One Model-Risk Engine, Three Regulators: RegTech Roll-Ups across Finance

A control-portability, diligence and integration framework for banking, insurance and securities acquisitions

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Regulatory-technology acquirers often expect one model-risk engine to serve banks, insurers and securities firms. The proposition is commercially attractive. A shared product can lower development cost, consolidate regulatory content, support cross-selling and create a larger recurring-revenue platform. The transaction risk lies in treating similar control language as proof that the underlying supervisory duties, evidence and operating decisions are interchangeable.

This paper develops a control-portability framework for RegTech roll-ups across regulated finance. It distinguishes the reusable product core from the sector overlays required for banking, insurance and securities workflows. The reusable core includes inventories, taxonomies, workflow orchestration, evidence records, issue management, testing infrastructure and audit history. Sector overlays preserve the legal perimeter, materiality tests, accountable roles, validation depth, customer-protection duties and supervisory reporting required for each use case.

The framework draws on the Federal Reserve's revised model-risk-management guidance, the Basel Committee's Core Principles, the Financial Stability Board's work on artificial intelligence in finance, the IAIS Application Paper on AI supervision, EIOPA's AI-governance opinion, IOSCO's work on AI and machine learning, FINRA supervisory guidance, the NIST AI Risk Management Framework and the European Union AI Act. It converts those sources into transaction diligence, revenue-cohort testing and a sequenced integration plan.

The central conclusion is that a RegTech roll-up creates value when the buyer standardises control primitives while preserving sector-specific decisions and accountability. A common interface or regulatory library provides limited evidence of portability. Durable value depends on a governed common data model, traceable evidence, configurable but constrained workflows, reliable integrations and revenue cohorts that demonstrate paid reuse. The buyer should price verified portability and retained contribution rather than claimed regulatory coverage. All numerical examples in this paper are hypothetical management assumptions prepared solely to demonstrate the framework.

JEL Classification: G21, G22, G23, G34, K22, L86, O32

Keywords: model risk, RegTech roll-up, financial regulation, control portability, banking, insurance, securities, M&A, integration

This paper is an educational and structural analysis prepared for research purposes. It is not investment advice, an offer, or a solicitation. It contains no client information.

1. DEFINE THE TRANSACTION DECISION

The investment committee must decide whether a target supplies a reusable model-risk engine or a collection of sector products that happen to share a brand. This distinction affects the price, integration budget, cross-selling thesis and liability allocation. The decision should be stated before confirmatory diligence begins.

The acquisition thesis should identify the reusable capabilities, the sector decisions that remain local, the customers expected to adopt additional modules and the evidence required to credit those benefits. Management should also state which synergies depend on product work, regulatory interpretation, customer consent or migration. A transaction model that records those dependencies is more useful than a broad claim that financial regulation is converging.

2. ESTABLISH THE COMMON CONTROL PROBLEM

Banks, insurers and securities firms all use models and automated systems to support consequential decisions. They need inventories, ownership, change control, testing, monitoring, issue management and evidence. This shared need creates a credible product core.

The operating context remains different. A bank may focus on credit, capital, liquidity and financial-crime decisions. An insurer may focus on underwriting, pricing, reserving and claims. A broker-dealer or asset manager may focus on trading, suitability, communications and market conduct. A shared engine must represent these differences without forcing every sector into the same approval logic.

3. USE CURRENT SUPERVISORY BASELINES

The Federal Reserve issued revised model-risk-management guidance in April 2026. It superseded earlier supervisory letters and retained a risk-based approach tied to a model's inherent risk, exposure, purpose and use. The revised guidance matters to diligence because it places proportionality inside governance rather than treating every model identically.

The Basel Committee's 2024 Core Principles address governance, risk management, operational resilience, technology and third-party dependencies. The IAIS 2025 Application Paper applies existing insurance principles to AI through proportionality, accountability, robustness, transparency, fairness and redress. Securities-sector expectations draw on supervisory controls, books and records, market rules, customer duties and model-risk practices. The common engine should support all three baselines while recording which authority and decision each control serves.

4. TREAT PORTABILITY AS AN EVIDENCE CLAIM

Portability means that a control component can move across sectors without losing its purpose, ownership, evidence or enforceability. It is an evidence claim. Product demonstrations and control counts do not establish it.

The buyer should sample live customer workflows in each sector. It should trace a governed object from inventory through classification, validation, approval, monitoring and issue closure. The same test should identify which data fields, roles, thresholds and artefacts change by sector. The output is a portability map grounded in deployed behaviour.

