Cash in Sight, Rights in Hand: Valuing Real-Time Treasury AI

A Valuation and Transaction Framework for Cash Visibility, Forecasting, Payment Control, Liquidity Decisions and Transferable Rights

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Real-time treasury technology promises a continuously updated view of cash, more accurate forecasts, lower idle balances, better funding decisions and faster intervention. Artificial intelligence can classify transactions, predict receipts and payments, identify anomalies, recommend transfers and support liquidity decisions across accounts, entities, currencies and banking relationships. The economic value of those capabilities depends on a harder question: whether the buyer acquires reliable sight of cash and enforceable rights to use data, issue instructions and operate the decision system after closing.

This paper develops a valuation and transaction framework for acquisitions of treasury-management platforms, cash-forecasting applications, bank-connectivity providers, payment-control systems and AI-enabled liquidity businesses. It separates observation from prediction, recommendation from approval, and approval from execution. It connects bank and enterprise data, ISO 20022 messages, account permissions, model governance, payment controls, reconciliation, fraud management, operational resilience and customer adoption to unit economics and enterprise value.

The analysis draws on current material from the Bank for International Settlements, Committee on Payments and Market Infrastructures, Basel Committee, Financial Stability Board, Federal Reserve, European Central Bank, Bank of England, European Banking Authority, European Union institutions, United Kingdom regulators, standards organisations and accounting and valuation bodies.[1][2][3][4][5][6][7][8][9][10] These sources establish that instant settlement and richer data can improve speed and visibility while liquidity, operational, cyber, fraud and governance obligations remain material.

A hypothetical acquisition illustrates a platform connected to 180 legal entities, 1,600 bank accounts and USD 8.0 billion of annual payment value. Every customer count, account count, payment value, forecast error, basis point, adoption rate, cost, multiple, probability and valuation amount is a management assumption created solely to demonstrate the method. None is a forecast, market benchmark or valuation opinion.

The paper concludes that a buyer should pay for real-time treasury AI only when the target proves current and reconciled cash coverage, stable forecast performance, controlled payment authority, lawful and transferable data rights, resilient connectivity, measurable customer adoption and durable contribution. Six figures and seven tables translate this conclusion into a diligence plan, unit-economic bridge, forecast-vintage analysis, model-governance test, valuation method, transaction protections and a 180-day programme.

Payments, banking, data, artificial-intelligence, competition, tax, accounting and corporate-law requirements vary by jurisdiction. Qualified specialists should determine the rules, permissions and contractual rights applicable to a specific transaction or operating model.

JEL Classification: G21, G23, G32, G34, O33

Keywords: treasury, artificial intelligence, cash forecasting, real-time payments, liquidity, payment controls, valuation, mergers and acquisitions

JEL Classification: G21, G23, G32, G34, O33

Keywords: treasury, artificial intelligence, cash forecasting, real-time payments, liquidity, payment controls, valuation, mergers and acquisitions

1. Define the acquisition decision

The investment question is whether the platform converts timely cash information into transferable, controlled and collected economics. A buyer may see impressive dashboards, high transaction volumes and machine-learning claims. Those observations do not prove that account coverage is complete, balances reconcile, forecasts remain accurate through stress, recommendations improve funding outcomes or the combined firm can exercise the same data and payment rights after closing.

The board should define the product perimeter before debating a revenue multiple. A target may aggregate balances, classify transactions, forecast cash, recommend funding, initiate payments, screen beneficiaries, route approvals or execute transfers through bank and payment-system connections. Each layer has different dependencies, permissions, liabilities and switching costs. A product that only observes cash can be valuable, while its valuation should not include execution economics unless the right and capability to transact are evidenced.

The acquisition thesis should be expressed as a testable chain. More complete and timely data should improve the forecast. A better forecast should reduce avoidable buffers, emergency borrowing, overdrafts, failed payments or manual work. Those benefits should reconcile to customer retention, pricing and collected contribution after connectivity, model, security, support, fraud, insurance and regulatory costs. Diligence should identify evidence capable of disproving every link.

Value should be divided into existing evidenced contribution, protected value that depends on transferable rights and service continuity, improvement value with funded actions, and future option value. Forecast expansion into autonomous transfers, embedded working capital or cross-border optimisation should remain outside the central case until authority, control and customer evidence support it.

Investment-committee evidence pack

The investment committee should receive one reconciled evidence pack. It should define legal entities, active customers, connected accounts, banks, currencies, message types, data latency, forecast vintages, payment authority, revenue, direct cost, incidents, losses, models, vendors and service levels using consistent dates and populations. Each valuation assumption should have an owner, evidence source and falsification test.

Sampling should precede management curation. The buyer can combine random customers with high-value, multi-bank, multi-currency, recently onboarded, highly automated, loss-affected and churned cohorts. For each sample it should trace selected balances, transactions, forecasts, recommendations, approvals, payments and reconciliations to source records. Failed imports, rejected payments and abandoned workflows belong in the population.

The decision paper should state which value survives a change of control. Bank consents, customer authorisations, API credentials, data-processing rights, model licences and cloud contracts can determine whether the service continues. A strong historical product can lose value if the acquirer cannot lawfully receive the data, renew the connection or operate the payment mandate.

2. Separate visibility, intelligence and transaction authority

Treasury automation is a stack. Visibility gathers balances and transactions. Intelligence classifies flows, forecasts positions and proposes actions. Orchestration routes approvals and instructions. Execution transmits an authorised payment or liquidity transfer. Reconciliation confirms settlement and updates the ledger. These layers should be valued separately.

Visibility value depends on coverage, timeliness and reconciliation. A dashboard that refreshes quickly from a subset of accounts may look real-time while missing material cash. Screen scraping, host-to-host files, SWIFT messages, open-banking interfaces and direct APIs can carry different data fields, frequencies and contractual rights. The buyer should measure economic coverage, not connection count.

Intelligence value depends on decision performance. Forecasting should be assessed by horizon, entity, currency, flow class and business condition. A model may predict regular payroll accurately and fail on tax, acquisitions, margin calls or concentrated customer receipts. Aggregate error can conceal offsetting mistakes that still cause local liquidity shortages.

Transaction authority changes the risk perimeter. A recommendation can be reviewed; an executed transfer can create immediate and potentially irreversible loss. The system needs authenticated users, segregated duties, approved beneficiaries, limits, sanctions and fraud controls, exception routes, confirmation and audit evidence. Rights to analyse data do not automatically include rights to initiate payment.

Table 1. Treasury capability and rights matrix

Layer	Principal function	Required right or evidence	Main valuation exposure
visibility	aggregate balances and transactions	customer mandate, bank access and reconciliation	incomplete or stale cash coverage
classification	identify flow type and counterparty	lawful data use and labelled history	weak forecast inputs and manual cost
forecasting	estimate future positions	model rights, vintages and outcome history	unstable decision quality
recommendation	propose transfer, funding or investment	policy logic and explainable rationale	unsuitable or uneconomic action
approval	apply authority and segregation	mandate, role, limit and authentication evidence	unauthorised instruction risk
execution	transmit payment or liquidity order	bank and scheme permission, security procedure	fraud, finality and operational loss
reconciliation	confirm settlement and ledger state	complete status and accounting data	false cash position and control failure

Proposed framework; transaction-specific legal and regulatory analysis remains required.

3. Reconstruct the cash evidence chain

Cash in sight is an evidence chain rather than a screen value. It starts with an identified legal entity and account. It connects the bank-reported balance, available funds, pending items and value date to imported transactions, enterprise records and intercompany positions. It then feeds the forecast, proposed action, approval, execution, settlement confirmation, accounting entry and realised liquidity outcome.

The chain must preserve content, time and provenance. A closing ledger balance may differ from available cash because of holds, sweeps, overdraft facilities, uncleared items or cut-off rules. A real-time payment can settle while the enterprise ledger remains unchanged. An API timestamp can show when data was

received without proving when the underlying position became effective. The buyer should define each cash measure and reconcile it to an authoritative source.

