FINANCIAL® DATA EXCHANGE

Recap of FDX Webinar: Agentic AI and Open Finance | *July 16, 2026*

Speakers:

- Kevin Feltes (CEO, FDX) + Jean-Paul LaClair (SVP Product, FDX)
- Matt Setless (Head of Product–Open Banking, JPMorgan Chase)
- Tom Carpenter (SVP, Global Services Policy, Mastercard)
- Sean Estrada (Head of Standards Advocacy, Stripe)
- Nishant Kaushik (CTO, FIDO Alliance)
- Bill Fisher (Security Engineer, NIST NCCoE)
- Ankit Agarwal (CTO, Skyfire)
- Nate Soffio (Head of Reusable & Agentic Identity Products, Prove)

Executive Summary

AI is increasingly being adopted by companies who access or transmit user-permissioned financial data, with a growing share of consumers now using AI to manage their finances. Frontier AI labs, aggregators, fintechs, and banks are all now touching permissioned data using AI tools — creating signifiant opportunities but also new risks. FDX gathered input via an industry Call For Input (30 submissions, 500+ pages) and held dozens of meetings with banks, tech companies, and others. In a webinar drawing 700+ registrants, FDX shared key preliminary themes from their findings.

Why this topic matters by stakeholder

- **Banks/credit unions:** Need to manage new credential-based non-human access, retain ability to trace downstream data flows, retain trust and confidence in consumer consent flows as AI agents access data
- **Data access platforms:** Continued role fostering trust and facilitating recipient onboarding and traceability as both they and their clients (including banks / fintechs) adopt this technology
- **Data recipients (fintechs, FIs):** Building AI-driven experiences but facing new accountability expectations from others to deploy in a way that enables traceability, data minimization, appropriate consent, and trust
- **Everyone:** Open Finance is interdependent: one party's choices affect the customer experiences for others

Seven new problem spaces identified, introduced by AI:

1. Consent for fluid, purpose-shifting agents
2. Human-in-the-loop principles, and how traceable consent / intent is captured
3. Interparty risk management & security (prompt injection, agent logging)
4. Data storage, revocation, and model training use
5. Governed "front door" access vs. agent-powered scraping using credentials
6. Read vs. write/act capabilities
7. Agent identity and accountability

Areas where API / technology integrations (and standards) may need to evolve include: risk management practices; improved API semantics for AI readability, agent identity/KYA frameworks, bounded delegation mandates, amended data minimization principles, consumer disclosures / UX guidelines, and a “responsible agent operator” framework

Panelist themes: AI breaks the deterministic assumptions behind legacy API security; standards should govern what crosses trust boundaries; "don't go backwards" on existing consumer protections; industry engagement at FDX can help shape incentives and mutual understanding, which matter greatly alongside modernized technical standards.

Next steps: Industry members should stay tuned for FDX synthesis paper coming soon; engage with AI Task Force this fall to drive evolving standards; share your company’s priorities and insights to input@financialdataexchange.org

Detailed Recap

1. What's happening: AI adoption in user-permissioned financial data sharing

1. **AI tools are being rapidly adopted by companies and consumers to** access, sharing, and processing user-permissioned data… i.e., in the "Open Finance" ecosystem
    1. Multiple studies suggest **a large and growing share of Americans are now using AI in some way to manage their finances**, and over half **expect fintech apps to use AI**.
    2. On the supply side, many leading companies—including AI frontier model companies, data access platforms, fintechs, and financial institutions—are now using AI to consume, process, or transmit user-permissioned financial data.
    3. The upside is a opportunity to improve financial health and serve consumers; the tension is that **privacy, control, and security** remain the biggest drivers of consumer trust.
2. **These new uses of AI creates a host of new important opportunities, risks, and open questions** for companies involved in user-permissioned data sharing (including data holders, and also recipients of the data)
3. **FDX has collected extensive industry feedback on this topic in recent months,** including through an industry-wide "Call For Input" that yielded 30 written submissions (>500 pages) and meetings with dozens more companies.
4. This webinar provided a **read-out on key focus areas** for banks, data aggregators, consumer advocates, technology providers, and other thought leaders in Open Finance. **Over 700 people registered** for this event, signaling the broad interest in this topic; this was inclusive of banks, credit unions, data access platforms, fintechs, AI builders, and regulators.

