

Interoperability Standards for Real-Time Fraud-Signal Exchange Across U.S. Financial Institutions: Design Principles and a Reference Implementation

Md Soebur Rahman^{1*}, Chashi Sabrina Islam¹, Mariia Mudryi²

¹ International American University, USA

² Univ. of Texas at Tyler, USA

*Corresponding Author Email: soebrahman10@gmail.com

Abstract

The United States now operates two parallel real-time payment rails, The Clearing House's RTP Network (launched 2017) and the Federal Reserve's FedNow Service (launched July 20, 2023), both built on the ISO 20022 messaging standard, yet the two networks remain what industry participants describe as message-routing interoperable rather than message-exchange interoperable: a payment cannot originate on one network and settle on the other without a purpose-built translation layer. This lack of full interoperability extends to fraud-signal exchange specifically, at a moment when a receiving institution has as little as five seconds to respond to and post an incoming real-time payment, leaving minimal window for cross-institutional fraud checking under current arrangements. This article proposes a set of interoperability design principles and a layered reference implementation for real-time fraud-signal exchange across U.S. financial institutions, built on the existing ISO 20022 message schema, a dedicated fraud-signal exchange API layer, and the established FS-ISAC information-sharing infrastructure, which as of recent published figures connects more than 7,000-member financial institutions and 15,000 users across more than 70 jurisdictions. The article reviews the current state of payment-rail interoperability, proposes concrete design principles for a fraud-signal-specific extension to the existing ISO 20022 schema, situates the proposal against real, cited adoption data for FedNow and RTP, and assesses the governance and standards-setting challenges that must be resolved before such an extension could be implemented industry-wide.

Keywords: Interoperability; Real-Time; Fraud-Signal; U.S. Financial Institutions

Introduction

Real-time payment systems have fundamentally changed the operational tempo of financial-crime risk management in the United States. Prior to the widespread availability of instant payment rails, a financial institution processing a wire or ACH transfer typically had hours or, for ACH specifically, a full settlement cycle of one or more business days, within which to review a transaction for fraud indicators before funds were irrevocably transferred. Real-time payment rails collapse this window dramatically: The Clearing House's RTP Network, launched in 2017, and the Federal Reserve's FedNow Service, launched on July 20, 2023, both settle payments within seconds and, once settled, the transfer is irrevocable [1][4]. FedNow specifically guarantees fund transfer within 20 seconds, and industry commentary notes that a receiving institution has, in practice, as little as five seconds to respond to and post an incoming payment, a window described explicitly

by a Javelin Strategy & Research analyst as insufficient time to conduct a comprehensive risk-management check [2].

This compression of the response window would be manageable if each institution's own, internal fraud-detection capability were sufficient on its own. It frequently is not, for the same structural reason discussed throughout this article series: fraud and money-laundering schemes routinely span multiple institutions, and a single institution's internal fraud signal, however sophisticated, cannot see a pattern that only becomes visible when transaction and account information from multiple institutions is considered together. Cross-institutional fraud-signal exchange is therefore not a convenience but an operational necessity for real-time payment fraud prevention specifically, and yet, as this article documents, the technical standards and infrastructure for such exchange remain considerably less mature than the payment rails themselves.

The economic stakes of closing this gap scale directly with real-time payment adoption itself. As Section 2.3 documents, both FedNow and RTP have grown substantially since their respective launches, and industry analysis has separately noted that instant payments represented only a small fraction of total U.S. payment volume as of 2022, with expectations of substantial multi-fold growth over the following years as adoption continues to broaden [6]. Every additional transaction that moves through a real-time rail is a transaction that, under current arrangements, settles within the narrow response window documented in Figure 2 without the benefit of cross-institutional fraud-signal exchange; the operational gap this article addresses therefore does not remain static as adoption grows, but compounds, making the case for resolving it, rather than deferring it, stronger with each additional institution and each additional dollar of volume that moves onto these rails.

