THE ADAPTIVE COMPLIANCE SERIES

VOLUME IV

THE MESSAGE IS THE CONTROL

How Payment Data Is Becoming the First Line of Financial-Crime Defence

MICHEAL SHEEHY

Revised Edition · July 2026

Copyright and Publication Note

This manuscript is an original work of analysis and opinion written for compliance, payments, risk, technology and business leaders. It is intended to advance discussion about the design of effective financial-crime controls. It does not constitute legal, regulatory or professional advice.

The examples used throughout are illustrative or composite. They are designed to explain operating models and control principles without describing confidential customer, employer or investigative information.

Contents

Foreword

Financial-crime compliance has spent much of the past twenty years looking at payments from the outside.

The payment moved. A message recorded the instruction. A monitoring system later examined selected attributes and decided whether the transaction resembled a rule, scenario or learned pattern. If it did, an investigator attempted to reconstruct the commercial reality behind it.

This model was understandable when payment messages were sparse, infrastructures were fragmented and processing time created a natural interval between initiation and settlement. But the conditions supporting it are disappearing. Payments are becoming faster, more interconnected and increasingly available around the clock. At the same time, payment messages are becoming richer, more structured and more capable of carrying information about the people, organisations and purposes behind the transfer.

The significance of that change is easy to underestimate.

The move to richer payment data is often described as a technology migration: new schemas, new fields, new validation rules and a long programme of implementation meetings. Anyone who has survived such a programme may reasonably feel that “new era of payments” is a rather glamorous description for several years of mapping spreadsheets.

Yet the deeper change is not technical. It is architectural.

When the information accompanying a payment becomes structured, reliable and available at the moment of decision, it can do more than explain the payment. It can help control it.

That possibility connects directly to the argument developed throughout the Adaptive Compliance series. Compliance effectiveness is not defined by the volume of controls an organisation can document. It is defined by how quickly the organisation can turn information into a sound decision, learn from the outcome and improve what happens next.

Payment data sits at the centre of that learning system.

This volume explores what changes when the message is no longer treated as a passive transport envelope. It examines the opportunity created by structured standards, the regulatory move toward enforceable payment transparency, the additional attribution complexity created by virtual-account models, the importance of preserving ultimate-party identity and the role of context in sanctions, fraud and anti-money-laundering decisions. It also considers the governance required when artificial intelligence begins interpreting payment narratives and recommending action in real time.

The argument is not that every field should become mandatory, every payment should carry a biography or every institution should share every piece of customer information. More data can create privacy risk, operational friction and false confidence. A beautifully structured lie remains a lie.

The argument is that payment information should be designed, governed and measured as part of the control environment.

The industry has spent billions teaching monitoring systems to infer facts that the payment message could have carried in the first place. The next stage of financial-crime control begins by closing that gap.

Introduction

The Information Was Always Part of the Payment

Every payment tells a story.

There is a payer, a beneficiary, a reason, a relationship and a route. There is usually an economic event behind the instruction: an invoice, payroll, a marketplace sale, a family remittance, a refund, an investment, a property purchase or the movement of money between accounts controlled by the same party.

Traditional payment messages captured only fragments of that story. Compliance systems were then asked to decide whether the payment was suspicious by examining the fragments that happened to survive the journey.

The result was a peculiar operating model. Institutions possessed detailed customer information in onboarding platforms, commercial information in product systems, device intelligence in fraud tools, sanctions data in screening platforms and transaction history in monitoring systems. The payment message passed between them carrying far less context than the organisation already knew.

We called the resulting alerts “false positives,” as if the principal problem were the alert. Often the real problem was missing information.

The transition to ISO 20022 and other richer messaging arrangements creates an opportunity to reconsider that model. ISO 20022 does not solve financial crime. It provides a common language capable of carrying structured information. Whether that capacity becomes a control depends on the quality of the data, the governance attached to it and the decisions the institution designs around it.

Payment transparency is becoming enforceable. Regulation (EU) 2023/1113 has applied since 30 December 2024, requiring specified payer and payee information to accompany transfers of funds. From 1 May 2025, the Bank of England mandated purpose codes and Legal Entity Identifiers for defined CHAPS payments. FATF revised Recommendation 16 in June 2025 and published the related assessment methodology in October. From 14 November 2026, Swift will no longer support fully unstructured addresses in CBPR+ payment messages, with CHAPS aligned to reject them. The revised FATF requirements must take effect by the end of 2030.

The migration, however, is not the destination.

There is a meaningful difference between a payment message that is technically valid and one that supports a good financial-crime decision. A field can be populated with incomplete information. A structured address can belong to the wrong party. A purpose code can describe the category of a

payment without explaining the commercial relationship. An account identifier can lead to a pooled structure that obscures the ultimate user.

The control objective must therefore be broader than completeness.

The payment-information chain should answer five questions:

1. Who are the real parties behind the payment?
2. Why is the payment taking place?
3. Is the route consistent with the parties and purpose?
4. Does the activity fit what is known and expected?
5. What decision should be made before the value becomes irretrievable?

These questions connect onboarding, account issuance, payment initiation, messaging, screening, monitoring and investigation. They also expose why functional separation has become a barrier to effectiveness. KYC teams validate identity. Product teams configure accounts. Operations teams manage payment exceptions. Fraud teams analyse devices and social engineering. Sanctions teams screen parties. AML teams monitor patterns. Each function sees part of the story.

The payment message can become the point at which those parts meet.

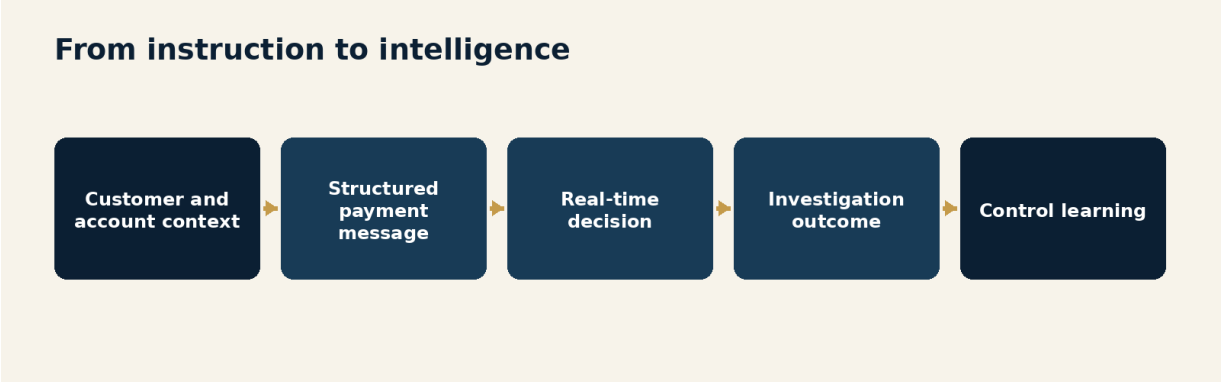

Figure 0.1 · From instruction to intelligence

The opportunity is not to turn the message into an overloaded file containing every fact an institution holds. It is to create a governed connection between the message and the information necessary for the decision. Sometimes that information will travel within the payment. Sometimes it will be resolved internally from an account identifier, token or reference. Sometimes it will be retrieved through an authorised network service.

What matters is that the control can reach the context when the decision must be made.

This volume develops that proposition through eight chapters. It begins with the historical separation between payment operations and compliance. It then examines structured data, enforceable payment transparency, ultimate-party identity, contextual decisioning, upstream screening and artificial intelligence. It concludes with a maturity model for turning payment information into an adaptive control system.

The message is becoming the control. The question is whether institutions will design it that way.

CHAPTER 1

We Treated Payment Data as Plumbing

The industry designed messages to move money and controls to explain it afterwards.

> **EXECUTIVE INSIGHT: The weakness in many monitoring systems is not the sophistication of the detection model. It is the thinness of the information the model receives.**

The historical bargain

Payment infrastructure was built to achieve a demanding but relatively narrow objective: transfer value accurately, securely and predictably between institutions.

Messaging standards therefore focused on the information required to execute and reconcile the instruction. Compliance was layered around that infrastructure. Names were screened against sanctions lists. Transactions were evaluated against scenarios. Investigators gathered additional information when something looked unusual.

This separation created an implicit bargain. Payment systems would move money efficiently; compliance systems would interpret risk separately.