2. Why this matters now

- **For Data Providers (e.g., banks & credit unions):** Many are focused on how to continue maintaining customer protection and trust. This includes identifying and responding to new forms of credential-based non-human access; retaining the ability to remediate problems that arise if an AI took an action or made a mistake; managing downstream risk and understanding where data goes; and supporting customers who want to engage through new channels. Because consumers often still look to their financial institution when something goes wrong, even if that failure happened in an agent or downstream tool, traceability and responsible risk management are key focus areas for many financial institutions.
- **For Data Access Platforms**: Many data access platforms see a (continued) critical role for themselves to intermediate trust, onboarding, and data flows that maintain the trust of consumers and counterparties as both they and a growing share of their clients increasingly use AI to access or process financial data.
- **For Data recipients (fintechs, banks, big tech):** Amid rapid innovation and strong competition for customer engagement, many are eager to use AI to access and process permissioned financial data to build better experiences. That said, doing so may change the risks they carry and their counterparties' expectations around traceability and accountability. Building experiences on top of industry best practices can help to bolster customer trust and adoption for the new tools and experiences they build.

- **For all: Open Finance is an interdependent ecosystem**, and your customers' safety / privacy / experience when using AI to share or process their financial data will depend on the choices of upstream or downstream parties… what they allow or deny or require, what experiences they build, how they manage risk, and how they govern the data. Understanding the actions others are taking (peers, counterparties, infrastructure players) in this period of change is therefore highly important.
- **Additionally: The emerging regulatory framework** around data access and AI may be shaped by industry practices and standards that emerge over the next couple of years.

---

3. Foundations | Understanding "agentic AI in Open Finance"

1. **An AI system is a layered stack** — user surface, a "harness" that orchestrates tools and multi-agent loops, a tool layer (data access platforms, scrapers, MCPs), and the underlying LLM — and *each layer may be owned by a different company*, which matters for assigning responsibility. "We use AI agents" tells a counterparty almost nothing about the actual risk.
2. **Permissioned financial data can be sent or stored at many points in that stack** — conversation logs, transaction databases, financial memory stores, and potentially for model training — in raw or derived form, complicating any simple answer to "who holds the data."
3. **"Agentic" is a spectrum, not a binary,** and FDX is using it as a lens to find real industry need rather than a gate to rule use cases in or out.
4. **Companies should be thinking surgically about the different *use cases*** for AI being used for processing or accessing permissioned data, as each brings unique risks and considerations. **Three notable examples**:
    1. **"AI" Data Recipient:** A full-stack chatbot (e.g., ChatGPT) run by one company acting as the data recipient and pulling data through a governed API;
    2. **AI model as sub-processor**: A fintech/bank app sending raw or derived data to an external AI model acting as a sub-processor;
    3. **Scraping agent**: A personal "scraping"/browser agent logging into a bank site with the user's credentials. (Here, FDX noted that some major model providers already embed policies that refuse credential-based bank scraping. Ask certain popular tools to log into your bank and grab transactions, and they decline. That said, open-source tools can still do it, and banks report this traffic is rising and harder to detect than typical cloud-scale scraping.

---

4. Where technology implementations and standards may need to evolve

There was a broad convergence among companies who shared feedback with FDX on the key problems that need solving to enable firms to safely enable user-permissioned data sharing in the age of AI.

7 Problem Spaces were highlighted in the webinar, for industry to consider

1. **Consent for fluid apps / agents**
    - Past: Apps often anchored to fixed purpose. Now: AI tools may receive a broad "intent mandate"; have emergent / evolving purpose; hyper-personalized
    - Key questions: How's consent scope set, changed? How to prevent over-asking or hallucination of purpose? Ensuring purpose is clear to user?
2. **Human in the loop + traceable consent**
    - Past: Explicit consent (scope, duration, accts.) by consumer. Now: Consumer may delegate authority to Agent to complete consent (upfront, or change)
    - Key questions: When should a human be in the loop? How to verifiably record and/or transmit the consent or mandate a consumer gave the agent?
3. **Interparty risk management**
    - Past: Many banks, aggregators have sought traceability + accountability for others in chain
    - Now / key questions: Does AI adoption raise novel risks to manage in new ways? E.g., Prompt injection defenses; Agent-behavior logging; Derived data retention/use policy; AI model sub-processor relationships
4. **Storage, revocation, and model training**
    - Key questions: When raw or derived permissioned financial data is fed into model… Is the consumer aware? What data could be used for model training, and what are the risks? How and where to minimize raw data being fed to model? When access is revoked, can data be "deleted" in same way? Are counterparties aware and able to manage appropriate risks and disclosures?
5. **The "front door" vs. screen-scraping**
    - Past: Harder to maintain "screen-scraping" scripts and evade blocking at many banks
    - Now / key question: Making the governed path work better—for consumers and for implementers—than ungoverned credential-based access can prevent problems associated with scraping
6. **Read versus write / act**
    - Past: Most data providers (North America) have not built "write" capabilities for third parties to utilize on behalf of consumers [outside bespoke strategic partnerships]
    - Now: If AI agents increasingly attempt write actions, does this up the pressure on establishing standardized write access capabilities (and associated technical standards)
7. **Agent identity**
    - **Is incremental Agent Identity or role signaling needed to enable traceability, remediation, and trust?**
    - **PERSPECTIVE 1:** New signaling would be helpful to trace which agent acted or who the agent operator was to aid remediation or assigning accountability
    - **PERSPECTIVE 2:** Existing frameworks already solve for this well (contracts for API channel access; recipient registration / onboarding, and risk management practices) … including for agent operators. Agents have no independent standing

Solutions proposed, where API integrations and technical standards may warrant a change; these will be topics of further industry discussion at FDX in the weeks ahead. These included:

1. **Improved semantics** in APIs, for AI readability
2. **Refining what industry defines / treats as a "data recipient",** and other taxonomy
3. **Agent identity / KYA-to-KYC** – at registration or runtime; **requestor-type signaling** (e.g., agent vs. human)
4. **Bounded agent delegation** framework (e.g., purpose, "intent mandate", time)
5. **Principles** for data persistence, minimization, masking when data travels to models
6. **"Third-hop" provenance** / transparency metadata
7. **Consumer disclosures / UX guidelines** on agent role and AI model data flows
8. **"Responsible Agent Operator" framework**; expand risk mgmt. to agent-specific risks
9. **Security**: Step-up authentication, prompt injection
10. **Finer-grained access controls** — augmenting coarse OAuth scopes

---

5. Why and how to engage in the Standards conversation at FDX

1. **What's coming next:**
   - **Read the FDX synthesis paper, coming in the next few weeks,** which captures everything heard to date.
   - **FDX members:** join the standards work this fall (the AI Task Force is where input becomes prioritization and standard-setting).
   - **Everyone — member or not:** connect and share your thinking at input@financialdataexchange.org. Membership is not required to participate.
   - **Watch for future events**, including the annual FDX Summit.
2. **FDX's role will be specific and complementary:** profile other standards where it makes sense (e.g., for agent identity), while focusing on the **financial-data-specific layer** — consent and delegation, scoped data access, traceability, and downstream accountability. A concrete artifact on the table is a possible **"agent-safe FDX API profile."**
3. **Why your voice changes the outcome:** many of the debates on how to evolve standards are yet to be had… including regarding identity, consent, disclosures, security methods, write access, and other key topics.
4. **The role of standards.** Standards can deliver value across three dimensions: a shared mental model and taxonomy so parties describe similar activities with similar terms; reusable building blocks that independent companies can adopt over time (not necessarily all at once); and consensus best practices that let firms deploy with confidence, ideally flowing through to greater consumer trust.

6. Insights from panelists

How AI changes the equation for data sharing

1. **Agentic AI challenges the architectural framework many systems were built on.** Kaushik, CTO of the FIDO Alliance, notes that decades of security relied on predictability and deterministic controls — an API client granted a credential with a specific scope, and the provider honors that scope. AI agents may challenge that assumption.