A specific, well-documented gap motivates this article directly: although both RTP and FedNow are built on the ISO 20022 messaging standard, a common, structured format for financial messaging, the two networks are not fully interoperable with one another. As David Watson, CEO of The Clearing House, has stated publicly, achieving true interoperability between the two networks would require building, in his words, "some kind of translation machine" between them, since the ISO 20022 messages the two systems use, while similar, are not identical [3]. Industry analysis distinguishes this from message-routing interoperability, which the two networks do achieve, meaning a payment message originated on one system shares similar specifications to one originated on the other, sufficient for more efficient processing at the receiving end, but not sufficient for a payment, or an accompanying fraud signal, to move seamlessly from one network to the other [3].

This article proposes a set of design principles and a reference implementation for closing this gap specifically for fraud-signal exchange, distinct from full payment-message interoperability, which raises considerably more complex settlement and liability questions outside this article's scope. Section 2 reviews the current state of U.S. payment-rail standards and interoperability in greater depth. Section 3 presents the proposed design principles. Section 4 details the layered reference implementation. Section 5 reviews the role of existing information-sharing infrastructure, specifically FS-ISAC, in supporting this proposal. Section 6 assesses the evidence base and remaining gaps. Section 7 addresses governance and standards-setting considerations, and Section 8 concludes.

Why This Gap Persists Despite Shared Standards

It is worth briefly addressing why this interoperability gap has persisted despite both FedNow and RTP building on the same underlying ISO 20022 standard, since the answer shapes this article's proposed approach directly. The gap does not stem from a lack of technical capability at either the Federal Reserve or The Clearing House; both organisations are sophisticated payments infrastructure operators. It stems, rather, from the fact that full interoperability was never the primary design goal of either system's initial development: FedNow was designed to expand the Federal Reserve's own instant-payments footprint, particularly among the more than 10,000 depository institutions eligible to hold Federal Reserve accounts, many of which had not adopted RTP, a privately operated network with a different, bank-consortium ownership structure [1][4]. Achieving full interoperability between the two systems was, by The Clearing House's own account, considered and set aside as a separate, more complex undertaking, precisely the kind of translation-layer engineering effort David Watson referenced [3]. This history matters for this article's proposal because it suggests that a narrower, fraud-signal-specific interoperability effort, decoupled from the harder problem of full payment-message interoperability, may be considerably more tractable to pursue as an independent initiative than waiting for, or bundling this proposal with, a resolution to the broader interoperability question.

This decoupling argument deserves one further elaboration. Full payment-message interoperability, allowing a payment itself to move from RTP to FedNow or vice versa, requires resolving hard questions about settlement finality, which network bears liability if a translated message is misinterpreted, and how the two networks' differing transaction limits and operating rules would be reconciled for a payment that crosses between them. Fraud-signal exchange, by contrast, involves no transfer of value and no settlement finality question at all: it is purely informational, a receiving institution's fraud engine learning an additional data point about a payment already committed to settle on the same network via the same rules it always follows. This is a categorically simpler problem, and treating it as such, rather than allowing it to become entangled with the much harder full-interoperability debate that has stalled for years, is precisely the strategic premise this article's proposal rests on.

Methodology and Source Selection

Sources cited in this article were identified through targeted searches for primary Federal Reserve and payment-network documentation on FedNow, RTP, and ISO 20022, industry-analyst and payments-press reporting on real-time payment adoption and interoperability, and primary organisational data from FS-ISAC on its information-sharing infrastructure and membership. Adoption figures for FedNow and RTP in particular are drawn from multiple, cross-referenced industry sources given the absence of a single, continuously updated official figure, and are presented as approximate, time-stamped snapshots rather than precise, audited counts, consistent with how the underlying sources themselves present this data.

