Autonomous Commerce and the Rise of Digital Trust

Toshihiro Koda

10 April 2026

Iakyara vol.413

Nomura Research Institute, Ltd.

Toshihiro Koda
Pacific Branch Manager
Nomura Research Institute IT Solutions America, Inc.

Executive Summary

As generative AI has advanced, autonomous agents capable of negotiating, executing transactions, and even making payments are rapidly becoming a reality. To usher in a new era of autonomous commerce, companies are making progress on challenges such as verifying the scope of AI agents' authority. Following such companies' lead, financial institutions should swiftly pursue their own digital trust initiatives.

Shift in transaction execution from humans to agents

Generative AI has evolved from a tool that answers queries conversationally to agents that autonomously carry out users' intentions. AI agents are starting to increasingly shop on behalf of humans within preset constraints, handling everything from gathering information and comparing terms to making inquiries, negotiating and even subscribing to services.

This shift's significance goes beyond efficiency gains. Today, online transactions rely on login, authentication and consent. These steps substantiate the legitimacy of online transactions. However, the same may not be true in a world where AI agents act as intermediaries. The agentic AI era requires infrastructure that verifiably records and recounts who authorized an agent to act, how much authority the agent has and what the agent based its decisions upon. These questions fall within the domain of digital trust, which will become increasingly important as AI agents become more involved in commerce.

Autonomous commerce leading the way for AI agents

Autonomous commerce refers to a mode of transaction in which an AI agent is authorized by a user to transact on the user's behalf and in line with their intentions. In doing so, the agent typically handles everything from comparing products and negotiating with sellers to making purchases and payments.

With conventional e-commerce, users themselves select products on their device screens, add them to their virtual cart and pay for them. In autonomous commerce, this entire process is carried out by AI agents. Agents are first starting

to be deployed in commercial transactions because they tend to be recurring and high-frequency and are accordingly expected to yield a big payoff. At the same time, challenges around control are likely to arise, as such transactions involve multiple parties, including merchants, distributors, and payment processors. Consequently, autonomous commerce is set to be an early testing ground for digital trust.

For example, AI agents gather information on product availability and returnability, redeem coupons and place orders. In the event of a delivery delay, an AI agent may switch to an alternative supplier and track product returns and refunds. However, such changes may entail adjustments to delivery destinations or payment methods, adding complexity to tracking transaction status. Autonomous commerce therefore requires the ability to track agentic workflows from start to finish on a transaction-by-transaction basis and visibility into the basis of agents' decisions at every stage.

Elements of digital trust supporting autonomous commerce

To resolve such issues, digital trust must encompass five components: consent to the scope and duration of an agent's permissions; explicit delegation of authority to the agent; fraud prevention and authentication by means of, e.g., electronic signatures; auditing of records of the basis of agents' decisions; and end-to-end status tracking with transaction IDs (see diagram).

Private companies have started proposing protocols that meet such requirements. In January 2026, Google unveiled a Universal Commerce Protocol (UCP) co-

