Sarcouncil Journal of Engineering and Computer Sciences

ISSN(Online): 2945-3585

Volume- 05| Issue- 11| 2026

Research Article

Received: 01-10-2026 | Accepted: 25-10-2026 | Published: 06-11-2026

Semantic Interoperability in Payment Infrastructure Modernization: A Canonical Transformation Reference Architecture for ISO 2022 Migration

Chandrashekhar Kharche

Payments Technology.

Abstract: The global migration to the ISO 20022 standard, which concluded its coexistence period on the Swift network in November 2025, is often referred to as a format change. The thesis of this paper is that the idea that the migration is a format change is the reason for the stalled and low-value migrations of many organizations. Instead, the real problem is that of semantic interoperability between the systems that need to exchange payment data. In this paper, I will detail the design science approach to the problem, analyze the standards and cutovers of various market infrastructures and the operations of those infrastructures. In doing so, I will differentiate between the syntactic translation between different formats (MT and XML) and the semantic transformations between them. In contrast to the one-directional converters currently available, this paper will present a canonical transformation reference architecture based upon the Canonical Data Model and Message Translator integration patterns. Using this architecture, all incoming and outgoing messages are normalized into the Canonical Data Model prior to the generation of the target message format. Furthermore, I will outline the transformation pipeline, the validation gates, and the exception and repair process within that architecture. Finally, I will discuss the benefits of this proposed architecture, evaluating it in terms of the measurable figures of the operator, the published benchmarks of third-party organizations, and projections of the outcomes of implementing this model. These projections indicate a reduction in the number of investigations required for transactions of approximately 1,600 staff hours for an organization that processes one million transactions per month. In addition, the presented architecture is intended to be a durable model rather than a one-time cutover plan, as the harmonization of the various payment systems will continue through at least 2027.

Keywords: ISO 20022; semantic interoperability; canonical data model; payment message transformation; straight-through processing; financial market infrastructure; reference architecture.

INTRODUCTION

For two decades, the messaging formats used for cross-border and high-value payments were based on the Swift MT series, which relied on the use of a series of fixed fields and free text. These formats were relatively compact and easily understood by the banks and other financial institutions that used them. However, the legacy formats had a number of drawbacks. Specifically, they often truncated the names and addresses of the parties involved in the payments, they often included the remittance information in a field that was not easily read by machines, and they did not provide a way for those involved in the payment to express the reason for the payment. A number of reports on the cross-border payments industry have noted that one of the causes of the high costs and the delays that were often associated with those payments is the long chains of banks that were involved in the payments and the difficulties of complying with the regulations governing the payments [Committee on Payments and Market Infrastructures, 2018]. The G20 report on the cross-border payments industry recognized these issues and recommended the adoption of a harmonised version of the ISO 20022 standard as one of the building blocks for the future of cross-border payments [Board, F. S. 2020; Committee on Payments and Market Infrastructures, 2023]. As the cross-border payments industry has moved towards faster payments and has faced an increase in the number of regulations with which those payments must comply, the limitations of the legacy messaging formats have become a significant bottleneck in the industry. The ISO 20022 standard, which was first published in 2004, was developed as a way of addressing these issues by creating a standard metamodel and a data dictionary that all banks would use to create the XML-based messages that would be used to communicate the details of each payment [International Organization for Standardization, 2013]. The migration to the ISO 20022 standard will be one of the largest changes ever implemented in the financial industry.

The scale can be described in terms of the deadlines associated with the migration. The Fedwire Funds Service adopted ISO 20022 on 14 July 2025 [Federal Reserve Financial Services, 2025]. The CHIPS network migrated on 2024 [The Clearing House, 2024]. The Eurosystem T2 service and CHAPS migrated on 2023 [European Central Bank. 2023; European Central Bank, 2023]. The Swift network ended the coexistence period between MT and ISO 20022 on 22 November 2025 [SWIFT.2025; Cornejo, H. et al.,2026]. From this perspective, the migration is a

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

*Corresponding Author: Chandrashekhar Kharche

DOI- https://doi.org/10.5281/zenodo.24009332

project with a definite end date. This paper takes a different perspective. A format cutover is necessary but not sufficient for the migration. Those organizations that migrate with the idea that the project is merely to change one message format into another will end up with the same limitations in the new format as they had in the old format. For example, they might map a truncated name field into a field that is to contain a structured name, but leave the name field truncated. Or they might include the remittance text in a field that was intended to contain structured text.