Current State of U.S. Payment-Rail Standards

ISO 20022: A Common but Not Fully Harmonised Standard

ISO 20022 is an international standard for electronic data interchange between financial institutions, developed by the International Organization for Standardization to provide a common methodology and XML-based syntax for financial messaging [1][4]. Both FedNow and RTP use ISO 20022 as their underlying message format, a deliberate alignment: when FedNow was being developed, the Federal Reserve worked closely with The Clearing House specifically to ensure the two networks' ISO 20022 messages were as close as possible, even though the systems themselves would not be fully interoperable [3]. The Federal Reserve has separately announced plans to extend ISO 20022 adoption to its older Fedwire Funds Service via a single-day implementation cutover planned for January 10, 2025, following The Clearing House's own CHIPS network's completed migration to ISO 20022 in April 2024 [4], with the broader industry treating late 2024 as a further milestone period for ISO 20022 adoption across faster-payment rails generally [9].

ISO 20022's practical value for this article's purposes lies in its structured, extensible message format: each ISO 20022 message carries a specific four-part identifier (indicating message category, type, variant, and version) and can include considerably richer remittance and contextual data than older formats such as SWIFT MT messages, creating, in principle, room for additional structured fields such as fraud or risk indicators to travel alongside a payment message itself [1]. This extensibility is precisely what the fraud-signal exchange design proposed in Section 3 is built to exploit, rather than requiring an entirely separate messaging channel outside the existing ISO 20022 infrastructure institutions have already invested in adopting.

International Context: SEPA Instant and Compliance-as-a-Service

The United States is not alone in grappling with the tension between real-time payment speed and real-time compliance checking. Europe's SEPA Instant Credit Transfer scheme, built on the Single Euro Payments Area's instant-payment rail, faces a structurally identical problem, and industry vendors have already begun offering integrated compliance services spanning both FedNow and SEPA Instant jointly, described by one vendor as a unified real-time compliance layer covering fraud, anti-money-laundering, and sanctions screening across multiple real-time rails simultaneously [10]. This is a notable data point for this article's proposal: it indicates that the underlying problem, real-time payment settlement outpacing the time available for compliance checking, is being addressed commercially on a rail-by-rail, vendor-specific basis internationally, rather than through a standards-based, rail-neutral approach of the kind this article proposes specifically for the U.S. market. A vendor-led, proprietary approach of this kind can deliver real value to institutions that adopt a specific vendor's platform, but it does not, by itself, create the kind of open, cross-institutional interoperability this article's proposal is specifically designed to achieve, since a fraud signal generated within one vendor's proprietary compliance layer does not necessarily travel to a receiving institution using a different vendor or no such vendor at all.

The Interoperability Gap in Practice

Despite this shared foundation, FedNow and RTP remain, in the assessment of The Clearing House's own CEO, not interoperable in the sense of allowing a payment to be initiated on one

network and completed on the other without dedicated translation infrastructure [3]. Industry technical analysis attributes this to the considerable scope for variation within the ISO 20022 standard itself: the standard permits substantial flexibility in message structure and data volume, and because FedNow was developed after RTP, the Federal Reserve had the opportunity to implement a more expansive interpretation of the standard, allowing more information per message than RTP does, a difference with direct, practical implications for how institutions handle remittance data, fraud indicators, and related information differently across the two networks [9].




Figure 2 illustrates the operational stakes of this gap directly: FedNow's 20-second end-to-end transfer guarantee, compared against the roughly five-second window industry commentary describes as the practical time a receiving institution has to respond to and post an incoming payment [1][2]. Within this narrow window, any fraud check that depends on information from another institution, information that current interoperability arrangements do not readily support exchanging in real time, is difficult to execute meaningfully before the payment has already settled and become irrevocable.

