



WHITEPAPER
SEPTEMBER 2026

Trust for agentic commerce

Powering secure value exchange in a world where
AI agents shop, pay and act on our behalf



Executive summary



*Trust is what lets the
agentic economy grow.*

The ceiling of the agentic economy is not technology. It is trust. As commerce becomes more invisible, automated, and intelligent, people and businesses need to know they remain in control: that systems are working for them, not instead of them. AI agents are learning to shop, pay, and act on behalf of consumers, businesses, and machines. ICSC and McKinsey project that agentic commerce in US consumer retail could reach \$1 trillion in revenue by 2030, and Gartner expects AI agents to intermediate more than \$15 trillion in B2B spending by 2028. Growth at that scale arrives only if every party in the chain can say yes with confidence.

Mastercard's role is to help new forms of value move with confidence between people, businesses, and communities. In agentic commerce, that means building the trust layer for delegated action. Trust is not a sentiment. It is the ability to verify delegated authority: who is acting, on whose behalf, with what permission, inside what limits, and with what accountability.

The **Mastercard Agent Pay Trust Framework** answers with five connected capabilities, each producing tamper-resistant proof the ecosystem can check. **Identity** proves who is acting. **Intent** proves what was authorized. **Controls** prove the limits. **Trusted execution** proves value moved securely, designed for any rail. **Intelligence** proves risk stays visible as behavior evolves. Together, they turn Mastercard's technology, standards, intelligence, partnerships, and protections into a practical operating model for trusted agentic commerce.

Trust should not only be built into agentic commerce. It should be visible. Mastercard is introducing an **Agent Pay Ready mark**, giving consumers, businesses, and ecosystem participants a simple way to recognize when an agent or merchant is operating within the Agent Pay framework.

The timing matters. As Agent Pay evolves with new standards, intelligence capabilities, and participation requirements, trust must become more consistent, recognizable and scalable. Together, the Agent Pay Trust Framework and Agent Pay Ready mark provide a common foundation for trusted agentic commerce, helping the ecosystem grow with greater confidence, accountability and control.



Contents

04	Introduction
06	The framework at a glance
07	Identity
08	Intent
09	Controls
10	Trusted execution
11	Intelligence
12	A win for every participant
13	How we are building it
14	Conclusion
15	The portfolio behind the framework



Introduction

The agentic economy will be defined less by what agents can do than by what people and businesses trust them to do.

“

The ceiling of the agentic economy is not technology. It is trust.

E-commerce digitized the storefront. Mobile put those businesses in everyone's pocket. Agentic commerce is changing the buying interface again, shifting commerce from search and discovery to checkout and disputes. It is commerce where AI agents search, decide, transact, and resolve issues on your behalf, guided by limits you set. AI agents are already learning to buy, return, replenish, and pay for consumers, businesses, merchants, and machines. Businesses will run procurement, B2B payments, and whole commercial workflows through agents. Machines will settle with machines, on card rails, account-to-account, or in stablecoins. Delegation comes to every corner of commerce, and the role of a trusted network becomes even more essential on every rail that carries it.

The opportunity is enormous, for the largest enterprise and the smallest household alike. ICSC and McKinsey project that agentic commerce in US consumer retail could reach \$1 trillion in revenue by 2030, and Gartner expects AI agents to intermediate more than \$15 trillion in B2B spending by 2028. Numbers that size make a simple point: capability is not the constraint. The agentic economy will only scale at the speed of trust.

The central questions are simple: who is acting, on whose behalf, with what authority, under what limits, and with what accountability if something goes wrong? In agentic commerce, trust is not a sentiment. It is the ability to verify delegated authority. That authority attaches to each action an agent takes, not to the agent itself. An agent can be legitimate, its credential valid, and its action still unauthorized. A valid payment credential no longer proves a valid decision. Agentic commerce must also verify delegated authority: proof that this agent, acting for this person or business, was allowed to take this action. This is where Mastercard's role comes forward: powering secure value exchange with the technology, protections, and intelligence that keep people and businesses in control.

For six decades, we have helped parties who have never met transact with confidence. We are extending that trust model into the agentic world, where AI agents act for people, businesses, merchants, and machines. The ecosystem has to know an agent is legitimate, that its action was authorized, and that it stayed inside its limits. Merchants need to trust the demand, issuers need to read the risk, and when questions come later, the evidence has to be there. That is how Mastercard helps new forms of value move securely, simply, and at scale.



