

Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report

Response to Consultation

Tokenization Systems

1. Do you agree with the benefits and risks of AI adoption by financial institutions described in this report? Are there any substantive benefits or risks not covered?

Tokenization Systems agrees with the benefits and risks described in Section 1, and in particular finds the agentic-AI risks subsection well targeted. The report observes that for autonomous agent actions, "Overriding, redressing, or remediating these actions can be difficult or impossible for humans." In the measured payments context, that is not a contingency but an architectural property. The dominant agentic-payment settlement scheme provides no chargeback primitive, no holding period, and no protocol-layer dispute mechanism; finality is the settlement network's finality, and refunds, where they occur, are voluntary off-protocol actions.

Two substantive risk areas are not covered.

First, cross-operator composition. The report's risk framing is institution-internal: a financial institution adopts AI, governs it, and manages third-party dependencies through service relationships. The glossary makes this scope explicit, limiting "supply chain" to "the services under a third-party service relationship." A live agentic payment, however, is delivered by a composition chain of five operator classes (an agent runtime, an authorization or checkout protocol, a settlement facilitator, a token issuer, and a settlement network), and the layer where the measured market concentrates, the facilitator, typically has no contractual relationship with any financial institution. Accountability in such a chain degrades to the weakest operator in it: every operator can be individually defensible while the composed product delivers an irreversible payment, on the instruction of a non-human agent, through an unlicensed intermediary, to a counterparty that cannot be identified from public records. Institution-internal governance, however sound, does not reach this structure. Leading AI-governance frameworks are organized around an institution managing its own AI and treating third-party providers as supply-chain risks to that institution (National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," January 2023), so the binding accountability gap sits at a layer they do not reach (it is not the model operator); payments regulation does not reach it either (it is not a licensed money transmitter or bank). The Board has itself recognized, for crypto-asset markets, that providers combining multiple functions warrant comprehensive regulation addressing both the individual functions and the risks arising from their combination, including conflicts-of-

interest and separation requirements (Financial Stability Board, "High-level Recommendations for the Regulation, Supervision and Oversight of Crypto-asset Activities and Markets," July 2023, Recommendation 9); the agentic-payment facilitator is such a function-combining operator, and it currently sits outside that perimeter.

Second, the integrity of AI-adoption statistics in agent-driven markets. Agents make transaction volume arbitrarily cheap to manufacture, so unweighted activity counts select for whichever operator has the strongest incentive to manufacture them. In the measured market, a public subsidy program correlates with a rise in the dominant count pipeline's monthly output, from roughly 3.0 million payments in March 2026, when the rebate activated, to 5.04 million in May, produced by a fleet of 35 buyer wallets. Separately, widely cited dashboard figures diverge from corrected measurement by more than threefold for attribution reasons alone (a published May 2026 community figure of \$179,793 against a corrected \$652,882 on the same chain, with the gap explained by address-set coverage). This matters directly for the FSB's monitoring agenda: adoption indicators built on raw activity counts in agentic markets will systematically mismeasure adoption. The same public settlement records that make activity-count metrics manipulable also make the market independently measurable, which is a benefit the report could note; the risk and the benefit are two sides of the same data layer.

2. Are the sound practices sufficiently comprehensive and clear to enable financial institutions' responsible AI adoption?

The twelve sound practices are clear, well structured, and comprehensive for institution-internal AI adoption. Tokenization Systems proposes three additions, each compatible with the report's menu approach and proportionality principle, and each grounded in the measured evidence above.

Addition 1: composition-chain documentation (extends Sound Practices 3 and 12). Sound Practice 3's documentation inventory already includes "dependencies on other AI use cases, models, systems, or third-party AI providers." For AI use cases with authority to execute transactions, the practices could ask institutions to document the full operator chain through settlement, explicitly including operators with which the institution has no direct contractual relationship, and to record for each layer who holds the authority to halt, reverse, or remediate a transaction, on what timeline, and under what commitment. The measured market documents why the contractual scope is insufficient: the layer through which the majority of agentic-payment value routes is outside every service relationship the current "supply chain" definition reaches.