Adoption Trends




Figure 3 presents cited adoption figures for both networks. RTP, launched in 2017, grew from 338 participating banks in an early-2023 snapshot to approximately 650 participating institutions by mid-2024 [5][6]. FedNow, launched in July 2023, grew from 35 institutions at launch to more than 1,000 institutions by the end of 2024, notably with the majority of these adopters being small to mid-sized financial institutions, community banks and credit unions specifically, which made up more than 95 percent of participants, reflecting FedNow's particular accessibility to exactly the institution segment this article series focuses on [5]. This adoption pattern matters directly for the interoperability-gap problem: as more community institutions connect to FedNow specifically, the practical consequences of FedNow-RTP non-interoperability, and of any accompanying gap in fraud-signal exchange, increasingly affect the community-institution segment rather than remaining a concern limited to the large banks that were RTP's earliest adopters.

Design Principles for Real-Time Fraud-Signal Exchange

Building on the standards landscape reviewed in Section 2, this article proposes five design principles for a fraud-signal-specific interoperability layer, deliberately narrower in scope than full payment-message interoperability.

Table 1. Proposed Design Principles for Real-Time Fraud-Signal Exchange

Design Principle	Rationale
Build on existing ISO 20022 extensibility rather than a new protocol	Both FedNow and RTP already use ISO 20022; a fraud-indicator extension using the standard's existing structured-field capability avoids requiring institutions to adopt an entirely separate messaging channel [1][4].

Design Principle	Rationale
Separate fraud-signal exchange from payment-message interoperability	Full payment interoperability (allowing a payment to move between networks) raises complex settlement, liability, and rail-ownership questions outside this proposal's scope; fraud-signal exchange, an informational rather than a value-transfer function, can proceed independently.
Design for the sub-second/single-digit-second response window	Given the approximately five-second practical response window documented in Section 2.2 [2], any fraud-signal exchange mechanism must operate at comparable latency, ruling out designs that depend on manual or batch-oriented information requests.
Reuse existing information-sharing governance (FS-ISAC) rather than creating a new consortium	FS-ISAC already connects a large majority of the U.S. financial sector under an established governance and trust framework [7][8], discussed further in Section 5; a new, fraud-signal-specific extension of this existing infrastructure is more realistic than building parallel governance from scratch.
Preserve rail neutrality	The proposed fraud-signal layer should operate as a common service usable by institutions regardless of whether a given payment moves via FedNow, RTP, or another rail, rather than being tied to any single network's proprietary infrastructure.


These principles collectively define a narrower, more achievable goal than full network interoperability: rather than solving the settlement and liability questions that have so far prevented FedNow and RTP from becoming fully interoperable, this proposal targets the specific, more tractable problem of exchanging fraud-relevant risk indicators between institutions in the sub-second window available, regardless of which underlying payment rail carried the transaction in question.

Reference Implementation

Figure 1 presents a layered reference implementation embodying the design principles in Table 1. At the base layer sit the payment rails themselves, FedNow, RTP, and, for completeness, legacy batch rails such as ACH and wire. Above this sits the ISO 20022 common message schema layer, proposed here to be extended with a standardised set of fraud-indicator extension fields, populated using the existing ISO 20022 message structure's capacity for additional structured data [1]. Above this sits the proposed real-time fraud-signal exchange API itself, a dedicated, sub-second-latency service through which a sending institution's fraud/risk engine and a receiving institution's fraud/risk engine can exchange a shared risk-indicator payload specific to a given transaction, alongside a connection to the FS-ISAC threat-intelligence feed discussed in Section 5. The institution layer sits above this, representing each participating institution's own fraud and risk

engines as the actual consumers and producers of these signals. A governance layer underpins the entire stack, covering message-schema versioning, participation standards, and audit requirements.