Five components of digital trust supporting autonomous commerce

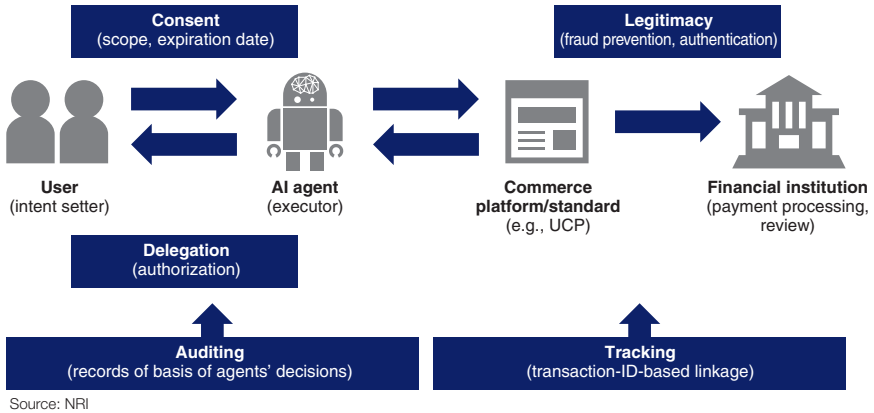

Source: NRI

developed with Shopify. UCP was designed as a common standard that can be used among companies throughout the process from searching for products through handling after-sale returns and inquiries. Before connecting to the UCP ecosystem, participating businesses publish a manifest of the UCP services and capabilities they have elected to support. Buyer consent is machine-verifiable and subject to user-imposed limitations with respect to purpose, scope and duration. UCP specifications include scoped delegation of authority and tracking with transaction IDs. Of the above five components of digital trust, UCP directly supports consent, delegation, auditing and status tracking. UCP's core capabilities include payments, account linking and order management also. UCP has been endorsed by major retailers including Walmart and Target.

In the payment space, Mastercard has proposed an authority verification and authentication protocol for purchases made by agents on behalf of cardholders through Mastercard Agent Pay. Visa's Trusted Agent Protocol aims to provide a common protocol for acceptance of agentic transactions by merchants and a standardized framework for recognizing legitimate agents and preventing access by malicious actors. Although Mastercard Agent Pay and Trusted Agent Protocol are both still in the proposal stage, payment is the point at which a transaction is consummated and linked to verification in the event of a dispute. Payment processors must adopt authentication and authority verification protocols that enable them to meet the requirements posed by payments from agents.

AI agents prospective spread to non-payment financial services

In the future, AI agents may also become involved in non-payment financial services. Potential applications include account operations and credit decisions at banks, suitability-based investment decisions at securities brokerages and insurer workflows ranging from policy enrollment to benefit claim processing. All of these require managing the scope of consent and authority in data form and recording the basis of agents' decisions in an auditable form. The five components of digital trust listed above are a useful lens for looking at how agentic AI might be deployed in the financial sector. However, financial institutions' businesses are diverse and have differing priorities and granularity requirements. One commonality across all financial services is a need for transaction-by-transaction verification infrastructure. A realistic way for financial institutions to proceed would be as follows.

First, define a common, transaction-ID-based key that ties together every task from authentication to payment/refund into a single workflow. Second, treat consent and authorization as structured data and establish a rule against processing any request unless authorization can be verified. Third, establish a standardized authentication procedure based on electronic signatures or tokens. Fourth, ensure that consent timestamps, records of the scope of agentic authorization, and verification results are retained as minimum required records for audit purposes. Fifth, deploy AI agents in phases, beginning with low-stakes tasks like processing of employee expense reports and then progressively expanding the scope of deployment while fixing any problems with the agents' performance.

Although agentic AI is still in its early stages, financial institutions that wait for standards and rules to be fully established may fall behind. Many of the required capabilities are adjacent to initiatives already underway, such as enhancements to open API platforms and authentication. Financial institutions can get started by utilizing such existing initiatives. As AI agents emerge as a new customer touchpoint, financial institutions that establish digital trust are more likely to be selected as transaction counterparties. How soon a financial institution achieves digital trust could be a key determinant of its future competitiveness.

about NRI

Founded in 1965, Nomura Research Institute (NRI) is a leading global provider of system solutions and consulting services with annual sales above $5.1 billion. NRI offers clients holistic support of all aspects of operations from back- to front-office, with NRI's research expertise and innovative solutions as well as understanding of operational challenges faced by financial services firms. The clients include broker-dealers, asset managers, banks and insurance providers. NRI has its offices globally including Tokyo, New York, London, Beijing and Sydney, with over 16,700 employees.

For more information, visit https://www.nri.com/en

Inquiries to : Business Planning Department, Financial Technology Solution Division
Nomura Research Institute, Ltd.
Otemachi Financial City Grand Cube,
1-9-2 Otemachi, Chiyoda-ku, Tokyo 100-0004, Japan
E-mail : kyara@nri.co.jp

https://www.nri.com/en/knowledge/publication/list.html#lakyara