FINANCIAL® DATA EXCHANGE

Agentic AI and Permissioned Financial Data:

What the Industry Told FDX

Synthesis of Industry Responses to FDX's Call for Input

November 2026

FINANCIAL® DATA EXCHANGE

Introduction

Agentic AI tools are increasingly being used to access consumer-permissioned financial data, bringing an array of critical questions for industry members to address.

From April through June 2026, FDX ran a **Call for Input** to seek broad industry input on

(1) how agentic AI is impacting consumer-permissioned financial data sharing, and

(2) how industry standards may need to evolve accordingly.

This document is a high-level summary of key themes that emerged.[1]

Around 30 organizations submitted comments, totaling about 500 pages. FDX also held meetings over the summer with dozens more organizations to hear their perspectives. FDX members are invited to review the detailed submissions here.

Two key overarching themes emerged.

First, AI tools and agents should be guided to use governed, interoperable, consumer-controlled channels when they enable people to connect their financial data, rather than ungoverned credential-based (screen-scraping) methods of access.

Second, important work is needed across a number of dimensions—including consent synchronization, semantics, security—to evolve the third-party API channels of yesterday to appropriately serve consumers using tools powered by Agentic AI.

To that end, several workstreams are underway at FDX— overseen by the FDX Agentic AI Task Force— to create new or adapted FDX technical standards that can help the industry safely enable interoperable data sharing with agentic AI tools and products. If you'd like to be engaged in those detailed conversations, contact FDXSupport@financialdataexchange.org.

In this synopsis paper, we'll summarize:

Where and how Agentic AI is being used today in Open Finance 3

Six scenarios in focus 3

1. An AI agent uses "screen-scraping" in lieu of APIs / governed channels 4
2. A consumer grants consent to an agent with broad or evolving purpose(s) 4
3. Permissioned financial data flows to a large language model (LLM) 5
4. An AI agent goes beyond "read only" to perform "write" actions at a financial institution 6
5. An Agentic AI tool seeks to integrate (e.g., with a bank) via an Agentic Tooling Interface (e.g., MCP) 6

[1] This document is not part of the FDX standard or an output of FDX's consensus-driven technical standard setting process, but rather a general summary. It does not represent a view on FDX's priorities or roadmap. Submissions are summarized, not quoted, and no respondent is named. If you want to know what a particular company argued, FDX members are invited to review written submissions [here].

6. An AI agent accesses an API but misinterprets the data retrieved 7
Three cross-cutting themes 7
What happens next at FDX 8
Getting engaged 8

Where and how Agentic AI is being used today in Open Finance

Numerous real-world examples were cited by respondents on where consumers are connecting their financial accounts to agentic AI tools today. Examples included:

- Large LLM providers offering personal finance features that enable consumers to connect their accounts through a data access platform
- Account connector tools registered in the directory of popular chatbots, enabling connectivity directly to select financial data holders
- Local agents and agentic browsers (sometimes open-source; sometimes hosted in virtual machines) that access data from a financial institution's website
- Fintech and banking apps (acting as data recipients) that leverage LLMs as back-end tools for processing user-permissioned data
- Registered Fintechs and Data Access Platforms using AI agents to assist in the process of fetching data from financial institutions, both via API and via credential-based methods.

The breadth of these use cases highlights the importance of anchoring any conversation around Agentic AI and Open Finance into specific **scenarios** and the problems they could create, before advancing to a discussion of where standards might help.

Six scenarios in focus

Six scenarios in particular received a lot of focus among the organizations we heard from in the Call for Input. These were:

1. An AI agent uses "screen-scraping" in lieu of APIs / governed channels
2. A consumer grants consent to an AI agent with broad or evolving purpose(s)
3. Permissioned financial data flows to a large language model (LLM)
4. An AI agent goes beyond "read only" to perform "write" actions at a financial institution
5. An Agentic AI Tool Seeks to Integrate (e.g., with a bank) via an Agentic Tooling Interface
6. An AI agent accesses an API but misinterprets the data retrieved

We'll take each of these scenarios in turn and highlight some of the suggestions that submitters provided on where standards might help.

1. An AI agent uses "screen-scraping" in lieu of APIs / governed channels

Many respondents focused on the potential risks and harms that may arise if agentic AI tools begin using credential-based ("screen-scraping") methods to access permissioned financial data, rather than using purpose-built channels (e.g., token-based, FDX-aligned APIs), where they exist.

The situation: An agent opens the bank's website inside the person's own browser, or using an agentic browser tool, in a live session using the consumer's saved credentials, and starts collecting data or performing other actions. To the data provider, that traffic may appear the same as the customer: same credentials, no agent identity, no scoped token, no receiving company or tool registered.

What we heard from respondents: Many respondents highlighted the risks this could pose, both to consumers and to financial institutions. For example, a bank may face increased account-takeover and unauthorized-transaction losses it can't assign to anyone, because the activity ran under the customer's own credentials and there's no way to establish which company acted or received the data. Meanwhile, bot defenses may not be able to distinguish an authorized agent, a real human login, and a malicious unauthorized bot. Many may block real customers in the process.

Many also highlighted that unless the governed path (e.g., a token-based API) is available and works better than the screen-scraping option, agents may route around it and use screen-scraping.

A majority of respondents urged industry participants to eliminate credential-based agent access, rather than build signaling around it. A smaller group argued for some form of new standards to be created (e.g., agent-attestation signaling) to make credential-based agentic access less risky, potentially as an interim measure (since many financial institutions still do not have dedicated APIs for third party data access).

At FDX, we believe that— in enabling permissioned financial data access—consumers and businesses stand to benefit long-term by using integration methods that support five core principles: Control, Access, Transparency, Traceability, and Security.

FDX's technical standards have been thoughtfully crafted over many years to enable implementers to deliver features that support those principles, in a broadly interoperable way. Those features include: consent synchronization; mechanisms for various parties to maintain consumer control dashboards; token security; a standardized path for consumer consent; methods for promoting data minimization; error handling; and more.

While technical standards and access methods will evolve over time, holding true to these five principles will allow end users to be 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.

2. A consumer grants consent to an agent with broad or evolving purpose(s)

The situation: Someone tells an agent to help manage their money. The agent decides what types of data it will ask for; either the agent or the consumer completes some version of the authorization flow; and the agentic tool collects the permissioned data.

Then later, the scope of data the agent seeks may change. For example, the agent decides it now wants additional categories of data, or two years of history across more accounts, pulled daily rather than once. That change may happen for different reasons: the person provided a new prompt; the model may be updated and now reason differently about the same unchanged instruction; or the agent (being non-deterministic) simply decided differently on Wednesday than it did on Monday.

What we heard from respondents: Many focused on the potential need for new mechanisms for capturing consent to handle these scenarios. The reasoning was fairly consistent: today's consent model assumes fixed software, one declared purpose, and a person present to start each action. Agentic AI tools may break these premises.

Respondents shared some common concerns. Ask for confirmation on every action, and people may approve reflexively without reading. Ask once for a broad consent upfront, and there is risk that consumers may not understand what they've authorized or may share more data than is needed. Others raised questions about whether data minimization would stop being a design principle and or become something enforced at runtime.

Respondents had differing suggestions on what should count as a new consent grant.

- Some suggested a machine-readable delegation object carrying purpose, scope, duration and renewal rules, evaluated on every call.
- Some suggested that today's standards for consent flows be extended with a purpose dimension, rather than replaced.
- Some suggested that the delegation of intent or consent to an agent be expressed as a signed credential (with the option for FDX to profile signed credential standards being developed by other standards bodies).
- Some suggested that FDX's existing token and consent structures already handle these new agentic scenarios adequately, if they are appropriately applied by implementers.
- There was not a broad agreement among respondents on when scope broadening should force a fresh authorization by the end user.

[While more discussion is underway at FDX to talk through these differing solution approaches, it would be fair to report that a large share of the respondents we've heard from to date believe ***some changes*** to consent scope management practices commonly used today would be needed to appropriately suit Agentic AI use cases.]

3. Permissioned financial data flows to a large language model (LLM)

Many respondents highlighted how numerous parties in the data-sharing chain may now be using LLMs in a variety of ways. While a person's permissioned financial data may be shared with a large language model, the circumstances under which that happens can vary widely. Many respondents focused on the risks associated with LLMs ingesting permissioned financial data, and a wide array of guardrails and transparency mechanisms were suggested.

The situation: A person's permissioned financial data is sent to a large language model. What are the role and responsibilities of that model—or the model provider—in the data sharing chain?

- Is the model a sub-processor working under an end user consent grant that already exists?
- Is the model a separate data recipient that needs its own, separate authorization?
- Is there transparency to others in the chain, including the consumer and the data provider, on which data is being sent to which model providers?
- How is this data being retained or used for training, and what contracts or user agreements exist to govern such use?

What we heard from respondents: Most respondents highlighted these as paramount questions for the industry to address.

Many respondents shared the view that responsibility for agent activity sits with the agent operator, not with the agent or the model, under the argument that an agent isn't a legal actor and can't carry duties.

Some respondents shared a view that new traceability mechanisms should be built to enable others in the data-sharing chain—including the data provider that issued the access token—to better understand when permissioned data is shared with large language models.

Solutions proposed to enable traceability included: provenance that travels with the data, signed receipts at each hop, expanding transparency around when LLMs are used as sub-processors, and frameworks to support scope that can narrow downstream but never widen. Others said the existing data controller-and-processor model already addresses scenarios where data is ingested by LLMs and can be leveraged by industry players—potentially with additional guidance—rather than trying to invent new frameworks or machinery to provide transparency. A few suggested making the chain cryptographically observable. Others provided suggestions for limitations on data retention, training use, or any other use beyond the one needed to serve the consumer's business purpose.

4. An AI agent goes beyond "read only" to perform "write" actions at a financial institution

The situation: A consumer directs an AI agent to make a payment, add a payee, or change a setting on their financial account. Or more broadly, the consumer directs an agent to accomplish a financial task or outcome, and the agent decides—either autonomously, or with human-in-the-loop approval— on specific actions to take with the consumer's financial accounts.

What we heard from respondents: Many respondents brought up these scenarios, and a large majority highlighted how "write actions" by agents were categorically different from a risk perspective than read-only data access.

Many industry members described significant barriers to financial institutions enabling agents to take actions on financial accounts, including liability assignment challenges, fraud risk, dispute resolution, and current regulatory frameworks.

Others proposed various suggestions for how the risks around agent-initiated or agent-executed actions could be managed. A few included: (i) risk classes for agent actions with named mandatory step-up triggers; (ii) write permissions granted and revoked separately from "read-only" permissions and never bundled together; (iii) a human confirmation requirement for anything irreversible, executed out of band so the agent can't approve its own action, and (iv) graduated autonomy tiers. Some focused on ways to reuse existing step-up mechanisms.

5. An Agentic AI tool seeks to integrate (e.g., with a bank) via an Agentic Tooling Interface (e.g., MCP)

The situation: Many AI platforms and tools are using new interface protocols that are purpose-built for AI agents (e.g., MCP; A2A) when integrating with data sources.

What we heard from respondents: Many parties across the Open Finance ecosystem—including Data Providers, Data Access Platforms, and Data Recipients—are now exploring whether and how they might build agentic tooling interfaces (e.g., MCP servers) to enable more seamless discovery and access by agents that call for permissioned financial data. Some have already established these new access methods. A range of companies we spoke with raised questions about how to ensure that data sharing

through these interfaces continues to reflect the features that FDX-based APIs have historically, including as pertains to data minimization, consumer control, security, and traceability.

FDX has a workstream underway presently to explore new or enhanced Agentic Tooling Interface standards or best practices that are tailored for the sharing of financial data.

6. An AI agent accesses an API but misinterprets the data retrieved

The situation: An AI agent retrieves data from an API. However, the data returned may not be interpreted correctly because the semantics are unclear. Whereas in the past a data retrieval script might be coded by an engineer who can call up their counterpart at the bank to understand what a given field (like "available balance") actually means or includes, an AI agent may infer the wrong answer.

What we heard from respondents: Several respondents share the perspective that data semantics—what financial data actually means—are increasingly important in the age of agentic AI. APIs provide an opportunity for clearer semantic accuracy than screen-scraping connections. Several respondents shared the view that API semantics need to be defined more clearly. Otherwise, misread data could harm consumers, degrade the quality of decisions made by consumer-facing tools, and raise provider servicing costs.

[FDX has a workstream underway to address semantic clarity in the FDX specification, to help address this potential problem for agents.]

Three cross-cutting themes

Across the various scenarios described above, a few cross-cutting themes emerged:

1. **Agent identity**. Many respondents highlighted the need for establishing a new way for distinguishing AI agents, human persons, and malicious bots. Several specific suggestions were made for ways to distinguish, bind, and provide broad visibility into the identities of: the agent, the operator responsible for it, and the person who delegated. Many respondents urged FDX to consider profiling emerging standards that address agent identity.
2. **Security:** Various respondents highlighted novel security risks that could arise as AI agents are deployed in the process of accessing financial data. Respondents encouraged FDX to consider ways to partner with other organizations—including OIDF and IETF—to describe and address these novel risks. Ideas shared included prompt injection defense mechanisms; agent-specific authorization scopes, and grant attenuation controls.
3. **Operator accountability**: Many respondents held to the position that an AI agents cannot take responsibility, and therefore someone else has to. Respondents suggested various approaches to make agent operators identifiable and provide for greater accountability. These included: registration of agent operators; mechanisms for capturing verified legal identity; incremental audit logging tooling; and frameworks to support incident response and consumer redress for agent operators.

FINANCIAL® DATA EXCHANGE

What happens next at FDX

The input gathered through this process is shaping the next phase of work at FDX.

Two workstreams now underway, as of September 2026, are:

- **Agent tooling interfaces**. FDX is developing an explainer on how agent-facing interfaces such as MCP servers relate to FDX-aligned APIs, so that members evaluating these interfaces can see where existing FDX controls apply and where questions remain open.
- **API semantics**. FDX is working to define more precisely what financial data elements mean within the FDX specification, so that an agent interpreting an API response reaches the same conclusion an experienced engineer would.

Agent identity, data recipient registration, consent, and security drew more comments in this process than any other topics. Further scoping on the relevant work for these issues is currently underway.

FDX will look to partner with or profile other emerging standards where it makes sense. Many respondents encouraged FDX to profile mature outside work, not to reinvent it. Respondents pointed at OAuth, OIDC, FAPI, FIDO and active IETF work, and said FDX should define only the financial-data layer on top.

Agent-tailored API profile: Many respondents packaged their recommendations the same way, encouraging FDX to establish a constrained, opt-in "agent-safe API profile" that an operator meeting a high bar could use. Support for that container was broad, though not unanimous.

Getting engaged

If you're not already engaged with FDX on this work, contact us at FDXSupport@financialdataexchange.org to learn how to get plugged in. Additionally, visit here to learn more about the upcoming FDX Summit happening March 13 to 15, 2027, where these topics will be explored further.