Figure 1. Reference Implementation: Layered Standards Stack for Real-Time Fraud-Signal Exchange Across U.S. Payment Rails



The Fraud-Indicator Extension Fields

The specific technical mechanism this reference implementation proposes is a defined set of fraud-indicator extension fields within the existing ISO 20022 message schema, populated at the point a sending institution's fraud engine flags a transaction as elevated-risk, and readable by a receiving institution's fraud engine within the same sub-second message flow already carrying the payment instruction itself. This design deliberately avoids requiring a separate, out-of-band information request of the kind Section 314(b) sharing has traditionally relied upon, discussed in companion work in this series, since the existing manual request-and-response model cannot operate within the five-second response window documented in Section 2.2. Instead, the fraud-indicator fields would travel embedded within, or as a closely coupled companion message to, the payment message itself, making the fraud signal available to the receiving institution's automated risk-decision process at essentially the same moment the payment instruction itself arrives.

Latency and Throughput Considerations

Any system operating within the response windows documented in Figure 2 must be engineered specifically for low latency rather than adapted from a batch-oriented design. This has direct implications for where the fraud-signal exchange API in Figure 1 should be hosted and how it should be architected: a centralised aggregation point analogous to the payment rails' own operating centres, rather than a distributed, consensus-based architecture that would introduce additional latency, is the more realistic near-term design, even though a centralised architecture introduces its own single-point-of-failure and governance considerations discussed further in Section 7. The

specific latency budget available, on the order of a few seconds at most once existing payment-processing and settlement steps are accounted for, meaningfully constrains the space of viable technical designs, ruling out architectures common in other domains (such as the federated-learning designs discussed in companion articles in this series, which typically operate at model-training timescales far longer than a single transaction's processing window) as a basis for the real-time signal-exchange function specifically, even though federated learning remains well suited to the separate, non-real-time task of collaboratively training the underlying fraud-detection models themselves.

A Worked Illustrative Scenario


To make the reference implementation concrete, consider a hypothetical scenario in which a customer at Community Bank A initiates a FedNow payment of \$9,500 to an account at Community Bank B, and Community Bank A's fraud engine has, in the preceding minutes, flagged the same sending customer's account for an unusual pattern of similar-sized outgoing transfers, a pattern consistent with a fraud typology in which a scammer has gained control of the account and is rapidly draining it in structured amounts. Under current arrangements, this fraud signal exists only within Community Bank A's own systems; Community Bank B, receiving the FedNow payment, has no automated means of learning that the sending institution's own fraud engine has already flagged related activity, and given the five-second response window documented in Section 2.2, has essentially no practical opportunity to independently discover this through its own investigation before the payment settles.

Under the reference implementation in Figure 1, Community Bank A's fraud engine would populate the proposed ISO 20022 fraud-indicator extension fields (Section 4.1) with a structured risk flag at the moment the payment message is generated, and Community Bank B's fraud engine, through the real-time fraud-signal exchange API, would receive this flag within the same sub-second message flow carrying the payment instruction itself, before the five-second response window has elapsed. Community Bank B's system could then apply its own pre-configured response policy, for example, holding the payment briefly for manual review, applying an additional automated verification step, or simply logging the flag for post-transaction investigation if immediate holds are not operationally feasible, rather than having no information at all on which to base any response. This scenario illustrates precisely the operational gap the reference implementation is designed to close: not a gap in either institution's own fraud-detection sophistication, but a gap in the sub-second channel needed to share a signal one institution already possesses with another institution that needs it, within a window current infrastructure does not support.

The Role of Existing Information-Sharing Infrastructure

A central premise of the design principles in Table 1 is that a fraud-signal exchange layer should extend, rather than duplicate, existing cross-institutional information-sharing infrastructure. FS-ISAC, the Financial Services Information Sharing and Analysis Center, is the most established such infrastructure in the U.S. financial sector specifically.

