FINANCIAL DATA EXCHANGE

AGENTIC AI & OPEN FINANCE

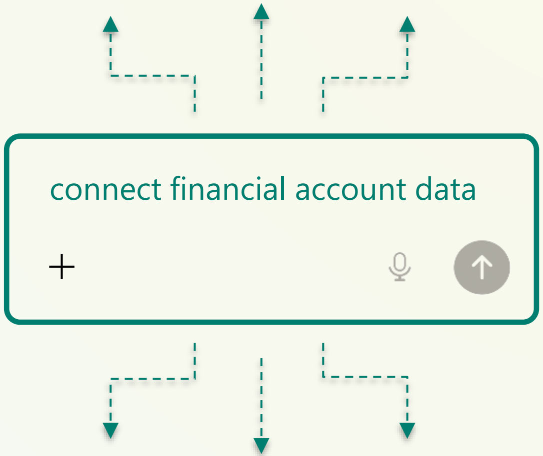

Focus

AI

AI + FINANCIAL SERVICES

AI + OPEN FINANCE

AI + OPEN FINANCE + TECHNICAL STANDARDS

Consent
Security
User experience
Access
Traceability
Onboarding & registration
Risk management
Semantics

Why this is important now?

1.

Adoption of AI
by companies and individuals in user-permissioned data sharing is accelerating

Why this is important now?

2.
That introduces new opportunities, risks, and questions

Consumers:
control, privacy, security?

Human in the loop?

Agent identity, traceability and trust

Intra-party risk management?

The “front door” vs. screen-scraping?

Why this is important now?

3.

Serving customers well requires collaborating across an interdependent ecosystem

CONSUMERS

FINTECHS / BIG TECH

BANKS, CREDIT UNIONS

DATA ACCESS PLATFORMS

POLICYMAKERS

AI BUILDERS

Industry standards can help

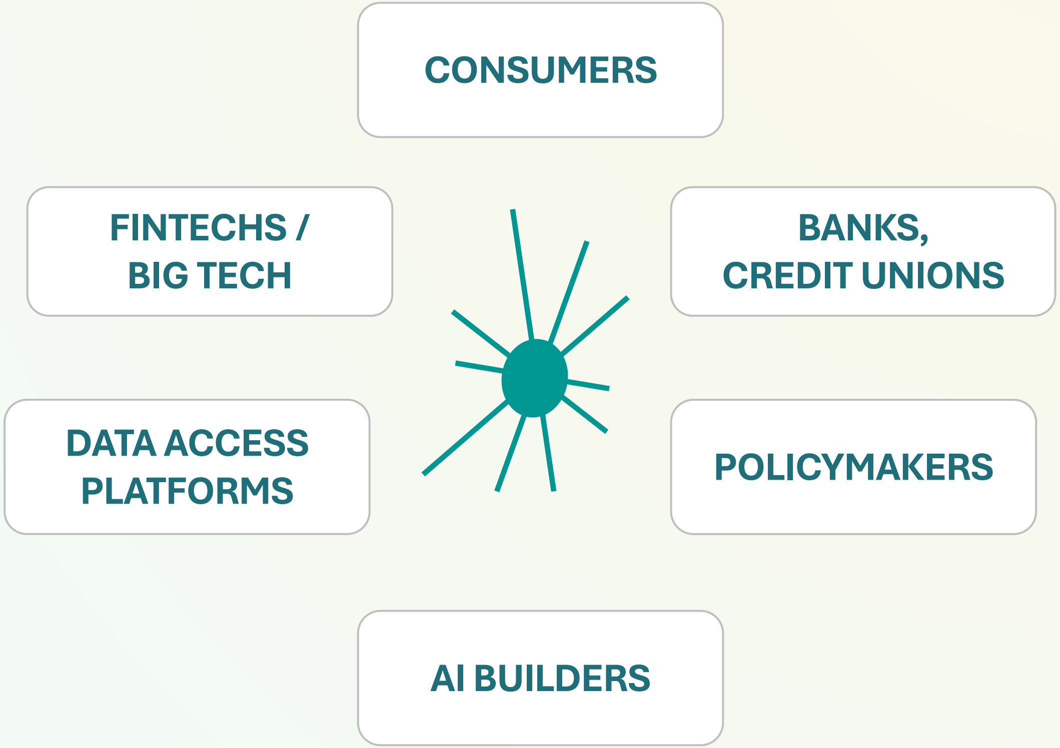

4 foundational points
that the FDX Call for Input surfaced

1. An AI tool or product can involve multiple layers ... from multiple companies

EXAMPLES [illustrative]

