ConvPayMAS: Conversational Payment Multi-Agent System with Agent-to-Agent Protocol and Three-Mandate Verification

Demonstration Track

Joon Kiat CHUA*
Singapore Management University
Singapore, Singapore
jk.chua.2023@msc.smu.edu.sg

Chen Hao TSE*
Mastercard
Singapore, Singapore
alvin.tse@mastercard.com

Donghao HUANG
Mastercard
Arlington, Virginia, United States
donghao.huang@mastercard.com

Zhaoxia WANG
Singapore Management University
Singapore, Singapore
zxwang@smu.edu.sg

ABSTRACT

Agentic commerce promises end-to-end shopping experiences driven by collaborating LLM agents, but deployment is constrained by fragmented checkout flows and strict payment-security and compliance requirements (e.g., PCI). We present ConvPayMAS, an LLM-based conversational payment multi-agent system that encapsulates payment execution behind a *Conversational Supervisor Agent (CSA)*, allowing shopping and merchant agents to invoke payment without handling sensitive card data. ConvPayMAS interoperates via Google's *Agent2Agent (A2A)* protocol for trusted handshakes and secure messaging, and implements Google's *Agent Payment Protocol (AP2)* using a cryptographic chain of three mandates - Intent, Cart, and Payment - to enable verifiable authorization and dispute resolution. Payment capabilities are exposed through an MCP tools server supporting sessions, wallet operations, and real-time authorization via Mastercard infrastructure. We demonstrate end-to-end conversational checkout, live mandate verification, card-security workflows, multi-agent coordination, and an autonomous travel-research scenario with pre-authorization.

KEYWORDS

LLM Multi-Agent Systems; Agentic Payments; Conversational Agents

ACM Reference Format:
Joon Kiat CHUA, Chen Hao TSE, Donghao HUANG, and Zhaoxia WANG. 2026. ConvPayMAS: Conversational Payment Multi-Agent System with Agent-to-Agent Protocol and Three-Mandate Verification: Demonstration Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/YDMY4904

1 INTRODUCTION

Large Language Models (LLMs) increasingly power multi-agent systems (MAS) across domains [5, 12, 13]. A particularly promising direction is *agentic payments*, where LLM agents negotiate and execute transactions on users' behalf [7, 9, 14]. This shifts commerce from manual marketplace browsing to an ecosystem of collaborating agents that deliver an end-to-end shopping and payment experience. Adoption, however, is constrained by technical and regulatory frictions in existing payment rails - including fraud detection, authorization workflows, and compliance controls - which make integration non-trivial [2, 7]. The challenges are:

(1) **Regulatory and PCI compliance.** Payment data handling is governed by PCI Security Standards [10, 11]. In agentic commerce, merchants and Payment Service Providers (PSPs) remain accountable for PCI obligations, complicating integrations where agents must enable payment without exposing sensitive cardholder data.
(2) **Lack of an end-to-end agentic payment framework.** Despite progress in LLM-based MAS [5, 12, 13], a practical reference architecture for agentic payment remains missing. Industry proposals still require substantial integration with existing rails [9, 14], slowing adoption.
(3) **Fragmented user experience.** Without native payment integrations, agents often require users to exit the conversational interface to complete checkout in external web forms.

To address these challenges, we propose the ***LLM-Based Conversational Payment Multi-Agent System*** **(ConvPayMAS)**, which abstracts payment-rail complexity within a specialized multi-agent architecture. In this demonstration[1], we show how ConvPayMAS interoperates with other agents in an agentic payment environment to enable practical conversational commerce. Our key contributions are firstly, a **single entry point for payment.** ConvPayMAS exposes a *Conversational Supervisor Agent* (CSA) as a unified interface via standardized protocols, enabling external agents to access Mastercard payment functions without integrating multiple APIs. Next, we **isolate PCI's scope**, where sensitive payment handling occurs within ConvPayMAS, allowing external agents to delegate payment to the CSA without touching sensitive cardholder data. Lastly, we create an **end-to-end conversational flow**, as ConvPayMAS keeps the payment journey within the conversational loop, reducing chat-to-checkout context switching from product discovery through payment confirmation.

*Both authors contributed equally to this research.