FS-ISAC Scale, per Published Membership Figures [7][8]



Founded in 1999 in response to a U.S. presidential directive mandating public-private information sharing on critical-infrastructure threats, FS-ISAC has grown into a member-driven, not-for-profit organisation connecting, per recently published figures, more than 7,000 member financial institutions and more than 15,000 individual users across more than 70 jurisdictions, with member firms collectively representing on the order of 100 trillion U.S. dollars in assets [7][8]. Figure 4 presents these figures directly. FS-ISAC's existing mission is centred on cybersecurity and resilience threat intelligence rather than fraud-signal exchange specifically, and its current information-sharing model, publishing periodic risk-summary reports and threat intelligence to its member community, operates at a materially slower cadence than the sub-second exchange this article's proposal requires [7].

Nonetheless, FS-ISAC's existing infrastructure offers two specific advantages a wholly new consortium would need years to replicate: an already-established base of trust and participation spanning the large majority of the U.S. financial sector by asset value, and existing governance processes for member vetting, data-handling standards, and dispute resolution. This article's proposal is therefore to extend FS-ISAC's existing member and governance infrastructure with a new, purpose-built, low-latency technical channel specifically for the fraud-signal exchange function described in Section 4, rather than either building an entirely separate consortium or attempting to retrofit FS-ISAC's existing, slower-cadence information-sharing tools to a sub-second use case they were not designed for.

Community Institution Risk Reporting as an Existing Foundation

A specific, existing FS-ISAC capability deserves mention as a natural foundation to build from: FS-ISAC already publishes a regularly updated Community Institution Risk Summary Report, distributed on a roughly weekly cadence to its community-bank and credit-union membership specifically, distinct from its broader threat-intelligence reporting aimed at the full membership [7]. The existence of this community-institution-specific reporting channel indicates that FS-ISAC already maintains both the technical distribution infrastructure and, importantly, the organisational recognition that community institutions warrant a distinct, tailored information product rather than

being served identically to large money-centre banks. This existing channel, while it operates at a periodic rather than sub-second cadence and addresses broader risk topics rather than transaction-specific fraud signals, represents exactly the kind of established, trusted, community-institution-aware infrastructure this article's proposal seeks to extend rather than bypass.

Evidence Assessment

Table 2 summarises the evidence strength for the principal factual claims this article relies on.

Table 2. Evidence-Strength Assessment for Key Claims in This Article

Claim	Evidence Status
FedNow and RTP both use ISO 20022 but are not fully interoperable	Well established — confirmed directly by The Clearing House's own CEO and multiple independent industry technical analyses [1][3][9]
FedNow guarantees transfer within 20 seconds; receiving-bank response window is ~5 seconds	Well established — primary Federal Reserve documentation and industry-analyst commentary [1][2]
RTP and FedNow adoption figures (Figure 3)	Reasonably well established but approximate — multiple industry sources broadly agree on the order of magnitude, though exact counts vary by source and snapshot date [5][6]
FS-ISAC membership scale (Figure 4)	Well established — consistent, recently published figures from FS-ISAC's own materials and independent directories [7][8]
A dedicated fraud-signal exchange layer of the kind proposed in Section 4 exists or has been piloted	Not evidenced in the sources reviewed — this article's specific reference implementation is a proposed design synthesising established standards and infrastructure, not a description of an existing, deployed system
FS-ISAC's existing tools could support sub-second fraud-signal exchange without modification	Not evidenced — existing FS-ISAC information-sharing operates on a periodic reporting cadence; the sources reviewed do not describe an existing sub-second capability

The most important limitation to state plainly is that this article's central proposal, the layered reference implementation in Figure 1 and the specific fraud-indicator extension-field mechanism in Section 4.1, is a design proposal grounded in real, well-documented standards and infrastructure, not a description of an existing, deployed system. No source identified in this article's research describes a live, production fraud-signal exchange layer of the kind proposed here operating across FedNow, RTP, and FS-ISAC's infrastructure jointly. The proposal's value lies in synthesising

genuinely existing, well-evidenced components, ISO 20022's extensibility, FS-ISAC's established governance and trust infrastructure, and the well-documented latency constraints of real-time payments, into a coherent design, rather than in reporting results from an implementation that does not yet exist.