Layer	Examples
Product / channel / UI layer The app you see. Where prompts, interaction, settings live.	**"AI" app:** ChatGPT, Perplexity, Claude Code **Bank app:** SoFi, Wells Fargo, North Island CU, ... **"Fintech" app:** Monarch, Quickbooks, Origin, ...
Tool layer Takes action / accesses data; write to database	**API:** [Data aggregators...] **Scraping:** Playwright, Puppeteer, Browserbase **MCP to API:** Era, Truthifi
Harness / orchestration layer governs how model interacts with tools, multi-step loops	**From "AI" co.:** Claude or OpenAI Agent SDK **Open-source**: Langchain, CrewAI **Commercial**: Relevance, Lindy, Stack AI
Model layer (LLM) Predicts the next output. No memory, cannot "act" alone	**Frontier**: GPT, Gemini, Opus, Grok, ... **Open source**: Llama **Internal**: custom bank / aggregator / fintech LLM

SIMPLIFIED MODEL *not comprehensive, ranked, or endorsements*

2. Permissioned financial data can be stored at multiple layers

Product / channel / UI layer

The app you see. Where prompts, interaction, settings live.

Tool layer

Takes action / accesses data; write to database

Harness / orchestration layer

governs how model interacts with tools, multi-step loops

Model layer (LLM)

Predicts the next output. No memory, cannot "act" alone

Where raw or derived financial data *may* be stored

Conversation logs

Transactions databases

Financial memory store

... used for model training

3. "Agentic" is a spectrum

LESS AGENTIC ⟷	MORE AGENTIC
Answer question / prompt	Take action
Single step	Multi-step
Human-in-the-loop always	More autonomous
Explicit instructions	Goal-directed execution
No tool use	Extensive tool use
No memory	Persistent memory

4. Clarifying the *use case* can help aligning on solutions

Use case 1: Chatbot accessing my financial data via API (directly, or indirectly)

Product / Channel	ChatGPT.com
Tools	ChatGPT "Connectors"
Harness	OpenAI policies / orchestration code
Model	GPT-5.5 model (OpenAI)

Data Provider

Data Access Platform (Aggregator)

SIMPLIFIED MODEL

Use case 2: Fintech using permissioned data, with AI as sub-processor

Product / Channel	Fintech mobile app
Tools	Orchestration code to pass data to the AI model, enforce safety, etc.
Harness	
Model	Sonnet / Opus (Anthropic)

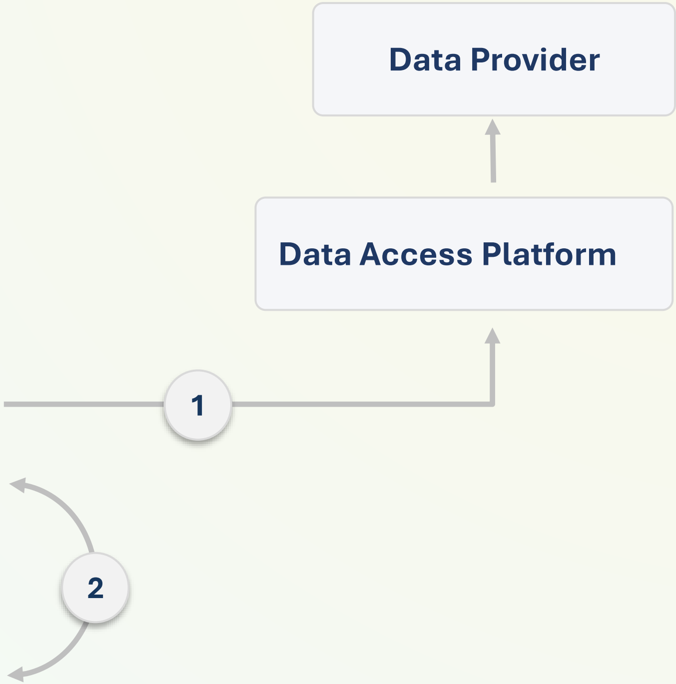

SIMPLIFIED MODEL

Use case 3: Scraping Agent

Product / Channel	OpenClaw + WhatsApp
Tools	Puppeteer or Playwright
Harness	OpenClaw
Model	Open-source model, or from OpenAI, Google, DeepSeek, …

Data Provider

What's happening today:

- Many major models have **policies to inhibit** this use case
- Possible with open-source tools
- Some banks report **seeing this traffic**, but differentiating it from "real humans" is harder