The Federal Reserve describes FedNow as a 24x7x365 service that clears and settles near-real-time transfers and includes a liquidity-management capability.[3] The European Central Bank describes TIPS as a 24x7x365 platform settling instant payments in central bank money, with liquidity transfers and ISO 20022 messages.[6][7] Continuous infrastructure changes the treasury day. Weekends and holidays become operating periods, and controls designed around a daily bank file can become stale before the next opening.

The transaction team should choose representative days and reconstruct them minute by minute. Tests should include ordinary operations, payroll, tax, debt service, a large receipt, a failed connection, a fraudulent instruction, a currency shortage and a market disruption. The objective is to determine when the platform knew, what it predicted, what it recommended, who authorised action and what settled.

Coverage and reconciliation protocol

Coverage should be measured by economic exposure. The denominator can include average and peak cash, payment value, forecasted obligations and material legal entities. Account-count coverage is a secondary measure because many low-value accounts can conceal one missing concentration account.

Reconciliation should distinguish imported, matched, classified, forecasted and settled transactions. Each stage needs an exception population. Management reporting that excludes rejected or unmatched records can overstate straight-through processing and understate support cost.

Evidence should be retained at source-record level. The platform should preserve message identifiers, timestamps, model version, input snapshot, recommendation, approval, bank response and final status. A buyer that cannot reproduce historical decisions cannot reliably validate performance or investigate loss.

Figure 1. Proposed cash-to-outcome evidence chain

The chain separates information, decision, authority, settlement and realised economics.

4. Measure real-time coverage honestly

The label real-time can describe several different clocks. A bank may make data available continuously while the platform polls every fifteen minutes. A platform may ingest instantly and refresh the user interface later. A payment may settle in seconds while the accounting system posts overnight. Valuation should follow the slowest component needed for the decision.

The buyer should construct a latency distribution from source event to actionable state. Median latency is insufficient because treasury loss often sits in the tail. Measures should include the ninety-fifth and ninety-ninth percentiles, maximum outage, stale-account rate and duration until reconciliation. Results should be segmented by bank, connection, currency, geography and time period.

Completeness also matters. ISO 20022 can provide structured and richer payment data, while implementation and field use vary. CPMI work on harmonisation recognises that consistent data requirements support cross-border payments.[8] The platform should show which fields arrive, which are mapped, which are discarded and which models depend on them. A message standard does not guarantee semantic consistency.

Coverage economics include onboarding and maintenance. Each new bank or enterprise system can require security review, certificates, mappings, testing and exception handling. A high gross margin calculated before connection operations can be misleading. The buyer should allocate recurring connectivity and data-quality cost to cohorts and test whether margins improve with scale.

Real-time implementation diligence

The buyer should obtain a complete connection inventory and reconcile it to active revenue. For each connection it should record institution, legal entity, account population, interface, protocol, message version, authentication method, refresh frequency, operating window, data fields, service owner, certificate expiry, incident history and termination terms. The inventory should identify connections marketed as live while dependent on batch files or manual intervention.

Raw logs should support a latency analysis. The team should select an ordinary period, a month-end, a weekend and an incident period. It should calculate time from bank event to ingestion, normalisation, model availability and user presentation. Missing observations should remain visible. The same exercise should test whether forecasts and alerts were recomputed after late data arrived.

Displayed balances should be compared with bank statements and available-funds information. Differences need a cause code: timing, hold, sweep, overdraft, pending item, currency conversion, duplicate, missing transaction or mapping error. Management should show how users are warned when a position is incomplete. A timestamp without a materiality assessment can create false confidence.

5. Test forecast vintages rather than one accuracy number

A cash forecast is useful only in relation to a decision horizon. Same-day liquidity, seven-day funding, monthly working capital and annual planning require different inputs and tolerances. The buyer should reconstruct forecast vintages: the estimate made at each prior date for the same future cash position.

Error should be measured using several lenses. Absolute error shows magnitude. Percentage error becomes unstable near zero. Directional error identifies whether the platform repeatedly overstates or understates cash. Quantile loss can test whether stated confidence ranges are calibrated. Liquidity-weighted error assigns greater importance to shortfalls that trigger borrowing, failed payments or covenant pressure.

The analysis should separate predictable and judgemental flows. Payroll, rent and contracted debt may be schedule-driven. Customer receipts, tax, acquisitions, dividends and exceptional capex may depend on business events. An AI model can improve recurring-flow classification while a structured human input process remains essential for material one-off events.

Back-testing must use the information available at the forecast date. Reconstructed forecasts that include later invoices or settlement outcomes create leakage. The buyer should preserve input snapshots and model versions, then compare original forecasts with actual bank and ledger outcomes.

Figure 2. Hypothetical forecast-vintage error curve

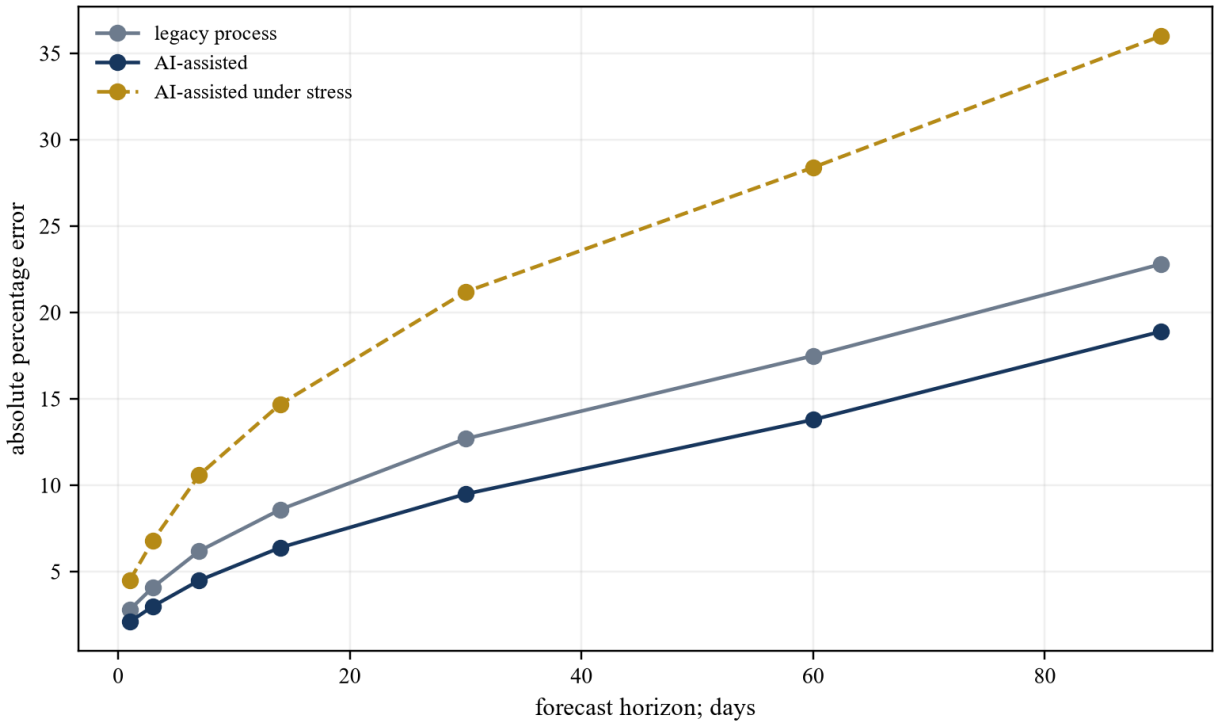

Wholly hypothetical percentage error; the figure is methodological and is not a market benchmark.

Forecast-governance evidence

The model inventory should identify purpose, owner, version, features, training period, validation, limits and downstream decisions. Forecast accuracy should be linked to adoption. A technically strong forecast creates little value if treasury teams override it, export results to spreadsheets or cannot explain it to approvers.

Override analysis needs context. Frequent overrides can show weak model quality, missing event data or user distrust. Rare overrides can show good performance or automation bias. The buyer should test who overrides, why, whether the change improves the outcome and whether lessons return to the model and process.