The contribution of this paper is to reframe the problem as one of semantic interoperability and to provide an architecture to address this problem. The paper makes four claims regarding ISO 20022 and how it should be implemented in the payments industry. The first claim is that the value of ISO 20022 lies in the meaning of the data, not in the markup of the data. The second claim is that the canonical data model is the best way to decouple source data formats from target data formats. The third claim is that a disciplined data transformation pipeline will create straight-through processing of data. The fourth claim is that the architecture of data transformation should be maintained and updated as the market evolves. The remainder of this paper is organized as follows. Section 2 discusses the materials and method used for this paper. Section 3 discusses the proposed architecture of data transformation and the quantified benefits of this architecture. Section 4 discusses the design rationale for this proposed architecture, as well as the risks and limitations of this architecture. Section 5 concludes the paper.

What ISO 20022 Actually Standardizes

ISO 20022 is frequently called a message format, but the standard is more precisely a recipe for producing message formats. Its metamodel separates three layers: business concepts (a creditor, a settlement amount, a purpose), logical message definitions derived from those concepts, and physical representations, most commonly XML [Bray, T. *et al*., 2008], that are exchanged on the wire [International Organization for Standardization, 2013]. Messages are organized into business domains, and the payments clearing and settlement domain supplies the message definitions that most migrations touch, such as the pacs.008 customer credit transfer, the pacs.009 financial institution transfer, and the camt cash management reports [SWIFT, 2024]. Because every element traces back to a shared dictionary, two institutions that both implement a message definition correctly should, in principle, agree on what each field means.

That qualification, in principle, is where most of the difficulty lives. The standard constrains structure tightly and meaning more loosely. Optionality is pervasive, local market practice narrows the standard in different ways in different jurisdictions, and reference data such as Legal Entity Identifiers, bank identifier codes, and purpose codes must be populated correctly for the structure to carry meaning. The academic literature on corporate adoption reaches the same conclusion from the practitioner side: Constantino, São Mamede, and Mira da Silva find that the success factors for corporations are organizational and data-quality factors, not the mechanics of generating XML [Constantino, J. *et al*., 2024].

The Migration Landscape

Market infrastructures have followed two main strategies for migration. The big bang cutover strategy requires all participants to move to the new infrastructure on a single date, which minimizes the burden of migration for the participants but introduces risk to the migration process. Alternatively, phased migration distributes the risk of migration over a period of time but introduces the need to maintain two systems in parallel to facilitate the migration. According to the Committee on Payments and Market Infrastructures, the choice of migration strategy depends on factors such as the size of the participant community, the readiness of the participants for migration, and the readiness of the vendors [Cornejo, H. *et al*., 2026]. Additionally, phased migration introduces the need to manage data truncation and the presence of default values in systems that have not yet been upgraded to handle the new data formats [J.P. Morgan. 2024]. Table 1 contains examples of cutovers that have occurred in various jurisdictions, which highlights the unevenness of the migration of market infrastructures even after the Swift coexistence window has passed.

Table 1. Representative ISO 20022 cutovers for high-value and instant payment systems, by strategy.

System (jurisdiction)	Strategy	Year	Note
T2 / Eurosystem (EU)	Big bang	2023	RTGS consolidation; Euro1 same year
CHAPS (UK)	Big bang	2023	Bank of England high-value rail

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

Lynx (Canada)	Big bang	2023	High-value replacement system
CHIPS (US)	Big bang	2024	Private high-value clearing
Fedwire Funds (US)	Big bang	2025	FAIM format retired at go-live
Swift CBPR+ (cross-border)	Phased	2022 to 2025	MT and MX coexistence ended Nov 2025
CIPS (China), RENTAS (Malaysia)	Phased	2023 to 2024	RENTAS used a centralised translator

*Compiled from CPMI Brief No. 11 (2026) [*Cornejo, H. *et al*., 2026*] and public infrastructure announcements* [Federal Reserve Financial Services, 2025; The Clearing House, 2024; European Central Bank. 2023; European Central Bank, 2023]. *Years denote the ISO 20022 go-live for the high-value or cross-border service noted.*

Problem Statement: Translation Is Not Transformation

Interoperability has a number of different layers. Systems are considered syntactically interoperable when one system can parse the data that is provided to it by the other system. They are considered semantically interoperable when both systems have an understanding of the meaning of the data that is being passed between them. A one-directional converter that converts an MT103 into a pacs.008 message will achieve the first of these, but not the second. Such a message will be able to validate against the schema that is defined for the new message format, but the meaning of the message will not be able to be conveyed beyond what was in the original message. Elements that could not be conveyed in the original message will be either empty, filled with a default value, or filled with text that cannot be used by the new system.