Table 1. Shared model-risk engine and sector-specific overlays

Control component	Reusable core	Banking overlay	Insurance overlay	Securities overlay
Inventory	Identifiers, ownership, versions, status and dependencies	Credit, capital, liquidity and financial-crime materiality	Underwriting, pricing, reserving, claims and policyholder impact	Trading, portfolio, communications, suitability and market-conduct use
Risk classification	Configurable taxonomy and scoring workflow	Prudential exposure and model-risk tier	Policyholder impact, actuarial judgement and fairness	Investor harm, market integrity and supervisory-control relevance
Validation	Test plans, evidence store, findings and approvals	Independent validation aligned with model use and materiality	Actuarial, conduct and governance review appropriate to use	Supervisory testing, market-rule checks and customer-duty review
Monitoring	Metrics, thresholds, alerts and issue routing	Performance, drift, capital and portfolio effects	Pricing, claims, discrimination and consumer outcomes	Trading behaviour, suitability, communications and execution outcomes
Accountability	Role model, attestations and escalation	Board, risk, business and validation functions	Board, senior management, actuarial, conduct and product functions	Supervisory principals, compliance, risk and business owners
Evidence and reporting	Versioned records, audit trail and export	Examiner-ready model and control evidence	Supervisor and policyholder-outcome evidence	Books, records and supervisory-control evidence

Reuse should be credited only where deployed evidence supports it.

5. DEFINE THE REUSABLE PRODUCT CORE

The reusable core contains technical and governance primitives. It should register governed systems, store versions, map obligations, assign roles, route work, preserve evidence, record decisions, monitor thresholds and manage exceptions. These functions can serve many regulated activities.

The core should avoid embedding one sector's assumptions into every workflow. Fields, states and permissions need explicit extension points. The platform should preserve a stable record even when an overlay adds a sector term, test or approver. This architecture supports reuse without erasing regulatory context.

6. PRESERVE THE BANKING OVERLAY

Banking model risk often connects directly to credit decisions, capital, liquidity, stress testing, fraud and financial crime. The model's purpose, exposure, use and potential adverse consequences determine the depth of governance. A banking overlay therefore needs materiality logic, independent challenge, outcome analysis and a clear relationship to enterprise risk management.

Diligence should test whether the target's workflow supports model limitations, compensating controls, overrides and use restrictions. A generic approval status is inadequate when an institution must understand how a limitation affects a portfolio or prudential decision. The overlay should preserve the responsible business owner and the independent review record.

7. PRESERVE THE INSURANCE OVERLAY

Insurance systems affect underwriting, pricing, reserving, claims, fraud detection and customer interactions. The IAIS and EIOPA emphasise governance, accountability, proportionality, robustness, transparency, fairness and redress. These concerns can require evidence that differs from bank validation.

The overlay should identify policyholder impact, product scope, actuarial judgement, protected characteristics, distribution effects and complaint or redress paths. It should preserve the insurer's responsibility when a vendor supplies a model. A shared testing engine may execute common tests, while the acceptance criteria and accountable decision remain tied to the insurance use case.

8. PRESERVE THE SECURITIES OVERLAY

Securities firms use models in trading, portfolio management, surveillance, customer communications, suitability, research and operations. FINRA identifies model risk, data governance, privacy, cybersecurity, outsourcing, books and records and supervisory controls as relevant considerations for AI-based applications.

The overlay should connect automated activity to written supervisory procedures, applicable market rules, customer obligations and escalation. Trading use cases need controls for volatility, unusual market conditions and unintended behaviour. Communications and advisory use cases need review, record preservation and customer-context evidence. One workflow engine can support these controls when the sector logic remains explicit.

9. BUILD A REGULATORY MAP THAT CAN CHANGE

A regulatory map should connect authorities, obligations, control objectives, evidence and effective dates. It should distinguish binding requirements, supervisory guidance, standards and internal policy. The distinction matters because a target may market broad coverage while providing only a generic control statement.

The buyer should test how the platform processes change. A source update should trigger review, mapping changes, affected-customer identification and a governed release. Historical mappings should remain visible. The cost and staffing required to maintain the map should appear in the operating model.

10. CREATE A COMMON CONTROL ONTOLOGY

The ontology defines reusable concepts such as governed system, model, use case, owner, risk, control, test, evidence, decision, exception and incident. It should support relationships without forcing a one-to-one mapping. A model can support several uses; one control can address several risks; one decision can depend on several tests.

Sector terms should extend the common ontology. A bank's model tier, an insurer's policyholder-impact assessment and a securities firm's supervisory procedure may share workflow primitives while retaining their own meaning. Diligence should inspect the data model and deployed configuration rather than relying on a slide describing a unified taxonomy.

11. TEST INVENTORY INTEGRITY

The inventory is the foundation for every downstream control. It should identify the governed object, versions, owners, purpose, users, data, vendors, deployment status and material dependencies. Duplicate records, stale ownership and uncertain boundaries weaken the control record.

The buyer should reconcile samples to source systems. Banking samples may come from model inventories and application portfolios. Insurance samples may come from actuarial, product and claims systems. Securities samples may come from trading, surveillance and communications systems. Coverage should be stated as a measured result, with management estimates identified as such.

12. VERIFY CLASSIFICATION LOGIC

Classification determines which controls apply. A common engine can provide questionnaires, scoring, decision trees and approvals. The thresholds and legal significance may differ by jurisdiction and sector.

Diligence should reproduce classifications for representative high, medium and low-risk cases. It should test missing data, conflicting answers, overrides and later changes. The platform should record the rule version and reviewer. A score without a traceable rationale creates operational risk and weakens portability.