[1] Demo video URL: https://vimeo.com/1152577656

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/YDMY4904

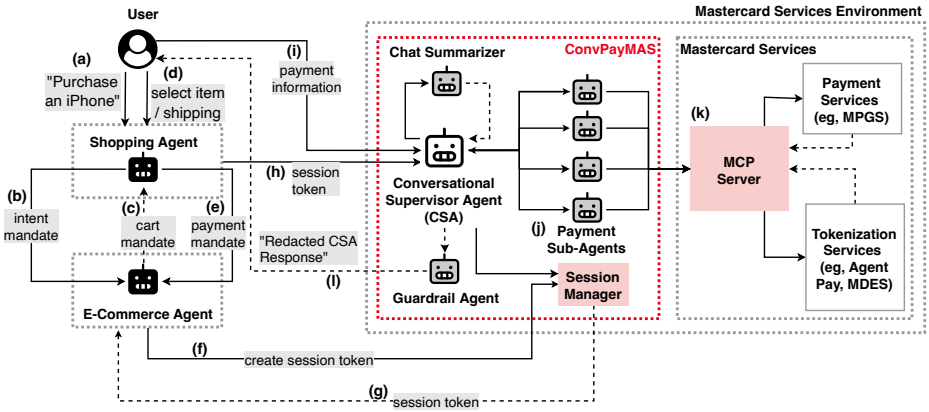

Figure 1: Overview of the agentic payment ecosystem with ConvPayMAS.

2 METHODOLOGY AND SYSTEM DESIGN

This section describes the agentic payment architecture (Fig. 1) and the design of ConvPayMAS. The system comprises three top-level agents that communicate via standardized protocols: **(1) ConvPayMAS**, a secure server-side agent providing a conversational interface for payment processing; **(2) Shopper Agent**, a user-side assistant that interprets intent and navigates marketplaces; and **(3) E-commerce Agent**, a merchant-side agent that handles catalog queries, cart operations, and payment initiation via ConvPayMAS.

2.1 ConvPayMAS Design

ConvPayMAS adopts a hierarchical modular architecture for deterministic payment execution, orchestrated with LangGraph [8]. **Conversational Supervisor Agent (CSA).** The CSA is the sole entry point for external agents. It parses natural-language requests and delegates execution to internal sub-agents.

Payment Sub-Agents. ConvPayMAS employs five specialized sub-agents (Fig. 1f) using **role-based protocols** [5, 6, 13]: **List Card** retrieves masked PANs from the wallet; **Create Card** accepts encrypted card data (AES-256-CBC) and initiates OTP; **Validate Card** verifies OTP to activate cards; **Charge Card** executes transactions; **Recommender** suggests optimal cards for the transaction. A **Chat Summarizer** compresses conversation history within bounded context windows. While the current implementation focuses on card payment, ConvPayMAS is designed to be extensible to additional payment rails (e.g., bank-account transfers or stablecoin-based payment) by swapping or adding function-specific sub-agents and MCP tools behind the same conversational interface.

MCP Server Tools. ConvPayMAS exposes MCP [1] tools for card registration, session management, and payment operations, enabling secure tokenization, real-time authorization, and wallet management through Mastercard's payment infrastructure.

2.2 Security and Protocol Design

Agent-to-Agent Protocol - We use Google's Agent2Agent (A2A) protocol [3] for trusted and secure inter-agent communication. Messages are protected with AES-256-CBC using session keys, and sessions are managed with JWT (HS256). **Three-Mandate Verification -** We implement Google's Agent Payment Protocol (AP2) [4] using three mandates: an *Intent Mandate* capturing purchase intent; a merchant-signed *Cart Mandate*; and a *Payment Mandate* binding user authorization to the cart via a hash chain for verifiable dispute resolution. **Card Data Security -** Card data (PAN, CVV, expiry) is encrypted client-side with AES-256-CBC. The MCP server issues per-session keys. PANs are validated with a Luhn check and expiry dates are verified before wallet storage. A Guardrail Agent redacts PCI/PII from conversation logs before any external transmission.

2.3 Payment Flow and Live Demonstration

The demonstration presents the below transaction flow live, with scenarios for mandate verification, card security, and multi-agent coordination. The transaction journey proceeds through four phases:

Phase 1 – Intent & Discovery. The user expresses intent to the Shopper Agent (e.g., "Purchase an iPhone") (Fig. 1a). The Shopper Agent creates an *IntentMandate* and sends it to the E-commerce Agent via A2A streaming (Fig. 1b). The merchant searches inventory and returns a signed *CartMandate* (Fig. 1c) with items, prices, and shipping options—digitally signed with HMAC-SHA256. If there are price modifications, a new *CartMandate* must be signed.