The Committee on Payments and Market Infrastructures noted that these converters are useful as a contingency measure for a short period of time, but that using them as a permanent solution is costly and will limit the benefits that can be obtained from the conversion of messages. Some of the intermediaries that were using these converters began to include additional messages in their systems to compensate for the loss of data that resulted from the use of the converters [Cornejo, H. *et al*., 2026]. The lesson that can be drawn from this is that the benefit of a standard is obtained only when the meaning of the data is reconstructed and enriched in the new format. If an address field is to be populated with a purpose code, then that code must be derived and validated for the message that is to be sent. This paper will examine how to build a system that achieves semantic interoperability between the systems that use different payment message formats.

MATERIALS AND METHODS

Research Design

The study adopts a design science research methodology, in which an artifact is constructed and evaluated against objectives [Hevner, A. R. *et al*., 2004]. The artifact of this study is a reference architecture and a transformation pipeline. The study is conducted in four steps. First, the problem is identified through the analysis of documentations (Section 1). Second, four objectives are derived from the problem: to preserve and enrich the meaning of each instruction (O1); to decouple source formats from target rails (O2); to maximize the straight-through processing while bounding the exception handling (O3); and to absorb the continuing change in the market without the need for a second migration (O4). Third, the architecture is designed by composing established enterprise integration patterns to the objectives. Fourth, the design is evaluated analytically by tracing each objective to the element that satisfies the objective and quantitatively using a parametric cost model with inputs obtained from published benchmarks.

Materials

The materials are documents that are publicly available. These documents include the ISO 20022 metamodel and the related standards [International Organization for Standardization, 2013; Bray, T. *et al*., 2008; SWIFT, 2024]; the analyses and requirements from the Committee on Payments and Market Infrastructures and the Financial Stability Board [Committee on Payments and Market Infrastructures, 2018; Board, F. S. 2020; Committee on Payments and Market Infrastructures, 2023; Cornejo, H. *et al*., 2026]; the migration announcements and reports from the market infrastructures and the central banks [Federal Reserve Financial Services, 2025; The Clearing House, 2024; European Central Bank. 2023; European Central Bank, 2023; Bank of Japan, 2022]; the publications from Swift and from the industry that describe the operational effects and the transition practices [SWIFT, 2025; J.P. Morgan. 2024; SWIFT, 2024; U.S. Faster Payments Council, 2023]; the research articles on the adoption of these practices within the corporate industry that have been published in peer-reviewed

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

journals [Constantino, J. et al., 2024]; the literature on enterprise integration and enterprise architecture [Hophe, G., & Woolf, B. 2003; Linthicum, D. S. 2000; Fowler, M. 2012]; and the standards for the metadata registries and the financial ontologies that are used to create the semantic grounding for these practices [ISO, I. 2015; Council, E. D. M. 2014]. These documents were selected for this study because they were issued by the standards-setting bodies, the central banks, the market infrastructures, or the peer-reviewed journals, or because they are the canonical references for the integration practices that were examined in this study. No data that is proprietary to any individual firm or the industry as a whole was used, and no figures in this study were taken from any unnamed institution.

Design Basis: The Canonical Data Model

The design response to the problem in Section 1.3 is to stop mapping formats to each other and instead map every format to a single shared model of meaning. This is the Canonical Data Model pattern from the enterprise integration literature: rather than build a translator for each pair of formats, which grows with the square of the number of formats, each format is translated once into and once out of a common internal representation by a Message Translator [Hophe, G., & Woolf, B. 2003; Linthicum, D. S. 2000]. The canonical model becomes the single place where the meaning of a payment is defined, where reference data is attached, and where validation rules live, in the manner of a layered domain model [Fowler, M. 2012]. Source and target formats then vary independently, which matters precisely because ISO 20022 market practice is still changing: a new message version or a new jurisdictional profile becomes an edge change at one adapter rather than a change rippling across every point-to-point mapping. Semantic grounding for the canonical model draws on shared financial ontologies and metadata registry discipline so that an element such as *creditor* means the same thing to every system that reads it [International Organization for Standardization, 2013; ISO, I. 2015; Council, E. D. M. 2014].