For a time, that bargain was workable. Payments moved through slower infrastructures, many products operated during defined windows and institutions could intervene through manual queues. Information lost from one message could sometimes be recovered through requests between banks or questions to the customer.

But the bargain carried a structural weakness: the institution responsible for deciding whether the payment was appropriate often lacked the information necessary to understand it.

Compliance became an exercise in reconstruction

Consider a payment containing a payer name, beneficiary name, amount, currency, account details and a short free-text reference.

Those fields may be enough to send the money. They are rarely enough to explain the economic event.

Is the payer buying a service, paying a supplier, moving its own funds or sending money on behalf of another party? Is the beneficiary the ultimate recipient or an intermediary collecting for thousands of users? Does the reference identify an invoice, describe the payment honestly or contain whatever text happened to fit?

Monitoring systems learned to work around these gaps. They used thresholds, velocity, geographies, counterparties and historical patterns as proxies for meaning. Investigators then searched across customer files, account notes and previous cases to assemble the context that did not travel with the payment.

This was less a single control than a chain of compensating controls.

The industry responded by improving the intelligence applied after the message was created. Rules became more complex. Machine-learning models found relationships across larger datasets. Case-management systems assembled additional evidence. These advances mattered, but they often left the underlying information architecture unchanged.

We built increasingly sophisticated telescopes and continued pointing them through a dirty window.

Faster payments expose the weakness

The separation between message and control becomes more consequential as payment speed increases.

When funds move in seconds, “detect and investigate” may mean “understand and document after the loss.” Fraud proceeds can be dispersed. Sanctions exposure can become irrevocable. Mule networks can move value through several accounts before a conventional alert reaches an investigator.

The answer cannot be to make every payment wait for a human review. That would defeat the purpose of modern payment systems and generate a queue large enough to qualify as its own national infrastructure.

The answer is to move more intelligence into the decision path.

That requires controls capable of distinguishing between three activities:

- **Execution:** Is the instruction technically valid and capable of settlement?
- **Permission:** Is the payment allowed under law, policy and product design?
- **Confidence:** Is there sufficient evidence that the parties, purpose and behaviour are consistent with legitimate activity?

Traditional messages were primarily designed for execution. Adaptive payment controls must support all three.

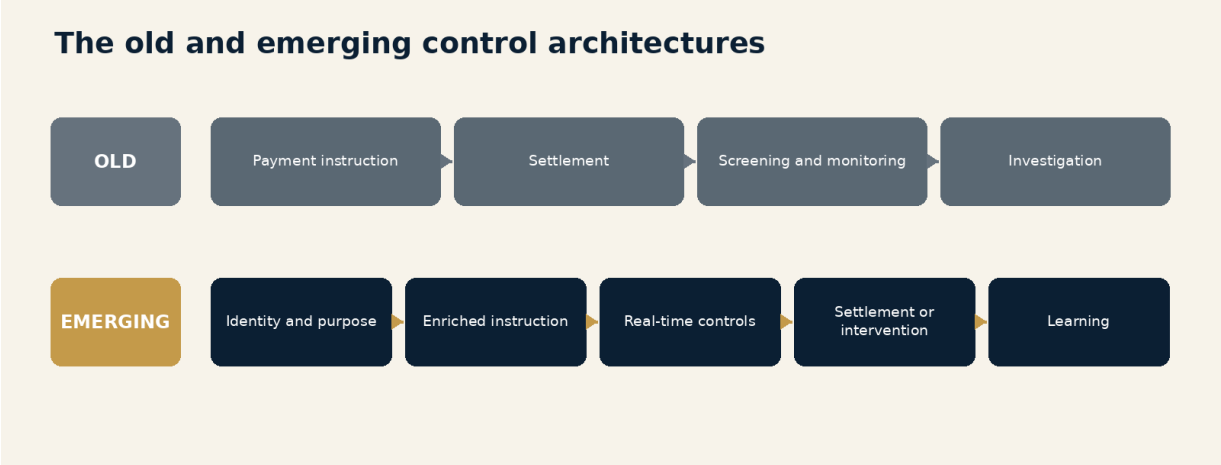

Figure 1.1 · The old and emerging architectures

Data fragmentation is a control problem

Most institutions do not lack information. They lack connected information.

Customer identity may sit in a lifecycle-management platform. Beneficial ownership may be stored in documents. Expected activity may exist in an onboarding questionnaire. Device and session data may remain in a fraud platform. The commercial purpose of a virtual account may live in a product configuration table. Payment details may be stored in a separate ledger.

Each system can be correct while the overall decision remains incomplete.

This is why data architecture belongs within financial-crime governance. If the organisation cannot connect what it knows about the customer with what it sees in the payment, the control environment is limited by design.

Data lineage is therefore not merely a technology concern. Leaders should know:

- where each critical payment field originates;
- whether the field is supplied, verified, inferred or defaulted;
- which transformations occur across the payment chain;
- whether truncation or translation changes meaning;
- which controls rely on the field; and
- who is accountable when the information is missing or wrong.

From principle to practice

The first step is to create a payment-information inventory rather than another list of monitoring rules.

For each major payment flow, map the information available at six points:

1. customer onboarding;
2. account or virtual-account issuance;
3. payment initiation;
4. message construction;
5. screening and decisioning; and
6. investigation and reporting.

The gaps will often be revealing. Information collected during onboarding may never reach screening. A product may permit an “on behalf of” use case without reliably capturing the ultimate party. A monitoring model may depend on a purpose field populated through a default value.

The objective is not maximum data. It is decision-relevant data with known provenance.

Common failure modes

- Treating a messaging migration as a technology programme with no financial-crime owner.
- Measuring field population without measuring accuracy or usefulness.
- Assuming screening can compensate indefinitely for missing ultimate-party information.
- Preserving separate fraud, sanctions and AML data models that describe the same payment differently.
- Enriching messages after settlement when the relevant decision should have occurred before it.

Questions for Leaders

1. What information does our monitoring system routinely attempt to infer because the payment does not state it?
2. Which customer or product facts are available internally but absent from the decision path?
3. Where is payment data truncated, translated, defaulted or overwritten?
4. Which controls would materially improve if purpose and ultimate-party information were reliable?
5. Are we modernising the message, or merely replacing its format?

CHAPTER 2

Structured Data Changes the Control Equation

A common language creates control capacity, but only governance turns that capacity into effectiveness.

> **EXECUTIVE INSIGHT: ISO 20022 is not a financial-crime control. It is an information architecture from which better controls can be built.**

The difference between text and structure

Legacy payment messages often compressed important information into limited or free-text fields. The same fact could appear in several formats, positions or abbreviations. Screening tools compensated through fuzzy matching, transliteration logic and extensive exception handling.

Structured data changes that equation. It separates parties, accounts, addresses, agents, remittance information and purpose into defined elements. That makes the information easier for machines to validate, compare and preserve.

The compliance value is not simply that more characters are available. Structure allows the institution to understand what a piece of information claims to represent.

A country in a country field can be validated as a country. A legal-entity identifier can be evaluated as an identifier. An ultimate creditor can be distinguished from the institution receiving the immediate instruction. A purpose code can be assessed against the customer, product and transaction.

That semantic clarity creates the foundation for better controls.

A valid message can still be a poor control

Technical validity is a low bar.

A required field can be populated with placeholder text. A structured address can be copied from an intermediary. A purpose code can be selected for operational convenience. An organisation name can be accurate while the underlying end user remains undisclosed.

Institutions therefore need at least four measures of payment-data quality:

- **Completeness:** Is the required information present?

- **Conformity:** Does it follow the required format and permitted values?
- **Accuracy:** Does it correspond to the real party, account or purpose?
- **Decision utility:** Does it improve the control decision for which it is used?

Many programmes measure the first two because systems can calculate them easily. The latter two require comparison with customer records, investigations, external evidence and outcomes.

The difference matters. A field that is 99 per cent populated but routinely wrong is not a control achievement. It is an efficiently distributed error.