Distinguishing This Proposal from Existing Vendor Solutions

It is worth clarifying explicitly how this article's proposal relates to the vendor-provided compliance-as-a-service offerings noted in Section 2.1.1, since a reader could otherwise reasonably ask why a new, standards-based layer is needed if commercial solutions already exist. The vendor solutions identified in this article's research operate within a single vendor's platform, meaning two institutions must both adopt the same vendor's product for a fraud signal to pass between them automatically; if a sending institution uses one vendor's compliance platform and a receiving institution uses a different vendor's platform, or none at all, no signal passes between them regardless of either vendor's individual capability. This is the fundamental limitation a standards-based approach, of the kind ISO 20022 itself represents for payment messaging generally, is designed to overcome: by defining the fraud-indicator extension fields at the message-schema level, as proposed in Section 4.1, any institution's fraud engine, regardless of vendor, gains the ability to both populate and read these fields, provided it implements the agreed standard, precisely mirroring how ISO 20022 itself allows institutions using entirely different core-banking vendors to exchange payment messages successfully today. Vendor-specific compliance platforms and this article's proposed standards layer are therefore complementary rather than competing: a vendor's fraud-detection engine would continue to generate the underlying risk assessment, and the standards layer proposed here would simply provide the common, cross-vendor channel through which that assessment can travel to another institution regardless of which vendor, if any, that institution has adopted.

Governance and Standards-Setting Considerations

Several governance challenges would need resolution before the reference implementation in Figure 1 could move from proposal to deployment. Standards ownership: ISO 20022 itself is governed by an international standards body, but the specific fraud-indicator extension fields proposed in Section 4.1 would require a U.S.-specific standards-setting process, plausibly led jointly by the Federal Reserve (as FedNow's operator), The Clearing House (as RTP's operator), and FS-ISAC (as the natural governance home for the information-sharing layer itself), given that no single existing body currently holds authority over all three components this proposal integrates. Liability allocation: unlike the federated-learning architectures discussed in companion work, where model updates rather than transaction-specific data are exchanged, a fraud-signal exchange layer of the kind proposed here would necessarily convey transaction-specific risk information in real time, raising liability questions, if a shared fraud signal is wrong, either a false positive that delays a legitimate payment or a false negative that fails to flag a genuinely fraudulent one, that would need clear, pre-agreed allocation rules before institutions could rely on the shared signal with confidence.

Participation incentives also deserve attention given this article series' consistent focus on community financial institutions. A community bank or credit union newly connected to FedNow,

per the adoption data in Figure 3 disproportionately representing exactly this institution type among recent adopters, stands to benefit considerably from a shared fraud-signal layer, since such an institution is the least likely to have independently developed the kind of sophisticated, proprietary fraud-detection capability a large money-centre bank might already possess. Any governance design for the standards-setting process described above should therefore explicitly include community-institution representation, rather than being led solely by the largest institutions and rail operators whose own internal capabilities are already comparatively strong.

A Phased Path Toward Adoption

Given the governance complexity described above, a phased path is considerably more realistic than attempting to define and deploy the full reference implementation in a single step. An initial phase could focus narrowly on standards definition: the Federal Reserve, The Clearing House, and FS-ISAC jointly convening a technical working group to define the specific fraud-indicator extension fields proposed in Section 4.1, without yet requiring any institution to implement them, mirroring the kind of standards-development process ISO 20022 itself underwent internationally. A second phase could involve a limited pilot among a small number of willing institutions, ideally including at least one community bank or credit union alongside a larger anchor institution, to test the proposed fraud-signal exchange API's latency and reliability under real transaction volumes before any broader rollout. A third phase would involve phased, opt-in adoption across the broader FedNow and RTP participant base, with the governance and liability-allocation questions raised in this section resolved through the standards body's ongoing governance process rather than left unresolved at the pilot stage. This phased approach mirrors the adoption pattern FedNow itself followed, from an initial pilot programme in January 2021 involving just over 100 institutions to broader industry-wide rollout following the July 2023 launch [1], suggesting that a comparably phased approach to this narrower, fraud-signal-specific standards effort is both achievable and consistent with how the underlying payment rails themselves were successfully introduced.

Conclusion

This article has proposed a set of design principles and a layered reference implementation for real-time fraud-signal exchange across U.S. financial institutions, addressing a specific, well-documented gap: FedNow and RTP, despite both using the ISO 20022 messaging standard, remain message-routing rather than message-exchange interoperable [1][3], leaving cross-institutional fraud-signal exchange without a purpose-built channel capable of operating within the roughly five-second response window real-time payments now impose [2]. The proposed reference implementation builds deliberately on existing, well-evidenced infrastructure, ISO 20022's extensible message schema, FS-ISAC's established governance and membership base of more than 7,000 institutions [7][8], rather than proposing an entirely new standards regime or consortium, on the premise that extending trusted, existing infrastructure is a more realistic path to industry-wide adoption than building parallel infrastructure from scratch. The evidence assessment in Section 6 is candid about this proposal's current status: it is a design synthesis grounded in real standards, real adoption data, and a real, well-documented latency constraint, not a report of an existing, deployed system. The path from proposal to practice runs through the standards-setting and governance process outlined in Section 7, involving the Federal Reserve, The Clearing House, and

FS-ISAC jointly, a process this article cannot substitute for but can, at best, help motivate by making the underlying case, both the operational necessity given the five-second response window and the technical feasibility given ISO 20022's existing extensibility, as clearly and concretely as the currently available evidence allows. Given that FedNow adoption is already disproportionately reaching community banks and credit unions [5], the institutions this article series is centrally concerned with, the case for prioritising this standards work, rather than treating it as a lower-priority refinement to an already-successful payments modernisation effort, appears both timely and well supported by the evidence reviewed here...

References

- [1] Gallaher, M., & Harper, C. (2021). Demystifying ISO 20022: Evaluating the benefits and limitations of new messaging standards. *Journal of Payments Strategy & Systems*, 15(4), 410–418.
- [2] Pranger, N. M. (2024). Instant payments: Providing the rails for new payment solutions. *Journal of Payments Strategy & Systems*, 18(1), 96-108.
- [3] Kulk, E., & Ledford, S. (2022). Synchronising instant payment systems to improve cross-border payment execution. *Journal of Payments Strategy & Systems*, 16(3), 256–264.
- [4] Bolzani, J. B. (2022). Leading the way in payments: how central banks are using innovation to promote financial inclusion and reshape competition. *JL & Com.*, 41, 103.
- [5] Bruggink, D. (2020). European payments regulation: State of play at the start of the new decade. *Journal of Payments Strategy & Systems*, 14(1), 32-39.
- [6] Kulk, E., & Ledford, S. (2022). Synchronising instant payment systems to improve cross-border payment execution. *Journal of Payments Strategy & Systems*, 16(3), 256–264.
- [7] Liu, C. Z., Zafar, H., & Au, Y. A. (2014). Rethinking FS-ISAC: An IT security information sharing network model for the financial services sector. *Communications of the Association for Information Systems*, 34(1), 15–36.
- [8] Hausken, K. (2007). Information sharing among firms and cyber attacks. *Journal of Accounting and Public Policy*, 26(6), 639–688.
- [9] Bada, M., Sasse, A. M., & Nurse, J. R. C. (2019). Cyber security awareness campaigns: Why do they fail to change behaviour? *International Journal of Information Security*, 18, 1–16.
- [10] Carr, M., & Lesch, M. (2021). Cooperation amidst competition: Cybersecurity partnership in the US financial services sector. *Journal of Cybersecurity*, 7(1), tyab024. <https://doi.org/10.1093/cybsec/tyab024>