“
*Five capabilities, one
promise: delegated action
the whole ecosystem can
trust.*

With Agent Pay, Mastercard is extending the trust that powers global digital commerce into agentic commerce—creating a trust layer that makes delegated actions safe, recognizable, governable, and controlled, unlocking what's next.

- 1 **Identity.** Know the agent, the human operator or business behind it, and the merchant, consumer, business, or machine involved, and tell legitimate actors apart from bots, spoofed identities, and malicious automation.
- 2 **Intent.** Prove what the user or business authorized through Verifiable Intent, and keep the records, audit trails, and dispute evidence that make delegation answerable.
- 3 **Controls.** Define what an agent is allowed to do, where it can act, how much it can spend, and under what conditions.
- 4 **Trusted execution.** Tokenization, authentication, and acceptance standards to protect credentials and make agent-led transactions recognizable and governable.
- 5 **Intelligence.** Apply fraud detection, behavioral analytics, risk signals, and threat intelligence to assess agent activity, surface suspicious behavior, and protect known and unknown agents alike.

Five capabilities, one connected system: the **Mastercard Agent Pay Trust Framework**. Each capability produces a proof the ecosystem can check. Together, those proofs create the foundation for recognizable trust signals, including the Agent Pay Ready mark for participants that meet Agent Pay program requirements.

Our role is not only enabling agents to make payments. It is enabling the agentic economy to grow, by giving every party a reason to trust agent-led action.

For merchants, that means recognizing legitimate agent-led demand. For issuers, richer context and clearer risk signals. For acquirers and PSPs, acceptance with consistent rules. For AI platforms, a trusted connection to commerce. For businesses and consumers, the confidence to delegate, with control, protection, and recourse.

The chapters ahead walk this framework through two journeys: a birthday gift that buys itself, guided by a household AI assistant, and a production line that never stops, kept supplied by a manufacturer's sourcing agent. One is a simple card purchase; the other settles machine-to-machine across card rails, account-to-account rails, and at times, stablecoin rails. Interfaces, actors, and rails will keep multiplying. The need for trust will remain constant. That is the trust this paper is about: how agent-led action in commerce can be verified, bounded, and answered for, not how the AI is built.





The Mastercard Agent Pay Trust Framework: five proofs that let every party say yes to delegated action.

Every delegated action raises the same question: who is acting, on whose behalf, with what authority, under what limits, and with what accountability if something goes wrong? The framework answers with five proofs. Before an agent acts, identity provides who was involved, intent proves what was authorized, and controls prove the limits. As value moves, trusted execution proves it moved securely. After and all around, intelligence proves the ecosystem can still see risk even as behavior changes. Miss any one of them and the secure default is refusal: merchants block good agents, issuers decline good transactions, and delegation stalls. With all five in place, the response can be proportionate instead of binary: approve, narrow the limits, refer back for confirmation, or decline as the risk requires. Accountability is not a sixth proof. It is what the five leave behind: evidence created before and during the action is what makes resolution possible after it. Five proofs, one system, designed for every rail—and built to power secure value exchange between people, businesses, and communities as commerce becomes more intelligent.

01

Identity

Who is acting, and on whose behalf?




The first question in every journey is the oldest one in commerce: who am I dealing with, and can they be trusted?

Start with the birthday gift. Your AI assistant sets out to find something under \$75, and within seconds it is talking to storefronts you have never visited. At the same moment, the sourcing agent is contacting suppliers its company has never bought from. Every one of those introductions carries the oldest question in commerce, now asked at machine speed: who am I dealing with, and can they be trusted? The question mirrors back at the merchant too, because a good agent and a hostile bot look almost identical at first contact.

This is why identity comes first, built as infrastructure, not a login. **Agent Portal and Registry** gives developers a front door to register agents, so your AI assistant arrives with credentials instead of a plea. Mastercard Credential Service helps establish every party in the chain with Know Your Agent, Know your Merchant and Know Your Customers (**KYA/KYM/KYC**). Verification is also how a stranger becomes a counterparty: the supplier the sourcing agent has never bought from is checked here first, then added to the approved list its spending limits point to. A **Merchant Product Catalog** grounds discovery in data merchants stand behind, a structured view of real merchants and real products instead of scraped guesswork.

The introduction goes both ways: the agent proves who it is, and the merchant proves it is who it claims to be. That is the first proof of delegated authority: who is acting, settled before anything else happens.