Table 2. Forecast-vintage diligence catalogue

Test	Segmentation	Evidence	Valuation relevance
absolute error	horizon, entity, currency and flow	original vintage and actual outcome	decision quality and retention
directional bias	normal and stress periods	signed error distribution	buffer and funding cost
interval calibration	forecast confidence band	frequency inside stated range	reliability of scenario use
tail error	largest shortfalls and excesses	event reconstruction	loss and liquidity exposure
data leakage	feature availability by timestamp	immutable input snapshot	validity of claimed performance
override value	user, reason and outcome	before-and-after decision	adoption and human-control quality
model drift	period and business change	stability and revalidation record	maintenance cost and durability

Proposed tests; thresholds should reflect the customer's decisions and risk appetite.

6. Connect prediction to economic outcome

Forecast accuracy is an intermediate metric. Economic value appears when a decision changes: cash is concentrated, borrowing is reduced, deposits are placed appropriately, foreign exchange is funded, a payment is rescheduled, a facility is drawn in time or manual work is avoided. The buyer should identify the counterfactual for each claimed benefit.

Idle-cash reduction needs care. A lower balance may reflect improved forecasting, business contraction, changed risk appetite or a treasury policy change. The platform should show matched cohorts or decision-level evidence connecting its recommendation to released cash without increasing failures or emergency funding.

Interest benefit should use actual rates, balances and days rather than a headline annual percentage applied to all cash. Borrowing avoided should exclude undrawn facilities that remained necessary for resilience. Working-capital benefit should not be attributed to the forecast engine when commercial teams changed payment terms or collections.

Manual-efficiency claims should reconcile to process activity and cost. Fewer spreadsheets or touches can support value, while control tasks may move into model validation, exception handling or connection support. The full operating model should be compared before and after deployment.

Benefit-attribution protocol

Each material benefit should have a baseline, intervention, outcome, counterfactual and evidence owner. For released cash, the baseline can be the policy buffer before deployment; the intervention the model-supported change; the outcome the actual balance and funding position; and the counterfactual the balance required under the previous process. The analysis should record concurrent policy and business changes.

Matched cohorts can strengthen attribution. Customers or entities with similar scale, volatility and banking complexity can be compared across adoption periods. Where selection bias remains, the valuation should use a conservative range. Customers who adopt most deeply may already have stronger treasury functions and better data.

Benefits should reconcile to financial records. Interest saved should connect to facilities and statements. Fees avoided should connect to bank charges. Labour savings should connect to roles, capacity or outsourced cost. Loss avoided requires an evidenced event and credible counterfactual. Vendor estimates can support a hypothesis while collected outcomes provide stronger evidence.

7. Treat payment rights as a core intangible asset

The right to see an account, analyse its data, initiate an instruction and execute a payment can arise from different contracts and technical credentials. Customer terms, bank agreements, scheme rules, data-protection law, power of attorney, user roles and security procedures may all be relevant. The buyer should map each right to the entity that holds it and test change-of-control consequences.

Credentials do not equal permission. An API token may technically access an account while contractual use is limited to a named customer or purpose. Historical data may be retained for service delivery but unavailable for model training or buyer integration. A model licence may permit hosted inference while prohibiting transfer of weights or use outside the current cloud environment.

The rights register should include data provenance, controller and processor roles, permitted purpose, retention, localisation, sub-processors, bank consent, termination, portability and audit evidence. Rights should be connected to revenue cohorts so the valuation can identify cash flows at risk.

Payment authority requires stronger evidence. The buyer should inspect signing rules, mandates, limits, beneficiary controls, dual authorisation, emergency access, certificate ownership and revocation. It should

test whether the platform can continue operating if a founder, bank sponsor or third-party integrator leaves.

Rights diligence and consent map

The transaction team should create a contract-to-capability map. For each material customer and bank, it should identify the contracted service, data classes, permitted processing, payment role, intellectual-property allocation, sub-processing, audit right, service level, liability, termination, assignment and change-of-control provision. The map should link directly to revenue and contribution.

Consent risk should be quantified. The team should identify contracts requiring prior consent, notice, credential replacement or re-papering. It should estimate timing, customer effort and the economic effect of refusal or delay. A consent plan should name relationship owners and follow the transaction's confidentiality and communication plan.

Intellectual-property diligence should trace employee and contractor assignments, open-source components, training data, third-party models, code repositories and deployment artefacts. The buyer should be able to build and operate the service without undocumented personal knowledge or credentials.

The valuation should use a rights waterfall. Fully transferable capabilities enter the central case. Capabilities requiring routine notice can enter with implementation cost. Material consents can receive probability weighting or contingent consideration. Capabilities that cannot transfer should be valued through replacement cost and delay, with dependent synergies removed.

Table 3. Transferable-rights register

Asset or capability	Evidence of right	Change-of-control test	Value response if incomplete
bank account data	customer mandate and bank terms	consent, notice and credential reissue	defer connected revenue value
enterprise data	integration contract and purpose	buyer access and migration right	exclude dependent model benefit
historical model data	provenance and lawful basis	continued training and validation use	reduce model and option value
forecast model	ownership, licence and dependencies	transfer, hosting and modification rights	replacement cost and delay
payment initiation	mandate, role and limit evidence	bank and scheme acceptance	exclude execution premium
cloud and security service	contract, controls and exit plan	assignment and continuity	resilience and migration deduction
customer workflow	product terms and audit record	continuation without re-papering	retention and implementation reserve

Proposed register; enforceability remains subject to contract and applicable law.

8. Govern model recommendations and autonomous action

Treasury AI can classify, predict, optimise and generate explanations. These functions should not share one control standard. A classifier affects data quality. A forecast affects a future view. An optimiser recommends allocation or funding. An agent that initiates action can move money. Materiality rises as the system gains authority and as reversibility falls.

The buyer should test the full decision system: data transformations, forecast model, liquidity policy, constraints, objective function, recommendation logic, approval workflow and payment interface. A

statistically accurate forecast can still produce a poor action if limits are missing, costs are stale or the objective rewards yield without preserving operational cash.

Deterministic controls should bound probabilistic components. Account ownership, approved beneficiary, legal-entity authority, payment limit, sanctions result, available balance and segregation of duties should not depend on a language model following a prompt. Generative models can summarise evidence or support investigation while control-critical checks remain versioned, testable and reproducible.

The 2025 BIS study on AI agents for cash management reports experimental evidence that a general-purpose model could preserve buffers, prioritise payments and balance liquidity cost against delay in simulated high-value payment scenarios.[1] The study also identifies the need for safeguards, human oversight and further research. Transaction valuation should therefore distinguish demonstrated experimental capability from production evidence in the target's own operating environment.

Model validation and action tiers

Validation should cover conceptual soundness, data lineage, implementation, performance, stability, explainability, security and use within the treasury process. Independence requires competent challenge and authority to restrict use. A vendor report can support diligence while it does not replace buyer testing on representative target data.

Action tiers can define increasing authority. Tier one observes and explains. Tier two forecasts. Tier three recommends. Tier four prepares an instruction for human approval. Tier five executes within pre-approved limits. Each tier should have evidence requirements, limits, monitoring, incident response and a clear accountable owner.

Decision-system test catalogue

Test cases should include ordinary and boundary conditions. Examples include incomplete balances, contradictory enterprise and bank records, delayed receipt, duplicate invoice, beneficiary change, unusual time, new device, currency shortfall, facility limit, weekend settlement, payment-rail outage and model-service failure. Expected behaviour may be a forecast, warning, restricted recommendation, enhanced approval or stopped action.

Explanations should match the actual decision logic. Generated text that sounds plausible while omitting a binding constraint creates control risk. The audit record should show inputs, calculations, constraints, versions, recommendation, human action and outcome. Reproduction should not depend on a mutable external service without retained evidence.