13. ASSESS VALIDATION ARCHITECTURE

Validation includes conceptual soundness, data quality, implementation checks, performance testing, outcomes analysis and ongoing monitoring. The appropriate combination depends on the model and its use. A shared engine can manage test plans and evidence without prescribing one universal validation standard.

The buyer should distinguish native tests, third-party tools, professional services and customer work. It should inspect execution logs, reproducibility, exceptions and review. Product revenue should receive software economics only where the capability is repeatable and supported within the contracted service.

14. TEST EVIDENCE LINEAGE

Evidence lineage records the source, system version, date, method, author, reviewer and later change for each artefact. A file upload can satisfy a checklist while failing to establish lineage. The platform should preserve the relationship between the evidence and the decision it supports.

The buyer should sample automated feeds, test results, tickets, documents and manual attestations. It should inspect access controls, timestamps, retention, deletion and export. Generated content should be identified and reviewed by an accountable person before it enters the official record.

15. EXAMINE WORKFLOW CONSTRAINTS

Configuration enables sector reuse, but unlimited configuration can weaken controls. The engine should distinguish required steps, optional steps and local extensions. It should warn or block configurations that remove segregation, evidence or approval required by the policy design.

Diligence should compare standard templates with deployed workflows. It should identify bespoke code, manual workarounds and customer-specific dependencies. The integration plan should preserve valid local controls until a tested common workflow is ready.

16. PRESERVE ACCOUNTABLE DECISIONS

The platform may recommend, route and document a decision. The regulated firm remains responsible for the decision. The control record should identify the owner, reviewer, approver, basis, conditions and review date.

Accountability cannot be collapsed into a generic role named approver. Banking, insurance and securities organisations assign responsibility differently. The product should support these structures while preserving independence and escalation. Acquisition integration should not replace local accountable roles merely to achieve a common configuration.

17. REVIEW THIRD-PARTY DEPENDENCIES

Financial institutions increasingly depend on cloud, data, model and specialist technology providers. The FSB identifies third-party dependencies and service-provider concentration as potential AI-related vulnerabilities. A RegTech platform can reduce fragmentation while creating its own concentration point.

The buyer should maintain a bill of materials for software, AI services, data and regulatory content. Contracts should support intended use, transfer and customer commitments. Exit, substitution and service-continuity plans should be tested. Dependency economics should include usage cost, monitoring, assurance and migration.

18. TEST SECURITY AND RESILIENCE

The platform stores sensitive model details, customer data, control findings and supervisory evidence. Security and resilience therefore affect both product value and regulatory acceptance. Diligence should assess identity, privileged access, encryption, tenant isolation, vulnerability management, recovery and incident response.

The buyer should also test integration failures. A connector that stops collecting evidence may leave a dashboard unchanged. The platform should identify stale data, incomplete runs and failed exports. Recovery objectives should reflect the customer workflow and contractual commitments.

19. REVIEW DATA RIGHTS AND CONFIDENTIALITY

The transaction should establish who owns customer data, derived data, benchmarks and generated output. Contracts should define permitted use, retention, deletion and cross-border transfer. Training or product improvement using customer information requires a clear contractual and governance basis.

The buyer should inspect whether data is commingled across customers and whether anonymisation claims are supported. Confidential supervisory material may require stronger restrictions. Consent or notification requirements can affect integration and cross-selling, so they should be reflected in the deal plan.

20. RECONSTRUCT PRODUCT CAPABILITY

Management's capability map should reconcile with code, deployed configuration, service records and customer use. Functions presented in demonstrations may depend on manual analysts or future roadmaps. Diligence should identify native, licensed, partner-delivered and manual components.

The reconstruction should also identify duplicate capabilities across targets. A roll-up may contain several inventories, workflow engines and content libraries. The buyer needs a rationalisation decision for each component, supported by customer impact and migration cost.

21. BUILD THE DILIGENCE EVIDENCE MATRIX

The evidence matrix links every value claim to a test, source and decision. It should be completed before the buyer gives credit for portability, revenue quality or synergy. Open evidence should reduce value or create a condition rather than becoming an undocumented assumption.

Table 2. Diligence evidence matrix for a cross-sector RegTech acquisition

Value claim	Evidence to obtain	Test	Failure signal	Transaction response
Shared control engine	Architecture, data model, code ownership and deployed configurations	Trace one workflow in each sector through the same core objects	Separate code bases or manual reconciliation	Reclassify as product consolidation and fund migration
Regulatory coverage	Source library, mapping method, reviewers and change logs	Reproduce a sample obligation-to-control mapping	Generic statements, missing rationale or stale effective dates	Limit coverage claim and price maintenance cost
Embedded use	Workflow logs, governed objects, decisions and evidence volume	Reconcile active use to contracts and renewals	Logins without completed decisions	Exclude shallow adoption from cross-sell case
Recurring revenue	Contracts, invoices, credits, collections and general ledger	Rebuild customer-level recurring revenue	Services, pilots or uncollected amounts included	Adjust recurring revenue and contribution margin
Cross-sector portability	Live sector workflows, configuration history and support tickets	Compare common and local components	Bespoke code, fragile integrations or control loss	Apply portability discount and staged integration
Synergy	Cost base, duplicate functions, roadmap and customer obligations	Assign owner, timing, cost and dependency	Benefit requires unapproved migration or customer consent	Delay credit and use milestone-based value

The matrix identifies the minimum record required before a value claim enters the transaction model.