Addition 2: a deployment-decision evidence package for execution authority (extends Sound Practices 5 and 10). Sound Practice 10 already contemplates "placing controls on AI agents executing financial transactions especially with customers' funds." The practices could go one step further and name a minimum evidence package an institution assembles before any AI use case receives execution authority (authority to move funds or bind the institution): evidence that deployed capability matches capability claims; a system of record sufficient to reconstruct each transaction's authorization context (the instruction, the mandate, the spend limits), not merely its settlement record; explicit documentation of finality and reversal authority at each layer of the chain; the discontinuation posture of each operator depended upon; pre-agreed incident-handling triggers and communication timelines; and at least one

viable exit path per critical dependency, or the documented absence of one. Tokenization Systems specifies an institutional diligence structure of this kind in its Minimum Viable Equivalence Pack framework for tokenized products (<https://tokenization.systems/papers/mvep/>), which it is extending to AI infrastructure. The measured market illustrates the stakes of the record item specifically: the settlement record that exists on public blockchains is the record no party chose to keep, while the authorization-context record that disputes would require is the record no party is currently obligated to keep.

Addition 3: system-level monitoring of AI-transaction dependencies (extends Sound Practice 12). Sound Practice 12 names "supply chain and concentration risks," and its body asks institutions to monitor concentration "at the individual financial institution (including group level)." Institution-level registers cannot observe system-level concentration: no individual institution's third-party inventory reveals that the settlement layer it depends on carries an HHI of 6,009, or that the count-dominant and value-dominant operators differ. Where AI-transaction infrastructure settles on public records, system-level concentration is directly measurable today, by institutions, by industry bodies, or by authorities. The practices could acknowledge measured system-level concentration of shared AI-payment dependencies as a complement to institution-level registers, together with continuous monitoring of each dependency's regulatory posture (licensing status, evaluation transparency, intervention capability, compliance infrastructure, and geographic restrictions), which provides the ongoing companion to point-in-time onboarding diligence. Tokenization Systems publishes a continuous index of this kind in "Routing the Dollar" (<https://tokenization.systems/papers/routing-the-dollar/>) and is extending the instrument to AI operators.

3. Do the sound practices strike an appropriate balance between managing risks relating to all forms of AI, and addressing some of the risks relating to emerging and new complex forms of AI, such as GenAI and agentic AI?

The balance is appropriate for a menu of practices, and the report's introductory note that certain practices "can be more relevant for specific AI technologies, such as in response to risks stemming from the autonomy of AI agents" is the right device. The measured evidence supports two calibrations.

First, agentic AI in payments is present-tense, not emerging. The measured market settled 8.87 million agentic payments in May 2026 on public rails, with both qualifications stated plainly: the dollar layer is pre-material (0.0000376 percent of USDC transfer value on the same rails), and the count layer is dominated by subsidized pipelines. Companies are forming around this infrastructure at 5.8 times the prior-year rate (81 agentic-payments companies founded in full-year 2025 against 466 founded in 2026 through May, 96 percent at pre-seed stage, per Verda Ventures data). The practices will meet this market mid-formation, which argues for treating the agentic-AI provisions as operational text rather than forward-looking annex material.

Second, Sound Practice 10's autonomy factor would benefit from distinguishing advisory autonomy from execution authority. An agent that recommends and an agent that irreversibly transfers value differ in kind, not degree: the former's errors are reviewable, the latter's are final at broadcast. The report's own controls bullet (human approval thresholds,

"restricting direct agent access to payment systems") implicitly recognizes the distinction; the measured market is precisely direct agent access to payment systems, already operating outside the institutions the bullet addresses. Making the distinction explicit in Practice 10 (and in the glossary, per Question 7) would let proportionality scale with the property that actually changes the risk class.

4. Are the sound practices sufficiently flexible to accommodate and address newer types of AI and responsible adoption over time?