Figure 3. Proposed authority ladder for treasury AI

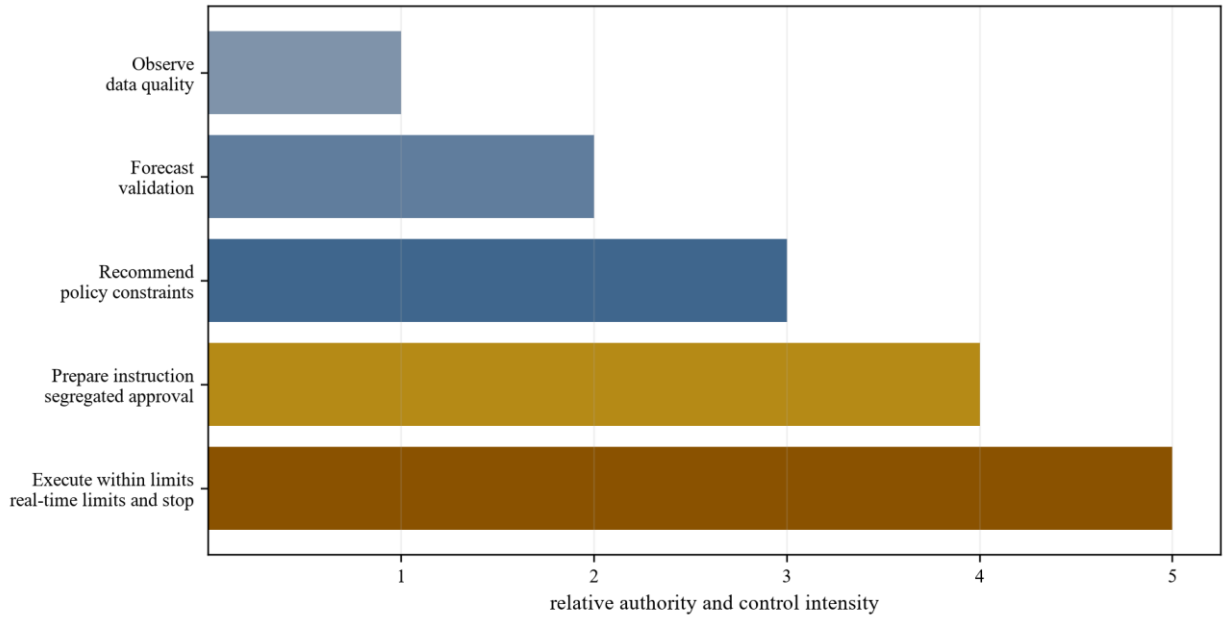

Control intensity should increase with authority, materiality and irreversibility.

9. Price fraud and payment loss into the model

Faster settlement compresses the time available to detect and stop fraud. A strong treasury platform should combine beneficiary verification, authentication, behavioural analytics, device and session evidence, transaction monitoring, sanctions controls, limits and human escalation. The buyer should inspect how these controls interact rather than counting features.

Loss data should be reconciled from alert to final economic outcome. Gross attempted value, prevented value, executed value, recovered value, customer reimbursement, insurance recovery and net loss are different measures. Alert precision, investigation time and false-positive cost affect both customer experience and operating margin.

The CPMI has identified fraud as a priority in cross-border fast payments and describes manipulation of payers, stolen credentials and altered instructions as relevant fraud forms.[9] Controls therefore need to cover authorised push-payment scenarios as well as account compromise. A technically authenticated payment can still result from deception.

The buyer should examine model and policy performance during change. New banks, payment rails, customer segments, currencies and user interfaces can alter fraud patterns. Integration can weaken established controls if identity, beneficiary history or device information is lost. Transaction protections should address known losses, open claims, control gaps and unseasoned cohorts.

Fraud-diligence work programme

The diligence team should reconcile alerts, cases, instructions, settlements, complaints, reimbursements, recoveries and insurance. Populations should use stable identifiers so a loss cannot disappear when it moves between operational and accounting systems. The analysis should include near misses because they reveal exposure without waiting for realised loss.

Control testing should cover enrolment and change. A legitimate user can be compromised after onboarding, and an approved beneficiary can be altered. Tests should inspect credential reset, device binding, privileged administration, beneficiary creation, limit change, approval routing and emergency access. Dual approval is ineffective if one administrator can change both beneficiary and approver populations.

Model metrics should be connected to investigation capacity. High recall with excessive false positives can delay payments or cause analysts to bypass alerts. Precision can appear strong if the target investigates only selected cases. The buyer should review sampling, queue ageing, escalation and quality assurance.

The economic model should include expected loss, investigation cost, reimbursement, insurance premium, deductible, coverage limit and uninsured scenarios. Past low loss may reflect a small or low-risk population. Expansion into execution, new geographies or higher payment limits should be treated as a new risk cohort until evidence seasons.

10. Measure intraday liquidity and buffer value

Real-time treasury AI can create value by reducing uncertainty around when cash is required, while it cannot abolish liquidity risk. Payment systems and firms need sufficient resources to meet obligations as they fall due. CPMI-IOSCO principles emphasise ongoing measurement and monitoring of settlement and funding flows, including intraday liquidity.[2]

The buyer should distinguish operational cash, precautionary buffer, trapped cash, regulatory liquidity, collateral, restricted balances and investable surplus. Releasing one category may be feasible while another remains unavailable. Currency and legal-entity constraints can prevent group cash from meeting a local obligation.

Liquidity benefit should be measured against service resilience. A platform that reduces buffers by assuming continuous connectivity may increase loss when a bank, cloud provider or payment rail fails. Stress tests should include delayed receipts, concentrated outflows, market closure, unavailable credit, currency disruption, fraud hold and data outage.

The decision engine should make its cost function visible. Delaying a payment can save liquidity and damage a supplier relationship. Drawing a facility can preserve settlement and incur fees. Investing surplus can add yield and reduce immediate access. The board should know which costs and limits the optimiser uses and who can change them.

Liquidity-scenario reconstruction

The buyer should reconstruct a full operating day for selected entities and currencies. Opening available cash, committed inflows, expected outflows, collateral, facilities and limits should reconcile to actual messages and statements. The analysis should show which obligations were time-critical and which could be delayed without contractual or commercial harm.

Intraday positions need more than end-of-day evidence. A firm can finish positive after experiencing a material shortfall. The team should calculate peak usage, minimum available balance, duration below policy buffer, facility draw timing and payment queue. It should compare the target's recommendation with the action taken and realised result.

Cross-entity optimisation should respect law, tax, covenant and operational constraints. Cash pooling, intercompany loans, notional structures and guarantees can have consequences beyond yield. The platform should represent restrictions explicitly and escalate exceptions to qualified decision-makers.

Stress liquidity should remain conservative. The value case can recognise verified reductions in avoidable buffer while preserving resources for plausible shocks. Benefit should not be created by assuming that a facility, market or payment system is available precisely when the scenario tests its absence.

Table 4. Intraday-liquidity evidence and valuation response

Claim	Required test	Economic measure	Valuation treatment
lower idle cash	matched entity and period	verified average released balance	capitalise only durable after-control benefit

Claim	Required test	Economic measure	Valuation treatment
fewer emergency draws	original forecast and facility record	avoided fees and interest	adjust for facility availability cost
fewer payment failures	complete instruction population	loss, fee and disruption avoided	use mature observed cohorts
improved cash concentration	legal-entity and currency test	usable cash transferred	exclude trapped or restricted balances
better intraday timing	timestamp reconstruction	overdraft and delay cost	test tail days and stress periods
higher investment yield	executed placement and maturity	collected net yield	deduct risk, liquidity and operating cost

Proposed framework; liquidity policies and constraints are institution-specific.

11. Test ISO 20022 and semantic data quality

ISO 20022 creates a common message framework and can carry richer structured information. Value depends on implementation. Banks and payment systems can populate fields differently, truncate data, map legacy formats or apply local usage rules. The platform needs a semantic layer that preserves provenance and exposes uncertainty.

The buyer should inspect the canonical data model, mapping rules, version control and rejection handling. It should select common and unusual messages, then trace fields from source through normalisation, model features, user display and export. Null, default and inferred values should remain distinguishable.

