Agentic AI, Open Finance, & Technical Standards

Call for Input

April 2026

FINANCIAL® DATA EXCHANGE

Introduction

As Agentic AI transforms numerous use cases across financial services, the potential for both benefit and harm to consumers and businesses is profound, and safe adoption will require responsible industry collaboration.

As the leading technical standards body in North America focused on "Open Finance", the **Financial Data Exchange** (FDX) is convening a cross-industry effort to promote safety and interoperability when agentic AI is used to facilitate **user-permissioned financial data sharing**.

We invite industry stakeholders to engage in three ways:

1. **Share your input** on the Five Question Areas described below by Friday, May 29, 2026. Please email your feedback to input@financialdataexchange.org [1]

2. **Review the accompanying Exploration Brief** - This presents an initial framing of the landscape, use cases, and novel questions to solve. This Brief is borne out of the early work of FDX's AI Exploratory Task Force.

3. **Get involved with FDX's discussions** on Agentic AI & Open Finance with other industry leaders, if you are not already. [Contact FDX for more information.]

Our goal

In the months ahead, FDX expects to issue updated technical standards and/or industry guidelines—shaped by the input we received—on the safe usage of Agentic AI in the sharing of permissioned financial data.

As AI technology gains widespread adoption, our aim is to help the industry moves forward—and not backward—in building out a durable and interoperable infrastructure layer for data-sharing connections that promotes widespread trust—among consumers, regulators, financial institutions, software providers, and others—while fostering innovation and efficiency.

[1] Any contributions by individuals or organizations who are not FDX members are subject to the Non-Member Contributor IP Rights Agreement linked here.[1]

FINANCIAL® DATA EXCHANGE

Why industry collaboration is needed now on the use of AI in user-permissioned financial data sharing

Starting point: Permissioned financial data sharing in North America today is built on a relatively mature foundation. Over the past several years, FDX and its members have established a robust technical framework for structured APIs, token-based authentication via OAuth, scoped consent management, and data recipient registration to govern how consumer financial data flows between parties. More than 114 million customer accounts today are connected through token-based APIs following the FDX API standard. These mechanisms – anchored by FDX's core principles of Control, Access, Transparency, Traceability, and Security – provide a strong technical foundation for extending to new patterns of data access.

The impact of AI agents: Now, AI-powered agents are introducing new patterns into this ecosystem. Intelligent software – sometimes acting alongside a user, sometimes acting on their behalf without real-time supervision – is beginning to intermediate the relationship between consumers and their financial data. These agents may operate asynchronously, persist across sessions, chain together multiple tools and data sources, and make autonomous decisions about what data to access and when. Unlike the applications that came before them, their purpose can be fluid: the same agent might help a consumer budget today, shop for a better mortgage tomorrow, and optimize tax strategy next quarter.

These new patterns challenge assumptions that are embedded in current consent and privacy frameworks. For example, current consent models generally assume a static application, a single well-defined purpose, and user-initiated actions. When an agent's future actions are not fully known at the time consent is given, and when outputs can be repurposed downstream in ways a user may not have contemplated, the meaning of "consent" itself changes.

This industry change is not just theoretical. Financial institutions, fintechs, and AI platform operators are already piloting or deploying agents that play various roles in accessing consumer financial data – and the pace is accelerating.

In many ways, FDX's current standards already help to support safe and transparent data integrations, including when Agentic AI is involved. Our focus is on the gaps where agentic AI introduces new challenges that current mechanisms do not fully address – particularly around consent delegation and purpose, agent identification, data minimization, and accountability.

There is an important and timely role for tailored standards or industry guidance in this space. FDX was created to help enable safe, secure, and interoperable user-permissioned financial data sharing. Where other organizations are building new standards, protocols, and infrastructure to support other agentic use cases (see Appendix), FDX sees an opportunity at the technical standards layer specific to user-permissioned financial data sharing – defining how consent and traceability should work when agents intermediate financial data sharing flows.

As the leading technical standards body for user-permissioned financial data sharing in North America, FDX seeks to lead this conversation – collaborating with other emerging standards and frameworks where they apply, and adding specific technical standards or guidance where gaps remain. Work is actively underway at organizations including NIST, the OpenID Foundation, and the Linux Foundation on adjacent dimensions of agent identity, security, and infrastructure. FDX seeks to complement rather than duplicate these efforts.

Aligning new patterns with five principles

In defining and building new or adapted access methods and technical standards that utilize AI, we believe that consumers and businesses stand to benefit long-term by maintaining a focus on five core principles.

When these principles are met, end users are empowered to better understand, leverage, and benefit from their financial data in a secure and reliable manner. Meanwhile, companies who share or collect permissioned financial data will be empowered to maintain the trust that consumers and regulators expect from companies offering financial services.

CONTROL

Account owners should be able to permission their financial data using intuitive navigation and be provided information that allows informed decision making. Account owners should be able to easily and intuitively grant, modify and revoke access to their financial data.

ACCESS

Account owners should have access to their data and the ability to determine which parties will have access to their data. Connections between parties should be convenient, secure, efficient, and avoid unnecessary steps. Authorized parties should only have access for the purposes for which user's consent was provided.

TRANSPARENCY

Individuals should know what data is being shared and how, when, for what purpose, and under what terms their permissioned data is used. Only data that is required to provide such services should be shared.

TRACEABILITY

All data transfers should be traceable. Account owners should have a complete view of all parties that are involved in the data-sharing flow.

SECURITY

All parties should ensure the safety and privacy of data during access and transport and when that data is at rest.

Call for Stakeholder Input

Who this Call for Input is for. Anyone involved in or affected by AI agents accessing consumer permissioned financial data, including: financial institutions, data aggregators, fintechs, AI platform operators, identity and security standards bodies, and consumer advocates. If you are not an FDX Member, we still welcome your input.

What is FDX: Financial Data Exchange (FDX) is a nonprofit standards body that seeks to unify the financial industry around a common, interoperable, royalty-free standard for the secure access of user-permissioned financial data. Founded in 2018, FDX has roughly 200 member organizations. More than 114 million customer accounts are connected through the FDX API standard today.

How to respond: Send email feedback to input@financialdataexchange.org. You may respond to any or all of the questions, share additional use cases or challenges we should consider, or provide broader feedback on FDX's role in this space.

1. *Agent Identification and Delegated Access*: How should agents identify themselves when accessing consumer-permissioned financial data?

A. For API access, can existing OAuth/OIDC mechanisms be extended with agent-type metadata (e.g., extending FDX's existing API headers)?

B. For credential-based access, is a new class of signaling needed?

C. What does "Know Your Agent" mean operationally in the context of consumer-permissioned financial data?

D. What minimum requirements should apply to an organization operating an agent that accesses consumer financial data – is there a "responsible agent operator" framework that FDX could define?

Helpful information could include: input from identity/security practitioners; analysis of existing data recipient requirements and what "minimum viable trust" looks like for agent operators; technical analysis of OAuth flows under agentic access patterns; NIST NCCoE findings

2. *Consent and Delegation*: What does meaningful consent look like when the consenting party is software acting over extended periods?

A. How should technical standards evolve to enable delegation and maintain user control when agentic AI is involved in data sharing process?

B. To allow for time-bounded or purpose-bounded delegation, what new specific features or principles could be introduced into technical standards, if any?

C. What role should step-up authentication play for high-risk agent actions?

D. When an agent acts autonomously, whose consent is being exercised – the user's, the agent's, or the platform operator's?

Helpful information could include: Consumer research on understanding of delegated consent; analysis of dynamic consent models; legal/regulatory perspectives on "continuing consent" across jurisdictions.

3. ***Downstream Accountability and Data Provenance*:** Who is accountable when data flows through agentic chains to parties not registered under existing frameworks?

A. When an agent accesses data via FDX APIs and forwards it to a third-party service for processing, does the existing data recipient framework extend to this "third hop"?

B. How should FDX encourage architectural patterns that minimize raw data exposure (e.g., structured tool calls, on-device processing) while maintaining accountability when data does leave the controlled chain? *(Note: This question is informational; FDX does not currently anticipate defining liability frameworks.)*

Helpful information could include: Architectural analysis of common agentic data flows; legal opinions on data controller/processor obligations in chained agent scenarios.

4. ***Scope, Data Minimization, and Revocation:*** How should scope frameworks and data handling evolve for agents whose purpose is fluid?

A. Given that FDX defines scopes at the data-cluster level, what (if anything) about agents requires evolving how we define and match scopes?

B. Should consent be transparent about AI use and fluid purpose, rather than trying to enumerate all possible purposes in advance?

C. What should "revocation" mean when an agent has already processed, summarized, or stored data in its memory or reasoning context?

D. Is there a practical mechanism for revoking derived access? And is this topic better suited for industry standards or for bilateral contracts and regulation?

Helpful information could include: Analysis of existing scope taxonomies; prototyping of adaptive scope models; technical analysis of agent memory architectures; legal precedent on data deletion in derived/processed contexts.

5. ***External Standards Alignment*:** How should FDX relate to the emerging cross-industry agent standards landscape?

A. How should FDX relate to emerging standards from NIST, OpenID/FAPI, the Linux Foundation (MCP, Agentic AI Foundation), and agent commerce protocols (Visa TAP, Stripe ACP, Google A2A)? *[See Appendix Table 1 in the Exploration Brief]*

B. Where should FDX consume or profile existing standards versus defining its own?

C. Is there a case for FDX to define an "agent-safe" API profile – a constrained set of endpoints, scopes, and behaviors designed specifically for agentic access patterns?

D. What would make FDX's standards and guidance compelling for AI platform operators to implement?