Evaluation Method and Evidence Tiers

Benefits are categorized into three distinct tiers of evidence so that it is possible to make a statement as to the strength of each. The first tier consists of figures that can be measured directly within the operator's own environment once the architecture is live. The second tier consists of third-party benchmarks with citations. The third tier consists of modeled projections based on Tier B benchmarks, which are labeled accordingly. We do not attribute any of the measured figures to any one institution because the reality is that many of the benefits are still being realized.

The Tier C efficiency model is parametric, allowing the model to be re-run with the operator's own numbers. If the monthly instruction volume is represented by V, the enquiry rate by r, and the mean investigation time in minutes by t, the monthly investigation effort in staff hours is H = V * r * t / 60. If the validation gate automatically classifies and repairs a share of the enquiries represented by a, the effort avoided is H * a hours, and the cost avoided is H * a * c per month, where c is the loaded operations cost. The figures for r and t are obtained from published operations data [SWIFT, 2024], a is set at a conservative percentage below the auto-sorting ceiling published in the same document, and c is an assumed cost figure for illustration purposes.

RESULTS

The Canonical Transformation Reference Architecture

Figure 1 illustrates the reference architecture. It is depicted as a stack of responsibilities with a single canonical model at its core. Payments enter from a variety of channels. They are then ingested and transformed into the canonical model, which includes the enrichment and validation of those payments. Finally, those validated payments are generated into the appropriate target message and routed to the appropriate payment rail. A governance and observability spine runs the height of the stack, as those functions are cross-cutting throughout the stack.

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

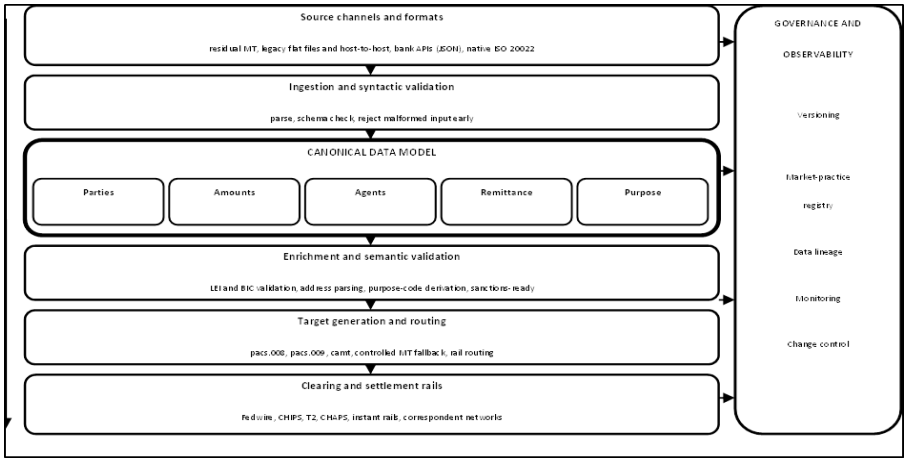


Figure 1. Canonical transformation reference architecture. Heterogeneous source formats are normalized once into a single canonical data model, enriched with validated reference data, and only then generated into target ISO 20022 messages and routed to clearing rails. A governance and observability spine spans all layers; the canonical core decouples source formats from target rails so each can evolve independently.

Layer Responsibilities

The source and channel layer accepts all the inputs that the estate needs to support, including MT traffic, files, batches, APIs, and ISO 20022 messages. The ingestion and syntactic validation layer takes these inputs and checks them against their schemas. Any input that is malformed is rejected before it reaches the core of the system. The canonical data model is the core component of the system and represents the data in a way that is independent of the format of the messages.

The enrichment and semantic validation layer adds value to the structured data. This layer validates the parties to the transaction, the LEIs and bank identifier codes, the addresses, the purpose codes, and the fields that are used for sanctions screening. The target generation and routing layer creates the message that is needed for the target system and routes the payment to the correct destination. The settlement rail layer represents the systems to which the payments can be made. The canonical data model and the enrichment layer satisfy objective O1. The adapter structure satisfies objective O2, and the governance spine satisfies objective O4.