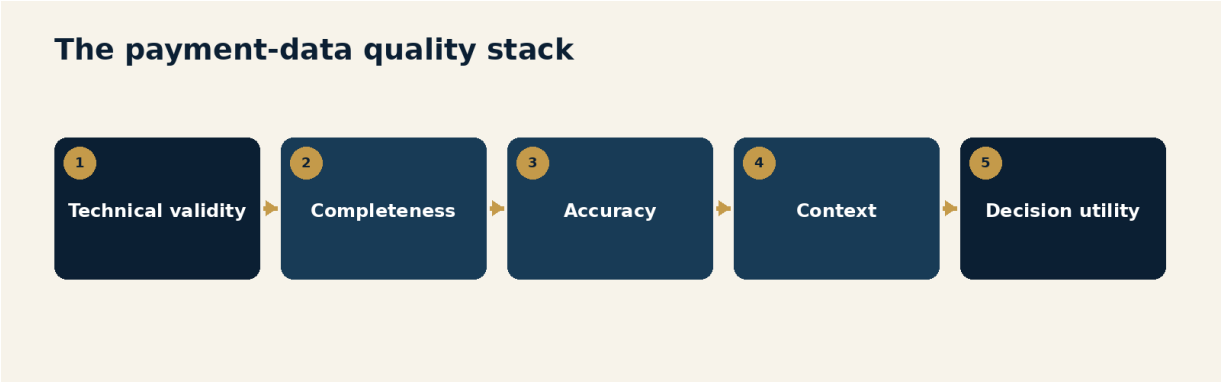

Figure 2.1 · The payment-data quality stack

Data provenance becomes essential

As payment messages grow richer, institutions must know where the information came from.

Some data is supplied directly by the customer. Some is retrieved from verified onboarding records. Some is generated by the product. Some is added by an intermediary. Some is inferred by a model.

Those sources do not have equal evidential weight.

A customer-entered description of purpose should not be treated like a verified account-ownership record. An inferred industry code should not silently replace a customer's declared business activity. An AI-generated summary should not be mistaken for a factual attribute unless the underlying evidence is available.

Payment-data governance should therefore preserve three things:

1. the value;
2. the source of the value; and
3. the confidence or verification status attached to it.

Without provenance, enrichment can create false confidence. The message appears more complete while the institution becomes less clear about which facts it actually knows.

Harmonisation and the risk of local interpretation

Common standards create interoperability only when institutions use them consistently.

The CPMI's harmonised ISO 20022 data requirements are intended to reduce differences across payment systems and jurisdictions. This matters because local implementation choices can recreate the fragmentation the standard was designed to solve.

One institution may use a field for the ultimate party. Another may repeat the immediate party. One market may mandate a purpose code. Another may leave it optional. Intermediaries may translate between internal formats and the network message, dropping detail in the process.

The control environment must therefore examine the entire chain, not simply the format at the point of origin.

Institutions should test whether information survives:

- channel and API ingestion;
- internal ledger transformation;
- message creation;
- intermediary processing;
- network translation;
- beneficiary receipt; and
- returns, recalls and investigations.

From fields to control hypotheses

The best way to create compliance value from structured data is to begin with a decision, not a field.

Instead of asking, "How can we use purpose codes?" ask:

"Which decisions are currently weak because payment purpose is unavailable or unreliable?"

Instead of asking, "What should we do with ultimate-party elements?" ask:

"Where do intermediary or platform models prevent us from seeing the real payer or beneficiary?"

This creates a control hypothesis that can be tested. For example:

> **IF VERIFIED ULTIMATE-BENEFICIARY INFORMATION IS CONSISTENTLY AVAILABLE, SANCTIONS SCREENING SHOULD IDENTIFY RELEVANT EXPOSURE EARLIER AND REDUCE INVESTIGATIONS CAUSED BY SCREENING AN INTERMEDIARY AS THOUGH IT WERE THE END RECIPIENT.:**

The institution can then measure whether the field improved detection, reduced ambiguity or changed intervention outcomes.

From principle to practice

Create a payment-data council with representation from payments, product, operations, financial crime, privacy, engineering and data governance. Its mandate should be practical:

- define critical payment-data elements;
- agree ownership and acceptable sources;
- establish quality thresholds;
- approve permitted defaults and transformations;
- identify controls dependent on each element;
- review outcome evidence; and
- prioritise remediation according to decision impact.

This is not another governance forum created to admire a dashboard. It should own decisions about how payment information is collected, preserved and used.

Common failure modes

- Declaring success when the new message passes schema validation.
- Treating optional fields as irrelevant without assessing their control value.
- Creating enrichment that cannot be traced to a source.
- Measuring the sender’s message without testing what the receiver obtains.
- Mandating data without considering privacy, localisation or proportionality.

Questions for Leaders

1. Which payment fields are critical to sanctions, fraud and AML decisions?
2. Do we know the source and verification status of each field?
3. How much information is lost through internal or external translation?
4. Are data-quality metrics connected to control outcomes?
5. Who can require remediation when a technically valid field is operationally useless?

CHAPTER 3

Payment Transparency Is Becoming Enforceable

The industry is moving from permission to carry richer data toward an obligation to capture it, preserve it, test it and act on it.

> **EXECUTIVE INSIGHT:** A populated field is not evidence of transparency. Institutions increasingly need to show that the correct information survived the payment chain and influenced a defensible decision.

From capability to obligation

For years, richer payment data was discussed mainly as an opportunity. ISO 20022 could carry more structured information. Purpose codes could explain why a payment was being made. Ultimate-party fields could distinguish the economic actor from an intermediary. Legal Entity Identifiers could improve the identification of organisations.

The language was usually permissive: institutions were encouraged to populate the fields, explore the benefits and prepare for future adoption.

That period is ending. Regulators, standard setters and payment infrastructures are converting richer data from a technical capability into an operating obligation. The expectation is no longer simply that an institution can send a modern message. It is that the institution captures the right information, preserves it through the chain, verifies what matters and can demonstrate how the information affected the control outcome.

This changes the compliance conversation. A successful migration programme cannot end with schema validation, message acceptance or a dashboard showing that fields are populated. Those measures prove that the plumbing works. They do not prove that the control works.

The timetable is no longer theoretical

The direction of travel can be seen across several overlapping requirements:

- 30 December 2024 - Regulation (EU) 2023/1113 began applying. It requires specified information on payers and payees to accompany transfers of funds and extends related transparency obligations to certain crypto-asset transfers.

- 1 May 2025 - The Bank of England mandated purpose codes for defined CHAPS payments and Legal Entity Identifiers for payments between financial institutions. Non-compliance is addressed through the CHAPS rulebook, senior-management attestation and remediation planning.
- June and October 2025 - FATF revised Recommendation 16 and then published the methodology that will be used to assess compliance with the revised standard in mutual evaluations.
- 14 November 2026 - Swift will stop supporting fully unstructured postal addresses in CBPR+ payment messages. At a minimum, town and country must be supplied in designated fields; invalid instructions may be rejected before entering the network. CHAPS is aligned to reject fully unstructured addresses from the same period.
- By the end of 2030 - FATF expects the revised Recommendation 16 requirements to have taken effect.

These milestones do not create one identical global rulebook. They do create a consistent strategic message: payment information is becoming part of the enforceable control environment.

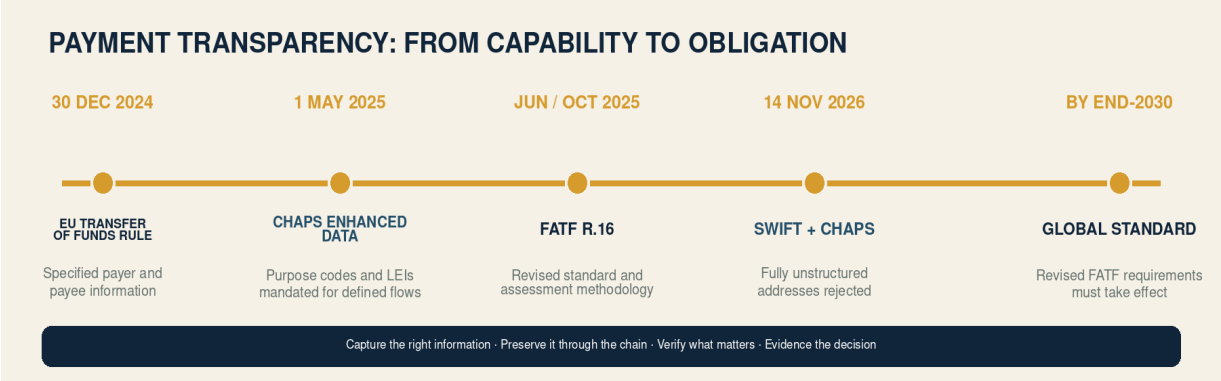

Figure 3.1 · Payment transparency: from capability to obligation

FATF changed the supervisory question

The 2025 revision of Recommendation 16 is important not only because it changes required information. It clarifies responsibility across the payment chain.

Under the revised standard, the payment chain begins with the financial institution that receives the customer instruction. FATF also clarifies responsibility for including required information and ensuring that it remains unchanged, standardises information requirements for certain peer-to-peer cross-border payments above USD or EUR 1,000 and requires tools intended to protect against fraud and error, including verification of recipient information.

The accompanying methodology matters just as much. It gives assessors a framework for examining how the revised requirements have been implemented. That moves the question beyond whether a jurisdiction or institution has reproduced the right words in law, policy or procedure.

The supervisory question becomes more practical: Did the correct parties enter the chain? Was the information preserved? Did verification prevent fraud or error? Could investigators obtain what they needed? What happened when the information was incomplete, inconsistent or wrong?

The standards do not make every participant omniscient. They make gaps, alterations and responsibility more visible.

Capture, preserve, verify and evidence

A mature payment-transparency control can be understood through four connected duties:

- Capture - obtain the required and decision-relevant information from the customer instruction or another governed source. The obligation begins with knowing what the field represents and who supplied it.
- Preserve - ensure information remains associated with the correct party as the instruction passes through channels, internal systems, intermediaries, translations, returns and recalls.
- Verify - apply the checks required by law, scheme rules and risk. Verification may include confirming recipient account information, validating identifiers and reconciling inconsistencies before value moves.
- Evidence - retain enough information to explain the source, transformation, control response, human involvement and outcome.

The four duties are interdependent. Verification cannot repair information that was never captured. Preservation of the wrong party is not transparency. Evidence that a field was populated is not evidence that it was useful.

A completed field is not evidence

Technical compliance is the beginning of the test, not the end. Payment information should pass through five levels of control quality:

- Technical validity - does the message conform to the relevant schema and network rules?
- Completeness - are the required fields present for the payment type and scenario?

- Accuracy - does the information describe the correct party, account, address, purpose or relationship?
- Context - can the institution understand the information in relation to the customer, product, corridor and payment journey?
- Decision utility - did the information change screening, verification, routing, enrichment, investigation, intervention or rejection?

A field can be perfectly formatted and perfectly wrong. Compliance teams have met this species before.

The evidence of effectiveness is therefore not the existence of the field. It is the difference the information made. Institutions should be able to show what was prevented, clarified, accelerated or escalated because the data was available and trustworthy.

Virtual accounts add another transparency layer

Virtual accounts deserve a specific caution, but they are not the central argument. A virtual identifier, the underlying customer or user and the settlement account may not be the same party or object.

- Visible identifier - what the payer or payment message may show.
- Underlying party - the customer or user associated with that identifier.
- Settlement account - where the funds ultimately settle.

The European Banking Authority's May 2024 report found divergent virtual-IBAN models and differences in how authorities interpreted and applied regulatory requirements. It highlighted issues involving money-laundering and terrorist-financing risk, authorisation, passporting, consumer and depositor protection and regulatory arbitrage.

The relevant compliance question is attribution: can the payment be linked to the correct underlying party throughout the chain, and can each institution understand which party it is expected to identify, screen or investigate?

Virtual accounts should therefore be treated as an additional transparency challenge within the payment architecture, not as a substitute for it. The EBA report is the sensible starting point for a deeper assessment of the particular model.

What supervisors can reasonably ask

As transparency requirements become more specific, institutions should expect questions that cross compliance, payments, product, operations and technology:

- Which legal, regulatory, scheme and network requirements apply to each material payment flow?
- Which information must be captured, and which additional fields are critical to the institution's own control decisions?
- Who owns each field, source, transformation and exception?
- How has the institution tested that party, address, purpose and identifier data survive the complete journey?
- What happens when information is missing, inconsistent, unverified or technically invalid?
- What evidence demonstrates that the information improved fraud, sanctions or AML outcomes?

A policy response alone will be insufficient. The evidence sits in messages, decision records, exception handling, remediation plans, testing results and actual payment outcomes.

From deadline management to control architecture

The safest response is not to treat each milestone as a separate compliance project. Institutions need one payment-transparency architecture capable of absorbing different requirements without rebuilding the control every time a rule, field or market changes.

That architecture should connect:

- an obligation-to-field map showing what each rule requires;
- a critical-data inventory showing source, provenance and owner;
- end-to-end journey testing across channels, products, intermediaries and returns;
- decision logic explaining which information affects screening, verification and intervention;
- exception rules for repair, enrichment, hold, rejection and escalation; and
- outcome measures showing whether the data improved the control.

A 2030 implementation date is not an invitation to discover in 2029 that a customer channel stores the address in one free-text field last redesigned when fax machines had opinions. The architecture work begins with current systems, current data and current payment journeys.

From principle to practice

Select one material cross-border payment flow and create a payment-transparency readiness record:

- identify the applicable obligations and implementation dates;

- map the customer instruction through every internal and external transformation;
- compare required fields with the information the institution actually captures;
- test whether immediate and ultimate parties remain correctly attributed;
- record the verification and recipient-checking controls;
- define what happens when information is missing or contradictory;
- measure which decisions changed because of the information; and
- assign an accountable owner and remediation date for every material gap.

The objective is not a beautiful inventory. It is a payment flow that can withstand operational failure, supervisory challenge and the occasional customer who enters 'N/A' with astonishing confidence.

Common failure modes

- Treating ISO 20022 adoption as proof that payment transparency has been achieved.
- Mapping regulatory obligations to policies without mapping them to fields, systems and decisions.
- Measuring field population without testing accuracy, provenance or preservation.
- Allowing intermediaries or internal transformations to replace the underlying party with the visible account holder.
- Preparing for the 2030 FATF deadline without addressing the 2026 network requirements.
- Documenting exceptions without changing the decision, remediation or product architecture.

Questions for Leaders

- Which payment-transparency obligations apply to our most material flows today?
- Can we prove that required information survives from customer instruction to beneficiary receipt?
- Which populated fields are technically valid but operationally unreliable?
- Can our controls resolve a virtual identifier to the correct underlying party?
- What decisions change when richer information is available?
- What evidence would satisfy a supervisor that the control works in practice?

CHAPTER 4

Identity Must Travel With the Money

Payment transparency depends on preserving the real parties through every layer of the transaction.

> **EXECUTIVE INSIGHT: An account name identifies the visible account holder. It does not always identify the person for whom the payment is being made.**

Immediate parties and ultimate parties

Modern payment chains involve more than a sender and receiver.

A marketplace may collect for a merchant. A payroll provider may pay employees for a corporate client. A payment platform may receive funds into virtual accounts on behalf of customers. A correspondent bank may process an instruction for another institution. A digital wallet may sit between the payer and the final beneficiary.

In these structures, the immediate debtor and creditor may not be the ultimate economic parties.

That distinction has always mattered to financial-crime compliance. Richer payment messages make it possible to represent it more clearly.

ISO 20022 includes elements capable of distinguishing immediate and ultimate parties in appropriate "on behalf of" scenarios. The Bank of England's enhanced-data requirements for CHAPS, for example, describe how ultimate debtor and ultimate creditor elements should be used when another party sits behind the immediate debtor or creditor.

The principle is simple: the payment chain should not replace the real party with the intermediary merely because the intermediary holds the account.

Identity is more than a name

Names remain important, particularly for sanctions screening. But a name alone is a weak identity architecture.

Different people share names. Organisations trade under brands. Characters are transliterated. Intermediaries populate their own address. Fraudsters use mule accounts with legitimate account-holder names.

A stronger payment identity can combine:

- verified name;
- account identifier;
- date or place of birth where appropriate;
- registered address;
- legal-entity identifier or registration number;
- country of incorporation or residence;
- role in the payment;
- relationship to the immediate account holder; and
- verification source.

Not all fields should travel in every payment. Privacy, proportionality and law matter. But the control should be able to resolve the visible party to a reliable identity.

Figure 4.1 · The identity chain

Identity must survive transformation

Payment data can degrade as it moves.

A platform may capture an underlying customer but fail to include it in the outbound message. An intermediary may translate the message into another format. A receiving institution may store only selected fields. A return payment may not preserve the original party information.

The result is identity decay: each participant sees a technically processable payment with less context than the participant before it.

Institutions should test identity continuity through real payment journeys. Synthetic test cases should include:

- direct account-to-account payments;
- on-behalf-of payments;
- virtual accounts;
- pooled and omnibus structures;
- cross-border intermediary chains;
- returns and recalls; and
- non-Latin names and addresses.

The test should compare what the initiator knows, what it sends, what each intermediary receives and what the beneficiary institution ultimately stores.

Verification and confidence

Payment identity should carry an evidential status.

Was the name entered by the payer? Retrieved from a verified customer profile? Confirmed through an account-name service? Supplied by an intermediary? Inferred from a narrative?

Controls can use these distinctions. A mismatch involving a verified beneficiary identity may warrant different treatment from a mismatch against unverified payer-entered text. An ultimate party supplied by a regulated partner under a governed arrangement may carry different confidence from a free-text reference.

This creates a move from binary data to confidence-aware data.

The payment does not merely say, “This is the beneficiary.” It allows the control to understand, “This beneficiary identity was verified through this source at this time.”

Transparency and privacy are joint design requirements

More identity data is not automatically better.

Payment information crosses borders, institutions and legal regimes. Unnecessary personal data can create privacy, security and misuse risk. Excessive requirements can also exclude customers whose local documents or address conventions do not fit a standard model.

The design question is therefore not “How much identity can we attach?” It is “What is the minimum reliable identity necessary for the participants to fulfil their responsibilities?”

That requires collaboration between compliance, privacy, security, legal and product teams. Privacy should not be invoked as a reason to preserve opacity, and compliance should not be invoked as a reason to collect everything.

From principle to practice

Create an ultimate-party transparency policy for products involving platforms, virtual accounts, pooled structures or intermediaries.

The policy should define:

- when immediate and ultimate parties differ;
- which party information must be captured;
- how it is verified;
- what travels externally;
- what can be resolved internally;
- how confidence is represented;
- how privacy requirements are applied; and
- when insufficient transparency prevents the payment or relationship.

The policy should be tested against actual messages, not only contractual descriptions.

Common failure modes

- Screening only the immediate account holder.
- Repeating the same party in immediate and ultimate fields.
- Assuming a regulated intermediary always knows its underlying user.
- Losing party information during translation or return processing.
- Treating unverified payer-entered names as verified identity.
- Collecting excessive identity data without a defined control purpose.

Questions for Leaders

1. Where do our products move money on behalf of another party?
2. Can sanctions and AML controls see the ultimate parties?
3. Which identity attributes are verified, supplied or inferred?
4. Does identity survive the complete payment journey?
5. Where does privacy law limit sharing, and what alternative control provides equivalent transparency?

CHAPTER 5

Context Is the Missing Control

The same payment can be ordinary or suspicious depending on the relationship and purpose behind it.

> **EXECUTIVE INSIGHT: Transaction monitoring frequently produces ambiguity because it observes movement without understanding the economic event.**

The limits of isolated attributes

Amount, country, counterparty and frequency are useful signals. None explains a payment on its own.

A large international transfer may represent a suspicious movement of funds, a property purchase, an intercompany treasury transfer or payment for a legitimate shipment. A first-time beneficiary may be a mule, a new supplier or an employee receiving expenses. Multiple payments of similar value may be structuring or the normal output of a marketplace.

Controls become stronger when they can compare the transaction with its declared and observed context.

Context includes:

- the purpose of the payment;
- the relationship between the parties;
- the customer's business model;
- the product and account configuration;
- expected geography and currency;
- invoice or remittance information;
- prior behaviour;
- device and initiation signals; and
- the wider network of related payments.

The objective is not to explain away unusual activity. It is to distinguish risk-relevant inconsistency from novelty.

Purpose should be testable

Purpose codes and remittance information can improve transparency, but only if they influence the decision.

A purpose code should be evaluated against what the institution knows. A payroll payment initiated from a business account should have characteristics consistent with payroll. A property-related payment should align with the customer, amount and counterparties. A transfer marked as internal should involve accounts under common control.

When purpose conflicts with behaviour, the conflict is a signal.

Examples include:

- “salary” payments sent to many unrelated businesses;
- “own account” transfers where ownership cannot be matched;
- “refunds” materially exceeding original purchases;
- “supplier” payments inconsistent with the customer’s industry;
- personal remittances routed through a commercial collection structure; and
- virtual accounts used for purposes outside the approved programme.

The control value comes from validation, not from accepting the label.

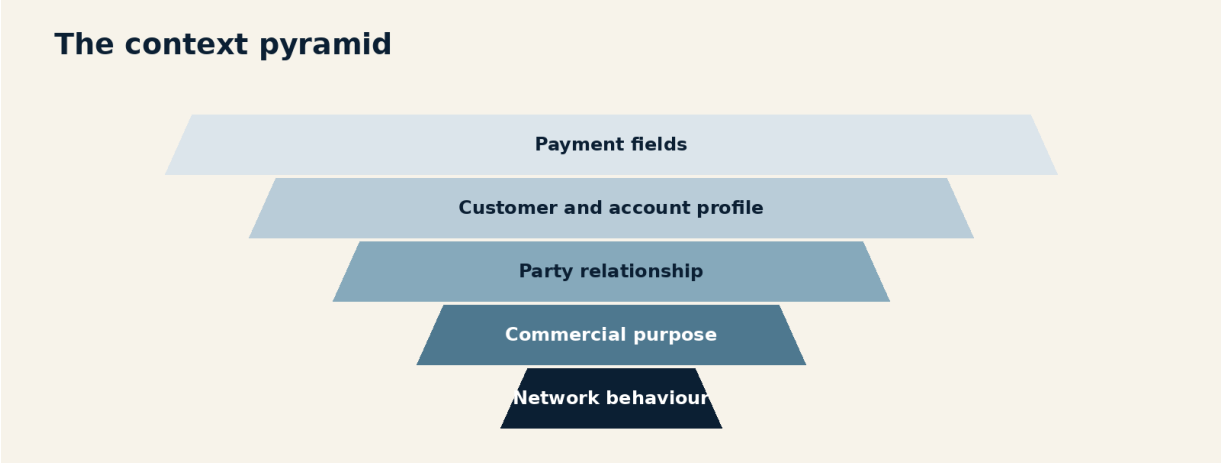

Figure 5.1 · The context pyramid

Context can reduce false positives without weakening controls

False-positive reduction is often framed as a tuning exercise: change thresholds, suppress known parties or add exceptions.

Context offers a different approach.

If a control knows that two accounts share verified ownership, it can evaluate a transfer differently without simply exempting it. If a virtual account is linked to a specific invoice and payer, the institution can distinguish the expected payment from unrelated credits. If a merchant's business model explains a seasonal increase, the model can incorporate the expected pattern while still monitoring counterparties and deviations.

This is more defensible than broad suppression because the decision is based on evidence.

It also improves investigations. Instead of presenting an analyst with an alert and asking them to search for context, the system can present the transaction, relevant customer facts, relationship evidence, purpose, inconsistencies and the reason the control escalated.

Context has a shelf life

Customer expectations change.

A business enters a new market. Ownership changes. A platform launches a new product. A customer begins using an account for a different purpose. Economic disruption changes payment patterns.

Context must therefore be refreshed through observed behaviour and customer events. Static onboarding declarations should not become permanent truth.

Adaptive controls treat expected activity as a living model:

1. begin with declared purpose and expected behaviour;
2. compare with early actual activity;
3. identify material differences;
4. seek explanation or evidence where necessary;
5. update the profile through governed decisions; and
6. retain the history of what changed and why.

This avoids two extremes: treating every deviation as suspicious or silently rewriting expectations to match whatever the customer does.

The shared-context problem

No single institution sees every part of a cross-border payment.

The originator institution knows the payer. The beneficiary institution may know the recipient. An intermediary sees the route. A platform may understand the commercial purpose. A fraud service may detect device compromise.

Better outcomes require governed forms of collaboration: structured payment data, requests for information, network utilities, shared typologies and public-private partnerships.

But collaboration must be designed. Institutions need clarity on which data can be shared, for what purpose, under which legal basis and with what assurance of accuracy.

The Financial Stability Board has identified divergence across payment, AML, sanctions and privacy data frameworks as a source of cross-border friction. The challenge is not merely to collect context. It is to make necessary context interoperable without turning the payment chain into an uncontrolled exchange of customer data.

From principle to practice

For each high-risk payment flow, define a minimum context standard:

- verified party identity;
- relationship between payer and beneficiary;
- payment purpose;
- expected account use;
- relevant invoice or transaction reference;
- geography and currency logic;
- initiation channel and device confidence; and
- programme or platform relationship.

Then identify which context is required to execute, which is required to control, which can be retrieved on demand and which should not be shared.

Common failure modes

- Treating customer-supplied purpose as verified fact.

- Suppressing alerts for expected activity without monitoring changes.
- Keeping product and commercial context outside the financial-crime data model.
- Requiring every institution to collect the same information independently.
- Sharing data without provenance, purpose limitation or quality controls.

Questions for Leaders

1. Which alerts exist primarily because the control lacks context?
2. Can we validate payment purpose against customer and counterparty information?
3. How are expected-activity profiles updated?
4. What context does another participant hold that would materially improve our decision?
5. What information should not travel with the payment, and how can it be retrieved safely when needed?

CHAPTER 6

Screening Moves Upstream

The best alert is often the payment that was safely redirected, paused or prevented before settlement.

> **EXECUTIVE INSIGHT: Faster payments require earlier decisions. A control that operates after irrevocable settlement may be informative without being preventive.**

From retrospective detection to intervention

Traditional monitoring is frequently post-transaction. The payment settles, activity is aggregated and an alert is generated according to a schedule.

That remains important for pattern detection and regulatory reporting. But it is insufficient for payment environments in which funds can move and disperse immediately.

Upstream control begins before payment initiation.

Customer permissions, account configuration, virtual-account restrictions, beneficiary setup and transaction limits all influence which payments can occur. The payment message then supplies information for screening and real-time decisioning. Post-transaction monitoring examines the wider pattern and feeds learning back into earlier stages.

The control system is therefore a continuum:

1. **Design-time controls** define permitted products and use cases.
2. **Issuance controls** determine who receives accounts and capabilities.
3. **Pre-initiation controls** evaluate access, device, beneficiary and intent.
4. **In-flight controls** screen parties, message data and risk signals.
5. **Post-event controls** identify patterns and emerging typologies.
6. **Feedback controls** change permissions, models and data requirements.

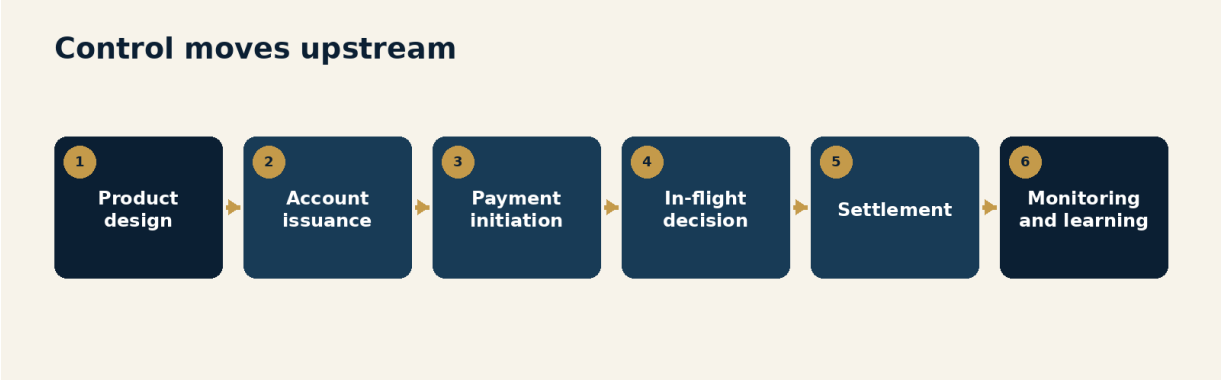

Figure 6.1 · Control moves upstream

The payment message as a decision surface

An enriched message can support multiple actions, not merely approve or decline.

Depending on risk, law and scheme rules, a control might:

- allow the payment;
- request missing or corrected information;
- require stronger customer authentication;
- reduce the permitted amount;
- route the payment for review;
- hold or delay execution where legally permitted;
- reject the instruction;
- block or freeze property as required;
- restrict the account or beneficiary; or
- permit the payment while increasing subsequent monitoring.

This range matters because binary controls create pressure toward either excessive friction or excessive tolerance.

Adaptive decisioning uses proportional responses tied to confidence and consequence.

Sanctions, fraud and AML see the same payment differently

A sanctions control asks whether a prohibited party, jurisdiction or interest is involved. A fraud control asks whether the instruction was authorised and whether the beneficiary or behaviour indicates deception. An AML control asks whether the payment forms part of suspicious activity.

These objectives differ, but the information overlaps.

A name mismatch may indicate sanctions evasion, account takeover or innocent data quality. A newly created beneficiary in a high-risk corridor may be a fraud signal, an AML signal or both. A virtual account receiving unrelated third-party payments may indicate misuse of the programme, mule activity or laundering.

Separate tools can remain necessary, but the decision architecture should reconcile their outputs.

The payment should not be approved by one control, rejected by another and queued indefinitely by a third while each system remains blissfully unaware of the others.

Latency becomes a control metric

Institutions traditionally measure alert volumes, false positives and investigation time. Faster payments require additional measures:

- time from risk signal to decision;
- percentage of relevant decisions made before settlement;
- time to propagate restrictions across channels and accounts;
- time to incorporate confirmed fraud or sanctions intelligence;
- value prevented or recovered;
- customer friction created by intervention; and
- quality of explanations for automated decisions.

The objective is not the fastest possible decision. It is a sound decision within the time available.

Designing safe fallback

Real-time controls will occasionally lack information or experience a service failure.

The fallback design is part of the control. Institutions should determine:

- which payments may proceed under degraded conditions;

- which require a hold or alternative screening path;
- how queues are prioritised;
- how long decisions remain valid;
- how manual overrides are governed; and
- how missed controls are identified and remediated.

“The screening service was unavailable” is an incident description, not a risk decision.

From principle to practice

Map each material financial-crime decision to the earliest point at which sufficient information exists.

If the risk can be controlled at product design, do not wait for monitoring. If it can be controlled at account issuance, do not wait for the first payment. If it requires network behaviour, preserve the post-transaction control but feed the outcome back into permissions and real-time decisioning.

Create a decision record containing:

- information considered;
- control or model version;
- confidence and material uncertainty;
- action taken;
- legal or policy basis;
- human involvement;
- customer impact; and
- subsequent outcome.

This creates the evidence needed to learn.

Common failure modes

- Applying real-time technology to a decision that still depends on manual data retrieval.
- Treating all uncertainty as a decline.
- Allowing product permissions and compliance controls to conflict.

- Measuring screening latency without measuring decision quality.
- Building no governed fallback for outages or incomplete data.
- Failing to feed confirmed outcomes back into upstream controls.

Questions for Leaders

1. Which financial-crime decisions currently occur after funds are irretrievable?
2. What earlier control could prevent or contain the risk?
3. Can sanctions, fraud and AML outputs be reconciled in one payment decision?
4. What happens when critical data or screening services are unavailable?
5. Do we measure prevention and intervention, or only alerts and investigations?

CHAPTER 7

AI Learns to Read the Payment

Artificial intelligence can connect the fields, relationships and history around a payment, but authority must remain governed.

> **EXECUTIVE INSIGHT: The value of AI is not that it can read more payment data. It is that it can identify which parts of the story do not agree.**

From matching to interpretation

Traditional screening compares selected fields with lists, rules or thresholds. Machine learning expands the pattern recognition. Generative and agentic systems add another capability: they can assemble information, describe relationships and explain inconsistencies in language that investigators and decision-makers can understand.

For a payment, an AI system might compare:

- the parties named in the message;
- verified customer and beneficial-owner information;
- virtual-account purpose and permissions;
- historical counterparties;
- invoice or remittance details;
- device and initiation data;
- sanctions and adverse intelligence;
- network connections; and
- similar confirmed cases.

The system can then explain that the payment is unusual not simply because it is large, but because the stated purpose conflicts with the customer's business, the beneficiary is new, the virtual account was issued for collections rather than outbound transfers and related accounts show similar routing.

That is closer to reasoning about the payment than matching a single field.

The explanation must remain evidence-based

AI-generated narratives can be persuasive. Persuasion is not proof.

Every material statement should be traceable to evidence. The user should be able to distinguish:

- facts retrieved from verified records;
- information supplied by the customer;
- external intelligence;
- model inferences;
- unresolved uncertainty; and
- the AI system's summary.

This is particularly important when AI compresses a complex payment journey into a short explanation. The summary may be useful while omitting contradictory evidence or overstating confidence.

The system should show its work without requiring the investigator to become a data engineer.

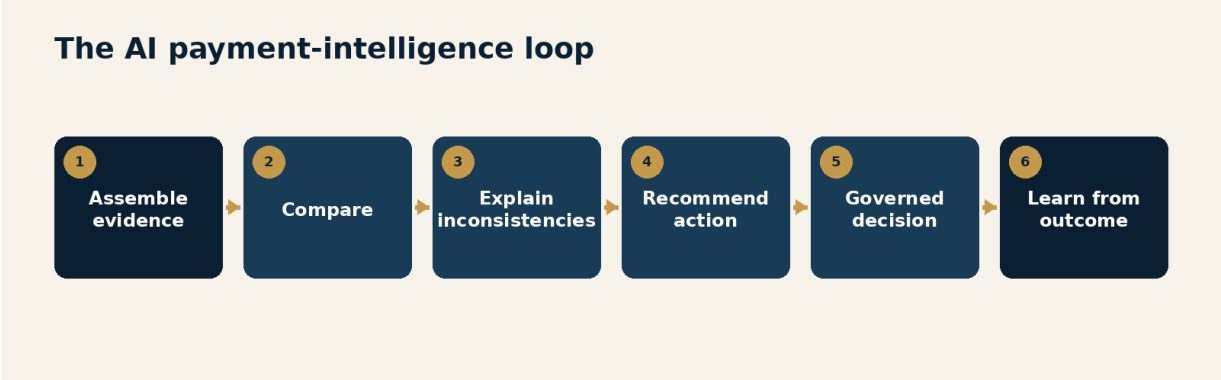

Figure 7.1 · The AI payment-intelligence loop

Authority should follow consequence

Not every AI use case requires the same oversight.

An AI system that organises evidence for an investigator carries different risk from one that recommends closing an account. A system that requests a missing address differs from one that blocks property. A model that prioritises a queue differs from one that executes a payment decision.

Governance should classify use cases according to:

- decision authority;

- financial and customer consequence;
- legal significance;
- reversibility;
- data sensitivity;
- model uncertainty; and
- availability of effective human review.

The phrase “human in the loop” is not enough. Leaders should know which human, reviewing what evidence, within what time, with what authority and under what workload.

AI agents and the message

The next stage is not limited to AI analysing payments. AI agents may initiate them.

An agent could select a supplier, agree terms, create an instruction and authorise payment within delegated limits. This creates new identity questions:

- Which person or organisation authorised the agent?
- What scope and spending limits apply?
- How is the agent identified in the payment information?
- Can the receiving institution distinguish agent-initiated activity?
- What happens when the agent’s action conflicts with the principal’s intent?
- Who is responsible for monitoring changes in the agent’s behaviour?

The payment message may need to represent both the legal principal and the initiating agent or service. Otherwise, institutions may see a valid customer payment without understanding how the decision was made.

Agentic payments make the principle of provenance even more important. The control must know not only who owns the account, but what actor initiated the instruction and under which authority.

Continuous model management

AI operating on payment data will drift as customers, products, fraud methods, sanctions risks and message quality change.

Model management should monitor:

- changes in field availability and population;
- changes in customer and corridor behaviour;
- performance across languages and geographies;
- explanation quality;
- false-negative and false-positive indicators;
- investigator agreement and override;
- customer impact;
- incidents and near misses; and
- performance after material product or policy changes.

Validation remains important, but periodic validation cannot be the only mechanism. A model interpreting live payment behaviour begins encountering new conditions immediately.

From principle to practice

Begin with bounded use cases where AI improves information quality and human decisioning:

- detecting missing or inconsistent message information;
- assembling payment and customer context;
- explaining why multiple signals conflict;
- prioritising cases;
- drafting targeted requests for information; and
- identifying clusters for investigator review.

Expand authority only when evidence demonstrates reliable performance, effective oversight and controlled consequences.

Every AI-supported decision should retain:

- the model and version;
- the input evidence;
- the output and explanation;

- the action taken;
- human review or override;
- the final outcome; and
- feedback used for improvement.

Common failure modes

- Treating an eloquent explanation as a correct explanation.
- Hiding source uncertainty inside a single risk score.
- Giving AI authority before establishing outcome evidence.
- Using “human review” as a label without workload or decision design.
- Failing to capture that an AI agent initiated the payment.
- Validating the model while ignoring deterioration in payment-data quality.

Questions for Leaders

1. Which AI statements can be traced directly to evidence?
2. Where does the system infer rather than know?
3. What decisions can the AI recommend, and which can it execute?
4. How will an agent-initiated payment identify both agent and principal?
5. What evidence would cause us to reduce or withdraw model authority?

CHAPTER 8

The Message Becomes the Control

The mature organisation turns payment information into a continuously improving system of prevention, decision and learning.

> **EXECUTIVE INSIGHT: Rich payment data creates strategic value only when it changes decisions, improves outcomes and teaches the organisation what to do next.**

From message migration to control transformation

The industry has reached an important milestone. Major payment infrastructures have adopted richer structured standards. Cross-border transparency expectations are strengthening. Institutions can carry more information and process it more consistently than legacy messages allowed.

But migration alone does not create effectiveness.

The strategic opportunity is to connect five capabilities:

1. **Identity:** reliable understanding of the parties and their roles.
2. **Purpose:** evidence of the economic reason for the payment.
3. **Context:** customer, account, product and network information.
4. **Decisioning:** proportionate action within the available time.
5. **Learning:** outcomes that improve data, permissions and controls.

When these capabilities operate together, the payment message becomes an active control surface.

The Payment Message Control Maturity Model

Institutions can assess maturity across five stages.

Stage 1: Transport

The message carries the minimum information necessary to execute and reconcile the payment. Financial-crime controls operate largely outside the message and after initiation.

Success is measured through processing accuracy, alert completion and operational compliance.

Stage 2: Structured

The institution adopts defined fields and common standards. Data is machine-readable, but quality remains focused on technical validity and population.

Success is measured through schema compliance, field completion and reduced repair.

Stage 3: Enriched

Payment data is connected with verified identity, underlying-party attribution, purpose and customer context. Provenance and confidence are visible.

Success is measured through improved attribution, fewer information gaps and better investigative decisions.

Stage 4: Decisioning

Structured and enriched information supports real-time sanctions, fraud and AML action. Controls can apply proportionate interventions before or during execution.

Success is measured through prevention, decision quality, latency, explainability and customer impact.

Stage 5: Adaptive

Investigation outcomes, confirmed fraud, sanctions intelligence, regulatory change and model performance continuously improve message requirements, account permissions and decision logic.

Success is measured through learning speed, resilience and control effectiveness.

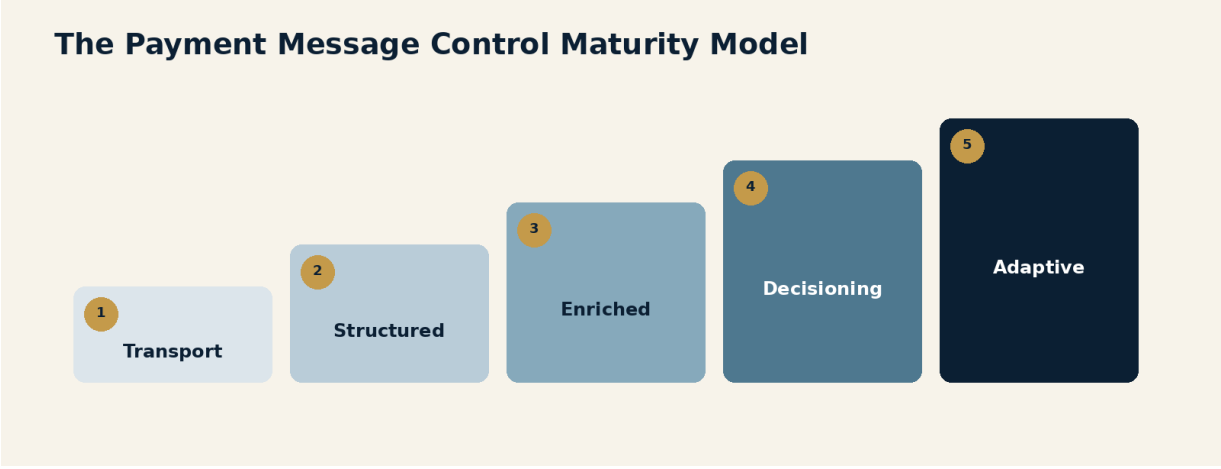

Figure 8.1 · The Payment Message Control Maturity Model

Maturity is uneven by design

An institution will not sit neatly at one stage.

Sanctions screening may operate in real time while AML monitoring remains post-event. Underlying-party data may be highly reliable in one product while another relies on sparse references. One payment rail may preserve ultimate-party information while another truncates it.

The maturity model should therefore be applied by payment flow and capability.

Leaders should identify:

- the current stage;
- the target stage justified by risk and use;
- the information or governance gap;
- the outcome expected from improvement; and
- the evidence that will demonstrate progress.

Not every payment requires Stage 5 sophistication. The goal is appropriate maturity, not technological maximalism.

The operating model

The message becomes the control only when ownership crosses traditional boundaries.

Payments teams own execution and network requirements. Product teams own customer journeys and capabilities. Compliance defines risk outcomes and legal obligations. Fraud teams understand initiation and deception. Data teams govern information. Engineering builds decision paths. Operations manages exceptions. Privacy and security protect the data.

The operating model should include:

- a payment-data council;
- a cross-functional control owner for each material flow;
- a critical-data-element inventory;
- decision and model governance;
- end-to-end testing;
- outcome-based metrics;
- incident learning; and
- a roadmap connecting regulatory change with product and infrastructure change.

Measuring what the message creates

Traditional migration metrics remain necessary: message acceptance, repair rates, field population and operational stability.

Control metrics should go further:

- verified ultimate-party coverage;
- underlying-party attribution across virtual, pooled and intermediary structures;
- purpose-data accuracy;
- information preserved end to end;
- decisions made before settlement;
- preventable losses or exposures avoided;
- reduction in investigation time due to available context;
- overrides and decision reversals;
- customer friction;
- time from confirmed typology to control change; and
- percentage of control changes supported by outcome evidence.

These measures shift the conversation from “Did we populate the field?” to “Did the information improve the decision?”

Trust travels with the payment

Payment innovation is frequently described through speed, cost and availability.

Those benefits matter. But the ability to move money faster is valuable only when institutions can trust the parties, purpose and controls surrounding the movement.

Richer messages can reduce ambiguity. Strong attribution can preserve the real parties through virtual, pooled and intermediary structures. Upstream screening can prevent harm. AI can assemble and explain evidence. None of these operates safely without governance, provenance and learning.

The message becomes the control when it carries or unlocks the information required to decide, when that decision occurs in time to matter and when the outcome improves the next payment.

That is adaptive compliance applied to the movement of money.

From principle to practice

Leaders can begin with four actions:

1. Select one material payment flow and map its information end to end.
2. Identify the three information gaps creating the greatest control ambiguity.
3. Move one material decision to the earliest safe point in the lifecycle.
4. Establish an outcome measure and use it to govern the next change.

The work does not require a grand transformation programme before it can begin. In fact, any transformation programme requiring its own transformation programme should probably be treated as an emerging risk.

Common failure modes

- Declaring the payment-message programme complete at technical migration.
- Pursuing maximum data without a defined decision purpose.
- Building real-time decisions without fallback, explanation or appeal.
- Treating virtual-account, KYC, fraud, sanctions and AML data as separate domains.
- Measuring activity rather than effectiveness.
- Learning from incidents without changing the information or decision architecture.

Questions for Leaders

1. At what maturity stage is each material payment flow?
2. Which information gap creates the greatest current risk?
3. What decisions can be moved earlier without disproportionate friction?
4. Do our metrics demonstrate better outcomes or simply more complete messages?
5. How quickly does a confirmed case change what happens to the next payment?

Toolkit

The Payment Intelligence Operating Model

This toolkit translates the volume's central argument into four practical artefacts.

1. Payment-Information Inventory

For each material flow, document:

- product and payment rail;
- customer and end-user structure;
- account and virtual-account model;
- immediate and ultimate parties;
- purpose and remittance information;
- source and verification status of critical fields;
- transformations across the chain;
- sanctions, fraud and AML controls;
- decision timing;
- outcome and feedback mechanisms; and
- accountable owner.

2. Payment Transparency Readiness Record

For each material payment flow, retain:

- applicable legal, regulatory, scheme and network requirements;
- originating customer instruction and capture channel;
- required and decision-critical party, address, purpose and identifier fields;
- source, verification status and provenance for each critical field;
- rules for preservation, translation and transformation;
- evidence of what downstream institutions receive;

- recipient-verification and fraud-prevention controls;
- repair, enrichment, hold, rejection and escalation rules;
- accountable owners and remediation dates;
- testing results and control exceptions;
- decision and customer outcomes; and
- approval, challenge and review dates.

3. Payment Decision Record

For automated or materially assisted decisions, capture:

- payment and customer identifiers;
- information considered;
- source and confidence;
- control or model version;
- sanctions, fraud and AML outputs;
- action and timing;
- legal or policy basis;
- human involvement and override;
- customer impact;
- investigation outcome; and
- feedback applied.

4. Executive Scorecard

The executive scorecard should balance five categories:

Information

- critical-field completeness and accuracy;
- verified ultimate-party coverage;
- underlying-party attribution across virtual, pooled and intermediary structures;

- end-to-end preservation.

Decision

- pre-settlement decision coverage;
- latency by risk and action;
- overrides and reversals;
- explanation quality.

Outcome

- prevented or contained exposure;
- confirmed fraud and suspicious-activity yield;
- customer friction;
- investigation time.

Resilience

- service degradation and fallback performance;
- time to propagate restrictions;
- data-quality incidents;
- control coverage during outages.

Learning

- time from outcome to control change;
- typology adoption;
- model and policy drift;
- repeat incidents.

References and Further Reading

1. Financial Action Task Force, “FATF updates Standards on Recommendation 16 on Payment Transparency,” 18 June 2025. Official source

2. Financial Action Task Force, *The FATF Recommendations*, updated October 2025, Recommendation 16. Official source

3. Committee on Payments and Market Infrastructures, *Harmonised ISO 20022 Data Requirements for Enhancing Cross-Border Payments*, updated February 2026. Official source

4. Committee on Payments and Market Infrastructures, *Navigating the ISO 20022 Migration Journey*, April 2026. Official source

5. Swift, “ISO 20022: A New Era for Global Payments,” 25 November 2025. Official source

6. Swift, “ISO 20022 Milestone for November 2026: Unstructured Addresses to Be Removed,” 25 March 2026. Official source

7. European Banking Authority, *Report on Virtual IBANs*, May 2024. Official source

8. Bank of England, “Mandating ISO 20022 Enhanced Data in CHAPS,” April 2024. Official source

9. Financial Stability Board, *Recommendations to Promote Alignment and Interoperability Across Data Frameworks Related to Cross-Border Payments*, December 2024. Official source

10. Wolfsberg Group, *Payment Transparency Standards*, October 2023. Official source

11. Wolfsberg Group, *Guidance on Payment Transparency - Roles and Responsibilities*, December 2024. Official source

12. Board of Governors of the Federal Reserve System, “Payment Innovation and the Information Value of ISO 20022,” remarks by Governor Christopher J. Waller, May 2024. Official source

13. European Union, Regulation (EU) 2023/1113 on information accompanying transfers of funds and certain crypto-assets, applicable from 30 December 2024. Official source

About the Author

Micheal Sheehy is a global financial-services and payments compliance executive with experience leading large-scale transformation across anti-money laundering, sanctions, customer lifecycle management, fraud, transaction monitoring, data, technology and model governance.

His work focuses on the future of adaptive compliance: how financial institutions can combine better information, modern technology, accountable governance and organisational learning to improve control effectiveness while enabling responsible growth.

He is the author of *The Adaptive Compliance Series*.

Contact: micheal@michealsheehy.com

Website: https://michealsheehy.com