22. REBUILD RECURRING REVENUE

Recurring revenue should be reconstructed from contracts, invoices, credits, cash and customer-level accounting. Subscription, usage, content, implementation, validation and managed-service revenue should be separated. Paid pilots remain separate until they convert to production.

Contract review should cover scope, term, renewal, termination, price protection, service credits, acceptance and data rights. A multi-year contract does not by itself establish recurring collected

revenue. Revenue quality depends on delivery obligations, customer adoption and contribution after support.

23. SEGMENT REVENUE BY SECTOR

Sector cohorts reveal whether one engine supports repeatable economics. Banking, insurance and securities customers may have different procurement cycles, implementation work, support demands and renewal patterns. Aggregated retention can hide a weak sector.

The buyer should measure gross retention, net retention, expansion, churn, collection and contribution margin by sector and product module. It should connect those measures to workflow depth. Expansion associated with more governed systems and recurring decisions is stronger evidence than an unrelated bundle concession.

Table 3. Illustrative revenue cohorts by regulated sector

Cohort	Opening recurring revenue USDm	Gross retention	Net retention	Services share	Contribution margin	Portability interpretation
Banking	12.0	93%	108%	18%	54%	Strong core adoption; validation support remains material
Insurance	7.0	89%	101%	27%	43%	Workflow reused; actuarial and conduct configuration raises delivery cost
Securities	5.0	84%	94%	31%	38%	Supervisory-control and integration work remains fragmented
Cross-sector customers	3.0	96%	119%	14%	61%	Highest evidence of paid reuse where modules share the same inventory and evidence layer

All values are hypothetical management assumptions and are included only to demonstrate the analysis.

24. MEASURE EMBEDDED USE

User counts and logins are weak adoption measures. The buyer should measure governed objects, completed classifications, validations, approvals, monitoring runs, exceptions and incidents. It should identify which roles participate and whether the record is used in internal audit or supervisory engagement.

Embedded use should be linked to renewal. A customer may maintain a contract because of a broader relationship even when the module is shallow. Cohorts with recurring decisions, integrated evidence and broader coverage provide stronger evidence that the platform forms part of the operating process.

25. CALCULATE PORTABILITY BY COMPONENT

Portability should be scored at component level. The buyer can assess data model, workflow, evidence, testing, integrations, reporting and regulatory content. Each score should reflect deployed reuse, configuration effort, sector exceptions and support cost.

The score is a decision aid rather than a substitute for judgement. A component with high technical reuse may still require sector-specific legal review. A component with moderate reuse may create value if the local overlay is repeatable. The transaction model should connect each score to cost and timing.

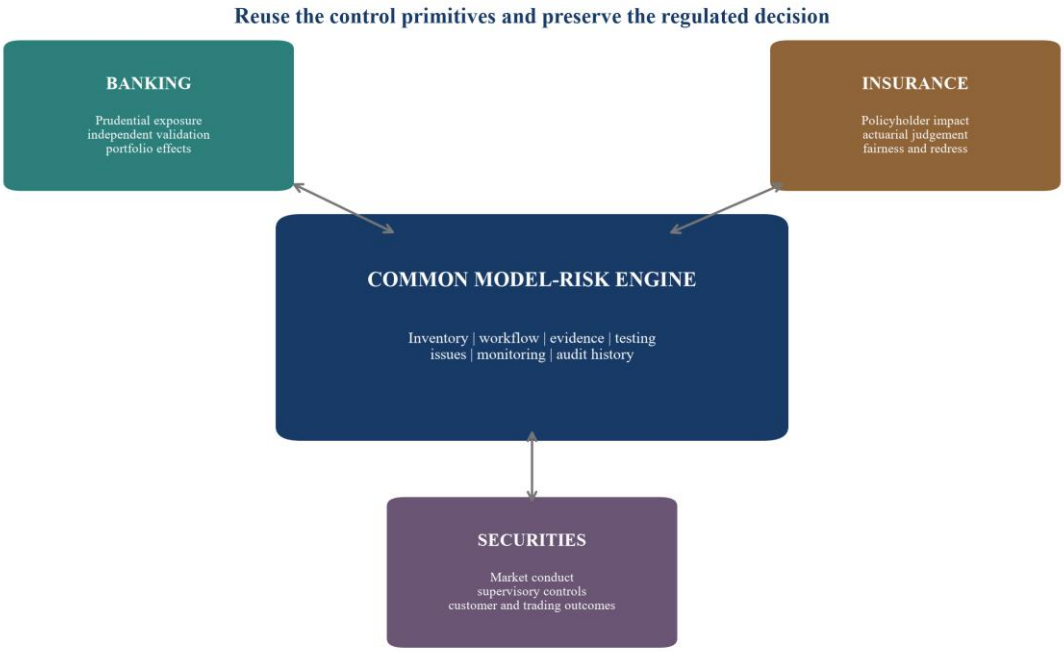

Figure 1. The common engine and three sector overlays

The common engine provides governed primitives; each sector retains its own decision logic and accountability.

26. ANALYSE CUSTOMER CONCENTRATION

RegTech targets often depend on a small number of large institutions. Concentration can support reference value while creating renewal and pricing risk. The buyer should measure revenue, contribution, receivables, support effort and roadmap dependence for each major customer.

The analysis should identify sector concentration and common dependencies. Several customers may rely on the same integration or regulatory content. A single technical failure or rule change can therefore affect more revenue than customer-level concentration suggests.

27. TEST CROSS-SELLING EVIDENCE

Cross-selling should be based on customer need, contractual permission, product readiness and a realistic sales process. An acquired bank customer does not automatically become an insurance or securities buyer. Many groups separate budgets, entities, risk functions and procurement.

The strongest evidence comes from paid multi-sector customers, qualified customer discussions and prior module expansion. The buyer should record the responsible executive, decision unit, required integration and expected sales cycle. Unverified pipeline should remain outside the base case.

28. MODEL INTEGRATION COSTS

Integration cost includes product rationalisation, data migration, connector work, security remediation, regulatory-content maintenance, customer support and change management. It also includes revenue risk during migration. These costs should be timed and assigned to accountable owners.

The buyer should preserve customer evidence and audit history. A technically efficient migration that loses lineage can damage the control record. Contracts may require notice, consent or continued support for the acquired platform. The plan should fund parallel operation where evidence or customer acceptance requires it.

29. PLAN PRODUCT RATIONALISATION

The combined group should decide which inventory, workflow, evidence store, testing layer and reporting engine will remain. The decision should use capability, scalability, security, customer adoption and migration evidence. Internal preference is not an adequate basis.

Rationalisation should sequence low-risk common services before sector decisions. Identity, observability and shared evidence formats may be consolidated early. Sector classification, validation and approval logic should move only after testing demonstrates that the target state preserves requirements and history.

30. PROTECT THE REGULATORY-CONTENT FUNCTION

Regulatory content requires qualified interpretation, source management, review and effective dating. A roll-up may expect to remove duplicate teams, yet broader coverage can increase the maintenance burden. The buyer should define ownership for each jurisdiction and sector.

The product should distinguish legal analysis, supervisory guidance, industry standards and internal controls. Customer-facing descriptions should state the product's role accurately. The integration budget should support continuing updates while the combined ontology and mapping process are tested.

31. USE A SEQUENCED INTEGRATION PLAN

The first phase should protect service continuity, evidence and customer obligations. The second phase should establish the common ontology, identity, security and integration standards. The third phase should migrate selected components based on verified portability. The final phase can rationalise duplicate products after customer acceptance.

Each migration should have entry criteria, reconciliation tests, rollback and approval. The integration committee should include product, engineering, regulatory, risk, security, customer and finance leaders. Sector owners should retain authority over their regulated decisions.

32. SET DAY ONE CONTROLS

Day One controls should govern privileged access, releases, incidents, customer communications, subcontractors, data transfer and regulatory content. The buyer should freeze uncontrolled changes during the transition and preserve access to historical records.

The team should also reconcile open issues, overdue controls and customer commitments. Material findings should have owners and due dates. Transaction close should not convert an acquired exception into an undocumented inherited risk.

33. DEFINE ONE-HUNDRED-DAY DECISIONS

Within one hundred days, management should approve the target architecture, common ontology, product roadmap, regulatory-content operating model and customer migration priorities. It should validate the opening recurring-revenue base and integration cost.

Management should also identify which cross-sell opportunities have customer evidence. Targets should separate customer discovery, qualified opportunities, contracted expansion and collected revenue. This protects capital-allocation decisions from optimistic pipeline treatment.

34. GOVERN MODEL CHANGES AFTER CLOSE

The combined platform may introduce new AI functions for classification, mapping, testing or evidence review. Each function should have a defined purpose, evaluation, human review and change record. The platform should preserve which version produced an output.

Changes that affect customer controls require impact assessment and communication. A product release can alter regulated decisions even when the interface changes little. Release governance should therefore connect code, content, configuration, test evidence and customer scope.

35. DESIGN THE MANAGEMENT INFORMATION

Management information should show product reliability, control completion, evidence freshness, issues, customer adoption, renewal and contribution margin. It should distinguish shared-core performance from sector-overlay performance. Aggregate green status can conceal a weak sector.

The board should review concentration, security, regulatory-content backlog, migration risk and material customer commitments. Measures should have definitions and source lineage. Management estimates should remain labelled until reconciled to records.

36. BUILD THE VALUATION BRIDGE

Valuation begins with verified recurring revenue and contribution. The buyer can then assess whether embedded use, retention, product reuse and growth support a premium. It should deduct required remediation, migration, duplicated cost and revenue at risk.

Portability should influence both revenue and cost. High reuse can improve cross-selling and development efficiency. Failed portability can require parallel products and sector specialists. The model should make these relationships explicit rather than applying a general RegTech multiple.

Table 4. Hypothetical transaction valuation bridge

Item	Illustrative value USDm	Evidence condition	Treatment
Enterprise value before confirmatory diligence	210	Indicative offer based on reported recurring revenue	Starting point only
Recurring-revenue quality adjustment	-22	Services, pilots, credits and collection differences	Deduct from base economics
Product-remediation and security cost	-12	Confirmed backlog and required control improvements	Deduct funded cost
Migration and parallel-operation cost	-16	Evidence-preserving integration plan	Deduct funded cost
Verified shared-engine premium	+18	Deployed reuse across all three sectors with retained contribution	Credit only after evidence
Probability-weighted cross-selling	+9	Qualified customer evidence and product readiness	Stage through milestones
Revenue at risk during migration	-14	Customers requiring consent or legacy support	Deduct expected exposure
Illustrative evidence-adjusted value	173	Sum of the hypothetical adjustments	Decision scenario, not a valuation opinion

All values are hypothetical management assumptions and do not represent market observations, forecasts or a specific company.

37. USE MILESTONE-BASED CONSIDERATION

Earn-outs, holdbacks or contingent value can align price with uncertain portability and cross-selling. Milestones should use observable definitions such as completed migrations, retained recurring revenue, collected module expansion and verified contribution margin. Product releases alone do not demonstrate customer value.

The mechanism should avoid incentives to delay necessary remediation or force unsuitable migrations. It should specify accounting, customer attribution, disputes and access to records. Regulatory compliance and customer outcomes should remain conditions rather than tradeable performance measures.

38. APPLY CAPITAL AND RISK GATES

The investment committee should use explicit gates. The product gate tests the common engine. The regulatory gate tests sector overlays and accountability. The commercial gate tests recurring revenue and paid reuse. The integration gate tests cost, evidence preservation and customer acceptance.

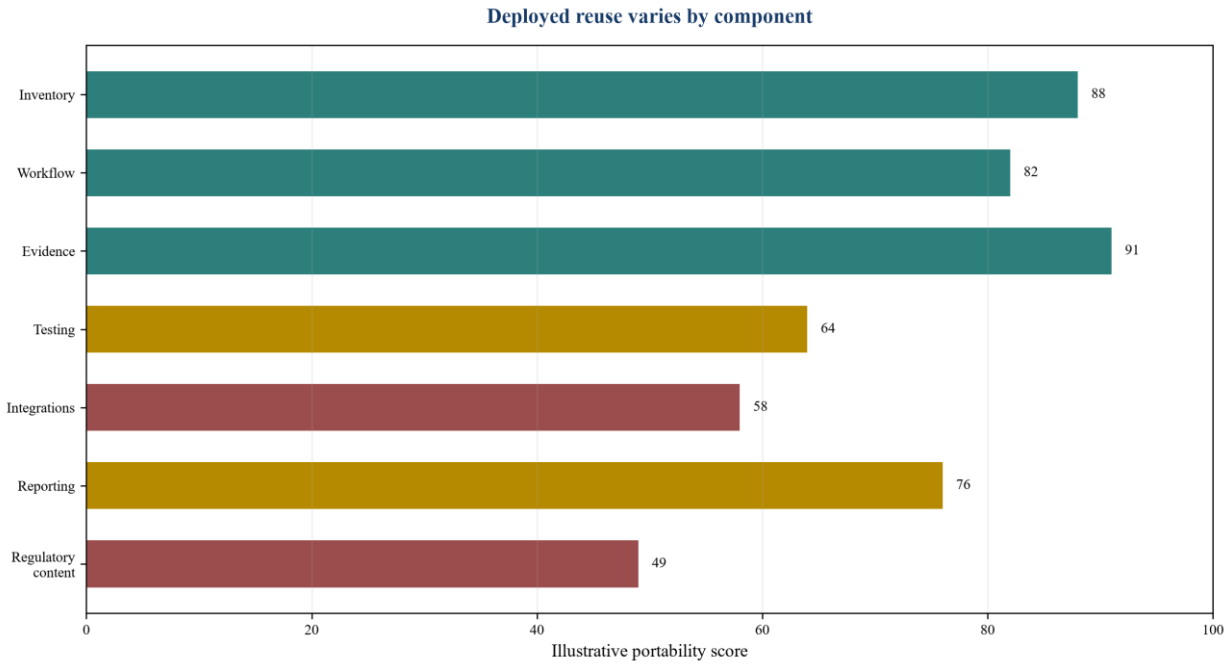

Figure 2. Portability score by control component

Scores are hypothetical management assumptions prepared only to illustrate prioritisation.

The committee should stop, reprice or condition the transaction when a critical gate remains unsupported. Open items need owners, evidence and deadlines. Approval should record which assumptions remain hypothetical and how the transaction protects against them.

39. RUN ADVERSE SCENARIOS

Adverse scenarios should test a major customer refusing migration, a supervisory objection, a material security incident, a core provider failure, a regulatory-content error and lower cross-selling. The model should show liquidity, covenant, integration and valuation effects.

The exercise should also test correlated failures. A provider outage can affect several sectors. A content error can propagate through a shared mapping. A common engine concentrates both capability and operational risk, so resilience should be valued alongside efficiency.

40. INTERPRET THE STRATEGIC VALUE

A cross-sector engine can create a valuable control record and a larger addressable market. The value comes from repeated use of reliable primitives, supported by local overlays and accountable decisions. Technical consolidation by itself does not establish the commercial result.

The strongest platform will show that customers use the same governed inventory, evidence and issue-management core for different regulated workflows. It will also show that sector specialists can maintain local obligations without fragmenting the system. This combination supports credible product economics.

41. APPLY THE DECISION FRAMEWORK

The buyer should first reconstruct the regulated use cases and common product core. It should then test live workflows, rebuild recurring revenue, score portability, fund integration and run adverse scenarios. The valuation bridge should credit only evidence that survives these steps.

The final investment paper should state the maximum value, required protections, closing conditions, Day One controls and one-hundred-day decisions. It should identify the remaining management assumptions. This produces a transaction decision that can be reviewed after close.

42. RUN THE BANKING DILIGENCE SAMPLE

The banking sample should include at least one model used in a prudential or customer decision, one vendor model and one automated process that management does not call a model. This range tests the inventory boundary and the target's ability to apply governance according to purpose and exposure. The buyer should select the sample independently from the complete inventory and reconcile it to business, risk and technology records.

For each item, the team should reconstruct inception, classification, validation, approval, limitations, monitoring and change. It should identify the business owner, independent reviewer and senior escalation route. Evidence should show how the institution responded when a threshold failed or a use changed. A clean record with no adverse events may indicate careful operation; it may also indicate incomplete issue capture. The team should examine source systems and committee materials before deciding.

The sample should test the revised Federal Reserve guidance's risk-based logic where it applies and comparable local requirements elsewhere. The platform should allow the institution to explain why the governance depth fits the model's risk. Templates can assist that explanation. They should not replace the institution's judgement or record of challenge.

43. RUN THE INSURANCE DILIGENCE SAMPLE

The insurance sample should span pricing or underwriting, claims and a customer-facing use. It should include a third-party system when the target claims vendor-governance capability. The team should trace the policyholder population, data, model purpose, outcome testing, human review, complaints and redress. It should also identify the actuarial and conduct roles involved in acceptance.

Fairness requires contextual assessment. The platform should record the affected population, relevant outcome measures, limitations and actions. A generic bias test with a green status does not explain whether the insurer considered product design, risk pooling, lawful differentiation or access. The record should preserve the evidence and reasoning used by accountable management.

The sample should also test changes. A pricing model may perform within aggregate thresholds while affecting a product or population differently. The overlay should support segmented monitoring and escalation. Where a vendor controls the model, the insurer still needs sufficient information to govern its use and outcomes. The product should help document that responsibility rather than transfer it rhetorically to the vendor.

44. RUN THE SECURITIES DILIGENCE SAMPLE

The securities sample should include a trading or surveillance model, a customer communication or advisory tool and an operational use. The team should connect each system to written supervisory procedures, recordkeeping, market or customer obligations and escalation. It should test unusual conditions, override, suspension and post-event review.

For trading, the record should show intended behaviour, limits, monitoring and response to market conditions outside the training or test set. For customer communications, it should show content review, basis, approval and preservation. For advisory or portfolio use, it should show how the system fits the firm's duties and how a responsible person reviews material output.

The buyer should inspect whether the target supplies configurable workflows or sector expertise through manual service. Both can have value, but their economics differ. Repeatable controls and supported integrations can earn software margins. Ongoing bespoke interpretation, surveillance tuning and supervisory review should be modelled with the required people and cost.

45. CONVERT FINDINGS INTO INTEGRATION WORK PACKAGES

Each diligence finding should become a work package with scope, owner, dependencies, evidence, cost and completion criteria. Product consolidation, security remediation, regulatory-content alignment, data migration, contract changes and customer acceptance should remain separate packages. This allows management to sequence work and prevents a broad synergy target from obscuring control dependencies.

A package should state the regulated decisions it can affect. It should identify the source and target records, reconciliation method and rollback. Historical evidence needs special treatment because regulators, auditors and customers may need to reproduce earlier decisions. Migration should preserve identifiers, versions, ownership, timestamps and relationships, or maintain an accessible archive with a tested retrieval path.

The finance work package should reconcile expenditure, operating cost and revenue risk. Savings should enter the model only after the related capability can be retired without weakening service or controls. Customer-facing product changes should include communication, consent where required and support. These disciplines connect the integration plan to the valuation bridge.

46. MAINTAIN A DECISION REGISTER

The transaction should maintain a decision register from diligence through integration. Each entry should state the decision, evidence, owner, date, assumptions, alternatives considered and review trigger. Important entries include product selection, regulatory mapping, customer migration, security exceptions, capital allocation and contingent consideration.

The register helps management separate facts from assumptions. It also supports later review when a source changes, a customer resists migration or a regulator raises a concern. An assumption should identify the evidence required to confirm it and the consequence if it is wrong. This is especially important for cross-selling, product portability and expected cost removal.

The register should avoid becoming a substitute for detailed records. It is an index to the underlying analysis, approvals and artefacts. The combined company can use the same discipline for material product changes after close. This creates continuity between transaction governance and the model-risk engine the buyer is acquiring.

47. TEST INTERNATIONAL SCALABILITY

International scalability requires more than translating controls or adding jurisdictions to a library. The buyer should identify the regulated entity, applicable authority, legal basis, data location, customer duty and supervisory record for each deployment. A global financial group

may use one platform while assigning responsibility to different legal entities and functions. The system should preserve that structure.

The common ontology can support cross-border reporting when terms and relationships are stable. Local overlays still need qualified review, effective dates and documented interpretations. The target should demonstrate how it handles conflicting or overlapping requirements, customer-specific policies and later regulatory change. A broad jurisdiction count has limited value without current mappings and deployed customers.

Commercial analysis should separate software localisation from professional interpretation. Language support, date formats and workflow configuration can be repeatable product features. Legal mapping, supervisory engagement and customer policy design may require specialists. The valuation model should reflect the actual delivery method, margin and update obligation for each market.

Data transfer and outsourcing requirements can also constrain a shared architecture. Diligence should examine hosting, subprocessors, access, encryption and customer options. The integration plan should avoid moving data or support activity until contractual and regulatory requirements are confirmed. International scale creates value when the operating model can maintain these controls without rebuilding the product for every customer.

48. AUDIT THE SYNERGY CASE AFTER CLOSE

The board should audit the synergy case at defined intervals after close. The review should compare the transaction model with actual migration cost, retained recurring revenue, collected cross-selling, product reliability and contribution margin. Differences should be explained by source evidence rather than revised narratives.

Product synergy should be measured through retired duplicate capability, reduced maintenance, common releases and verified customer use. Commercial synergy should be measured through contracted and collected revenue linked to the combined offering. Cost savings that shift work to customer support or regulatory-content teams should be recorded at net contribution.

The audit should also review control outcomes. A lower cost base does not support value when evidence quality, security, customer acceptance or supervisory readiness deteriorates. Material exceptions should trigger remediation and, where relevant, contingent-consideration provisions. This review closes the loop between the acquisition thesis and the operating record.

The same method supports future add-ons. Actual experience provides a better basis for portability estimates, integration schedules and pricing. Management can refine the diligence matrix while preserving the distinction between observed outcomes and assumptions for the next transaction.

49. CONCLUSION

RegTech roll-ups across finance can create operating leverage when one model-risk engine supports reliable inventories, workflows, evidence, testing, issue management and audit history. Banking, insurance and securities firms can share these primitives. Their materiality tests, legal duties, validation depth, accountable roles and customer protections remain specific to the regulated use.

The acquirer's task is to prove where reuse already works and fund the work required where it does not. The Control Portability Framework links that proof to diligence, valuation and integration. It protects the buyer from paying for a regulatory narrative while giving genuine shared infrastructure appropriate credit.

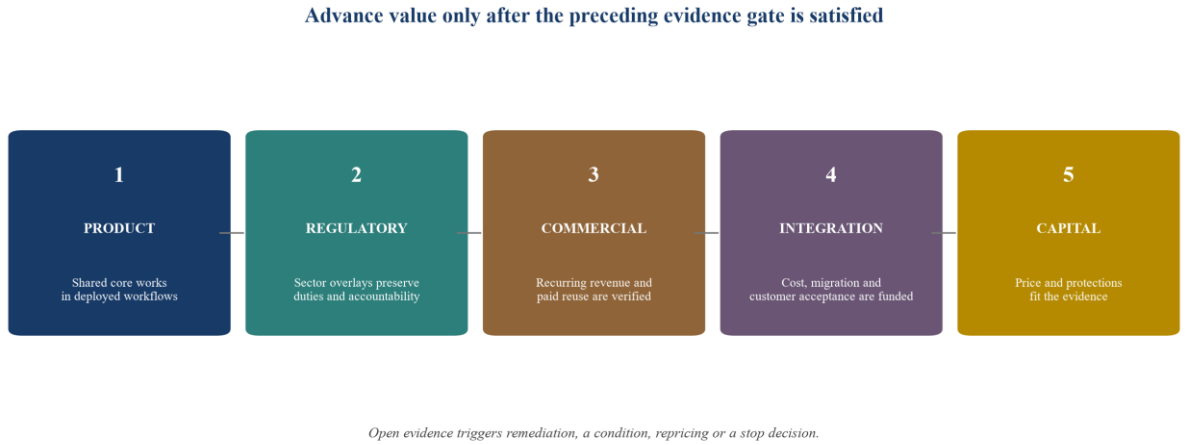

Figure 3. Capital, value and risk gates for a RegTech roll-up

Every gate requires evidence before the next value claim enters the transaction model.

REFERENCES

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