The Transformation Pipeline and the Exception Path

Architecture defines where the responsibilities live; the pipeline describes the movement of a single instruction through those responsibilities and satisfies objective O3. As shown in Figure 2, the pipeline includes two lanes. The upper lane is the straight-through path consisting of capture, parse, normalize into the canonical model, enrich and validate, generate the target message, and dispatch. The lower lane is the exception path, with the validation gate between the two lanes being the single most important design decision in the pipeline.

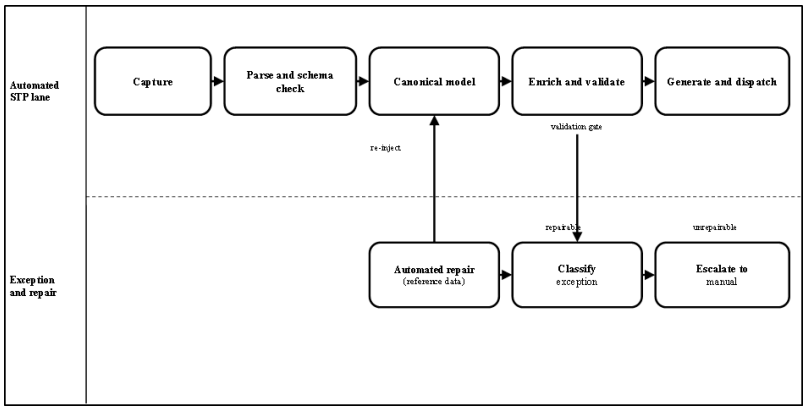


Figure 2. Message transformation pipeline. The automated straight-through path (upper lane) carries conformant instructions from capture to dispatch; a validation gate diverts only the instructions that fail syntactic or semantic checks into a bounded exception and repair lane (lower), which re-injects a corrected instruction rather than returning it to the sender.

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

A naive pipeline is one where a failure in the validation phase results in the rejection of the instruction, which is returned to the sender. This approach is inefficient as it leads to a maximum number of round trips, thus destroying the advantage that the standard was meant to provide in terms of speed. A better approach to the gate is one where invalid inputs are rejected as before, but those that are structurally valid but contain some semantic incompleteness (e.g., missing a purpose code, address that did not parse) are diverted into a repair stage. In this repair stage, the instruction is automatically enriched with the necessary information from reference data, before being re-injected into the canonical processing stage. Only those instructions that cannot be repaired after some defined number of attempts are sent to manual handling. This increases the straight-through rate of transactions while avoiding the creation of an exceptionally large number of exception transactions. The reason for bounding the repair stage is that an unbounded repair loop is simply a slower form of rejection.

It is as important to instrument the gate as it is to design it. Each diversion out of the main processing stage should be accompanied by a reason for the diversion. Each repair should be documented as to the changes made and the reference source from which the missing information was retrieved. This information forms the foundation for later investigations of those transactions. In the payments operations literature, a typical number of payments that require enquiries from the operations team is 2-5% of the total payments made in a day, with each enquiry requiring an average of 3-4 minutes to investigate [SWIFT, 2023]. A gate that can automatically repair the majority of these transactions would be the mechanism that enables the level of efficiency that will be discussed later in the report.

Quantified Benefits by Evidence Tier

Table 2 outlines each benefit and the corresponding architecture and tier of evidence. Four of the benefits have Tier B benchmarks, while interface maintenance relies on structural arguments from the integration literature to become Tier A for operators that count their adapters.

Table 2. Benefit areas by evidence tier.

Benefit area	Mechanism	Evidence tier	Indicative basis
Straight-through processing	Structured data removes truncation and parsing ambiguity	B	Up to 84% of investigation messages auto-sortable once ISO-native [SWIFT, 2024]
Exception handling cost	Automated classify-and-repair at the validation gate	B	2 to 5% enquiry rate; 3 to 4 min each [SWIFT, 2024]
Compliance screening	Clean structured parties and purpose codes reduce ambiguity	B	Richer data strengthens financial-crime controls [Cornejo, H. *et al*., 2026; Constantino, J. *et al*., 2024]
Reconciliation	End-to-end structured remittance and references	B	Faster reconciliation and working-capital gains [Cornejo, H. *et al*., 2026; U.S. Faster Payments Council, 2023]
Interface maintenance	Canonical core replaces n-squared mappings	A/C	Linear vs quadratic adapter growth [Hophe, G., & Woolf, B. 2003]

Tier A*: operator-measured once live. Tier B: published third-party benchmarks. Tier C: modeled projection. Sources as cited.*