2. **Separate intent, consent, and authorization — and structure the moment they convert.** Agarwal stated that intent-based mandates by consumers can be "dubious," because over a long agent conversation the semantics get fuzzy and the consumer's understanding may not match the recipient's. Soffio (Prove) reframes it as a progression: intent graduates to consent, and consent and authorization are "bedfellows." Open banking already shows what structured consent looks like — specific purposes, data fields, permissible use, duration, and revocation. The opportunity is to **structure the inflection point** where open-ended browsing/chatting flips into a deliberate, authorized action — so all parties in the chain stay safe without "tiptoeing around semantics." Practical implication: precise definitions and clear authorization UX aren't bureaucracy; they're what makes disputes resolvable.

3. **Need for shared understanding of an agent.** Nate Soffio (Prove) provided three metaphors:
    - **Shipping containers:** an agent is a standardized "box" of claims, authorization, permissions, and instructions. Be agnostic to who publishes the box (a merchant, a large lab's runtime) as long as you can verify its provenance when you need to.
    - **Three-legged races:** the consumer and the agent are bound together for a specific purpose; both must be trustable, and your verification/instructions have to be bound to the agent for that task.
    - **Permission slips:** something cryptographically signed travels in the box and can be checked when the agent acts.

Principles for ensuring new standards and technology governance mechanisms get adopted

1. **Work backward from the data holder.** Agarwal (Skyfire): start from what the company receiving the data access request (e.g., bank / data provider) needs to feel comfortable before acting on the request/instruction

2. **The role of standards.** Soffio's rule of thumb: a good standard should govern anything that crosses a trust boundary — credential formats, delegation objects, signature schemes, registry interfaces. Everything else — key custody, industry-specific risk engines, UX — should compete on the open market. Saying "we embrace these standards" makes emerging tech "less spooky to buyers."

3. **Reuse; don't fully rebuild.** Agarwal's statement to data holders: use the tech you've already built. A good standard identifies which protocols make agentic access safe (verified human principal, verified agent, clearly transmitted authorization) and lets the industry evolve as better protocols emerge — creating a learning flywheel.

4. **Fragmentation would increase costs for all.** Estrada warned that per-holder "trusted agent registries" could be problematic in practice and at scale, leading to a degraded consumer experience; Bill Fisher and Kaushik echoed that fragmentation slows everyone down.

5. **There is not going to be one architectural solution**, argued one panelist. There are already some services uplifting themselves to support AI. MCP brings a different architecture and different solutions. But what happens when services aren't ready to build that infrastructure — which could take significant time and cost? Companies are urgent to act, but they are also looking for ways to implement with what they have.

6. **“Don't go backwards"** is the first principle (per Carpenter): preserve a decade of gains in consumer control, security, transparency, and fair competition, and amend existing standards before building net-new ones.

7. **Getting the incentives right:** NIST's Fisher argued the solutions will be not just technical-standards-based but also market- and incentive-driven, and that as the financial sector faces common problems, sharing information and getting together to promote the "happy path" while standards evolve will be important. FIDO’s Kaushik returned to incentives as "critical," noting FIDO focuses not just on defining standards but on enablement and usage.

Navigating a rapidly changing ecosystem

8. **A panelist from a large bank raised the concept of "DARM" (downstream agent risk management)** as the agent-era complement to downstream data-recipient risk management. At the same time, downstream visibility may be one of the ***harder problems*** to solve as data now traverses multiple hops and orchestrators.

9. **Standards + policy landscape (Estrada).** He noted that legislators and regulators are actively learning alongside the technical community while standards are being built — "the right approach" — and argued the industry should let regulators reference open, interoperable industry standards rather than the two tracks moving independently.

10. **A warning against fear (Estrada).** His ask to the community: don't act from a **fear-based** posture reacting to headlines and risk discussions; instead, the industry succeeds by engaging across standards bodies and creating open, interoperable infrastructure that can be *referenced* by regulators — keeping the conversation going and sharing as much as possible.