Identity

Who is acting,
and on whose behalf?



What exactly was authorized?

“

When the box arrives wrong, the question is never whether the agent acted. It is what was actually asked.

The gift is chosen. But before money moves, and long after it has, one question decides whether delegation survives: what exactly did you authorize? Perhaps you said under \$75 and the agent found the right gift at the wrong price: \$90. Perhaps the sourcing agent committed to a delivery window the supplier reads differently. In human commerce these are conversations. In agentic commerce they have to be evidence. Without it, every disagreement becomes the agent’s word against everyone else’s, and delegation stops at the first dispute.

Verifiable Intent creates that evidence: an open-source, privacy-preserving cryptographic framework that produces tamper-resistant proof of exactly what a consumer or business approved, including the item, budget, and timing. Through *selective disclosure*, each authorized party sees only the information needed to verify the transaction, and nothing more. Paired with *Payment Passkeys*, which confirm that the approval came from the user. With Verifiable Intent, merchants are in control, and have an enforceable mechanism to automatically validate that the purchase the agent is trying to fulfill actually corresponds to what the consumer ordered.

When the sourcing agent commits to pay on delivery next quarter, **Agent Pay for Machines** turns that promise into a verifiable, auditable artifact. And **Insight Tokens** keep even personalization inside the bounds of what you consented to, so data is governed by the same principles consumers expect.

The proof keeps working after the moment passes. When the wrong size ships, **Agentic Disputes** uses Verifiable Intent’s embedded audit trail to show what was asked, provide clear evidence to resolve the dispute, and help banks automate resolution without manually reconstructing a chain of agent decisions. People delegate more when mistakes can be fixed. The agentic economy will scale at the speed of recourse, not only the speed of execution. Intent turns delegation from something asserted into something demonstrated, before, during, and after the action. The second proof is on the record: what was authorized.



03

Controls

What is the agent allowed to do?



Autonomy is not the absence of limits. It is freedom inside boundaries someone chose on purpose.

The record proves what was asked. The next question is what the agent may do at all. Would you hand a new AI assistant your wallet on day one? Neither would a CFO. Delegation starts narrow, a single errand, a modest budget, and earns its way wider. That pattern is not reluctance; it is how trust has always been built. Agentic commerce needs the same pattern by design, so a household can say yes to the gift without saying yes to everything, and a business can let an AI agent buy parts without letting it buy whatever it likes. Without enforceable limits, delegation does not stop. Trust does.

Agent Controls put those boundaries into the transaction flow itself: budgets, categories, time windows, and conditions that decide what an agent can do before it does it. Your AI assistant holds \$75 for one gift; the sourcing agent holds a quarterly envelope for approved suppliers. And when the actors are machines entirely, **Agent Pay for Machines** holds them to the same standard, verified and transacting inside defined permissions and settlement rules.

Controls are how confidence compounds. Every widened limit is a choice someone made on purpose, visible, enforceable, and reversible. The boundary always belongs to the person or business doing the delegating, and that is exactly why they keep delegating more. The third proof defines the boundaries: what the agent may do, and how much it may spend.



04

Trusted execution

Can the transaction itself be trusted?




The moment value moves is the moment trust is spent. Execution is where every promise upstream is kept.

The gift sits in a cart you never saw; the supplier's invoice is due the moment the parts arrive. Now the money moves. Everything the journey has established, who is acting, what was authorized, what is allowed, converges on a single moment that has to be safe whether it clears on a card, an account-to-account transfer, or a stablecoin.


Tokenization makes that moment safe, agentic payment infrastructure built on tokenization and agent identity, so transactions run securely at scale, and the real credentials are never shared. The **Acceptance Framework** gives merchants, acquirers, and issuers shared rules to recognize agent-led transactions and treat them as first-class rather than anomalies to decline. **Agent Connect** ties agent platforms, merchants, and payment providers into Mastercard's agentic services, and the **Agentic Toolkit** gives developers a direct way to build trusted execution into their own agents. Your gift clears without your card number ever changing hands; the supplier is paid on rails its bank already trusts.

And when no person is present at all, **machine payments** extend the same execution to verified machines settling with machines, from the production line reordering its own parts to autonomous procurement running through the night. New interfaces, new rails, new actors, and the same standards underneath: tokenized, authenticated, recognizable, governable. This is where six decades of building trust into commerce, from acceptance rules to authentication to tokenization, matter most. New rails will arrive; the trust logic travels. And the fourth proof travels with it: value moved, nothing exposed.