Worked Efficiency Illustration

The following illustration is for Tier C, which is a modeled projection from the parametric model defined in Section 2.4. A service with one million instructions per month (V) at an enquiry rate of four percent (r) results in forty thousand instructions that require investigation; at 3.5 minutes per investigation (t), the total effort required to investigate these is 2,333 hours per month [SWIFT, 2024]. A validation gate that automatically classifies and repairs 70 percent of these investigations (a), which is consistent with the up-to 84 percent rate of auto-sorting reported for fully ISO-native estates, would remove 1,633 hours per month, leaving only 700 hours for manual handling [SWIFT, 2024].

Assuming a cost of $40 per hour for operations, the 1,633 hours that are eliminated results in $65,320 in cost savings per month and $783,840 in cost savings per year for one service. These cost savings can be increased or decreased according to the serviceâ ™s volume of instructions per month and the number of instructions that can be

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

automatically repaired. Additionally, these savings do not include the cost savings that would result from the faster reconciliation of structured remittance data [U.S. Faster Payments Council, 2023], or from the cost savings that would result from the reduction of false positives from the screening of instructions.

DISCUSSION

Why the Canonical Core Earns Its Place

Placing a canonical model at the center of your architecture is not a free exercise. It introduces a layer to the system as well as a need to adopt a modeling discipline. The value of this addition comes from the way it changes the way in which your costs scale. If you use a point-to-point mapping between n formats, your number of mappings will scale with n squared. If you use a canonical model, however, the number of mappings will scale linearly with the number of formats [Hophe, G., & Woolf, B. 2003; Linthicum, D. S. 2000]. Each of these mappings is isolated from the others. If your standards are still being harmonized, your ability to introduce a new version of the standard as a localized adapter change will save you from having to migrate your entire data estate. Furthermore, using a canonical model allows you to locate all of your reference data lookups and validation rules in a single place. This makes it easier to audit your systems and ensure that your data is always subjected to the same screening procedures.

Interpreting the Benefit Case

The point of the worked illustration is not the precise number, which depends entirely on an operator's volume, enquiry rate, and cost base, but the structure of the case: the return comes from moving exceptions out of human queues, and that is an architecture and data-quality outcome, not a format outcome. This is consistent with the finding that success in adoption is driven by organizational and data-quality factors rather than by message generation [Constantino, J. *et al*., 2024], and with the observation that converters limit benefits when they are treated as permanent [Cornejo, H. *et al*., 2026].

Risks and Governance

The architecture has some failure modes that should be stated plainly. The canonical model is a single point of failure in that if it is modeled badly, all the adapters will inherit that error. Therefore, the canonical model must be governed as a shared asset with explicit ownership and change control. The enrichment of the canonical model is only as good as the reference data used to enrich it. If the reference data is stale or incorrect, it will introduce errors into the system, which is worse than having gaps in the model. Any automated repair of the canonical model should be bounded and fully logged. An automated repair that modifies a payment instruction is both an operational and compliance risk. The lineage of such a repair is essential to allow the system to be audited.

Limitations

Two limitations bound the claims above. The quantified benefits of Tier B and C benchmarks are offset by the Bank of Japan experience in which forty percent of the participants in one of the systems in the market had implemented ISO native core banking by the summer of 2025 but the false positive rates of the anti-money laundering systems of those participants had not changed [Cornejo, H. *et al*., 2026]. The interoperability of the systems of the various participants is another limitation because the markets are adopting the market practice updates at different rates [Cornejo, H. *et al*., 2026]. The architecture of the proposed system is a living architecture that can evolve with the market. The harmonization requirements will be maintained through at least 2027 and the governance spine in Figure 1 will allow for the absorption of a new version or profile of the architecture without requiring a second migration of the systems [Committee on Payments and Market Infrastructures, 2023; Cornejo, H. *et al*., 2026]. Future work will require validation of the model with the Tier A benchmarks from the live implementations of the system, specifically the baseline and post-implementation enquiry and repair success rates at the validation gate.