Structured remittance data can improve matching and forecasting. It can also contain personal or commercially sensitive information. Data minimisation, access, retention and security should follow purpose. The acquisition plan should identify which historical messages can migrate and whether the buyer can continue using them for analytics and model improvement.

Semantic quality has a direct support cost. Each bank-specific exception, manual mapping and unresolved field increases onboarding time and weakens automation. Unit economics should allocate this cost to cohorts rather than treating it as central research and development.

12. Rebuild unit economics after the control stack

Revenue can include subscription, account, entity, user, payment, transaction-value, implementation and premium-analytics fees. The buyer should reconcile contracted pricing to invoices, credits, collections and active usage. Annual recurring revenue should exclude non-recurring implementation and pass-through bank or network charges unless separately identified.

Direct cost should include bank connectivity, messaging, cloud, data, model inference, onboarding, mapping, customer support, payment operations, fraud investigation, security, compliance, insurance and loss. Sales commission and implementation subsidy should be matched to cohort economics. Costs often rise non-linearly as the platform acquires larger and more complex customers.

The hypothetical case has 180 customer entities, 1,600 connected accounts and USD 8.0 billion of annual payment value. Subscription and usage revenue is USD 18.0 million. Connectivity and data cost USD 2.4 million; cloud and model operations cost USD 1.6 million; onboarding and support cost USD 2.5 million; payment control, fraud and insurance cost USD 1.8 million; product, security and compliance operations cost USD 2.0 million. Contribution before central cost, tax and capital is USD 7.7 million.

Every amount is hypothetical. The example does not claim that scale, pricing or margin is achievable. Its purpose is to show that model and payment-control cost belong inside contribution rather than below a headline software margin.

Cohort profitability method

Cohorts should be segmented by customer size, bank count, geography, payment authority and onboarding period. Revenue retention alone can conceal costly connectivity or support. Contribution retention measures whether the economic relationship survives.

Implementation payback should use collected gross contribution. Capitalised implementation cost should not disappear from the acquisition model. The buyer should test whether onboarding effort falls with reusable connectors and mappings or rises as the product enters new banks and jurisdictions.

Revenue-quality and retention tests

The buyer should reconcile bookings, contracts, invoices, credits, collections and revenue recognition. Multi-year commitments should be assessed for termination, minimums, implementation dependencies and customer acceptance. Usage revenue should be separated from pass-through charges and volatile payment activity.

Retention should be presented by customer count, revenue and contribution. Gross revenue retention can remain high while expensive cohorts consume support and connectivity resources. Net retention can reflect price increases or payment volume rather than broader product adoption. Cohort bridges should explain expansion, contraction, churn, credits and cost movement.

Sales concentration should include channel and bank dependencies. Several customers acquired through one sponsor or enterprise platform may share a renewal risk. Pipeline value should remain outside the central case unless conversion evidence is mature and delivery capacity is funded.

Figure 4. Hypothetical annual contribution bridge

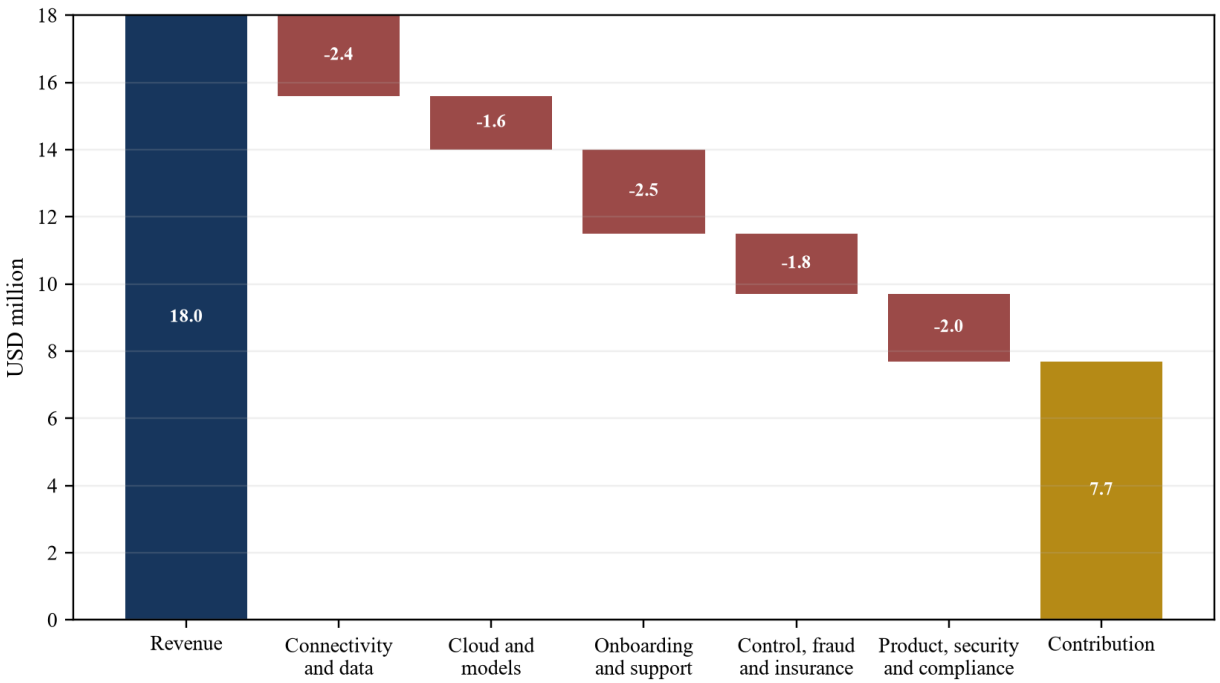

Wholly hypothetical USD millions; central cost, tax and capital remain outside the displayed contribution.

13. Measure adoption and decision quality together

Customer logins, connected accounts and payment volume show activity while they do not prove decision value. The buyer should measure whether treasury teams use forecasts, accept recommendations, complete approvals, reconcile exceptions and change funding behaviour. Adoption should be linked to outcome and contribution.

Shadow workflows matter. Customers may export a forecast and complete decisions in spreadsheets, messaging applications or bank portals. The platform can retain subscription revenue while lacking control of the economic workflow. Diligence should observe representative users and trace the complete process.

Adoption should be segmented by role. An analyst may use classification, a treasurer may use scenarios, a controller may approve payments and a chief financial officer may view liquidity. Loss of one critical role can reduce value even when monthly active users remain stable.

Product telemetry should respect customer rights and confidentiality. The buyer should confirm that analytics are lawfully collected and sufficiently precise for the intended conclusion. A click does not establish reliance, and absence of a click does not establish absence of value when information is delivered through an interface or API.

14. Test operational resilience and third parties

Real-time treasury depends on continuous systems. The critical path can include the customer's enterprise software, identity provider, connectivity vendor, payment network, bank, cloud platform, model service and support operation. The buyer should map dependencies and test failure at each boundary.

The Basel Committee's operational-resilience principles and third-party risk work emphasise governance, dependency management, incident response and continuity.[10][11] The European Union's Digital Operational Resilience Act establishes requirements around ICT risk, incidents, testing and third-party risk for covered financial entities.[12] Applicability depends on the target and service, but the operating evidence remains commercially relevant across transactions.

Service-level statistics should be reconstructed from raw monitoring and incidents. Contractual uptime can exclude maintenance and downstream bank failures. Mean availability can conceal a severe month-end outage. Recovery time should be tested for business service, data consistency and payment authority rather than infrastructure alone.

Exit plans need executable detail. The buyer should know how to export customer configurations, forecasts, approvals and audit records; replace a model or connectivity provider; revoke credentials; and continue critical payments. A plan without tested data and responsible owners provides weak valuation support.

15. Protect privacy, confidentiality and cyber security

Treasury data can reveal payroll, acquisitions, suppliers, financing, tax, distress and strategy. The buyer should map personal and corporate confidential information, processing purposes, locations, access, retention and onward sharing. Change of control and model-training use require specific review.

Cyber diligence should focus on the money-moving path. Identity, privileged access, secrets, certificates, code deployment, beneficiary data, approval rules and bank connections require strong controls and logs. Penetration testing is one input; secure design, monitoring, incident handling and recovery provide broader evidence.

AI introduces additional attack surfaces through prompts, training data, model endpoints and generated explanations. Control-critical values should be protected from untrusted text. The system should prevent a

payment instruction, beneficiary or policy limit from being changed through a conversational interface without deterministic validation and proper authority.

Data segregation should survive acquisition integration. Combining customer data sets can create attractive analytics and new restrictions. Synergy should remain excluded until lawful purpose, access, security and customer commitments support the proposed use.

16. Value the platform by evidence layer

A single revenue multiple can hide the reasons value exists. The buyer should triangulate discounted cash flow, comparable-company and transaction evidence, replacement cost, customer-cohort economics and scenario value. Each method should use consistent revenue, contribution, rights and risk assumptions.

The valuation can be organised into five layers. Layer one is collected stand-alone contribution. Layer two is protected value from transferable contracts, rights, connectivity and customer continuity. Layer three is evidenced improvement from funded operational actions. Layer four is buyer-specific synergy. Layer five is option value from new authority, products or geographies. Confidence and discounting should fall as evidence weakens.

Intangible assets require careful identification. Customer relationships, technology, data, contracts, licences and trade names can have different lives and transfer conditions. IFRS 3 and IAS 38 provide accounting frameworks for business combinations and identifiable intangible assets.[48][49] Purchase-price allocation does not itself determine investment value, but it can expose assumptions about separability, useful life and economic benefit.

Table 5. Evidence-layer valuation architecture

Layer	Evidence threshold	Valuation method	Typical protection
collected contribution	invoices, cash and direct cost reconcile	DCF and cohort economics	ordinary warranties
protected continuity	contracts, rights and service survive close	retention-adjusted DCF	consent conditions and covenant
evidenced improvement	funded action and measured baseline	probability-weighted benefit	completion funding and milestones
buyer synergy	named integration owner and capacity	buyer-specific NPV	excluded from seller consideration
option value	authority and market evidence remain incomplete	staged real-option analysis	contingent consideration

Proposed architecture; amounts and weights remain transaction-specific.

17. Apply a rights-and-control discount transparently

The valuation committee should avoid one undifferentiated risk premium. Specific deductions can reflect missing bank consents, weak data provenance, non-transferable model rights, forecast instability, payment-control gaps, fraud exposure, customer concentration, resilience weakness and integration cost.

The hypothetical bridge begins with enterprise value of USD 110 million supported by stand-alone contribution and market assumptions. Verified distribution and working-capital opportunities add USD 14 million and USD 9 million. Incomplete bank and data rights reduce value by USD 8 million; forecast and model uncertainty by USD 6 million; payment-control and fraud exposure by USD 7 million; resilience and integration requirements by USD 5 million. The resulting illustrative value is USD 107 million.

Every amount is hypothetical. The bridge demonstrates method and is not a valuation opinion. A specific transaction requires buyer returns, capital structure, tax, market evidence and legal analysis.

Figure 5. Hypothetical rights-and-control valuation bridge

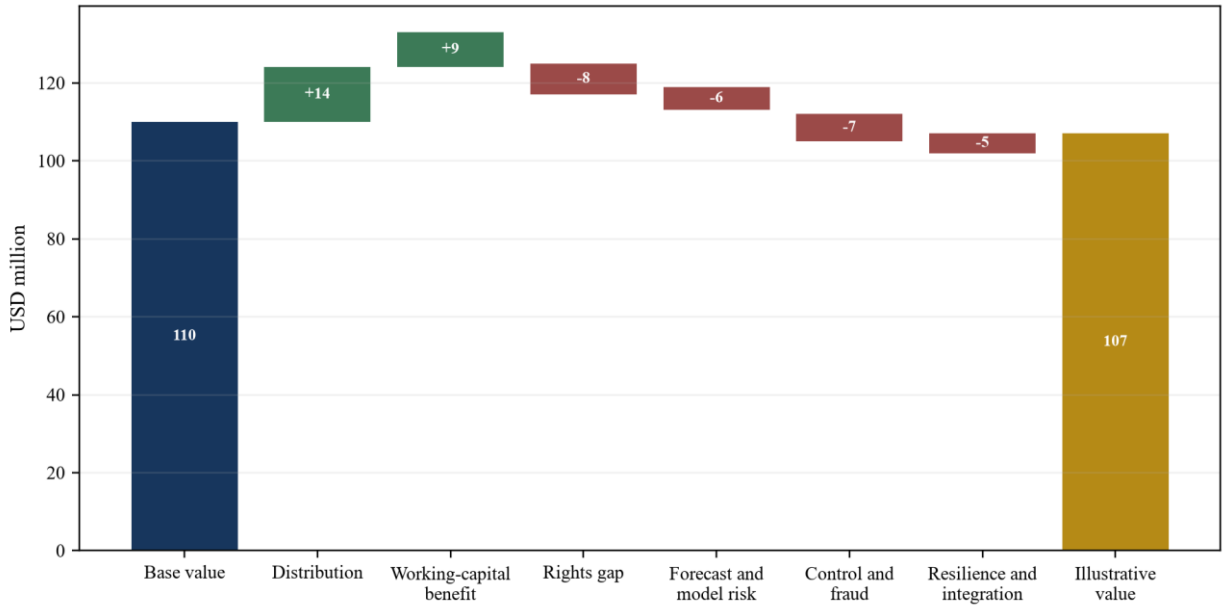

Wholly hypothetical USD millions; the bridge is methodological and is not a valuation opinion.

18. Test sensitivities and downside cases

Sensitivity should expose variables that determine value. Customer retention, connected-account coverage, forecast performance, implementation effort, payment adoption, loss, support productivity and vendor cost can change the outcome. The model should avoid assuming that all variables move favourably together.

Downside cases should include loss of a major bank connection, customer consent renewal, model underperformance, fraud loss, cloud outage, increased insurance cost, slower onboarding and delayed permission to offer payment initiation. The board should see cash funding requirements as well as enterprise value.

Forecast benefit should be capped by addressable decisions. A customer with little cash variability may obtain workflow efficiency without releasing material liquidity. A complex group may have high theoretical benefit and low adoption because authority is decentralised. Cohort evidence should inform penetration and benefit assumptions.

Table 6. Hypothetical contribution sensitivity

Net revenue; USDm	Control cost USD 4.5m	Control cost USD 5.5m	Control cost USD 6.5m	Control cost USD 7.5m
15.0	6.6	5.6	4.6	3.6
17.0	8.6	7.6	6.6	5.6
19.0	10.6	9.6	8.6	7.6
21.0	12.6	11.6	10.6	9.6

Wholly hypothetical annual USD millions; no cell is a forecast or market benchmark.

19. Translate evidence into transaction protections

Transaction documents should allocate identified uncertainty. Representations can address customer and bank contracts, data rights, payment mandates, model ownership, source code, intellectual property, security procedures, losses, incidents, regulatory correspondence, vendors and financial metrics. Definitions should match the diligence data.

Conditions can require bank or customer consent, transfer of critical licences, successful credential reissue, delivery of reproducible forecast vintages, closure of a material security issue or funding of a loss reserve. Interim covenants should govern model, connection, security, pricing and payment-authority changes between signing and closing.

Escrow, indemnity, retention and insurance should match enforceable exposure. Contingent consideration can be tied to retained contribution, connected economic coverage, forecast performance on seasoned cohorts, verified payment adoption and successful transfer of rights. Gross payment volume alone can reward risky or low-margin activity.

The buyer should preserve scope options. A payment-execution product can be delayed while visibility and forecasting transfer. A jurisdiction or bank connection can be carved out. A customer cohort can remain on existing infrastructure until consent and control tests pass. Purchase agreement and integration plan should use the same evidence gates.

Table 7. Evidence-to-protection matrix

Evidence gap	Price response	Protection	Release evidence
incomplete bank consent	defer connected revenue value	consent condition and covenant	accepted transfer and working connection
uncertain historical-data rights	exclude dependent model benefit	representation and restricted use	lawful transfer and documented purpose
unseasoned forecast model	lower improvement probability	retention or earn-out	mature vintage performance
payment-control weakness	funded remediation deduction	condition, escrow and indemnity	tested limits, approval and recovery
unresolved fraud loss	reserve adjustment	specific indemnity	closed claim and paid outcome
critical vendor dependency	continuity deduction	assignment and exit covenant	consent and tested fallback
high implementation effort	cohort margin adjustment	completion funding	verified onboarding productivity

Proposed matrix; legal drafting and remedies remain transaction-specific.

20. Design integration around cash continuity

Integration can change every part of the evidence chain. Bank connections, credentials, account mappings, legal entities, approval rules, models, data stores, cloud services and customer support may move. The buyer should determine which changes require consent, retesting or customer action.

Cash continuity comes first. Customers need accurate balances, approved payments, statements, exception handling and support while systems change. The target should freeze unnecessary configuration changes, preserve logs and maintain an emergency operating route. Every migration exception should have a severity, owner, deadline and customer-impact assessment.

Data migration should reconcile at account, transaction and forecast level. Opening balances, unmatched items and payment status need explicit treatment. Duplicate and missing records can create false positions or repeated instructions. Migration tooling should be tested on representative and edge-case customers.

Model migration is a controlled change. The combined firm should compare old and new forecasts and recommendations on matched inputs, investigate differences, validate limits and monitor post-migration outcomes. A model that remains technically identical can behave differently after upstream mappings or customer populations change.

Synergy should be released after evidence. Removing support, security or payment-control capacity before replacement operations are proven can generate apparent savings and later loss. Board reporting should connect customer continuity, rights transfer, forecast quality, payment control, incidents, contribution and cash.

21. Establish governance and management information

One accountable executive should own the end-to-end service. Product, treasury, engineering, security, compliance, fraud, operations and customer support should share definitions for cash coverage, stale data, forecast error, override, payment incident, loss, recovery and contribution.

Board information should remain concise and traceable. A monthly pack can include economic coverage, latency tails, reconciliation exceptions, forecast vintages, override value, payment-control breaches, fraud outcomes, service availability, customer adoption, cohort contribution and remediation status. Each metric should have a defined population and source.

Limits should trigger action. A stale material account, breached payment limit, model drift, unusual beneficiary, unreconciled settlement or severe outage should route to named owners. Management should document restriction, override, recovery and closure.

Governance should cover vendors and models after closing. Contract renewals, model changes, API versions, bank certificate expiry and payment-scheme releases can affect continuity. A forward calendar and tested ownership reduce hidden operational cliffs.

22. Execute a 180-day programme

Days one to thirty should preserve cash visibility, payment authority, credentials, logs, model versions, customer and bank contracts, incident records and loss evidence. The buyer should establish governance, change restrictions and emergency paths. It should reconcile headline accounts, payment value, revenue and contribution to source records.

Days thirty to seventy should reconstruct representative cash days, forecast vintages and payments; measure economic coverage and latency; test authority and rights; and identify material gaps. High-risk autonomous actions should be restricted or routed to enhanced approval while evidence is incomplete.

Days seventy to one hundred and twenty should remediate priority mappings, models, security controls, vendor dependencies and consent requirements. Integration pilots should use reversible cohorts and matched outputs. Fraud, stress and outage scenarios should be exercised.

Days one hundred and twenty to one hundred and eighty should season forecast and payment outcomes, verify contribution and adoption, complete customer and bank migrations, and release contingent value only after defined gates pass. Remaining uncertainty should stay in reserves, escrow, delayed scope or lower forecast confidence.

Figure 6. Proposed evidence-gated 180-day programme

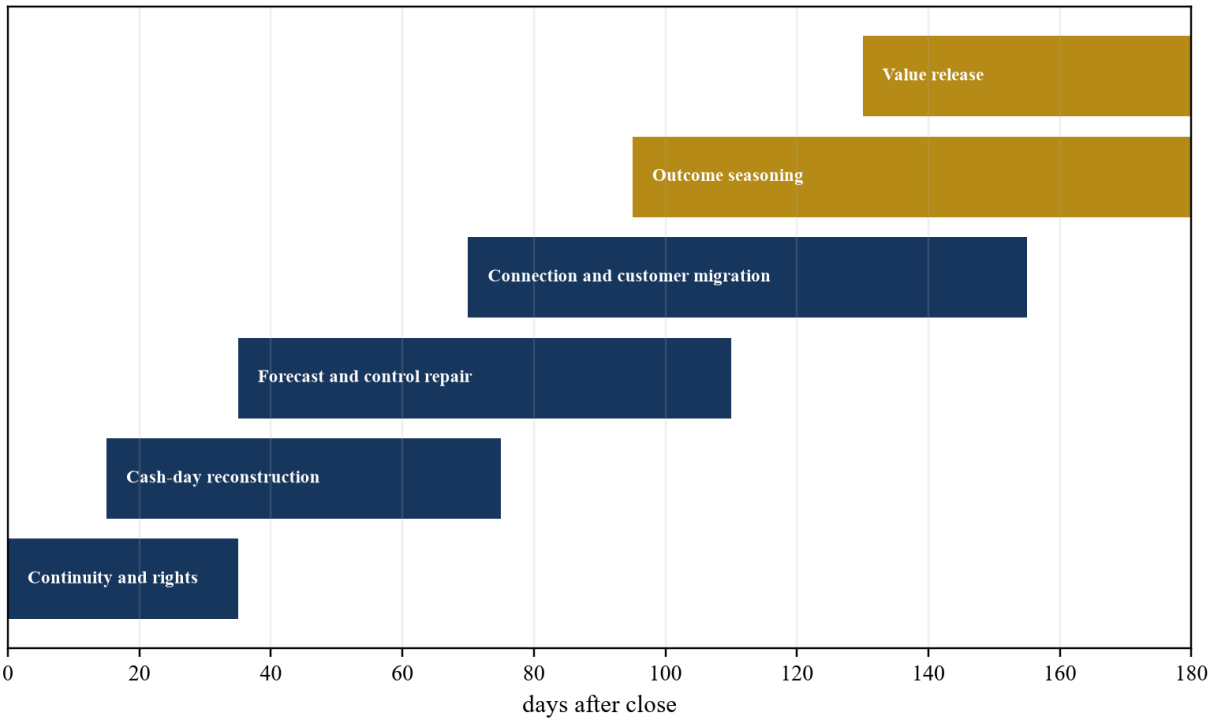

Timing should follow transaction, bank, customer, regulatory and technology constraints.

23. Decision and conclusion

Real-time treasury AI deserves value when it converts reliable cash information into better, controlled decisions and durable contribution. A modern interface, rich payment data and sophisticated model can support that result. The evidence chain must still connect authoritative balances, complete coverage, forecast vintages, policy constraints, accountable approval, settlement, reconciliation and realised economics.

The buyer should separate visibility from transaction authority, reconstruct historical cash days, test forecast performance by horizon and condition, and measure adoption at decision level. It should treat customer, bank, data, model and payment rights as core transaction assets. Fraud, liquidity, security, resilience and ongoing model governance belong inside operating economics.

Rights in hand means more than possession of software. It means the combined firm can lawfully obtain the data, use the model, operate the connection, instruct the bank, preserve the audit trail and serve the customer after closing. Missing rights can turn an apparently scalable platform into a costly re-papering and migration programme.

The resulting investment decision is practical. A premium is supportable when economic cash coverage reconciles, forecast performance is reproducible, actions remain within controlled authority, losses and incidents are transparent, customer adoption produces collected contribution, and contracts and permissions survive the transaction. Price protection, narrower scope, funded remediation or contingent value is appropriate when those conditions remain incomplete.

Frequently asked questions

Q: What is real-time treasury AI? A: It is software that combines timely bank and enterprise data with classification, forecasting, scenario analysis, recommendations or controlled payment workflows to support treasury decisions.

Q: Why separate analytics from transaction capability? A: Viewing and analysing cash require different rights and controls from initiating or executing a payment. Transaction authority introduces immediate fraud, liquidity, legal and operational exposure.

Q: Which metric best demonstrates value? A: No single metric is sufficient. The board should connect economic cash coverage, forecast performance, controlled adoption, realised funding or workflow benefit and collected contribution after full operating cost.

Q: How should forecast accuracy be tested? A: The buyer should preserve original forecast vintages and input snapshots, compare them with actual outcomes, segment error by horizon and flow, and examine bias, tail error, confidence calibration and stress periods.

Q: How should data and payment rights affect valuation? A: Revenue and benefit dependent on a right should receive full value only when the contract, permission, lawful purpose, change-of-control position and operational transfer are evidenced.

Q: How should fraud affect the acquisition model? A: Attempted, prevented, executed, recovered, reimbursed, insured and net losses should be reconciled by cohort. Control cost, investigation capacity and unresolved exposure belong in contribution and transaction protection.

Q: Which protections can address incomplete evidence? A: Conditions, covenants, representations, indemnities, escrow, reserves, retention, carve-outs and contingent consideration can be tied to consents, rights, forecast performance, control remediation and verified contribution.

Q: What should happen during the first 180 days? A: The buyer should preserve cash and payment continuity, reconstruct evidence, transfer rights, repair material data and controls, pilot migrations, season outcomes and release value only after defined gates pass.

References

[1] Bank for International Settlements, AI agents for cash management in payment systems https://www.bis.org/publications/working-paper-1310-ai-agents-cash-management-payment-systems

[2] CPMI-IOSCO, Principles for Financial Market Infrastructures https://www.bis.org/committees/cpmi/pfmi/overview

[3] Board of Governors of the Federal Reserve System, FedNow Service https://www.federalreserve.gov/paymentsystems/fednow_about.htm

[4] Board of Governors of the Federal Reserve System, FedNow frequently asked questions https://www.federalreserve.gov/paymentsystems/files/fednow_faq.pdf

[5] Board of Governors of the Federal Reserve System, Policy Statement on Payment System Risk https://www.federalreserve.gov/frrs/regulations/federal-reserve-policy-statement-on-payment-system-risk.htm

[6] European Central Bank, TARGET Instant Payment Settlement https://www.ecb.europa.eu/paym/target/tips/html/index.en.html

[7] European Central Bank, TARGET Annual Report 2023 https://www.ecb.europa.eu/press/targetar/html/ecb.targetar2023.en.html

[8] Committee on Payments and Market Infrastructures, ISO 20022 harmonisation and cross-border payments https://www.bis.org/committees/cpmi/cross-border-payments/engagement.htm

[9] Committee on Payments and Market Infrastructures, Enhancing cross-border payments: addressing fraud https://www.bis.org/cpmi/pietf/fraud_report_2026.pdf

[10] Basel Committee on Banking Supervision, Principles for operational resilience https://www.bis.org/bcbs/publ/d516.htm

[11] Basel Committee on Banking Supervision, Principles for the sound management of third-party risk https://www.bis.org/bcbs/publ/d605.htm

[12] European Union, Digital Operational Resilience Act https://eur-lex.europa.eu/eli/reg/2022/2554/oj

[13] European Union, Instant Payments Regulation https://eur-lex.europa.eu/eli/reg/2024/886/oj

[14] European Union, Artificial Intelligence Act https://eur-lex.europa.eu/eli/reg/2024/1689/oj

[15] European Union, General Data Protection Regulation https://eur-lex.europa.eu/eli/reg/2016/679/oj

[16] European Banking Authority, Guidelines on ICT and security risk management https://www.eba.europa.eu/regulation-and-policy/internal-governance/guidelines-ict-and-security-risk-management

[17] European Banking Authority, Guidelines on outsourcing arrangements https://www.eba.europa.eu/regulation-and-policy/internal-governance/guidelines-outsourcing-arrangements

[18] European Banking Authority, Payment services and electronic money https://www.eba.europa.eu/regulation-and-policy/payment-services-and-electronic-money

[19] European Central Bank, TIPS onboarding and liquidity management https://www.ecb.europa.eu/paym/target/tips/onboarding/html/index.en.html

[20] European Central Bank, TIPS user requirements https://www.ecb.europa.eu/paym/target/target-professional-use-documents-links/tips/shared/pdf/TIPS_User_Requirement_Documents_Version_R2025.JUN.pdf

[21] Bank of England, RTGS renewal programme https://www.bankofengland.co.uk/payment-and-settlement/rtgs-renewal-programme

[22] Bank of England, CHAPS and RTGS https://www.bankofengland.co.uk/payment-and-settlement/chaps

[23] Bank of England, Model risk management principles for banks https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-supervisory-statement

[24] Board of Governors of the Federal Reserve System, SR 11-7 model risk management https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm

[25] Board of Governors of the Federal Reserve System, Intraday credit policies https://www.federalreserve.gov/frrs/regulations/ii-federal-reserve-intraday-credit-policies.htm

[26] Board of Governors of the Federal Reserve System, Liquidity risk management https://www.federalreserve.gov/supervisionreg/topics/liquidity_risk.htm

[27] Financial Stability Board, Recommendations for achieving greater convergence in cyber incident reporting https://www.fsb.org/2023/04/recommendations-to-achieve-greater-convergence-in-cyber-incident-reporting-final-report/

[28] Financial Stability Board, Artificial intelligence and financial stability https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/

[29] CPMI-IOSCO, Guidance on cyber resilience for financial market infrastructures https://www.bis.org/cpmi/publ/d146.pdf

[30] CPMI, Linking fast payment systems across borders: governance and oversight https://www.bis.org/cpmi/publ/d223.pdf

[31] CPMI, Extending and aligning payment-system operating hours https://www.bis.org/cpmi/publ/d216.pdf

[32] CPMI, Harmonised ISO 20022 data requirements https://www.bis.org/cpmi/publ/d230.pdf

[33] International Organization for Standardization, ISO 20022 financial services messaging https://www.iso20022.org/

[34] National Institute of Standards and Technology, AI Risk Management Framework https://www.nist.gov/itl/ai-risk-management-framework

[35] National Institute of Standards and Technology, Cybersecurity Framework 2.0 https://www.nist.gov/cyberframework

[36] National Institute of Standards and Technology, Generative AI profile https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-generative-artificial-intelligence

[37] United Kingdom Information Commissioner's Office, AI and data protection https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/

[38] European Data Protection Board, Automated decision-making and profiling https://www.edpb.europa.eu/our-work-tools/our-documents/guidelines/guidelines-automated-individual-decision-making-and_en

[39] Financial Conduct Authority, Artificial intelligence approach https://www.fca.org.uk/firms/innovation/ai-approach

[40] Financial Conduct Authority, Operational resilience https://www.fca.org.uk/firms/operational-resilience

[41] Payment Systems Regulator, Authorised push payment fraud reimbursement https://www.psr.org.uk/our-work/app-scams/

[42] UK Finance, Confirmation of Payee https://www.wearepay.uk/what-we-do/overlay-services/confirmation-of-payee/

[43] United States Treasury, Financial sector adoption of cloud services https://home.treasury.gov/system/files/136/Treasury-Cloud-Report.pdf

[44] Office of the Comptroller of the Currency, Third-party relationships risk management https://www.occ.treas.gov/news-issuances/bulletins/2023/bulletin-2023-17.html

[45] Federal Financial Institutions Examination Council, Authentication and access guidance https://www.ffiec.gov/press/pr081121.htm

[46] International Organization of Securities Commissions, AI and machine learning by intermediaries and asset managers https://www.iosco.org/library/pubdocs/pdf/IOSCOPD684.pdf

[47] International Valuation Standards Council, International Valuation Standards https://ivsc.org/standards/

[48] IFRS Foundation, IFRS 3 Business Combinations https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/

[49] IFRS Foundation, IAS 38 Intangible Assets https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/

[50] Organisation for Economic Co-operation and Development, AI principles https://oecd.ai/en/ai-principles

About the Author

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

https://www.linkedin.com/in/ckadya/

https://www.matchpoint-partners.com/team/ck-adya.html

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.