Trusted execution

Can the transaction itself be trusted?



05

Intelligence

Can risk be seen as agent behavior evolves?



Most of the agents the ecosystem will meet have not been written yet. Intelligence is how trust keeps up.

The journeys do not end at settlement, because the ecosystem never stops moving. Tomorrow your AI assistant updates itself. Next month the sourcing agent negotiates with a supplier that did not exist this month. And somewhere, someone is building an agent designed to look exactly like the good ones. A trust layer that only knows yesterday's actors would fail on contact. Intelligence is the capability that reads risk in motion.

Agentic trust and intelligence services fuse identity, intent and risk signals so an agentic transaction is recognized and approved when it is legitimate. The decision is not only yes or no. A transaction can be approved as asked, approved within a narrower limit, referred back for confirmation, or declined. The response scales with the stakes: more scrutiny when consequences are high, less friction when they are not. A safe decline, explained and correctable, is a trust outcome too. Purpose-built **Fraud Models for Agentic** score the new kind of risk agent-led transactions carry, catching suspicious behavior without treating every agent as an anomaly. And **Recorded Future** safeguards watch beyond the transaction itself, tracking malicious actors and the infrastructure they build before either reaches a payment.

Intelligence is also how we assess risk when agents we have not onboarded enter the ecosystem. When an unknown agent interacts with a merchant, issuer, or processor, behavior and risk signals can help distinguish legitimate activity from malicious or anomalous behavior. This does not extend trust to the agent; it provides the visibility needed to evaluate its activity and respond appropriately. As agents and attack patterns evolve, intelligence must evolve with them. That is the fifth proof, and the one that never finishes: risk, kept visible even as it moves.

05

Intelligence

Can risk be seen as agent behavior evolves?



A win for every participant



A transaction only happens when every party has a reason to say yes.

Nothing moves unless everyone in the chain has enough trust to say yes. The framework is built to produce exactly that, on every side at once. Each yes enables the next: the consumer's control becomes the merchant's recognizable demand, the merchant's acceptance becomes the issuer's context, and every good approval makes the next delegation easier. It is how Mastercard extends its role beyond payments alone, connecting new forms of how we exchange value with confidence, control, and accountability.

Consumers and businesses get delegation they can live with: agents that act with greater confidence, control, protection, and recourse when something goes wrong.

Merchants get legitimate agent-led demand they can recognize and accept, a new front door to discovery, and confidence that what an agent brings is business worth serving.

Issuers, wallets, and payment providers get richer transaction context, with the identity, intent, and risk signals to tell an agentic purchase from ordinary e-commerce and say yes to the good activity.

Acquirers and PSPs get new forms of acceptance with clearer rules and signals, so agent-led volume arrives as opportunity, not anomaly.

AI platforms and developers get a connection to trusted commerce infrastructure, with fast, standard on-ramps instead of a payments project.

Two journeys, completed

The gift that bought itself. The AI assistant proved who it was, carried proof of what you approved, stayed inside its budget, and paid without exposing your card. The issuer recognized the purchase as agentic, saw it fit your pattern, and approved it without friction. **Verifiable Intent, Agent Controls, and Agent Pay** each did their quiet work, and the only thing you noticed was the gift arriving on time. Had anything gone wrong, the evidence to fix it was already on file.

The line that never stopped. The sourcing agent kept a factory supplied all quarter. **Agent Pay for Machines** verified every machine, **On Behalf Billing** recorded who each one acted for, and intelligence watched the pattern for drift. Suppliers were paid on time, on the rails they preferred, and nothing stepped outside the limits the manufacturer set.



How we are building it



*Trust at scale only holds
when it is open and shared.*

Agent Pay was announced in April 2025, ahead of every other payment network, with the United States as its first market. The milestones since span the major markets: Europe's first live agent-executed payment, with Santander; live, fully authorized transactions across Latin America and the Caribbean; Korea's first live agentic transactions; Singapore's first, with DBS and UOB; and Australia's first authenticated agentic transactions, with Commonwealth Bank. At the same time, we have expanded the ecosystem through collaborations with leading AI platforms and technology companies, including Google, Microsoft, OpenAI and others, helping connect trusted commerce capabilities to the next generation of digital experiences.

Agent Pay for Machines now extends the same trust to machines paying machines. Together, these proof points show Mastercard's operating-system role in action: bringing network scale, standards, partnerships, technology, protections, and intelligence to new ways of moving value. Our approach rests on three principles.

- 1 Consistent proof.** Verifiable Intent, co-developed with Google and contributed to the FIDO Alliance, is an open specification: proof of what a consumer approved is something any party to the transaction can verify, not a private claim locked to one company.
- 2 Known actors.** KYA, KYM, KYC, agentic trust and intelligence services, and purpose-built fraud models let the ecosystem recognize who is acting, whether registered or not; and approve legitimate activity.
- 3 Bounded autonomy.** Agent Controls define what an agent may do, and Agent Pay clears what it does without exposing credentials.

We also give the ecosystem tools to build and connect. **Agent Suite** lets enterprises embed agentic services into their operations, **Agent Suite for Merchants** does the same for merchant teams, and **Merchant Cloud** gives merchants and developers an environment to get commerce ready for agents. **Shopping Muse** and **AEO from Dynamic Yield** sharpen discovery and personalization on the storefront itself. **Agentic Spark** takes a team from idea to working prototype, and **Consulting for Agentic** brings the strategy to put it into production.

No one company can finish a trust layer alone. If you issue, acquire, build agents, or run commerce, the framework is open to build on; the portfolio appendix is the on-ramp. The work is not only to make agents smarter. It is to make the whole system trust what they do, so more people and businesses can participate in what comes next with confidence and control.



Conclusion

Every era of commerce began with a new form of trust. Cards let strangers trade without cash. E-commerce let them trade without ever meeting. The agentic economy asks for the biggest extension yet: trust in action itself, delegated to AI agents, on behalf of people, businesses, and machines, across every rail from cards to stablecoins. It also reinforces Mastercard's broader role: powering secure value exchange between people, businesses, and communities as new forms of value emerge.

The five proofs are how that trust gets built: identity that knows who is acting, intent that proves what was authorized, controls that set the boundaries of delegation, execution that never exposes a credential, and intelligence that learns as fast as behavior changes. Together they make delegated action safe, recognizable, governable, and scalable.

Three calls we are ready to make

Acceptance goes universal. Agent-led transactions will become routine at major digital storefronts sooner than most merchants are planning for, and demand will quietly route around the ones that are not ready.

Consumer leads, commercial scales. The steepest growth in agentic commerce will come from procurement, B2B, and machine-to-machine flows that never touch a checkout page.

Open proof wins. Evidence of authorization will consolidate on open, checkable standards, the way tokenization did, and proprietary claims will lose their audience.

This paper began with a gift bought well and a production line that never stopped. Neither is remarkable, and that is the point: when trust is built in, delegation becomes ordinary. Sixty years of turning strangers into trusted counterparties prepared us for exactly this work. We are not only enabling agents to make payments. We are helping the entire ecosystem trust agent-led action, from the first introduction to the final resolution, so people and businesses can unlock what's next. Because realizing potential is priceless.

For more information, please [contact your Mastercard representative](#).

The portfolio behind the framework

The offerings behind the Mastercard Agent Pay Trust Framework, in plain language, grouped by the capability each supports, plus the services that help the ecosystem build and connect. Every entry notes the participants it primarily serves.

● Identity

Agent Portal and Registry *Developers & AI providers*

Where developers and agent platforms discover Mastercard services, register agents, and start integrating.

KYA / KYM / KYC *Issuers & wallets, acquirers & PSPs, merchants*

Know Your Agent, Merchant, and Customer checks that establish identity and accountability for every party in a transaction.

Merchant Product Catalog *Merchants, developers & AI providers*

A structured, agent-readable view of a merchant's products, pricing, availability, and attributes, so agents can find a merchant, then recommend and buy accurately.

● Intent

Agentic Disputes *Merchants, issuers & wallets*

Dispute and chargeback tools that use agentic transaction data, intent evidence, and audit trails to resolve issues faster.

Insight Tokens *Merchants, developers & AI providers*

A secure, governed way for agents to use permissioned Mastercard insights, so personalization happens only with the consumer's consent.

Verifiable Intent *Whole ecosystem*

An open-source cryptographic framework that creates tamper-resistant proof of what a consumer or business actually approved.

● Controls

Agent Controls *Consumers, enterprises & machines, issuers & wallets*

Permission and policy controls that set what an agent can do, where it can transact, and under what limits.

On Behalf Billing *Enterprises & machines, merchants*

A billing model for transactions an agent initiates or manages for a consumer, a business, or a machine.

● Trusted execution

Acceptance Framework *Merchants, acquirers & PSPs, issuers & wallets*

The rules and expectations that let merchants, acquirers, and issuers recognize and accept agent-led transactions as first-class, not anomalies to decline.

Agent Connect *Acquirers & PSPs, developers & AI providers*

An integration layer that connects merchants, PSPs, wallets, issuers, and agent platforms into Mastercard's agentic services.

Tokenization *Merchants, issuers & wallets, developers & AI providers*

Mastercard's agentic payment infrastructure for secure, scalable agent transactions, using tokenization and agent identity instead of exposed credentials.

Agent Pay for Machines *Enterprises & machines, developers & AI providers*

Payment infrastructure for verified machines and agents to transact with defined permissions, controls, and settlement, including verifiable promise-to-pay artifacts that commit an agent to a future or conditional payment with auditable proof.



Agentic Toolkit *Developers & AI providers*

Developer tools and APIs for building agentic commerce, including Agent Pay and related workflows.

● **Intelligence**

Agentic trust and intelligence services *Issuers & wallets, acquirers & PSPs*

A decisioning layer that combines identity, intent, and risk signals so an issuer can recognize an agentic purchase, tell it apart from ordinary e-commerce, and approve the good ones.

Fraud Models for Agentic *Issuers & wallets, acquirers & PSPs*

Risk models built to spot suspicious agentic behavior and score the new kind of risk an agent-led transaction carries.

Recorded Future safeguards *Issuers & wallets, merchants*

Threat intelligence that identifies malicious actors, infrastructure, and emerging threats around agentic commerce.

● **Build and connect**

AEO from Dynamic Yield *Merchants*

AI-driven experience optimization that adapts content, offers, and recommendations to sharpen personalization and lift conversion.

Agent Suite *Enterprises & machines, developers & AI providers*

Configurable AI-agent services that let enterprises embed agentic AI into customer experiences, operations, and commerce.

Agent Suite for Merchants *Merchants*

Merchant-focused agents for conversational shopping and discovery, tied to inventory, promotions, margin, and brand voice.

Agentic Spark *Merchants, enterprises & machines*

A lightweight consulting accelerator to identify, prioritize, and prototype agentic commerce use cases.

Consulting for Agentic *Merchants, issuers & wallets, enterprises & machines*

Advisory and implementation services for agentic strategy, use cases, operating models, and go-to-market.

Merchant Cloud *Merchants, developers & AI providers*

A developer and merchant environment for commerce services, including tools to prepare merchants for agentic commerce.

Shopping Muse *Merchants, consumers*

A generative AI shopping assistant from Dynamic Yield that helps shoppers search, discover, and choose products from a retailer's catalog, in the retailer's own voice.

● **The portfolio by participant**

The same offerings, grouped by the ecosystem participant each primarily serves.

Merchants

Acceptance Framework, AEO from Dynamic Yield, Agent Pay, Agent Suite for Merchants, Agentic Disputes, Agentic Spark, Consulting for Agentic, Insight Tokens, KYA / KYM / KYC, Merchant Cloud, Merchant Product Catalog, On Behalf Billing, Recorded Future safeguards, Shopping Muse, Verifiable Intent.

Issuers, wallets, and payment providers

Acceptance Framework, Agent Controls, Agent Pay, Agentic Disputes, agentic trust and intelligence services, Consulting for Agentic, Fraud Models for Agentic, KYA / KYM / KYC, Recorded Future safeguards, Verifiable Intent.

Acquirers and PSPs

Acceptance Framework, Agent Connect, agentic trust and intelligence services, Fraud Models for Agentic, KYA / KYM / KYC, Verifiable Intent.

AI platforms and developers

Agent Connect, Agent Pay, Agent Pay for Machines, Agent Portal and Registry, Agent Suite, Agentic Toolkit, Insight Tokens, Merchant Cloud, Merchant Product Catalog, Verifiable Intent.

Product availability, features, and timing are subject to change. Any forward-looking statements involve assumptions and judgments, and there is no assurance they will be realized. Mastercard and the circles design are trademarks of Mastercard International Incorporated.