The practices are appropriately technology-neutral, and the proportionality and reassessment provisions (materiality reassessment at scale-up under Sound Practice 5) provide genuine flexibility. One structural addition would improve durability: attach record-keeping expectations to the evidence itself rather than to the current observability of any particular system design.

The measured market illustrates the hazard. Agentic payments are measurable today at per-payment granularity from public records; that is the only reason the measurements in this response exist. The protocol roadmap already includes a batch-settlement scheme (pre-funded escrow, off-chain vouchers, bulk claims) that removes per-payment public observability where adopted. The window in which independent public-record measurement of this market is possible at current fidelity may be temporary. Practices written against today's observability would age; practices that name the durable expectation (preservation of transaction-level records, including the authorization context described under Question 2, wherever AI executes transactions) remain binding across settlement designs. Sound Practice 10's "maintaining audit trails of agent transactions" bullet is the natural anchor; the addition is to make clear the expectation survives changes in settlement architecture that would otherwise extinguish the record.

5. Do the case studies in this report sufficiently highlight how different types of financial institutions can benefit from responsible AI adoption? Are there additional case studies for inclusion in the report? If so, please provide such case studies, particularly for nonbanks.

The case studies highlight institution-internal benefits well, across banks, insurers, and systemically important institutions. They do not yet illustrate the benefit case for AI capability that sits outside any one institution's governance, which is where a growing share of financial-sector AI value will be delivered and where the report's institution-internal frame does not reach. Tokenization Systems offers the following nonbank case study for inclusion, drawn from its June 2026 measurement memo, with the full memo, dataset, and reproduction queries available as supplementary material.

Case study offered: settlement infrastructure for agentic payments. Agentic payments on the x402 protocol settle through facilitators: nonbank services that receive an AI agent's signed payment authorization, broadcast the settlement transaction on a public blockchain, and pay the network fee. Facilitators decide what settles, when, on which chain, and at what fee, yet no licensing, registration, capital, or operational requirement attaches to the function. The full measurement memo is available at <https://tokenization.systems/papers/agentic-payments-control-layer/>.

First measurement of the layer (May 2026, three chains):

- concentration is above the 2,500 highly-concentrated threshold on every reading (count HHI 6,009 on Base, 4,203 on Solana; corrected volume HHI 4,040 and 4,915; single-operator service on Polygon);
- the count-dominant and value-dominant operators differ on two separate chains;
- the largest count contributor is a chain operator's official facilitator under a public gas-rebate program, producing 5.04 million monthly payments from 35 buyer wallets;
- the majority of corrected dollar volume routes through one proxy-settlement facilitator with 17 active buyers and a directly observable 0.99 percent facilitation fee, against roughly 14 basis points on an on-chain settlement-router path and a posted \$0.001 per transaction elsewhere;
- settlement is irreversible at the protocol layer; and
- none of the top 25 sellers by dollar volume resolves to a named operator from public records.

The structural lesson for the report: every sound practice in Sections 2 and 3 presumes an institution that can see and govern its AI deployment, while the live agentic-payments market routes through a layer no institution governs, most participants cannot identify, and no supervisory perimeter currently reaches. The case study gives Question 5's nonbank request concrete content: the relevant nonbanks are not only nonbank financial institutions adopting AI, but nonbank AI-payment infrastructure operators on whom financial institutions and their customers will increasingly depend.

6. Do the case studies in this report provide actionable insights for financial institutions in their responsible AI adoption?

The case studies are actionable for institution-internal adoption, and the strongest among them share a feature worth generalizing: they name the control point. The fraud-detection case study, for example, specifies the human-in-the-loop checkpoint (analyst review and approval of every agent-proposed detection rule before implementation), which tells a reader exactly where the institution's authority binds. Tokenization Systems suggests each case study identify (i) the operator chain involved in the use case and (ii) the layer at which each cited sound practice operates. For composed and agentic use cases this discipline does additional work: it makes explicit which control points sit inside the institution and which sit with outside operators, which is the information a board needs to apply Sound Practices 1 and 5 honestly. A case study that surfaces an uncontrolled layer is more actionable, not less, than one that does not.