CONCLUSION

The ISO 20022 migration is typically described in terms of its deadlines, most of which have already passed. Yet, to consider the migration finished would be a mistake. The migration cutover was the precondition of the migration, but the value of the migration is unlocked only if the migration is treated as a problem of semantic interoperability. To unlock the value of the ISO 20022 migration, the organization must implement a canonical transformation architecture that uses a single internal model at the core of the architecture, a pipeline that validates and classifies the data as they move through the pipeline, and a governance spine that allows the organization to accommodate the evolving ISO 20022 standard. The benefits of such an architecture are real, yet conditional upon

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

data quality, automated exception handling, and governance. The benefits of the ISO 20022 migration accrue to those organizations that are constructing their infrastructure for the decade after the migration cutover, not for the cutover itself.

Acknowledgments
This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. The author declares no conflict of interest.

REFERENCES

1. Committee on Payments and Market Infrastructures. "Cross-border retail payments." *Bank for International Settlements.* (2018).
2. Board, F. S. "Enhancing cross-border payments. Stage 3 roadmap." *Basel. https://www. fsb. org/wp-content/uploads/P131020-1. pdf* (2020).
3. Committee on Payments and Market Infrastructures. "Harmonised ISO 20022 data requirements for enhancing cross-border payments." *Bank for International Settlements*. (2023).
4. International Organization for Standardization.. "ISO 20022-1:2013, Financial services, Universal financial industry message scheme, Part 1: Metamodel. ISO, Geneva." (2013).
5. Federal Reserve Financial Services. "Fedwire Funds Service adoption of the ISO 20022 message format (go-live 14 July 2025)." *Federal Reserve, Washington, DC*. (2025).
6. The Clearing House. "CHIPS migration to the ISO 20022 message format." *The Clearing House, New York*. (2024).
7. European Central Bank. "T2 go-live and the migration to ISO 20022." *ECB, Frankfurt*. (2023).
8. Bank of England. "Migrating CHAPS to ISO 20022." *Bank of England, London*. (2023).
9. SWIFT. "ISO 20022 for financial institutions: navigating the end of coexistence." SWIFT, La Hulpe. (2025).
10. Cornejo, H., Fahrmann, S., Kanaguchi, T., Lammer, T., Mulaudzi, M., Orina, C. G., and Semorile, F. "The future of financial messaging: navigating the ISO 20022 migration journey. CPMI Brief No. 11." *Bank for International Settlements*. (2026).
11. Bray, T., Paoli, J., Sperberg-McQueen, C. M., Maler, E., and Yergeau, F.. "Extensible Markup Language (XML) 1.0 (Fifth Edition)." *W3C Recommendation*. (2008).
12. SWIFT. "ISO 20022 programme: message definitions for payments clearing and settlement (pacs) and cash management (camt)." *SWIFT, La Hulpe*. (2024).
13. Constantino, J., Sao Mamede, H., and Mira da Silva, M. "Adopting ISO 20022: Opportunities, Challenges, and Success Factors for Corporations in Payment Processing." *Emerging Science Journal*, 8.4 (2024).
14. J.P. Morgan. "Navigating the ISO 20022 message transition." *J.P. Morgan Payments*. (2024).
15. Hevner, A. R., March, S. T., Park, J., & Ram, S. "Design science in information systems research1." *MIS quarterly* 28.1 (2004): 75-106.
16. Bank of Japan. "Payment and Settlement Systems Report." *Payment and Settlement Systems Department, Bank of Japan, Tokyo.* (2022).
17. SWIFT. "Supercharge your payments business with ISO 20022." *SWIFT, La Hulpe*. (2024).
18. U.S. Faster Payments Council. "The value of ISO 20022 for U.S. B2B instant payments." *FPC, Washington, DC*. (2023).
19. Hophe, G., & Woolf, B. "Enterprise integration patterns." *Designing, Building, and Deploying Messaging Solutions. Boston et. al.: Addison-Wesley* (2003): 48.
20. Linthicum, D. S. "Enterprise application integration." *Addison-Wesley Professional*, (2000).
21. Fowler, M. "Patterns of enterprise application architecture." *Addison-Wesley*, (2012).
22. ISO, I. "Information Technology–Metadata Registries (MDR)–Part 1: Framework." (2015).
23. Council, E. D. M. "Financial Industry Business Ontology." *Online verfügbar unter http://www. edmcouncil. org/financialbusiness, zuletzt aktualisiert am* (2014).

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:
Kharche, C. "Semantic Interoperability in Payment Infrastructure Modernization: A Canonical Transformation Reference Architecture for ISO 20022 Migration" *Sarcouncil Journal of Engineering and Computer Sciences* 5.11 (2026): pp 1-8.

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher