ANT GROUP | DIGITAL TECHNOLOGIES

AGENTIC PAYMENTS IN WHOLESALE FINANCE

Exploring use cases
value potential
and conditions for responsible adoption

THOUGHT LEADERSHIP

Deutsche Bank and **Ant Digital Technologies**

October 2026

AGENTIC PAYMENTS IN WHOLESALE FINANCE

Exploring use cases, value potential, and conditions for responsible adoption

Deutsche Bank and **Ant Digital Technologies**

October 2026

Foreword I:

A banking perspective



Sabih Behzad
Head of Digital Assets and Currencies Transformation

Deutsche Bank

The next evolution in payments will not be defined only by faster rails or new forms of digital money.[1] It may also be shaped by software agents capable of interpreting objectives, assessing alternatives, and initiating payment-related actions. For banks, corporates, and financial institutions, this raises a fundamental question: how can selected payment-related decisions be delegated to software while authority, control, and accountability remain within the institution?

Wholesale payments already benefit from extensive automation, but many of the decisions surrounding them remain fragmented across systems, teams, and operating windows. Treasury, cross-border payments, and trade finance require institutions to interpret changing information, assess alternatives, and coordinate actions under demanding regulatory and operational conditions. Agentic capabilities could help address this decision layer by allowing selected payment-related choices to be delegated to software within explicitly defined limits.

Applying agentic capabilities to wholesale finance presents a different order of opportunity and complexity from retail commerce. Institutional transactions involve material values, multiple parties and systems, and decisions made under changing liquidity, market, and compliance conditions. The potential value lies not simply in executing payments automatically, but in improving how those decisions are made, coordinated, and controlled.

This paper focuses on three priority use cases. In **treasury liquidity management**, agents could anticipate liquidity shortfalls or surpluses and assess permitted responses across internal transfers, investment, external funding, and foreign exchange. In **cross-border payment orchestration**, agents could determine how best to fulfil a client's payment objective by coordinating foreign exchange (FX), liquidity or credit, service providers, timing, and settlement options. In **trade finance**, agents could interpret transaction information, identify discrepancies, and orchestrate documentary checks, control, and exception workflows before financing or payment is released. These are value hypotheses to be assessed, not established outcomes. Their potential depends on the structure and the quality of the underlying data, integration, and institutional control environment.

The opportunity should therefore not be understood as a pursuit of maximum autonomy. In wholesale finance, trust, predictability, legal certainty, and accountability are essential. Agents may provide judgement and orchestration, but independent controls must determine whether an action is authorised and may proceed. Authority must remain explicit, bounded, and revocable, while high-value, high-risk, and exceptional transactions should retain appropriate human involvement.

Banks have a particular role in making this model operational. They connect the client and account relationship with liquidity, regulated controls, payment execution, and recourse across an increasingly diverse set of payment and settlement rails. Agentic payments may expand how decisions are made and coordinated, but they do not diminish the obligations associated with moving institutional money.

Realising the potential of agentic payments will require collaboration among financial institutions, technology providers, payment and settlement infrastructures, policymakers, regulators, and clients. Progress will depend on common approaches to identity, mandates, controls, evidence, and interoperability that enable agentic capabilities to operate safely across institutional boundaries.

This paper is intended as a practical foundation for further dialogue. We invite industry participants to assess its value hypotheses, challenge its proposed approaches, and work collectively towards agentic payment models that combine innovation with the trust, resilience, and accountability expected of wholesale finance.

Foreword I: A banking perspective

Foreword II:

The technology perspective



Zhuoqun Bian
President of Blockchain Business

Ant Digital Technologies

Two technologies are converging to reshape financial markets. **Artificial intelligence (AI)** is moving from systems that generate information to agents that act: they coordinate services, negotiate terms, and initiate transactions. **Blockchain** is moving money and assets onto programmable ledgers, where value is verifiable, carries its own rules and settles continuously. Together, they are laying the foundations of the token economy – an on-chain economy in which software can hold, move, and settle value within rules that institutions define.

In this economy, **agentic payments** are not simply a new payment channel. They are a core layer of payment and settlement infrastructure: the point at which an agent's decision becomes a transfer of tokenised money or assets. How that layer is designed will shape the digital financial markets built on top of it.

This is the infrastructure Ant Digital Technologies builds: enterprise blockchain through AntChain, institutional-grade settlement through Jovay, and tokenisation and agent-to-agent payments through Anvita. That work has taught us that AI and blockchain cannot be designed separately. An agent is only as trustworthy as the assets it moves and the ledger on which it settles.

Realising this potential in wholesale finance depends on foundations that are neither optional nor purely technical. Agents need verifiable identity, bound to an accountable owner. Their authority must be explicit, machine-readable, bounded, and revocable. Decision-making, control, and execution must be separated, so that an agent may propose but cannot extend its own authority. Agents, systems, and ledgers must interoperate across institutional boundaries. And settlement must fit the transaction: rapid confirmation for high-frequency, low-value activity, and stronger finality where value is material.

Institutional adoption remains at an early stage. Performance, security, scalability, and control must be demonstrated through practical testing, and each extension of agent authority earned as evidence accumulates. Banks bring the obligations that come with moving money: accountability, compliance, resilience, and recourse. Technology providers bring an understanding of how agents, blockchains, and programmable assets behave at scale. Neither perspective alone describes what institutional adoption requires end to end.

The token economy will not be built by any single institution. We invite financial institutions, infrastructure providers, standards bodies, and regulators to build its payment and settlement layer with us, so that agent-payment models are open, trusted, and interoperable by design.

Authors and contributors

Authors

Ivy (Ming) Schuhmacher, Digital Assets and Currencies Transformation, Deutsche Bank

Jonathan Au-Yeung, Head of Product Solutions (Anvita), Ant Digital Technologies

Contributors

Deutsche Bank

- **Digital Assets and Currencies Transformation**: Thomas Brophy
- **Cash Management Product**: Marvin Ulbrich, Harrison Tan, Jeffrey Smeeton
- **Corporate Treasury Services:** John Olatunji, Hyesi Jun
- **Trade Finance:** Anup Parmar
- **Technology**: Peter Hung
- **Editorial Marketing:** Will Monroe, Desirée Buchholz

Ant Digital Technologies

- **Strategic Business and Solutions**: Bo Wang, Alexandra Wang
- **Web 3 Product**: Cobe Zhang, Edward Liu

CONTENTS

Executive summary 01

1. From automation to agency: what are agentic payments? 02

- 1.1. Introduction 02
- 1.2. Why do agentic payments matter for wholesale finance? 03
- 1.3. Definition 04

2. Unlocking value across wholesale finance 06

- 2.1. Treasury liquidity management 06
- 2.2. Cross-border payment orchestration 08
- 2.3. Trade finance 11
- 2.4. How tokenised settlement could enhance the value hypotheses 13
- 2.5. Institutional capabilities required to enable agentic payments 14

3. How wholesale agentic payments could work 15

- 3.1. Reference architecture 16
- 3.2. Recommended design choices for wholesale agentic payments 17
- 3.3. Settlement rails and assets 18
- 3.4. Institutional roles and responsibilities 19
- 3.5. Agent lifecycle and control gates 20
- 3.6. Case study: Anvita Flow as an implementation example 21

4. Conditions for responsible adoption 24

- 4.1. Regulatory and legal clarity 24
- 4.2. Industry standards and interoperability 24
- 4.3. Technology and operational resilience 25
- 4.4. Institutional control and risk management 26

5. A progressive pathway to institutional adoption 27

Executive summary

This paper defines an **agentic payment** as a payment in which a software agent interprets information and selects or initiates a transaction under an explicit, bounded, and revocable mandate delegated by a legally accountable institutional principal. Unlike deterministic automation, which follows pre-scripted instructions, an agent can respond to changing circumstances and select from permitted alternatives. The defining characteristic is therefore **delegated decision authority**, not automated execution alone.

Wholesale applications require particular consideration because institutional transactions involve material values, complex processes, and extensive regulatory and operational requirements. Agentic payments in wholesale finance must therefore combine adaptive decision-making with predictable execution, legal certainty, and institutional accountability.

The paper examines three priority use cases:

- **Treasury liquidity management**, where agents could anticipate liquidity shortfalls and surpluses and assess permitted responses across internal transfers, investment, external funding, and foreign exchange.
- **Cross-border payment orchestration**, where agents could evaluate interdependent decisions across FX conversion, payment rails, and execution timing, based on cost, speed, availability, regulatory requirements, and settlement risk.
- **Trade finance**, where agents could interpret transaction data and documents, identify discrepancies, orchestrate controls and exceptions, and prepare financing or payment instructions once the required conditions and approvals are satisfied.

The central value hypothesis is that agentic payments could improve decision quality and responsiveness, reduce manual hand-offs, increase straight-through processing (STP), and support more efficient use of liquidity and payment infrastructure. Tokenised settlement is not a prerequisite, but where it is available, continuous availability, atomic execution, and programmable transfer conditions could provide additional value. These benefits remain to be demonstrated through controlled implementation.

This paper applies a three-layer architecture comprising **intent and orchestration**, where the agent interprets objectives and proposes an action; **control and authorisation**, where independent safeguards determine whether the proposed action may proceed; and **settlement**, where authorised instructions are executed with finality. The core design principle is that the agent may provide judgement, while independent institutional controls provide permission. Authority must be explicit, bounded, and revocable; controls must assess current conditions at execution; and decisions and transactions must be auditable.

Responsible adoption will require legal and regulatory certainty, interoperable standards, operational resilience, and effective risk governance. Supporting capabilities include trusted data, machine-readable mandates, secure agent lifecycle management, immediate suspension, clear accountability and recourse, and sufficient evidence to reconstruct decisions, controls, and outcomes.

Near- to medium-term adoption should focus on **AI-assisted decision-making** and **bounded agentic execution**, rather than unconstrained autonomy. Institutions should begin with narrowly defined use cases, retain human approval for high-value, high-risk or exceptional transactions, and expand delegated authority only when evidence demonstrates business value, reliable performance, and effective control. The objective is not maximum autonomy, but the appropriate degree of agency for each use case.

1. From automation to agency: what are agentic payments?

1.1. Introduction

AI agents are extending payment automation beyond the execution of predefined instructions. They can interpret objectives, assess relevant information, and select from among permitted actions. This creates the potential to delegate selected payment-related decisions to software while retaining institutional authority, control, and accountability.

Most existing research focuses on retail commerce or e-commerce use cases. It has examined whether AI agents could purchase goods and services on a user's behalf,[2] whether broad user mandates provide valid authorisation for individual transactions,[3] and how agent identity, intent, liability, and control should be addressed.[4] Their practical implication is that agent-mediated payments require verifiable authority, bounded permissions, and an auditable connection between the institutional principal's mandate and the resulting transaction.

Two recent papers are particularly relevant to wholesale payments. The Bank for International Settlements (BIS) tested an AI agent in simulated wholesale real-time gross settlement (RTGS) scenarios and found that the agent could produce recommendations reflecting key intraday liquidity management practices, such as maintaining liquidity buffers, prioritising urgent payments, and balancing liquidity usage against settlement delays. The findings indicate potential for agentic decision-making support in wholesale payments, while also underlining the need for safeguards, oversight, and further research before production deployment.[5] The International Monetary Fund (IMF) has proposed separating an agent's adaptive decision-making from the deterministic layers responsible for authorisation, control, and settlement. The practical implication is that institutions should allow agents to act only within verifiable mandates, while retaining independent controls over credentials, execution, auditability, and human oversight.[6]

However, direct research on wholesale agentic payments remains limited. Existing studies do not yet provide an integrated assessment of where agentic decision-making could create value across wholesale finance activities, how agents should interact with banking and settlement infrastructure, or which governance and control structures would be required for institutional adoption.

This paper seeks to address that gap by identifying priority wholesale use cases, assessing their value potential, and proposing a progressive pathway towards responsible institutional adoption.

1.2. Why do agentic payments matter for wholesale finance?

Agent-mediated payments have already reached live use in retail. Mastercard completed its first live agentic transaction in Hong Kong in April 2026, and Visa enabled agent-initiated purchases at European merchants from July 2026.[7] In March 2026 Ant Digital Technologies launched Anvita Flow, an agent-to-agent network in which agents discover services, negotiate terms, and settle in real time; services are listed and certified in an Agent Store, and settlement uses the open x402 standard so that an agent can pay for a service inside the request itself, in denominations below one cent.[8]

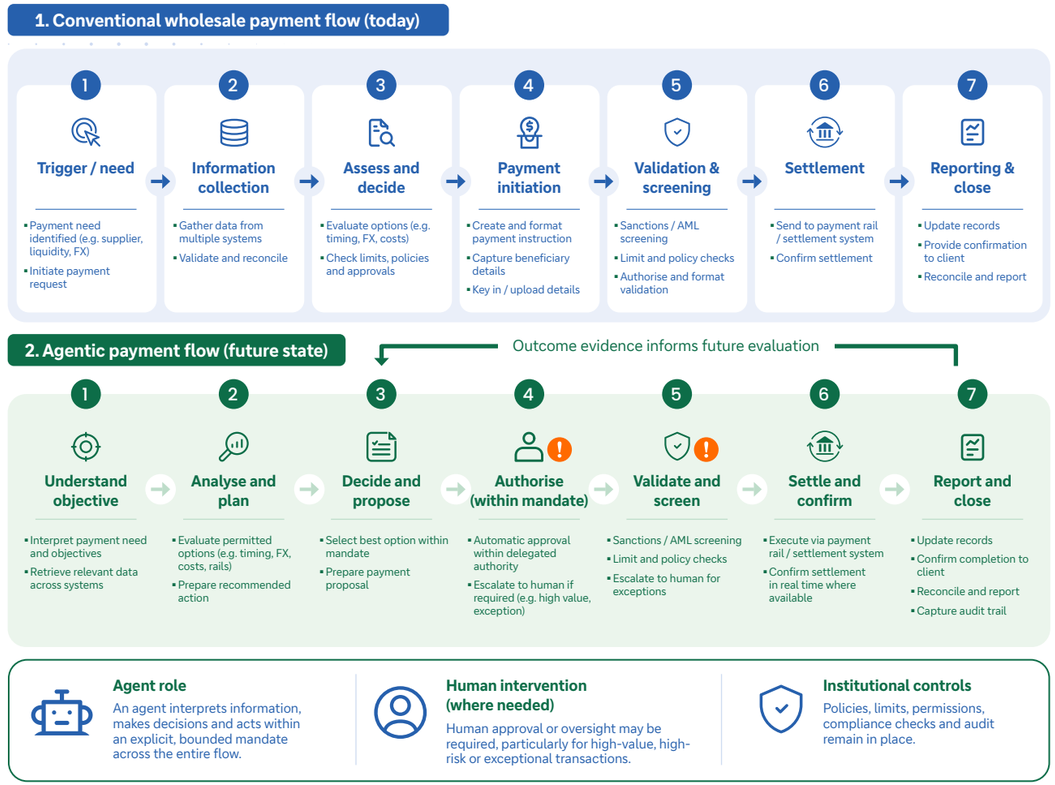
Although wholesale transactions differ materially from retail payments in their value, complexity, regulatory requirements, and risk profile, the same capability – an agent selecting and initiating a permitted payment within a defined mandate – can be applied in institutional settings. Wholesale activities across cash, liquidity, collateral, FX, and settlement involve repeated, data-intensive decisions that often depend on manual intervention, fragmented processes, and operating hours. Financial institutions and corporates already govern these activities through policies, delegated authorities, counterparty limits, and approval thresholds, which could provide a foundation for machine-readable agent mandates.

This opportunity is reinforced by the emergence of a multi-rail payments environment. As traditional correspondent banking networks increasingly coexist with instant payment systems and tokenised settlement arrangements, the challenge may shift from executing payments within individual networks to orchestrating value, liquidity, and risks across them. In this context, AI-enabled systems could support payment routing, liquidity management, and the selection of appropriate settlement paths across an increasingly diverse payments landscape (for more information read our *flow* whitepaper, 'Cross-border payments in 2040: five hypotheses for a multi-rail future').[9]

The tokenisation of money and assets strengthens this case. As deposits, money-market funds, bonds and other assets are represented on programmable ledgers, value can move continuously, settle atomically against its counterpart, and carry its own transfer conditions.[10] This creates a software-native settlement environment in which agents can operate beyond traditional processing hours, and the controls can be enforced directly at the point of value transfer. AI and blockchain are in this sense complementary. AI provides the capacity to interpret information and select actions, while tokenised infrastructure can provide programmable, verifiable, and continuously available settlement. Agentic payments sit at their intersection and could therefore become part of the payment and settlement infrastructure to support an emerging on-chain financial economy.

The following example graphic illustrates how introducing an agent into a conventional wholesale payment flow could streamline decision-making and execution, reduce manual hand-offs, and improve efficiency while preserving institutional controls and accountability.

Figure 1: Proposed wholesale payments model



1.3. Definition

For the purposes of this paper, **a wholesale agentic payment is a payment in which a software agent interprets relevant information, selects or initiates a transaction, and acts under an explicit and bounded mandate delegated by a legally accountable institutional principal and revocable in near real time.**

The defining feature of agentic action is **delegated decision authority**. Unlike programmable automation, which executes pre-scripted instructions when predefined conditions are met, an AI agent interprets relevant information and selects among permitted actions within a defined policy perimeter.

Payment activity can be positioned along an automation spectrum (see *Table 1*), based on the level of decision authority delegated to software:

Table 1: Definition of automation spectrum

Stage	Definition	Application status
1. Human-directed execution	A person evaluates the information, selects the action, and authorises the transaction	Established in both retail and wholesale finance
2. Deterministic automation	The system executes a pre-scripted action when predefined conditions are met, without interpreting the broader context	Widely used in retail and wholesale processes
3. AI-assisted decision-making	An AI system interprets information and recommends action, but a human or established institutional process approves the transaction	Emerging in wholesale applications
4. Bounded agentic execution	An agent evaluates permitted alternatives and selects or initiates an action independently within an explicit mandate, while deterministic controls govern authorisation and execution	Emerging in retail; under active exploration in wholesale finance
5. Higher-autonomy orchestration	An agent dynamically plans and coordinates multiple actions, potentially across counterparties, systems or agents, with human involvement primarily for exceptions and oversight	Longer-term potential

Existing retail applications increasingly combine AI-supported product discovery and decision-making with controlled payment initiation. In wholesale finance, deterministic automation is already widely embedded in payment and settlement processes, while the BIS study provides early experimental evidence that AI agents may support more contextual decisions, such as payment prioritisation and intraday liquidity management.

The near- to medium-term opportunity for wholesale finance therefore lies between **AI-assisted decision-making** and **bounded agentic execution**. Higher autonomy orchestration should remain a long-term possibility.

2. Unlocking value across wholesale finance

Wholesale finance includes activities across corporate banking, investment banking, and asset management. Agentic capabilities could support a wide range of activities, including collateral optimisation, securities financing, risk monitoring, investment research, and portfolio management.[11] [12]

This paper focuses specifically on agentic payments. The most credible near-term applications are those in which payment-related decisions are frequent, data-intensive, and currently fragmented across systems and teams. It therefore prioritises three corporate banking use cases in which payment-related decisions are integral to the underlying process:

1. **Treasury liquidity management;**
2. **Cross-border payment orchestration; and**
3. **Trade finance.**

For each use case, the paper also outlines how tokenised money and assets could strengthen the model where tokenised settlement is available.

2.1. Treasury liquidity management

Corporate treasury teams need to monitor balances, cash-flow forecasts, payment obligations, funding needs, currency positions, and operating cut-offs across multiple entities, jurisdictions, and systems. These activities ultimately require the movement of funds across accounts, entities and currencies to meet obligations and optimise liquidity. Fragmented information and manual coordination can delay decisions, increase precautionary liquidity buffers, and leave available cash underutilised. Furthermore, corporate treasury teams must also determine whether forecast changes require immediate action, continued monitoring, or an adjustment to existing liquidity buffers.

How AI can be leveraged:

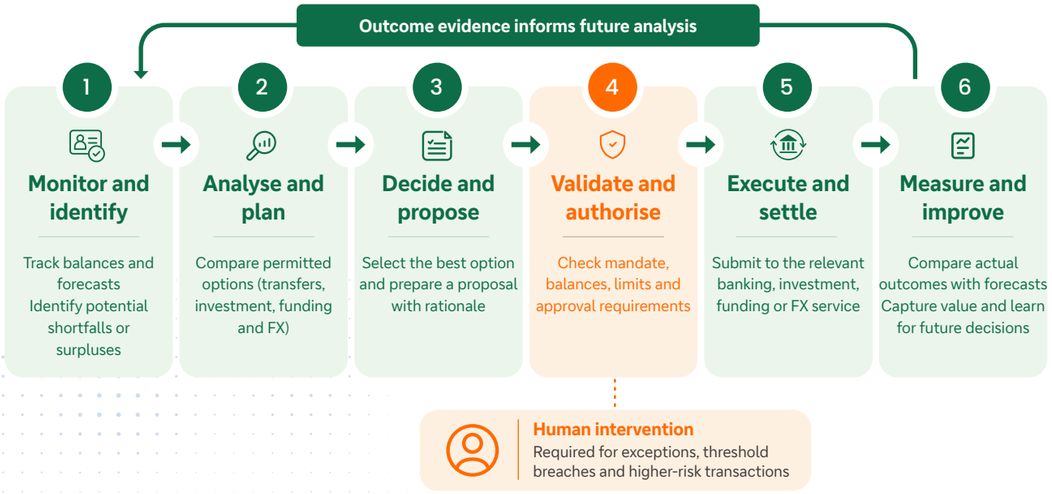
Two complementary but distinct AI capabilities could create value: 1) predictive forecasting and 2) agentic decision-making. Forecasting estimates future liquidity positions, while the agent interprets the resulting forecasts, determines whether action is required, and orchestrates permitted responses within its mandate. The agent could coordinate available treasury arrangements and banking services, including cash concentration, payments, investment, funding and FX, and assess them against cost, timing, liquidity impact, currency exposure, applicable requirements, and internal policy. In this model, the agent provides an orchestration layer across available options, while independent controls determine whether the proposed action is authorised and may proceed.

The following example illustrates how the use case could operate in practice.

An unexpected receipt creates surplus liquidity in one part of a corporate group, while another entity is using short-term funding. An agent could assess whether mobilising the surplus to reduce borrowing provides greater value than retaining or investing the surplus funds, subject to the corporate's liquidity requirements and policies. This example illustrates how an agent could move from identifying a change in liquidity position to evaluating the available options, proposing an action, and supporting its controlled execution, as outlined below.

1) **Monitor and identify**: The agent monitors current balances, expected cash flows, and upcoming obligations across permitted entities and identifies a material change in the projected closing position.
2) **Analyse and plan**: The agent determines whether the projected position requires action and compares permitted responses against the required liquidity buffer, funding cost, expected return, timing, currency exposure, and internal policy.
3) **Decide and propose**: The agent selects the preferred response and prepares a proposal with the supporting information, alternatives considered, and decision rationale. Where immediate action is not justified, the agent may recommend continued monitoring.
4) **Validate and authorise**: An independent policy engine verifies the agent's mandate, eligible entities and accounts, exposure and value limits, current balances, and approval requirements. Proposals that remain within delegated authority and satisfy all controls may proceed through STP; exceptions and threshold breaches are routed for human approval.
5) **Execute:** Once authorised, the transaction is submitted to the relevant banking, funding, or FX service for execution and settlement.
6) **Measure and improve**: Following settlement, the agent records the decision, approval, execution outcome and exceptions, and compares realised outcomes with forecasts and expected benefits. This enables treasury to monitor decision quality, quantify the value delivered, and support governed improvement of future analysis and recommendations. Importantly, the feedback loop does not allow the agent to alter its mandate, limits, or controls. Outcome data may inform future analysis and model improvement, subject to data-quality checks, performance monitoring, and institutional change control.

Figure 2: Illustrative agent-led workflow for treasury liquidity management



Value hypothesis

- Earlier identification of liquidity shortfalls and surpluses.
- More targeted alerts and faster evaluation of permitted treasury responses.
- More effective allocation and mobilisation of liquidity across entities, accounts, and currencies.
- Reduced routine intervention through bounded straight-through execution.
- Measurable improvement in cash utilisation, liquidity buffer calibration, and avoidable funding costs.

Use-case-specific inputs and dependencies

- Entity-level current and expected closing balances, cash-flow forecasts, and upcoming payment obligations.
- Visibility across permitted accounts, entities, currencies, and banking relationships.
- Current liquidity arrangements and available internal liquidity, funding, investment, and FX options.
- Treasury policies covering liquidity buffers, exposures, counterparties, permitted activities, and funding limits.
- Access to treasury management, investment, funding, and FX services.
- Outcome data comparing forecast and realised positions, transaction costs, funding usage, and liquidity buffer utilisation.

2.2. Cross-border payment orchestration

For the purposes of this use case, cross-border payment orchestration covers the execution of payments across currencies, jurisdictions, providers, and payment or settlement networks. It may include commercial, intercompany, and treasury-related payments. Aggregate liquidity management across entities and accounts remains within the treasury liquidity management use case in Section 2.1. Where a transaction also requires document examination, trade financing, guarantees, or verification of contractual trade conditions, those activities fall within the trade finance use case described in Section 2.3.

Cross-border payments require coordinated decisions on timing, currency conversion, liquidity or credit, service provider, and execution route. This complexity is likely to increase as traditional correspondent banking networks coexist with domestic and real-time payment systems, card networks, and emerging tokenised settlement platforms. Information, controls, credit availability, and transaction status may therefore be distributed across multiple systems and networks, creating additional fragmentation and operational hand-offs. Applicable requirements may also vary according to transaction purpose, corridor, currency, and the status of the entities involved, including residence status and any specific regulatory approvals, quotas, or reporting obligations. Effective orchestration will therefore require more than selecting the fastest or lowest-cost route. The agent must determine how best to fulfil the client's payment objective. This requires coordinated choices across FX, liquidity or credit, service providers, fees, timing, and available rails, while ensuring that the proposed approach remains permissible under applicable jurisdictional and client-specific arrangements.

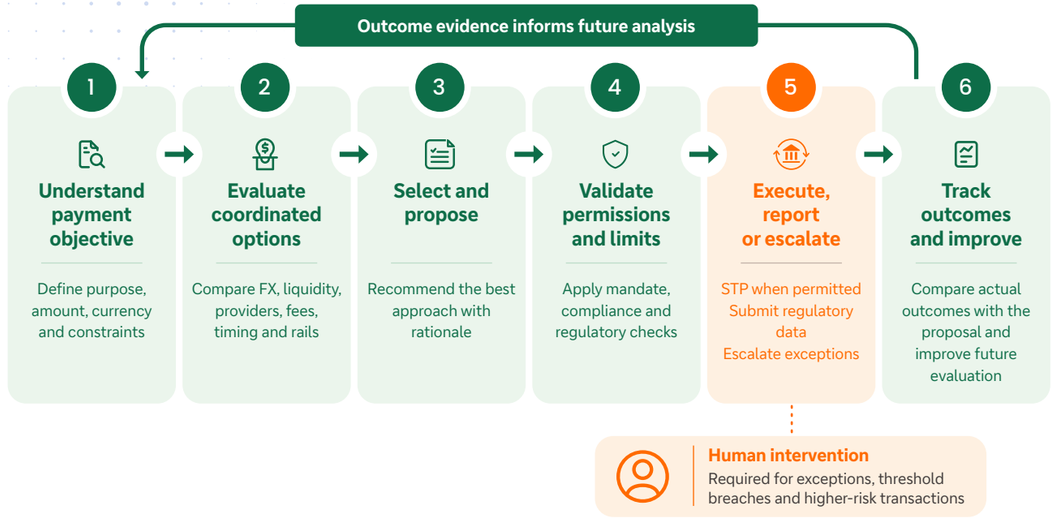
How AI can be leveraged:

An agent could orchestrate a cross-border payment by interpreting its commercial purpose and client requirements. It could classify the transaction type, identify the jurisdictions and entities involved, and determine the execution decisions and regulatory requirements that apply. Rather than applying a fixed routing rule, the agent could evaluate interdependent choices across FX booking and timing, liquidity or credit availability, service provider, transaction fees, payment rail, and settlement arrangement. It could assess permitted combinations against the client's objective and balancing cost, speed, availability, and settlement risk, within applicable corridor-, currency- and entity-specific requirements. The agent would then propose a preferred approach most likely to achieve the required outcome, supported by the alternatives considered and its rationale. Independent institutional controls would continue to determine whether the proposal is authorised and may proceed. Outcome evidence could inform future option evaluation and governed model improvement as providers, rails, and requirements evolve. Any resulting change to the agent's model, decision parameters, or operating perimeter would remain subject to institutional validation and change control.

An illustrative flow could operate as follows:

1) **Understand payment objective**: The agent interprets the payment's amount, currency, beneficiary, purpose, required delivery time, and relevant client constraints. It identifies the entities and jurisdictions involved and determines whether FX, liquidity, credit, documentation, or regulatory approval is required.
2) **Evaluate execution options**: The agent compares permitted combinations of FX booking and timing, liquidity or credit source, service provider, transaction fees, payment rail, and settlement arrangement. It evaluates how each combination performs against cost, speed, operational availability, compliance requirements, and settlement risk.
3) **Select and propose**: The agent selects the preferred approach that best meets the client's objective and prepares a proposal, supported by the relevant transaction data, alternatives considered, and decision rationale.
4) **Validate**: Independent control services validate the proposal against the applicable mandate, beneficiary information, sanctions status, account and currency permissions, pricing ranges, transaction limits, and approval requirements. They also check applicable residence restrictions, regulatory approvals, reporting obligations, and transaction or aggregate limits.
5) **Execute or escalate**: A proposal that remains within the delegated mandate and satisfies all applicable controls proceeds through STP. Where supported, required regulatory data or documentation is submitted through approved interfaces. Human intervention is required where a proposal falls outside the mandate, exceeds an approval threshold, or presents an exception relating to value, risk, pricing, compliance, or settlement conditions.
6) **Track and improve**: The agent tracks execution, settlement, reporting, and reconciliation. The transaction record captures the proposal, rationale, controls, approvals, execution status, and exceptions. Furthermore, actual cost, timing, settlement, and exception outcomes are retained and compared with the proposal to support monitoring and governed improvement of future decisions. As in the treasury use case in Section 2.1, outcome feedback informs future evaluation but does not change the agent's mandate or controls.

Figure 3: Illustrative agent-led workflow for cross-border payment orchestration



Value hypothesis

- More consistent optimisation of the end-to-end payment outcome.
- Better coordination of FX, liquidity or credit, providers, fees, timing, and settlement rails.
- Greater STP potential for transactions within mandate.
- Faster identification and escalation of exceptions requiring human intervention.
- Improved outcome visibility, supporting reconciliation, and future decision improvement.

Use-case-specific inputs and dependencies

- Complete payment instructions, including amount, currency, beneficiary, purpose, required delivery time, and client constraints.
- Current FX options, payment-rail coverage, processing times, and route availability.
- Available liquidity and credit sources, including relevant limits and costs.
- Entity and transaction classifications, including residence status and payment purpose.
- Corridor- and currency-specific requirements, including local formats, cut-offs, and restrictions.
- Current transaction, entity, and aggregate group limits.
- Connectivity to relevant bank, market-infrastructure, and regulatory data or reporting services, where available.
- Outcome data covering actual fees, FX execution, processing time, settlement status, and exceptions.

2.3. Trade finance

Trade finance is especially suited to AI because value is constrained less by payment execution than by the work required to interpret documents, validate conditions, coordinate controls, and resolve exceptions before financing or payment can be agreed and released. Today, information is often distributed across contracts, invoices, transport documents, client data files, emails, and bank systems. Repeated extraction, comparison, and hand-offs consume scarce specialist capacity, create inconsistent turnaround times, and make growth dependent on additional sales, operations, legal, and risk resources.

How AI can be leveraged:

The near-term opportunity is not a large autonomous agent making uncontrolled credit, compliance, or payment decisions. It is an AI-enabled operating layer that converts unstructured trade information into structured, traceable evidence; compares that evidence with transaction terms and policy; orchestrates the correct workflow across product, operations and control functions; and prepares offers and instructions only after independent controls and approvals are satisfied. This moves AI from isolated document extraction and processing to end-to-end process orchestration, while keeping sanctions, anti-money laundering (AML), credit, legal, and payment authorisation under authoritative control and retaining human judgement for material or uncertain exceptions.

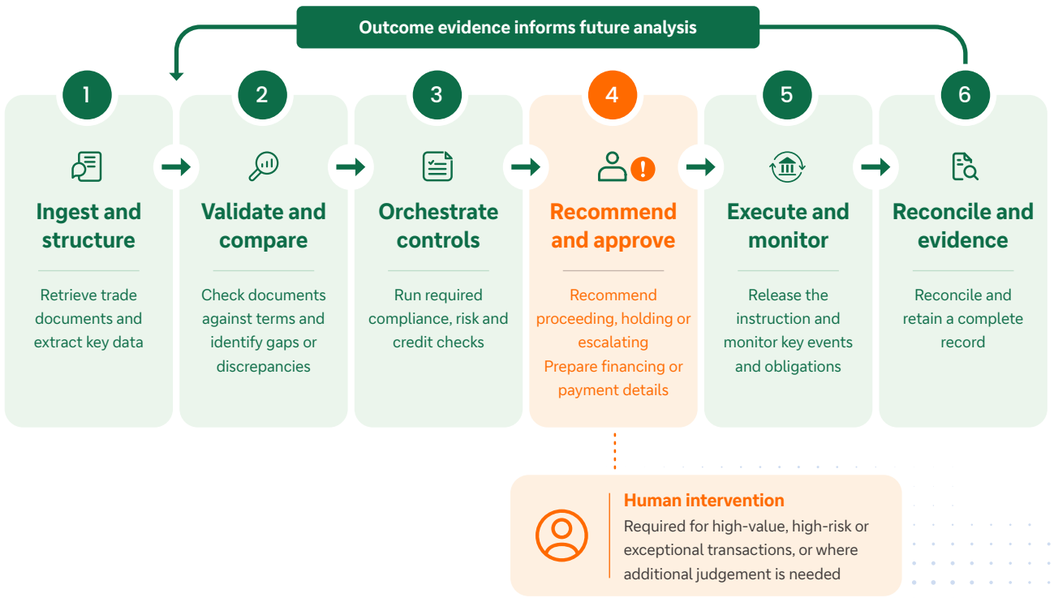
From use cases to scale. This direction is already visible in Deutsche Bank's Trade Finance & Lending portfolio. A reusable capability for converting emails, PDFs, Office files and images into structured data for downstream systems has been developed. AI-supported bank-guarantee text vetting has been rolled out in Germany with a broader roll-out planned thereafter. AI-based automation is targeting higher extraction accuracy and broader handling of complex documentary-credit structures. Tools and platforms to allow trade experts to progress their use cases to production are currently being piloted. Together, these and other such initiatives show a deliberate progression from reusable AI components and bounded pilots to production pathways and repeatable platform-level scaling.

Deutsche Bank Corporate Bank's AI agenda is centred on delivering concrete examples that deliver measurable value – in faster client turnaround, lower cost, and operational efficiency – and identified document processing as the primary trade-finance benefit areas. This framing is important: the strategic objective is not to accumulate proofs of concept, but to industrialise capabilities that can be reused across products, connect them to governed data and workflows, and scale them where impact can be evidenced.

An illustrative flow could operate as follows:

1) **Ingest and structure**: Agents retrieve permitted transaction records, client data, and trade documents, extract key fields, and link every output to its source.
2) **Validate and compare**: AI compares invoices, transport evidence, guarantees, credit terms and client data with approved transaction conditions, identifying omissions, discrepancies, and anomalous patterns.
3) **Orchestrate controls**: An agent invokes the required sanctions, AML, goods, vessels, fraud, credit, and policy checks. Authoritative services determine the control outcome; the agent consolidates evidence and routes exceptions to the accountable specialist.
4) **Recommend and approve**: AI recommends proceeding, holding, or escalating, explains the rationale and prepares the financing, guarantee, or payment details. Human approval remains mandatory where value, risk, uncertainty, or policy requires it.
5) **Execute and monitor**: Once independently authorised, the instruction is released through the approved channel. The agent tracks shipment events, obligations, payment status, and emerging exceptions against expected milestones.
6) **Reconcile and evidence**: The transaction is reconciled and closed with a complete record of source documents, extracted data, discrepancies, decisions, controls, approvals, communications, and settlement.

Figure 4: Illustrative agent-led workflow for trade finance



Value hypothesis

- Faster client turnaround by identifying documentary, data, and policy exceptions earlier.
- Greater processing capacity without proportionate increases in specialist operations, legal and risk resources.
- More consistent application of playbooks, policies, and transaction conditions, supported by source-linked evidence.
- Fewer manual hand-offs and better end-to-end status visibility across clients, product, operations, and control functions.
- A scalable path from assisted review to bounded orchestration and, where justified, faster release of financing or payment once verified conditions are satisfied.

Use-case-specific inputs and dependencies

- Controlled access to trade documents, transaction terms, and client instructions.
- Document extraction that preserves links to the original source material.
- Machine-readable contractual, documentary and financing conditions.
- Reliable external information on goods, vessels, shipment events, and receivable status, where relevant.
- Connectivity to trade-finance, lending, and document management workflows.

2.4. How tokenised settlement could enhance the value hypotheses

Tokenised settlement is not a prerequisite for agentic payments; however, continuous availability, atomic execution, and programmable transfer could provide additional value across all three use cases.

- In treasury liquidity management, tokenised deposits could support more timely liquidity repositioning outside traditional operating windows.
- In cross-border payment orchestration, tokenised settlement could improve execution certainty and enable interdependent transaction components to complete together or not at all, reducing settlement risk. Additionally, it could coordinate and enable higher payment velocity associated with more fragmented agentic payments and overcome cut-off time through its 24x7 availability.
- In trade finance, financing or payment could be released once independently verified conditions are met, such as compliant documents, an accepted shipment event, or confirmation of receivable eligibility. Tokenised invoices and receivables could also provide a verifiable record of ownership and financing status.

These capabilities would complement, rather than replace, the bank's mandate, approval, and control framework, and their application would depend on the relevant infrastructure, asset, jurisdiction, and use case.

2.5. Institutional capabilities required to enable agentic payments

The three use cases depend on a common institutional foundation. Agent intelligence alone is not sufficient: institutions must provide the processing, data, controls, connectivity, and resilience required to convert an agent's proposed action into an authorised and accountable transaction.

Five enabling capabilities are particularly important.

- **Extended-hours processing**: Real-time payment, liquidity, screening, and monitoring capabilities that can support agent-driven decisions and transactions beyond traditional operating windows, where required by the use case.
- **Trusted data foundation**: Timely, consolidated, and permissioned data across accounts, liquidity, payments, FX, counterparties, and transaction status, with clear quality and provenance.
- **Machine-readable authority and controls**: Explicit digital mandates defining what the agent may do, supported by independent policy, financial-crime, limit, and approval controls at execution.
- **Secure integration and execution**: Controlled connectivity across treasury, FX, trade-finance, payment, and settlement systems, with credentials and transaction signing kept outside the agent's reasoning environment.
- **Monitoring, evidence, and resilience**: End-to-end records of mandates, decisions, controls, and outcomes, supported by continuous monitoring, immediate suspension, and the ability to continue the underlying process without the agent.

3.

How wholesale agentic payments could work

In this paper, the **principal is the institution** on whose behalf the agent acts, and which retains legal and regulatory responsibility for transactions executed under the delegated mandate. The institution acts through an **authorised human owner** who defines the agent's objective, approves, and may revoke its mandate, and oversees its use. The human owner must be clearly distinguished from the agent: the agent acts only within delegated authority, cannot extend or alter its own mandate and does not bear legal or regulatory responsibility for the resulting transaction.

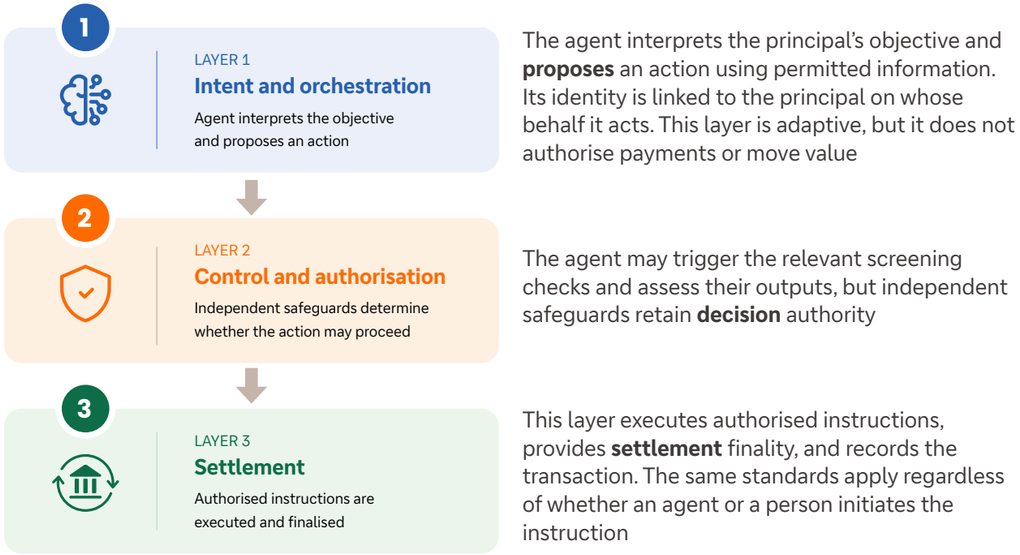
The authorised human owner anchors the accountability framework. Delegated authority must be traceable to a natural person who is empowered to grant it, review its use, approve exceptions and withdraw or suspend it. The human owner therefore connects the mandate and suspension rights described in Sections 3.4 and 3.5 with the named accountable ownership and human-approval thresholds set out in Section 4.4 and Chapter 5. **This does not transfer the institution's legal responsibility to the individual; rather, it establishes clear human ownership and oversight within the institution's governance framework.**

Not every individual should be eligible to perform this role. The authorised human owner should have formally delegated authority from the institution, an appropriate understanding of the agent's capabilities, limitations and controls, and sufficient experience to respond to unexpected behaviour or control failures. The role should include authority to restrict or suspend the agent and to invoke established manual or fallback processes where agentic execution is unavailable, repeatedly unsuccessful or operating outside expected parameters. **Where an agentic-payment system is classified as high-risk under the applicable AI framework, these arrangements should also meet the relevant statutory requirements for human oversight.** Article 26(2) of the EU AI Act requires deployers of high-risk AI systems to assign oversight to natural persons with the necessary competence, training, authority and support.[13]

3.1. Reference architecture

This paper applies the three-layer model proposed in the IMF's Note No. 2026/004, *How Agentic AI Will Reshape Payments* (the 'IMF Note'), which separates **intent and orchestration**, **control and authorisation**, and **settlement**.[14] Its value is that it draws the boundary in the place that matters to wholesale finance: it isolates where adaptive, probabilistic agent behaviour is permitted from where deterministic control over authorisation and payment finality must be preserved.

Figure 5: Three-layer reference architecture for wholesale agentic payments



The agent interprets the principal's objective and **proposes** an action using permitted information. Its identity is linked to the principal on whose behalf it acts. This layer is adaptive, but it does not authorise payments or move value

The agent may trigger the relevant screening checks and assess their outputs, but independent safeguards retain **decision** authority

This layer executes authorised instructions, provides **settlement** finality, and records the transaction. The same standards apply regardless of whether an agent or a person initiates the instruction

The central architectural principle is straightforward: the agent supplies judgement, while independent institutional safeguards supply permission. An agent should not be able to decide, approve, and execute a payment on its own.

In wholesale finance, a mandate alone is not sufficient. Account eligibility, sanctions status, credit exposure, liquidity, and pricing may change after authority has been granted. **The control layer must therefore assess each proposed transaction against current conditions at the point of execution.**

3.2. Recommended design choices for wholesale agentic payments

Design question	Choices	Recommended wholesale approach	Rationale
1. How should decision-making and execution be organised?	Functional separation vs. integrated agent	**Separate agent judgement** from validation and execution	An agent's decision should not move value without independent verification and authorisation
2. How should authority be represented?	Signed mandates vs. institution-controlled policy engines	Use an **explicit**, **bounded**, and **near-real-time revocable mandate** together with an **institution-controlled policy engine**	The mandate establishes delegated authority; the policy engine determines whether the transaction remains permissible under current conditions
3. When should human approval apply?	Human approval vs. bounded autonomous execution	Use **risk-based approval thresholds** and **bounded execution**	Human approval should be retained for high-value, high-risk, exceptional, or uncertain transactions rather than required uniformly
4. Where should adaptive reasoning be permitted?	Fixed workflow vs. adaptive agent reasoning	Use **adaptive reasoning** for analysis, selection, and exception handling, while preserving **deterministic authorisation and settlement**	Probabilistic outputs should not directly control irreversible actions or finality
5. How should evidence be retained?	Institutional records vs. shared or on-chain evidence	Keep detailed evidence in **controlled institutional records** and share only the **minimum proof required** across institutional boundaries	This supports traceability and auditability and dispute resolution while protecting confidential information
6. What agent model is appropriate initially?	Single-agent vs. multi-agent architecture	Begin with **one agent acting for one institutional principal**	Multi-agent operation requires cross-institutional standards for identity, authority and precedence that remain at an early stage

3.3. Settlement rails and assets

Agentic payments do not depend on a single settlement rail or form of money. Subject to the use case and jurisdiction, an agent could initiate payments through existing account-based infrastructure, including correspondent banking networks, RTGS systems, automated clearing houses, and instant-payment systems, as well as card networks or emerging tokenised platforms. Therefore, the settlement asset could be central bank money, commercial bank deposits, wholesale central bank digital currency, tokenised deposits or regulated stablecoins.[15]

Tokenised settlement assets nevertheless merit particular attention in an agentic context. Although they are not a prerequisite for agentic payments, three characteristics may be particularly relevant. Firstly, continuous availability can support transactions outside conventional operating windows and cut-off times. Secondly, atomic settlement can enable two interdependent legs, such as delivery and payment or dual-currency transfers, to complete together or not at all. Thirdly, programmability can allow eligibility, transfer conditions, dynamic pricing, and transaction limits to be enforced at the point where value moves.

Where these characteristics are relevant to the use case, tokenised deposits, regulated stablecoins and other tokenised assets could complement the control architecture described in Section 3.1 by extending selected safeguards into the settlement layer. Section 3.6 illustrates how this may be applied in practice.

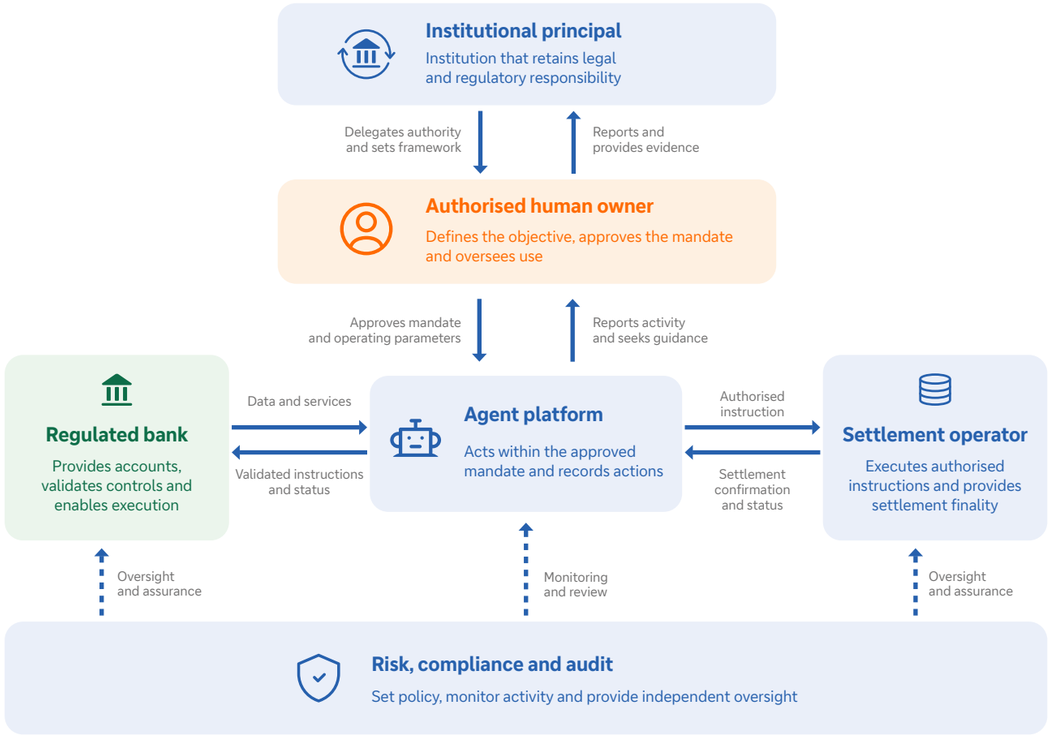
The IMF Note supports architecture in which agents can interact with different settlement arrangements, rather than being tied to one particular rail or asset. The BIS blueprint for the future monetary system highlights programmable platforms that combine central bank money, tokenised deposits, and tokenised assets, while acknowledging that multiple interoperable ledgers may coexist. Stablecoins represent another possibility, particularly for certain cross-border use cases, although the BIS Committee on Payments and Market Infrastructures report on stablecoins in cross-border payments emphasises regulatory, financial-stability and monetary-sovereignty considerations and thereby treats stablecoins as one future model among several.[16]

Where settlement takes place through a financial market infrastructure, existing requirements for governance, risk management, access, and settlement finality continue to apply.[17] **Agentic capabilities should be accommodated in the decision and authorisation layers without weakening the standards governing settlement. Where appropriate, programmable settlement controls may provide an additional safeguard at the point where value moves.**

3.4. Institutional roles and responsibilities

Agentic payment involves multiple participants whose responsibilities must be clearly allocated.

Figure 6: Roles and responsibilities for wholesale agentic payments



Within the proposed operating model, **the regulated bank's distinctive contribution** is to connect agent decision-making with the client's account, liquidity, and regulated payment capabilities. The bank validates the agent's authority, confirms that the transaction is appropriately funded and permitted, applies consistent controls across available payment and settlement rails, and retains the records required for reconciliation, investigation, and recourse. Other regulated institutions may perform some of these functions, depending on their permissions and operating model. The bank's particular value lies in combining them within an established client, deposit, and risk management relationship.

3.5. Agent lifecycle and control gates

Onboarding, suspension, and offboarding

- Customer onboarding establishes the institutional principal and the authorised human owner or owners permitted to approve and oversee agent mandates. **Agent onboarding** extends customer onboarding through know-your-agent[18] verification, linking the agent to the verified institutional principal, authorised human owner, responsible agent operator, approved mandate, and revocable credential.
- **Suspension** should be distinct from permanent offboarding. The institutional principal, acting through an authorised human owner, and executing institution should each be able to suspend an agent, mandate, or specific capability immediately and unilaterally. Suspension should prevent new activity while preserving the records required for investigation and review. The operating model should define suspension triggers, treatment of pending instructions, and conditions for reinstatement.
- **Offboarding** permanently withdraws the agent's mandates, credentials and access to systems, tools, and data. Customer offboarding should apply to every agent acting for that customer in compliance with relevant regulations including GDPR and data retention policies. The underlying process should remain operable without the agent so that suspension or removal does not create a continuity failure.

Control gates

- **Before deployment**: Assign an authorised human owner, approve the scope and decision boundaries, test rollback arrangements, and restrict the agent to the minimum data and tools required.
- **Before each decision**: Confirm that the mandate remains valid and that the agent is using approved, integrity-checked information sources.
- **Before each transaction**: Independently assess the proposed action against current limits, eligibility, sanctions status, exposure, and approval requirements.
- **During execution**: Keep signing keys outside the AI model, in secure modules that the agent may use but cannot extract or alter, with controls against duplicate or replayed instructions.
- **After execution**: Reconcile settlement and retain sufficient evidence to reconstruct the decision, rationale, controls, approvals, and resulting transactions.
- **Ongoing**: Monitor activity against expected behaviour and risk thresholds, with automated and manual mechanisms to restrict or suspend the agent following anomalies, mandate breaches, or control failures.

Two agent-specific risks require particular attention when agents can initiate transactions. Firstly, an agent may automatically retry an instruction if it does not receive confirmation, potentially creating duplicate payments. Payment systems should therefore recognise repeated instructions and prevent the same transaction from being executed more than once. Secondly, agents using similar models, data, and objectives may make similar decisions at the same time. At scale, this could amplify liquidity pressures or market movements across institutions. Potential mitigants include greater diversity in agent design, limits on the scope of delegated activity, and real-time monitoring for unusual or correlated behaviour.

3.6. Case study: Anvita Flow as an implementation example

Anvita Flow is Ant Digital Technologies' agent-to-agent payment platform.[19] The following case study describes how one existing implementation applies selected elements of the reference architecture and design principles described in Sections 3.1 and 3.5. It is provided as an implementation example, not as a recommendation of a particular platform, protocol, or settlement model. Figure 7 and the transaction flow in Section 3.6.2 are simplified illustrations of this implementation.

Institutional agentic payments raise three fundamental questions:

- Who is the agent, and on whose authority does it act?
- What is the agent permitted to do, and how are those permissions enforced?
- How can frequent, potentially low-value transactions be settled efficiently?

Anvita Flow addresses these questions through capabilities mapped to the intent and orchestration, control and authorisation, and settlement layers of reference architecture in Section 3.1.

3.6.1. Anvita Flow: Implementation overview

The implementation separates the agent's adaptive decision-making from the controls governing authority, credentials, and settlements. A steward agent acts for a verified institutional principal under a mandate approved by an authorised human owner and interacts with service agents within defined parameters. Delegated authority specifies boundaries or thresholds including permitted activities, counterparties, limits, frequency, and expiry, while signing credentials and settlement controls remain separate from the agent's reasoning environment.

Figure 7: Illustrative mapping of Anvita Flow to the three-layer reference architecture

Reference architecture	What the layer must guarantee	Anvita Flow implementation
Intent and orchestration	Adaptive reasoning only Proposals, not payments Agent bound to a principal	**Steward agent acts for principal;** service agents expose bank services Agent identity: KYC + KYB + KYA, bound to verified owner, on-chain ID
The boundary that matters - judgment above, permission below - the agent never holds the executing credential		
Control and authorisation	Deterministic Same instruction + conditions = same decision, always	**Scoped delegated authority:** EIP-7702, or a per-agent smart account Caps, payees, expiry; principal can revoke Agent key in secure module, not the model
Settlement	Finality and authoritative record Unchanged by agent origination	**Eligibility + limits enforced in the** settlement asset contract itself **Finality graded to value at risk:** ~100ms / ~1s attested / ~min proof

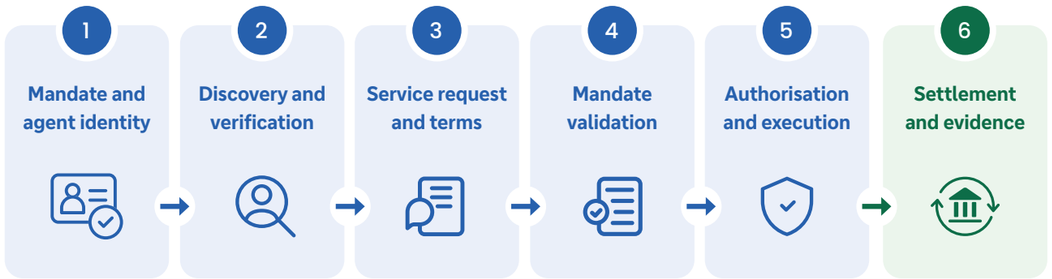
Anvita Flow components mapped to the neutral reference architecture. Settlement is x402-compatible; mandates align to the AP2 model.

The mapping preserves the central architecture principle described in Section 3.1: the agent may provide judgement and orchestration, but guardrails and independent controls determine whether an instruction is authorised and may proceed.

3.6.2. Illustrative agent-to-agent transaction flow

The following sequence outlines an example agentic transaction flow compatible with both x402 and AP2 protocols:

Figure 8: Illustrative agent-to-agent transaction flow



1) Mandate and agent identity

The institutional principal, acting through its authorised human owner, approves a signed mandate defining the steward agent's scope, spending limit, permitted counterparties, and expiry. The agent is linked to a verified institutional principal and authorised human owner through an admission process combining identity binding, customer and business verification, know-your-agent checks and controlled credential issuance.

Key design choice: Establish verified agent identity and explicit, bounded, and revocable authority before the agent can act.

2) Discovery and verification

The steward agent identifies a service agent and verifies its credentials, institutional principal, and accountable authority before requesting a service.

Key design choice: Require counterparties to verify the identity and authority of the agents with which they interact..

3) Service request and terms

The steward agent requests a service, and the service agent returns the applicable terms and payment requirement within the request-response exchange.

Key design choice: Use standardised protocols to support service discovery, structured payment requests, and interoperability between agents.

4) Mandate validation

The steward agent checks the payment requirement against its mandate. Where the final amount is not known in advance, the permitted authority is limited by a maximum spending threshold.

Key design choice: Validate every proposed transaction against the current mandate, limits, and permitted conditions before authorisation.

5) Authorisation and execution

The instruction is signed using a credential held separately from the agent's reasoning model. The institutional principal, through an authorised human owner or institutional credential service, can revoke the authority or rotate the credential, while independent controls verify eligibility and transaction limits before execution.

Key design choice: Separate decision-making, credential custody, and execution, with independent controls applied before value moves.

6) Settlement and evidence

Once the required controls have been satisfied, the transaction settles and the mandate, decision, control results, and outcome are retained as a traceable record.

Key design choice: Apply settlement assurance proportionate to the transaction's value and risk, supported by end-to-end traceability and auditability.

This flow is a simplified illustration of the implementation. The applicable identity, control, approval, and settlement arrangements would depend on the use case, institution, and jurisdiction.

4. Conditions for responsible adoption

Agentic payments introduce new possibilities for wholesale finance, but responsible adoption depends less on maximum agent capability than on the institution's ability to establish authority, maintain control, and show evidence of accountability. The design challenges are therefore to connect adaptive, probabilistic decision-making with payment infrastructure that requires predictability, auditability, and legal certainty.

4.1. Regulatory and legal clarity

Agentic payments are currently governed by existing AI, payments, financial crime, data protection, outsourcing, and operational resilience frameworks rather than through a dedicated regulatory regime. Institutions must therefore establish when an agent-generated instruction is legally attributable to its institutional principal, how delegated authority is defined and revoked, and how responsibility is allocated among the institutional principal, authorised human owner, agent operator, executing institution, and settlement operator.

Legal certainty is particularly important where agents interact with smart contracts, distributed ledgers, or tokenised settlement assets. In these arrangements, AI governance must be considered alongside ledger governance, automated execution, settlement finality, confidentiality, and liability.[20] Furthermore, institutions must also retain sufficient evidence to reconstruct the mandate, decision, controls, approvals, and resulting transactions.

For institutional adoption, agents should initially operate within narrow, machine-readable mandates, independent transaction controls, and defined human approval thresholds. Broader authority should be introduced progressively and only where legal attribution, contractual responsibility, suspension rights, evidence requirements and arrangements for errors, disputes and recourse are clearly established.

4.2. Industry standards and interoperability

Agentic payments will not scale through proprietary arrangements alone. An agent acting for an institutional principal must be able to establish its authority to a counterparty that has no prior relationship with it, and to do so without either party exposing its internal controls. Three areas of standards are converging on that problem, and an institution's interoperability position is a question of how it engages with them.

- **Authorisation**. AP2 expresses authority as cryptographically signed mandates specifying scope, limits, actor identity, and permitted conditions, and was contributed to the Fast Identity Online Alliance in 2026, placing it under neutral governance.[21]
- **Settlement.** The x402 standard carries a payment requirement inside an ordinary web exchange, enabling settlement at sub-cent values; stewardship passed to the Linux Foundation in 2026, and the founding membership spans the major card networks, cloud providers, and stablecoin issuers.
- **Identity.** On-chain agent registries provide portable identity and reputation that can be resolved across institutional boundaries, complementing rather than replacing regulated onboarding.

Two practical observations follow. Firstly, convergence is happening at the integration layer rather than through consolidation: a single protocol-agnostic integration point that supports several agent protocols simultaneously is now offered commercially, which suggests institutions will not be required to choose a single standard.[22] Secondly, existing financial messaging remains the interoperability backbone for wholesale, and structured transaction data continues to carry the originator, beneficiary and remittance information on which downstream controls depend. The realistic interoperability requirement is therefore that agent-layer standards map cleanly onto existing message standards (e.g. Swift) rather than displacing them.

For financial institutions, the near-term implication is to treat standards participation as a control activity rather than a technology choice: engagement determines whether an institution's own control expectations are reflected in the standards its counterparties will rely on.

4.3. Technology and operational resilience

Agentic payments introduce technology dependencies and failure modes that require specific treatment within existing resilience frameworks.[23] Existing requirements for incident management, resilience testing, business continuity, and third-party risk continue to apply. Four areas warrant particular attention in an agentic-payment context.

- **Non-determinism and change:** An agent's behaviour may be affected by changes to its model, data, purpose, or execution authority. A material change to any of these may alter the agent's operating perimeter and should trigger formal reassessment rather than being treated as routine maintenance. Resilience testing should therefore cover behavioural regression as well as availability.

- **Agent-specific failure modes**: Retry behaviour may produce duplicate instructions once an agent can initiate transactions, requiring idempotency and replay protection at the instruction layer. Information retrieved from external sources should be treated as untrusted because manipulated inputs may influence the agent's decisions.

- **Tokenised settlement dependencies:** Where agents transact using tokenised money or assets, the resilience of the end-to-end payment chain also depends on the underlying smart contracts, signing arrangements and ledger. Institutions should independently review and test relevant contracts, protect signing keys through secure infrastructure with tested recovery and rotation procedures, and understand the ledger's operating and finality model before executing transactions. These established disciplines become particularly important where an agent can initiate transfers within delegated authority.

- **Concentration and correlation:** Where multiple institutions use agents based on similar models, data and objectives, individually rational decisions may produce correlated market behaviour. This risk may require industry and supervisory monitoring in addition to institution-level controls.

Across all four areas, institutions should be able to suspend an agent immediately, continue the underlying process without it, and reconstruct the agent's actions and decision basis. These capabilities are essential to operational resilience, investigation, and recovery.

4.4. Institutional control and risk management

Agentic payments should be governed through the institution's existing risk management framework, with ownership, oversight and controls proportionate to each use case.[24] Decision-making authority may be delegated to an agent, but risk ownership and accountability remain with the institution. The management body retains overall responsibility for effective governance, clear allocation of responsibilities, and an appropriate control environment.

Each use case should have a named authorised human owner, a defined risk classification, and an approved operating perimeter. This should specify the agent's permitted activities and limits, applicable human-approval thresholds, and clear arrangements for escalation, restriction, and suspension. Material changes to the agent's purpose, mandate, model, data, or execution authority should trigger reassessment under the relevant change-control process. Where a tokenised settlement asset is used, the operating perimeter should also specify the eligible assets and ledgers and the asset-level controls on which the institution relies.

Human oversight should not depend solely on the agent's own explanation of its actions. Emerging research indicates that AI systems may produce misleading outputs, omit material information, or behave differently when subject to oversight. Even where this behaviour is not deliberately designed, it may prevent reviewers from forming an accurate view of the agent's decision or compliance with its mandate.[25]

Institutions should independently verify agent decisions against source data, transaction records and externally enforced controls. Testing and monitoring should identify material omissions, inconsistencies between explanations and observed actions, and attempts to avoid escalation. Uncertain, anomalous or high-risk cases should remain subject to independent human review.

The institution should be able to demonstrate how the agent was authorised, which controls were applied, which institution remains legally accountable for the outcome, and which authorised human owner oversaw the mandate. Ongoing management information should cover exceptions, overrides, rejected instructions, control failures, losses, and activity outside expected patterns. Arrangements for errors, disputes, and recourse should also be defined, including where an agent acts within its approved mandate but produces a commercially adverse outcome.

5.

A progressive pathway to institutional adoption

The near- to medium-term opportunity for wholesale finance lies between **AI-assisted decision-making** and **bounded agentic execution**, rather than unconstrained autonomy. Institutions should therefore introduce agentic payments progressively, beginning with narrowly defined use cases and expanding delegated authority only where evidence demonstrates business value, reliable performance, and effective control.

A practical adoption pathway comprises four stages:

1. **Assist**: The agent analyses information and recommends an action, while a person or established institutional process retains decision and execution authority.
2. **Orchestrate**: The agent prepares and coordinates the transaction across relevant systems, while independent controls and human approval determine whether execution may proceed.
3. **Execute within mandate**: The agent selects or initiates approved transactions within explicit limits. Human involvement is triggered by transaction value, risk, uncertainty, or exception.
4. **Scale against evidence**: The institution expands permitted products, counterparties, transaction values, or operating hours only after demonstrating reliable outcomes, effective controls, operational resilience, and sufficient legal certainty.

Progression should be use-case-specific and supported by clear entry, review and exit criteria. Each stage should define the agent's authority, transaction limits, approval requirements, and identify the institutional principal and authorised human owner. Existing lifecycle, change-control, suspension, and continuity requirements described in Sections 3 and 4 continue to apply. Higher-autonomy and multi-agent arrangements should remain longer-term possibilities, introduced only where the incremental value justifies the additional complexity and risk.

Appendix

A.1 Glossary

Term	Definition
AI agent	Software that interprets an objective, assesses relevant information, selects among permitted actions, and uses connected systems or tools to pursue that objective.
Agent operator	The person or entity responsible for operating, maintaining or supporting the agent's technical environment. The agent operator may be different from the authorised human owner and does not acquire the institutional principal's transaction authority solely by operating the agent.
Agentic payment	A payment in which a software agent selects or initiates a transaction within authority delegated by a legally accountable institutional principal acting through an authorised human owner.
Authorised human owner	A natural person formally authorised by the institutional principal to define the agent's objective, approve and revoke its mandate, oversee its use and manage or escalate exceptions.
Bounded agentic execution	An arrangement in which an agent may select or initiate an action only within an explicit mandate and subject to independent institutional controls.
Counterparty eligibility	The rules determining whether an institution, client or other party is permitted to participate in a transaction.
Deterministic automation	Automation that executes a predefined action when specified conditions are met, without interpreting a broader context or selecting among alternatives.
Exception handling	The process for identifying and resolving transactions that cannot proceed normally because of missing information, control failures, or other unexpected conditions.
Higher-autonomy orchestration	An arrangement in which an agent plans and coordinates multiple related actions across systems, transactions, or counterparties, with human involvement focused mainly on oversight and exceptions.
Institutional principal	The institution on whose behalf an agent acts, and which retains legal and regulatory responsibility for transactions executed under the delegated mandate.

Term	Definition
Machine-readable mandate	A digital representation of an agent's permitted activities, limits, conditions, and approval requirements that systems can interpret and enforce automatically.
Programmable payment	A payment executed automatically when predefined conditions are met. Unlike an agentic payment, it does not necessarily require interpretation or action selection by an AI agent.
Settlement rail	The infrastructure or network through which funds or assets are transferred, and settlement is completed.
Stablecoin	A digital token designed to maintain a stable value relative to a reference asset, typically an official currency such as the euro or US dollar.
Straight-through processing	End-to-end processing of a transaction with limited or no manual intervention or re-entry of information.
Tokenised asset	A digitally represented financial or real-world asset recorded and transferred through token-based infrastructure.
Tokenised deposit	A digital representation of a commercial bank deposit issued and recorded using token-based infrastructure.

Sources

1 Deutsche Bank (2026), 'Digital Money: A Perspective on Stablecoins, Tokenised Deposits, and CBDCs', *flow*

2 Restrepo Amariles, D., Charlotin, D. and He-Guelton, L. (2026), 'AI Agents in Payments: Applications, Risks and Regulations', European Journal of Risk Regulation

3 Remolina, N. (2026), 'Agentic Payments: When Is a Payment (Un)Authorised?', Singapore Management University School of Law Research Paper, forthcoming

4 Acharya, V. (2025), 'Secure Autonomous Agent Payments: Verifying Authenticity and Intent in a Trustless Environment', Journal of LATEX Class Files, VOL.1, No.1

5 Aldasoro, I. and Desai, A. (2025), 'AI Agents for Cash Management in Payment Systems', BIS Working Papers, No. 1310, Bank for International Settlements

6 Davidovic, S. and Tourpe, H. (2026), 'How Agentic AI Will Reshape Payments', IMF Note No. 2026/004, International Monetary Fund

7 World Economic Forum on Artificial Intelligence (2026)

8 Ant Digital Technologies (2026), 'Real Up Cannes Summit Unveils 'Anvita' to Power the Agent-to-Agent Economy,' PR Newswire

9 Deutsche Bank (2026), 'Cross-border Payments in 2040: Five Hypotheses for a Multi-Rail Future', *flow*

10 BIS Annual Economic Report (2025), Chapter III, 'The next-generation monetary and financial system'. Bank for International Settlements

11 Citi Institute (2025), 'AI in Investment Management: Beyond Efficiency Gains', Citi Technology & Innovation

12 Boston Consulting Group (2026), 'Rebuilding Asset Management for an AI-First World'. Global Asset Management Report 2026

13 'Regulation (EU) 2024/1689' (Artificial Intelligence Act), Article 26(2), requires deployers of high-risk AI systems to assign human oversight to natural persons with the necessary competence, training and authority, as well as the necessary support

14 Davidovic, S. and Tourpe, H. (2026), 'How Agentic AI Will Reshape Payments', IMF Note No. 2026/004, International Monetary Fund

15 Deutsche Bank Research Institute (2026), 'Tokenised Money and the Future of FX'. The paper reviews existing and emerging settlement arrangements, including corresponding banking, domestic payment systems, RTGS and CLS, alongside stablecoins, tokenised commercial bank money, tokenised central bank reserves, hybrid models, shared-ledger platforms, and synchronisation infrastructure

16 Committee on Payments and Market Infrastructures (CPMI) (2023), 'Considerations for the Use of Stablecoin Arrangements in Cross-Border Payments', CPMI Papers No. 220, Bank for International Settlements

17 Committee on Payments and Market Infrastructures and International Organisation of Securities Commissions (2012), 'Principles for Financial Market Infrastructures'. The IMF Note observes that where the settlement layer corresponds to a financial market infrastructure, its operation is subject to these principles, which set standards for governance, risk management, access, and settlement finality

18 World Economic Forum (2026), proposing a Know Your Agent framework alongside existing Know Your Customer requirements, 'AI agents could be worth $236 billion by 2034 – if we ensure they are the good kind'; the IMF Note observes that OAuth 2.0 and OpenID Connect are already used to support attested agent identity

19 Ant Digital Technologies (2026), 'Real Up Cannes Summit Unveils 'Anvita' to Power the Agent-to-Agent Economy'

20 Clifford Chance and Deutsche Bank (2025), 'The Convergence of AI and Distributed Ledger Technology: Opportunities and Risks'. This paper examines practical implementation considerations and the legal, commercial and technical challenges arising from the combined use of AI and DLT

21 FIDO Alliance (2026), 'FIDO Alliance to Develop Standards for Trusted AI Agent Interactions'; Google (2026), 'Agent Payments Protocol and the FIDO Alliance', announcing AP2 v0.2 and its contribution to the FIDO Alliance

22 Visa (2026), 'Visa opens the door to AI-driven shopping for businesses worldwide', provides a single integration point that is protocol-agnostic, supporting the Trusted Agent Protocol, x402, the Machine Payments Protocol, the Agentic Commerce Protocol, and the Universal Commerce Protocol through one integration

23 'Regulation (EU) 2022/2554' on digital operational resilience for the financial sector (DORA) applies to financial entities' information and communication technology arrangements, including incident management, resilience testing, and third-party risk

24 The three lines of defence should apply: first-line management owns and manages agentic payment risks and controls; second-line risk and compliance functions set standards, provide oversight and challenge; and third-line internal audit provides independent assurance

25 Park, P. S., Goldstein, S., O'Gara, A., Chen, M. and Hendrycks, D. (2024), 'AI Deception: A Survey of Examples, Risks, and Potential Solutions', Patterns, Vol. 5, No. 5, Article 100988

ANT GROUP | DIGITAL TECHNOLOGIES

AGENTIC PAYMENTS IN WHOLESALE FINANCE

Exploring use cases
value potential
and conditions for responsible adoption