

“Reaching merchants will always mean cards and bank rails - but agents at machine speed can benefit from onchain, where mandate, guardrails and settlement execute together instantly. The work ahead is making these rails interoperable, so agents can move between traditional and onchain seamlessly, whenever appropriate.”

— **Edwin Lau, Southeast Asia Lead, Monad Foundation**

3.3 The Role of Stablecoins

Stablecoins represent one component of the broader payment infrastructure for agentic commerce. Their relevance is tied to three characteristics discussed in the context of machine-driven transactions:

- **Rapid settlement:** Agents can operate across workflows where waiting for payment confirmation or settlement can introduce friction. Faster settlement can allow an agent to complete a transaction and continue its workflow without prolonged reconciliation.
- **Programmability:** Stablecoins can be incorporated into programmatic payment flows, allowing payment execution to form part of an automated workflow rather than requiring a separate manual payment step.
- **Support for small, frequent transactions:** Agentic and machine-to-machine activity may involve transactions at a frequency or value where conventional payment processes can become inefficient. Stablecoins can support these smaller, more frequent settlement requirements.

“The next thing is the speed of settlement, which is almost instantaneous, and the ability by design to be programmable. Those are the core strengths of what stablecoin related solution have enabled agents to do so much more.”

— **Tianwei Liu, CEO, StraitsX**

“Potential opportunities include agents paying for information, expert-curated workflows, and off-the-rack solutions”

— **Wilfred Ngia, Commercial Director, OKX SG**

These characteristics are particularly relevant to **machine-to-machine transactions**, where an agent may need to pay for data, APIs, computing resources or other digital services as part of an ongoing workflow.

However, stablecoin settlement alone does not solve the full payment problem. **Acceptance remains critical.** Agents may need to transact with merchants and services that operate through cards, wallets, bank transfers or local payment networks rather than directly accepting stablecoins.

Stablecoins should therefore be viewed not as a replacement for existing payment rails, but as **one settlement mechanism within an interoperable payment environment.**



“The convergence of stablecoins as a payment instrument and agentic AI opens up new vectors of commerce. For the first time we have a possibility where an agent can discover, decide, and settle in a single workflow in real time and with governance oversight. That's not an incremental improvement on existing rails. That's a fundamentally different relationship between commerce and money.”

— **Sriram Royam, Chief Operating Officer, Tazapay**

3.4 The Economics of Autonomous Transactions

Agentic commerce introduces another consideration into payment economics: **the cost of making the decision can become part of the cost of making the transaction.**

In conventional commerce, transaction economics are typically evaluated through factors such as payment fees, foreign exchange costs and settlement time. Autonomous agents introduce an additional variable: the computational resources required to evaluate options and determine whether and how to transact.

An agent may therefore make multiple decisions before completing a single economic action. It may search for alternatives, evaluate prices, access external services and reassess its decision based on changing information. Each step can carry a computational cost.

The resulting optimisation problem is broader than selecting the payment method with the lowest

transaction fee. It requires consideration of the **combined cost of computation, payment, foreign exchange and settlement.**

This creates a potential role for payment infrastructure in helping agents evaluate the overall efficiency of a transaction. Rather than optimising for a single payment metric, an agentic system may need to consider the cost and timing of the entire workflow.

As agentic commerce develops, this interaction between **compute economics and payment economics** may become an important consideration in the design of machine-to-machine transactions.

"We're all looking at it from an efficiency perspective on two spectrums: the cost of the transaction — transaction cost, micropayment, FX cost, settlement timing — and compute cost: how efficiently the agent itself can make that transaction."

— Tianwei Liu, CEO, StraitsX

Chapter 4

Trust, Governance, and Industry Priorities



As financial activity becomes increasingly agent-driven, the mechanisms used to establish authority, manage risk and assign accountability need to evolve alongside the technology.

The central challenge is not simply whether an AI agent can execute a transaction. It is whether the financial system can determine who the agent is acting for, what it is authorised to do, whether the transaction remains within those parameters, and who is accountable when something goes wrong.

The roundtable identified several areas where existing approaches may need to evolve, particularly as agentic payments move from controlled pilots towards broader production use.

4.1 Defining Authority and Liability

Human delegation introduces a fundamental governance question: **how should an agent's authority be defined and recognised across different systems?**

An agent may be instructed to act within specific parameters, such as a spending limit, transaction type or business purpose. However, if those parameters are defined differently across payment providers, networks or applications, the agent's authority may not be consistently recognised throughout the transaction chain.

"For an agent to act on your behalf, you need trust. For there to be trust, you need very controlled guardrails."

— Akhil Sadarangani,
Product Lead, Global
Self-serve Product, Stripe

Two related questions therefore require further development:

- **Mandate definition:** How should human or institutional authority be expressed in a form that different systems can understand and
- **Liability attribution:** When an agent acts outside its authorised intent, how should responsibility be determined across the parties involved in the transaction?

These questions become increasingly important as agents move from assisting with transactions to executing them autonomously. Clear mechanisms for defining authority and establishing accountability can provide the basis for greater confidence among financial institutions and enterprises.

The objective is not necessarily to establish a single technical implementation, but to enable **consistent recognition of authority and accountability across the transaction lifecycle.**

4.2 Identity, Access and Runtime Controls

Agentic payments introduce a new security requirement: financial systems need to distinguish not only between authorised and unauthorised users, but also between different agents acting on behalf of those users.

Several layers of control become relevant:

- **Agent identity and access**

Agents need mechanisms that allow financial infrastructure to establish which agent is requesting access and whether it is authorised to interact with a particular financial service or settlement environment.

- **Runtime controls**

Security cannot rely solely on reviewing transactions after they have occurred. Controls may also need to operate during execution, monitoring agent behaviour and intervening when activity falls outside established parameters.

- **Transaction-level controls**

Limits and conditions can be applied to the agent's financial activity, including parameters such as transaction amount, category or other defined restrictions.



“Increasing controls that involve Knowing your Agent, having identity and access controls for agents, are becoming common.”

— **Pramodh Rai,**
Co-Founder & CEO,
Cyber Sierra

The emphasis on runtime controls reflects a broader shift from **post-transaction monitoring towards controls that operate as part of the transaction itself.**

This is particularly important because the consequences of an autonomous agent making an incorrect or unauthorised decision can extend beyond a single transaction. An agent may be capable of initiating multiple actions rapidly, making the ability to detect and contain anomalous behaviour an important component of financial security.



As agentic payments become more autonomous, infrastructure must account for agent behaviour and the guardrails within which agents operate. Reliability, security and identity, runtime controls, and observability become increasingly important.

— **Naveen Gupta, Head of Business Development - Payments, APJ Leader, AWS**

4.3 Compliance by Design

Agentic payments also require compliance mechanisms that can operate at the speed of autonomous execution.

Traditional financial controls often rely on established processes for identifying customers, screening transactions and reviewing activity. When an agent can initiate transactions continuously and across multiple systems, controls need to be incorporated into the agent's operating environment rather than treated solely as a downstream review process.

This points towards a **compliance-by-design** approach in which relevant controls are incorporated into the transaction workflow itself.

Depending on the use case, these controls may include:

- spending and transaction limits;
- permitted transaction categories;
- geographic restrictions;

- identity and access requirements; and
- monitoring for activity outside established parameters.

The specific controls will vary according to the risk and nature of the use case. The underlying principle is that **compliance needs to remain effective even when transaction execution becomes increasingly autonomous.**

“None of the existing frameworks comprehensively address the current AI environment. Either a new framework has to develop, or the existing ones will have to evolve pretty fast. Otherwise, we're all going to be blocked in pilot environments.”

— **Pramodh Rai,**
Co-Founder & CEO,
Cyber Sierra

4.4 From Pilots to Production

The challenge for agentic payments is therefore not simply technological capability. It is establishing sufficient trust for autonomous systems to operate within regulated financial environments.

Highly controlled use cases may already be capable of operating in production, but broader adoption requires greater confidence in how agents are identified, authorised, monitored and held accountable.

This creates several priorities for the ecosystem:

1. **Develop clear mechanisms for delegated authority** so that agents can act within defined boundaries.
2. **Strengthen agent identity and access controls** so financial systems can distinguish authorised agents from unauthorised actors.
3. **Introduce runtime safeguards** that can monitor and intervene during autonomous execution.
4. **Clarify accountability and liability** when an agent acts outside its authorised parameters.
5. **Evolve existing compliance frameworks** to accommodate autonomous financial activity.

Together, these mechanisms can provide a foundation for moving from highly supervised experiments towards broader applications of agentic payments.

The objective is not to remove human accountability from financial systems. It is to **translate human and institutional authority into controls that remain effective when execution is delegated to machines.**

Chapter 5

Conclusion & Strategic Priorities

Agentic commerce represents a shift in how financial transactions may be initiated and executed. As AI agents take on a greater role in commercial decision-making, payment infrastructure will need to support autonomous execution while maintaining the controls, interoperability and accountability required by financial systems.

The discussion points towards a multi-rail environment in which agents can interact with different payment and settlement mechanisms according to the requirements of each transaction. Stablecoins may form one component of this infrastructure, particularly where rapid settlement, programmability and frequent transactions are relevant. At the same time, cards, bank transfers and local payment networks remain important to reaching the wider economy.

For the ecosystem to progress from experimentation towards broader production use, several priorities emerge:

1. **Enable interoperable payment infrastructure**
Connect agents to multiple payment and settlement mechanisms rather than designing around a single rail. Orchestration can allow the appropriate mechanism to be selected according to the transaction's requirements, including cost, speed and availability.
2. **Establish clear authority and accountability**
Develop mechanisms for defining what an agent is authorised to do, on whose behalf it acts, and how responsibility is determined when an agent operates outside those parameters.
3. **Build controls into autonomous execution**
Strengthen agent identity, access controls, transaction limits, monitoring and runtime safeguards so that financial controls remain effective as transaction execution becomes increasingly autonomous.
4. **Evolve standards and regulatory frameworks**
Existing security, compliance and governance frameworks may need to adapt to account for autonomous financial activity. Greater interoperability in areas such as agent identity, delegated authority and transaction controls can help reduce fragmentation as adoption develops.

The path towards agentic commerce is therefore not simply a question of enabling AI to make payments. It requires financial infrastructure that can connect **intent, authority, payment execution and accountability** across an increasingly interconnected ecosystem.

The opportunity for the industry is to build this infrastructure in a way that allows autonomous commerce to develop while maintaining the trust and safeguards required for real-world financial activity.