7. Are the definitions in the glossary clear and aligned with industry sound practices, including recent developments in AI?

The definitions are generally clear and well aligned with industry sound practices, including recent developments in AI. Two refinements are suggested, both consequences of the measured evidence.

First, distinguish execution authority within the agentic-AI definition. The current definition, "Systems designed to autonomously perform complex and extended tasks, often making decisions and taking actions with limited human oversight," accurately describes autonomy but does not distinguish advisory or analytical autonomy from execution authority: the authority to take actions that are final for third parties, such as transferring value or binding an institution. The report's own risk discussion and controls bullets treat that property as the operative one, and the measured market documents why: at the settlement layer the difference between recommending a payment and executing one is the difference between a reviewable output and an irreversible transfer. A two-part definition (agentic AI; agentic AI with execution authority) would let the sound practices and proportionality language reference the sharper category directly.

Second, reconsider the scope limitation in "supply chain," or add a companion term. The glossary limits supply chain to "the services under a third-party service relationship." As measured, the operative dependency chain for an agentic payment includes operators with no service relationship to the financial institution or to each other (an agent runtime, an authorization protocol, a facilitator, a token issuer, a settlement network). A companion term such as "composition chain," defined as the sequence of operators that jointly deliver an AI-enabled outcome, including operators outside any direct contractual relationship, would give Sound Practices 3 and 12 vocabulary that reaches the structure where, on present evidence, the binding accountability gaps sit.

8. Are there any comments you would like to make on this report? Please fill in.

Tokenization Systems submits this response to add a measured, present-tense case study to the FSB's responsible-AI work: live agentic payments. The submission's value is empirical and operational. It shows where autonomous financial actions already settle, who controls the settlement layer, which activity measures misstate adoption, and which control points sit outside institution-internal AI governance.

In June 2026 Tokenization Systems completed what it understands to be the first control-layer measurement of live agentic payments: facilitator-level concentration, organic-volume decomposition, settlement-property analysis, and composition-chain accountability mapping for the x402 agentic-payment protocol across the three blockchains that carry materially all of its activity (Base, Polygon, Solana). Agentic payments are the case in which the autonomy the report addresses is already executing irreversible financial transactions at scale, largely outside supervised financial institutions. The measurements are published in "The Control Layer of Agentic Payments" (<https://tokenization.systems/papers/agentic-payments-control-layer/>); every figure cited in this response reproduces from committed queries and scripts, and Tokenization Systems offers the full memo, dataset, and reproduction code as supplementary material.

Five measured facts frame the responses below.

1. The settlement layer of agentic payments is highly concentrated on every reading. May 2026 facilitator concentration measures at a Herfindahl-Hirschman Index (HHI) of 6,009 by payment count and 4,040 by corrected dollar volume on Base, 4,203 and 4,915 on Solana, and effective single-operator service on Polygon. Every reading exceeds the 2,500 threshold that competition practice treats as highly concentrated.

2. The count layer and the value layer have different dominant operators on two separate chains. On Base, one facilitator carries 75.6 percent of May payment counts while another carries 57.2 percent of corrected May dollars; on Solana the same flip recurs with different operators. Single-metric claims about this market are structurally fragile.

3. Headline transaction counts are dominated by subsidized, machine-generated pipelines. The largest single contributor to cross-chain counts is a chain operator's official facilitator, operating under a public gas-rebate program, which produced 5.04 million May payments (56.8 percent of cross-chain counts) from 35 buyer wallets. Buyer cardinality separates this regime from organic usage without any wash-trading heuristic: 26.7 payments per distinct buyer on Base versus roughly 144,000 on Polygon.

4. Measured at the value layer, the market is currently small. Corrected cross-chain May volume is \$940,323 across 8.87 million payments: 2.83 percent of all USDC transfer events on the three chains, and 0.0000376 percent of USDC transfer value.