Phase 2 - Selection & Payment Mandate: The user selects an item and shipping method (Fig. 1d). The Shopper Agent creates a *PaymentMandate* containing a JWT that references the Cart Mandate's authorization and the hashed payment contents.

Phase 3 - Session Creation: The E-commerce Agent receives the *PaymentMandate* (Fig. 1e) via A2A and validates all three mandates' cryptographic integrity—verifying JWT signatures, expiry times, audience claims, and hash chain. Upon validation, it creates a transaction session (Fig. 1f), returning a session token (Fig. 1g).

Phase 4 - Payment Execution: The user is handed off to ConvPayMAS with the session token (Fig. 1h, i). The CSA routes requests to appropriate sub-agents (Fig. 1j)—*List Card, Create Card, Validate Card, Charge Card*—which invoke MCP server tools for payment processing (Fig. 1k). Upon completion, the CSA returns confirmation via natural language (Fig. 1l).

3 ACKNOWLEDGMENTS

This work was supported by Mastercard Foundry R&D. The authors thank Anamika Thokal and Varuna Ektare from the Mastercard Foundry Pune R&D team for their assistance with the demonstration video and program management, respectively. This work was also supported by the EngD Program at Singapore Management University.

REFERENCES

[1] Anthropic. 2024. What is the Model Context Protocol (MCP)? https://modelcontextprotocol.io/docs/getting-started/intro.
[2] Lin Claudia, Taylor Ken, Zukowsky Rich, and Wrocherinsky Dalia. 2025. The Next AI Frontier: From Prompts to Purchases. https://www.dwt.com/blogs/artificial-intelligence-law-advisor/2025/10/agentic-ai-concerns-for-merchants-and-issuers. Accessed: 2025-12-30.
[3] Google. 2025. Agent2Agent (A2A) Protocol. https://a2a-protocol.org/latest/.
[4] Google. 2025. Powering AI commerce with the new Agent Payments Protocol (AP2). https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol.
[5] Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: a survey of progress and challenges. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence* (Jeju, Korea) *(IJCAI '24)*. Article 890, 10 pages. https://doi.org/10.24963/ijcai.2024/890
[6] Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. MetaGPT: Meta Programming for A Multi-Agent Collaborative Framework. In *The Twelfth International Conference on Learning Representations*. https://openreview.net/forum?id=VtmBAGCN7o
[7] Schumacher Katharina, Roberts Roger, and Giebel Katharina. 2025. The agentic commerce opportunity: How AI agents are ushering in a new era for consumers and merchants. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-agentic-commerce-opportunity-how-ai-agents-are-ushering-in-a-new-era-for-consumers-and-merchants. Accessed: 2025-12-30.
[8] LangChain. 2025. Balance agent control with agency. https://www.langchain.com/langgraph. Accessed: 2025-12-30.
[9] Mastercard. 2025. Mastercard Agent Pay - Powering the next frontier of commerce. https://www.mastercard.com/global/en/business/artificial-intelligence/mastercard-agent-pay.html. Accessed: 2025-12-30.
[10] PCI Security Standards Council. 2025. AI Principles: Securing the Use of AI in Payment Environments. https://blog.pcisecuritystandards.org/ai-principles-securing-the-use-of-ai-in-payment-environments. Accessed: 2025-12-30.
[11] PCI Security Standards Council. 2025. PCI Security Standards Overview. https://www.pcisecuritystandards.org/standards/. Accessed: 2025-12-30.
[12] Aske Plaat, Max van Duijn, Niki van Stein, Mike Preuss, Peter van der Putten, and Kees Joost Batenburg. 2025. Agentic large language models, a survey. *arXiv preprint arXiv:2503.23037* (2025).
[13] Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O'Sullivan, and Hoang D Nguyen. 2025. Multi-Agent Collaboration Mechanisms: A Survey of LLMs. *arXiv preprint arXiv:2501.06322* (2025).
[14] Visa. 2025. Enabling AI agents to buy securely and seamlessly. https://corporate.visa.com/en/products/intelligent-commerce.html. Accessed: 2025-12-30.