SIMPLIFIED MODEL

Call for Input: Agentic AI & Open Finance

FDX Members: See repository of submissions on Confluence

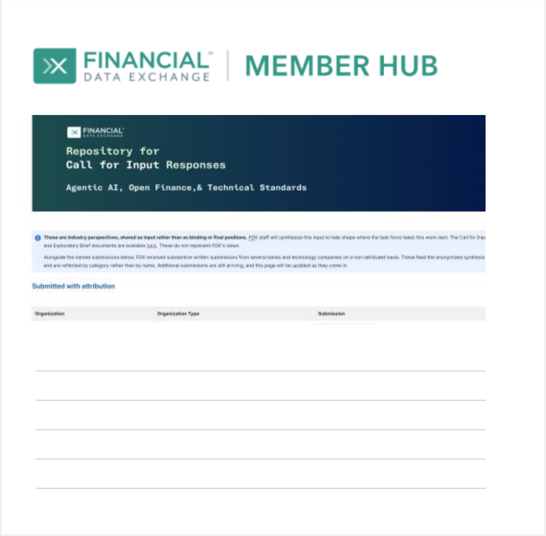

What we heard

CONVERGE

on the problems

DIVERGE

on the solutions

Some of the key problems raised by the industry...

1. Consent for fluid apps / agents
2. Human in the loop + traceable consent
3. Intra-party risk management
4. Storage, revocation, and model training
5. The "front door" vs. screen-scraping
6. Read versus write / act
7. AI / agent identity, traceability and trust

Use case 1: AI chatbot (e.g., ChatGPT) is the app, accessing via API

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

- 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)

- 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. Intra-party risk management

Past: Many banks, aggregators have sought traceability + accountability for others in chain

Now: 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

Use case 2: AI model as sub-processor behind an (e.g., fintech/bank) app

4. Storage, revocation, and model training

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?

FINANCIAL DATA EXCHANGE

Use case 3: Personal agent scraping a (e.g., bank) website

5. The “front door” vs. screen-scraping

Past: Harder to maintain “screen-scraping” scripts and evade blocking at many banks

Now: 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)

A topic with diverging views: Is incremental Agent Identity or role signaling needed to enable traceability, remediation, and trust?

PERSPECTIVE 1

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

PERSPECTIVE 2

New signaling would be helpful

to trace which agent acted or who the agent operator was

.. to aid remediation or assigning accountability

A peek at some of the **Solutions** proposed for standards

1. **Improved semantics** in spec for AI readability
2. **Refining taxonomy:** e.g., “data recipient”
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

What’s next from FDX

Where we are: the phases ahead

Phase 1 · DONE	Problem-space exploration; Call for Input; Advisory Group
Phase 2 · NOW	**Synthesis, scope and lanes, education**
Phase 3 · August-Sept	Prioritization and a possible consolidated working draft
Phase 4 · Oct onward	Formal standards and RFCs, through FDX consensus process

We are here

Agent-safe FDX API profile?

Prospective FDX focus areas, and where we may profile / partner

Agent communication and transport	MCP · A2A · AGENTS.md
Agent identity and authentication	OIDF · NIST · FIDO · IETF
Consent and delegation	**FDX**
Scoped data access and retrieval	**FDX APIs · FAPI**
Traceability and downstream accountability	**FDX**
Financial institution / data source	Banks · credit unions · brokerages

FDX's lane historically

The strong signal from the input: profile mature external work and define only the financial data sharing part.

We'll prioritize high-impact use cases without losing sight of others

AI CAN PLAY MANY ROLES IN OPEN FINANCE

Who builds/operates

model co, fintech, bank, device maker, vendor, aggregator

Via which channel

governed API, or "scraping" (credential based)

On whose behalf

consumer, data recipient, platform, multiple

Where in the chain

Direct vs. downstream recipient, sub-processor, ...

What action taken

suggest/decide consent scope, give consent, authenticate, retrieve data, process data, use data

Who touches data

product operator, model provider, vendor, ...

How autonomous

multi-step; human-in-the-loop; goal-directed; tool use; persistent memory

Aggregator in loop

if not, functions (e.g., data min., flow-downs) may need solving in other ways

How to stay involved

1. Read the FDX synthesis paper coming soon
2. For FDX Members: Standards work this fall
3. For all: Connect with us & share your thoughts
4. Watch for future events on this, including FDX Summit

input @ financialdataexchange.org

About FDX

OPEN MEMBERSHIP | ~200 ORGS TODAY

LEADING STANDARDS BODY FOR USER-PERMISSIONED FINANCIAL DATA SHARING