5. No layer of the composition chain that delivers an agentic payment currently holds the authority to reverse one, and the sell-side endpoints are pseudonymous even to full-data analysis. Settlement is push-payment and irreversible at the protocol layer; the facilitator function carries no licensing, registration, capital, or operational requirements; and none of the top 25 May sellers on Base resolves to a named operator from public surfaces.

No impropriety is implied by any figure in this response. The measured subsidy program is public, the proxy-settlement architecture discussed below is documented by its operator, and the measurement claims are about market structure, not conduct.

By background, Tokenization Systems is an independent research program on stablecoin, tokenization, and digital-asset market structure. It has published eight papers on SSRN, including a fifty-two-protocol governance-concentration audit accepted at a peer-reviewed journal, and has filed thirteen comment letters with six United States federal financial agencies, carrying sixty-seven recommendations in total. This is Tokenization Systems' first submission to an international standard-setting body.

Requested edits to the report:

1. Composition-chain documentation: ask institutions to document the full operator chain for AI use cases with execution authority, including non-contractual settlement layers and each layer's halt, reversal, and remediation authority. Anchor: Sound Practices 3 and 12. Reason: the measured settlement facilitator layer sits outside ordinary third-party service relationships.

2. Execution-authority evidence package: name a minimum evidence package before AI systems receive authority to move funds or bind the institution. Anchor: Sound Practices 5 and 10. Reason: reconstructible authorization-context records are needed before irreversible transactions occur.

3. System-level dependency monitoring: acknowledge measured system-level concentration of shared AI-payment dependencies as a complement to institution-level

registers. Anchor: Sound Practice 12. Reason: individual third-party inventories cannot reveal market-wide facilitator concentration.

4. Agentic-AI glossary: distinguish agentic AI generally from agentic AI with execution authority. Anchor: Glossary and Sound Practice 10. Reason: advisory autonomy and irreversible transaction authority differ in kind.

5. Supply-chain glossary: reconsider the service-relationship limit or add "composition chain" as a companion term. Anchor: Glossary, Sound Practices 3 and 12. Reason: agentic payments are delivered by operators that may lack direct contractual relationships with each other or with a financial institution.

These edits do not require the FSB to convert the report into a prescriptive standard; they clarify how the existing menu applies when AI systems can execute transactions.

Tokenization Systems appreciates the FSB's work on the consultation report, and in particular its early and specific attention to agentic AI. Tokenization Systems' offer is the same one it has made to national authorities: independent, reproducible measurement. The agentic-payments market is small at the value layer, highly concentrated at the control layer, forming faster than any accountability framework around it, and, for a limited window, fully measurable from public records. Tokenization Systems would welcome the opportunity to contribute the case study above, the underlying dataset, or further measurement to the FSB's continuing work, and intends for this response to be published.

Respectfully submitted,

Zach Zukowski

Tokenization Systems

zach@tokenization.systems

www.Tokenization.Systems

ORCID: 0009-0006-3642-2450

SSRN Author ID: 10386216

References:

Financial Stability Board (2026), "Sound Practices for Responsible Adoption of Artificial Intelligence (AI): Consultation report," 10 June 2026.

Zukowski, Z. (2026), "Minimum Viable Equivalence Pack: An Institutional Diligence Framework for Tokenized Products," SSRN working paper. Available at <https://tokenization.systems/papers/mvep/>

Zukowski, Z. (2026), "Routing the Dollar: Stablecoin Gateway Concentration and Dollar Transmission," SSRN working paper. Available at <https://tokenization.systems/papers/routing-the-dollar/>

Zukowski, Z. (2026), "The Control Layer of Agentic Payments: Facilitator Concentration, Subsidy Pipelines, and Settlement Accountability in x402," Tokenization Systems research memo, June 2026. Available at <https://tokenization.systems/papers/agentic-payments-control-layer/>