Algorithmic Deal Risk: Pricing Drift, Dependency and Change-of-Control Exposure

A Transaction Framework for Converting AI Operating Uncertainty into Price, Escrow, Covenants and Integration Controls

Authored by Chennakeshav Adya, Independent Researcher

September 2026

Abstract

Artificial-intelligence businesses can change between signing and closing without a conventional asset being sold, lost or visibly impaired. A foundation-model provider may alter price, capability or terms. Model performance may drift as customer behaviour and data change. A data licence may restrict assignment, training or post-acquisition use. A small group of engineers may hold essential operating knowledge. New regulatory duties may attach to a provider, deployer or downstream integrator. These exposures can affect revenue retention, gross margin, remediation investment, closing certainty and the buyer's ability to operate the acquired product.

This paper develops an algorithmic deal-risk framework for boards, investment committees, acquirers, sellers, sponsors, lenders and transaction teams. It translates model drift, upstream dependency, data and software rights, critical-person concentration and change-of-control obligations into diligence evidence, valuation scenarios and transaction protections. The framework separates an observed baseline from an underwritten case, identifies cash-flow pathways, assigns control owners and preserves measurable post-close tests.

Five original figures and five decision tables present an algorithmic dependency map, drift scenarios, a licence waterfall, a critical-person heat map and a transaction-protection term sheet. A worked example uses a hypothetical acquisition and analytical assumptions. All amounts, percentages, probabilities, multiples and scenarios are illustrative assumptions. The paper does not provide legal, regulatory, tax, accounting, cybersecurity, employment or investment advice and does not recommend a transaction, valuation, contractual term or technology architecture.

JEL Classification: G34, G32, K22, L86, O32

Keywords: artificial intelligence, mergers and acquisitions, model drift, change of control, data licensing, transaction pricing, escrow, covenants, technology due diligence, integration risk

1. Define algorithmic deal risk in cash-flow terms

Algorithmic deal risk is the possibility that the technology, rights, people or external services required to produce an AI-enabled customer outcome will change the cash flows underwritten in a transaction. The risk can arise before signing, between signing and closing, or after control transfers. It becomes financially relevant through lost customers, lower conversion, higher inference cost, service credits, remediation investment, delayed product releases, regulatory restrictions or a longer integration timetable.

A diligence finding should therefore state the affected operating capability, the evidence available, the mechanism of change, the cash-flow pathway, the time to detect, the time to remediate and the party able to control the outcome. Describing a model as high risk without this chain does not support a price or contract decision. Describing a product as accurate without

defining the tested population, period and threshold provides equally little protection.

The transaction team should establish one register that connects technical evidence to commercial and legal decisions. Each item receives an owner, materiality range, mitigation, residual exposure and post-close test. The register should distinguish conditions that affect enterprise value, equity value, closing certainty, integration cost and ongoing governance. A single issue may affect several of these categories through different mechanisms.

2. Establish the signed-state operating baseline

The buyer needs a reproducible picture of the product and economics at the measurement date. The baseline should identify exact model and prompt versions, routing rules, data snapshots, evaluation suites, infrastructure, human-review steps, upstream services, customer configurations and cost allocations. It should record the period over which performance, availability, revenue and margin were observed. This is the reference state for warranties, interim covenants, closing tests and later value measurement.

Management presentations often combine product versions or customer cohorts. A buyer should reconcile headline metrics to logs, invoices, customer records and controlled evaluations. If historical results cannot be reproduced, the buyer can treat the difference as a measurement limitation and build a wider valuation range. The objective is a defensible baseline, not a laboratory-perfect reconstruction of every past decision.

The baseline should be preserved through hashed artefacts, access-controlled exports and documented test procedures. The buyer and seller should agree which changes are ordinary operation and which require notice or consent before closing. A version change that appears technical can alter cost, output quality, explainability or compliance. Interim controls should therefore apply to the full operating system rather than the model name alone.

3. Map the complete algorithmic dependency chain

An AI product can depend on foundation-model APIs, cloud regions, vector databases, labelling vendors, security tools, open-source libraries, proprietary datasets, payment systems and named individuals. A dependency map should trace the customer outcome through every material component. It should identify the contractual counterparty, technical interface, ownership, jurisdiction, concentration, change rights, termination terms, replacement option and revenue exposed.

Runtime evidence is essential. Architecture diagrams may omit shadow services, fallback models and customer-specific routes. The buyer can compare code repositories, environment configuration, telemetry, network calls, invoices and incident records. Differences should be resolved before the dependency is scored. A low-cost component can remain highly material if it sits on every production request and cannot be replaced quickly.

Dependency should be measured through impact and substitutability. Impact considers revenue, margin, customers, obligations and safety. Substitutability considers technical equivalence, data portability, contractual permission, integration effort, evaluation time and customer approval. The combination determines whether the risk calls for price, a closing condition, escrow, a covenant, a funded replacement plan or continuing monitoring.

Figure 1. Algorithmic dependency and value-exposure map

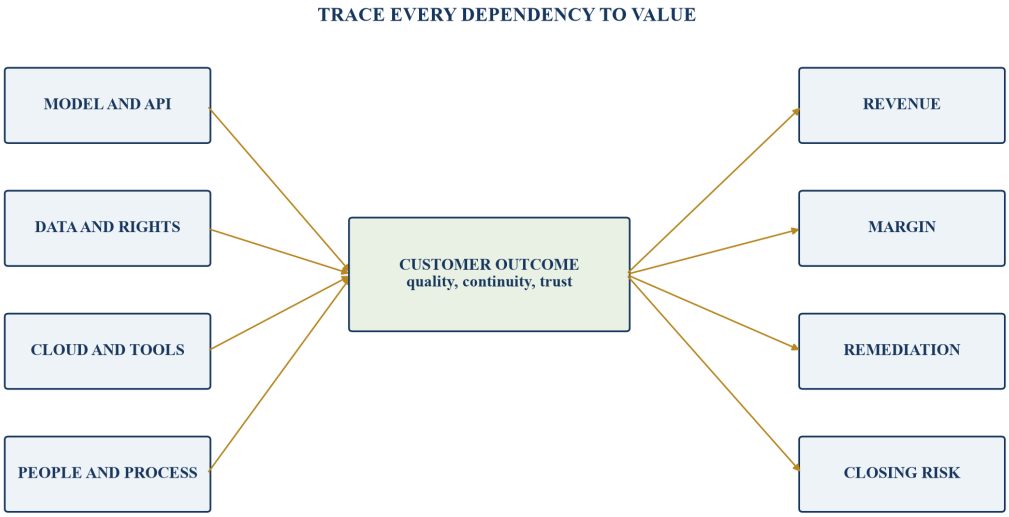

Each operating dependency is linked to the customer outcome, exposed cash flow and transaction response.

Table 1. Dependency evidence and transaction response

Dependency	Evidence	Value pathway	Potential response
foundation model or API	contract, version history, telemetry and invoices	capability, price, latency and availability	scenario, covenant, fallback and escrow
proprietary data	provenance, licence, consent, lineage and use logs	product quality, differentiation and compliance	condition, warranty, indemnity and isolation
cloud and specialist tool	agreement, architecture, region and exit plan	continuity, margin and security	consent, commitment, migration reserve and test
open-source component	inventory, licence, contribution and vulnerability record	distribution rights, remediation and release delay	remediation, warranty and holdback
critical person	access, decision rights, knowledge and succession	continuity, release capability and customer trust	retention, transition covenant and operating gate
customer-specific configuration	statement of work, acceptance and support record	retained revenue and service obligation	consent, communication and revenue protection

Materiality combines cash-flow impact, replacement time and control.

4. Separate model drift from model failure

Drift is a change in inputs, relationships, outputs or operating context that can move performance away from the approved range. Data drift changes the characteristics of inputs. Concept drift changes the relationship between inputs and desired outcomes. Performance drift changes measured quality. Operational drift can arise through routing, prompts, retrieval content, infrastructure or human review. These conditions can exist even when the underlying model file remains unchanged.

A transaction should distinguish normal variation from a threshold breach. The diligence team should identify metrics, segments, observation windows, confidence limits and business

consequences. A global average can hide a material decline for a regulated use case, language, customer cohort or high-value workflow. Monitoring should therefore follow the value and obligation at risk rather than a single model score.

Failure is an inability to satisfy an agreed requirement, such as availability, accuracy, safety or contractual service. Drift may create an early warning before failure. A buyer can value that warning only if the target has stable measurement, ownership and response procedures. Where monitoring is weak, the buyer may need an independent baseline, a funded observation period and wider downside scenarios.

5. Build drift scenarios around observed evidence

The valuation model should use a limited set of coherent drift scenarios. A base case can reflect observed performance and current remediation spend. A moderate case can model deterioration in a material segment, greater human review and delayed sales. A severe case can model customer restrictions, replacement of an upstream model, regulatory remediation or loss of a data source. Each scenario should state trigger, duration, response and recovery evidence.

Probability estimates should not imply precision unsupported by history. The buyer can use ranges and decision thresholds, then test valuation sensitivity to the probability and duration of each state. The central question is whether the purchase price and capital plan remain acceptable when operating evidence moves outside the underwritten range.

The model should avoid double counting. Lower revenue, higher operating cost, remediation investment and multiple compression may reflect the same root cause. The team should specify which cash-flow effects enter the forecast and which risks influence the discount rate or transaction terms. A documented bridge supports investment-committee scrutiny and post-close measurement.

Figure 2. Drift scenarios and cash-flow transmission

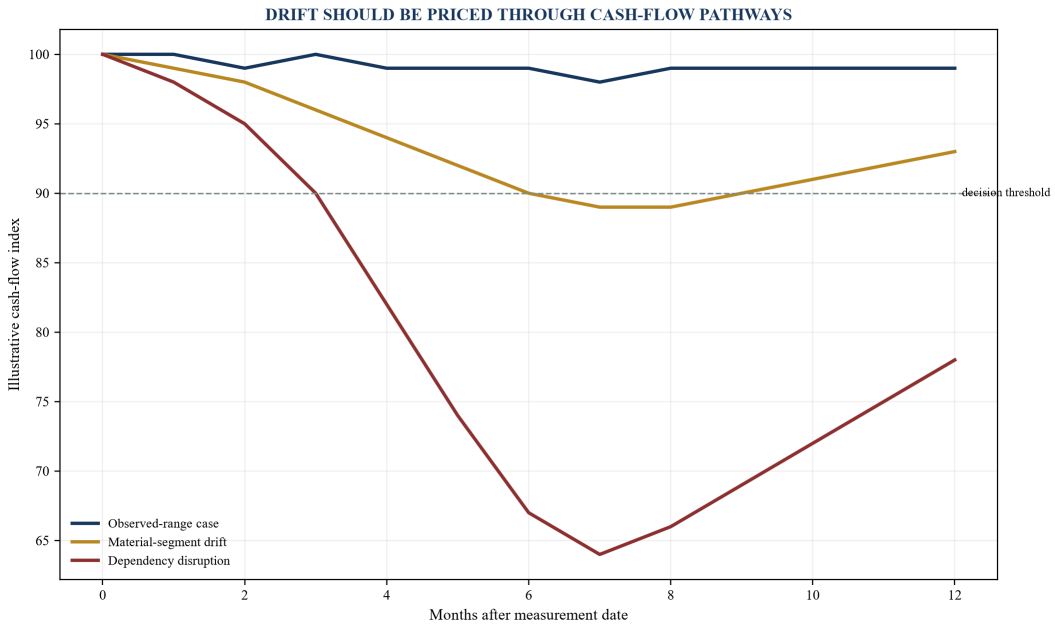

Illustrative scenarios show how operating movement can affect revenue, margin and remediation.

Table 2. Illustrative drift scenario record

Scenario	Trigger	Operating response	Valuation pathway	Protection test
observed range	metrics remain within validated limits	ordinary monitoring and maintenance	underwritten revenue, margin and investment	baseline reproduced at closing
segment deterioration	material cohort crosses threshold	investigate, retrain, route or increase review	lower conversion and higher service cost	cohort metric and remediation milestone
upstream model change	capability, price or terms move materially	test fallback, renegotiate or migrate	margin pressure, delay and replacement investment	provider notice and approved fallback
data restriction	source, consent or permitted use changes	isolate, delete, replace or limit features	reduced differentiation and remediation	verified rights and technical enforcement
regulatory restriction	use case requires new control or suspension	assess, document, constrain and validate	delayed revenue and compliance investment	specialist opinion and control acceptance

Values are analytical assumptions and require transaction-specific evidence.

6. Test benchmark stability and decision relevance

A benchmark can create false comfort when its data are stale, contaminated, narrow or weakly connected to the customer outcome. The buyer should reproduce material claims using locked environments, documented prompts, representative samples, repeated runs and clear scoring. It should examine variance and error distribution rather than relying only on a mean result.

Benchmark relevance should be traced to revenue and obligations. A coding score may not predict performance in a customer workflow that depends on retrieval, tools, permissions and human review. A safety benchmark may not cover the language or domain in which the product is sold. The diligence plan should include transaction-specific tests designed around customer use, failure cost and expected integration changes.

NIST's generative-AI profile emphasises that systems often rely on multiple third-party components and data sources, complicating attribution. This matters to a buyer because a reproduced result should identify the model, surrounding system and evidence owner. A claim that cannot be attributed cannot be protected reliably through a simple model warranty.

7. Identify upstream provider change rights

Foundation-model and infrastructure providers can change versions, deprecate endpoints, adjust rate limits, revise safety filters, alter data handling and reprice usage. The diligence team should read the operative agreements and service documentation, then reconcile them to the production configuration. Marketing descriptions and informal assurances should not substitute for contractual rights.

The analysis should identify notice periods, unilateral change provisions, service levels, liability limits, assignment restrictions, data-use terms, termination rights and available commitments. It should test whether the target receives enterprise terms or relies on a standard online agreement. A technically portable application may remain commercially locked in if comparable performance requires expensive redevelopment or customer revalidation.

The buyer should quantify replacement lead time and interim cash exposure. This includes engineering, evaluation, customer approval, security review, data migration, parallel run and possible revenue loss. A credible fallback is a tested operating capability, not the name of an alternative provider on a slide.

8. Price API and compute concentration

AI gross margin can change quickly when token price, context length, routing, caching, utilisation or human review changes. The buyer should reconstruct cost per successful customer outcome using invoices and telemetry. Unit cost should distinguish input, output, embeddings, retrieval, tools, storage, networking, reserved capacity, failed calls and review labour.

Concentration analysis should show spend and revenue exposed to each provider, model family, cloud region and capacity commitment. The team can model price increases, discounts expiring, volume growth, inefficient prompts and fallback use. Revenue-based pricing can hide cost sensitivity until usage rises or customer behaviour changes.

Transaction protection can combine a purchase-price scenario with operating controls. These may include an interim covenant limiting material provider changes, a closing test for continued access, a funded migration reserve, or an earn-out metric based on contribution margin rather than revenue alone. Contract specialists should draft and interpret all provisions.

9. Construct a rights and licence waterfall

The acquired product may include code, weights, prompts, data, content, labels, documentation, trademarks, patents, open-source components and customer-specific developments. Each asset can have different owners, licences and assignment rules. The rights waterfall should start with the customer feature and trace every required right through the chain of contributors and suppliers.

The analysis should distinguish ownership from permission. The target may own source code while lacking the right to train on customer data after control changes. A permissive software licence may impose notice requirements. A dataset may be available publicly while its collection, personal data or database rights remain constrained. Specialist advice is required for the relevant jurisdictions and facts.

The transaction team should connect gaps to remedies. A missing consent may be a closing condition. An ambiguous training right may call for a specific warranty, indemnity, remediation covenant or product limitation. A replaceable component may justify a quantified reserve. The response should reflect enforceability, materiality and the buyer's ability to verify compliance.

Figure 3. Rights and licence waterfall

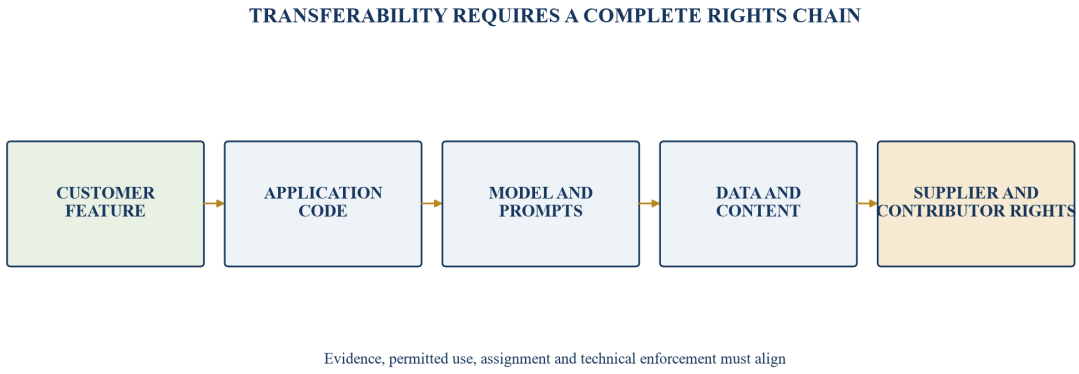

A customer feature is supported only when every material right in the chain is evidenced and transferable.

Table 3. Licence waterfall diligence record

Layer	Core evidence	Change-of-control question	Transaction response
customer feature	contract, specification and acceptance	may the buyer continue and change the service?	consent, communication and service covenant
application code	title, assignments and repository history	did employees and contractors assign rights?	confirmatory assignment, warranty and remediation
models and prompts	provider terms, ownership and version record	may the buyer access, modify and commercialise?	consent, fallback, escrow or use restriction
data and content	provenance, purpose, licence and retention	may data transfer and support intended post-close use?	condition, segregation, deletion or indemnity
open source	software bill, licences and notices	do distribution or source obligations affect the product?	cure, replace, disclose and warrant
supplier contribution	agreement, deliverables and termination	does control change trigger consent, price or termination?	condition, renegotiation and replacement reserve

Qualified specialists should interpret the operative rights and restrictions.

10. Examine training-data and output exposure

Training and fine-tuning data should be traced by source, purpose, jurisdiction, licence and retention. The buyer should identify whether the target created a model, modified a third-party model, or operates as a downstream provider. The EU's current general-purpose AI framework places documentation, copyright-policy and training-content-summary duties on relevant providers, with additional duties for models with systemic risk. Applicability requires fact-specific legal analysis.

Output exposure includes confidentiality, personal data, intellectual property, harmful content, discrimination and reliance. The target's contracts may allocate these risks differently across customers. Product controls such as filtering, retrieval permissions, review and logging should be tested against actual commitments. A policy document without technical enforcement has limited transaction value.

The buyer should distinguish a remediable documentation gap from an operating model that depends on a prohibited or unlicensed use. The first may require cost and time. The second can affect the product's lawful operating perimeter and terminal value. Transaction terms should reflect that difference.

11. Allocate provider and deployer responsibilities

AI regulation increasingly distinguishes the entity that develops or places a model or system on the market from the entity that deploys it. An acquisition can change branding, integration, purpose, distribution and control. Those changes may alter the parties' responsibilities. The diligence team should map each product and jurisdiction before assigning a generic compliance label.

The map should identify the relevant system, intended purpose, users, affected persons, risk classification, documentation, human oversight, monitoring, incident reporting and downstream information. It should also identify which entity can produce the evidence. A contractual allocation does not remove a statutory obligation, although it may support cooperation and recourse between parties.

The buyer can connect missing capability to closing and integration decisions. If the target cannot supply information required by downstream customers, revenue retention may be exposed. If the buyer's planned integration creates a new provider role, the integration budget should include the resulting assessment, documentation and controls.

12. Review UAE financial-services AI obligations

The Central Bank of the UAE issued responsible-AI guidance for licensed financial institutions in February 2026. The guidance focuses on consumer protection, transparency, bias, accountability, explainability, data privacy and responsible use. A transaction involving a regulated financial institution or its material technology provider should examine how these principles apply to the specific product and contractual chain.

The buyer should test governance, customer disclosure, model approval, data protection, outcome monitoring, complaints, human review and third-party oversight. It should identify which records support management's statements and whether controls operate across outsourced models. The analysis may affect customer consents, product changes, integration sequencing and ongoing assurance cost.

Regulatory interpretation belongs to qualified counsel and relevant specialists. The transaction model can nevertheless capture observable remediation work, affected revenue, decision dates and capital requirements. This creates a bridge between compliance evidence and valuation without treating a legal conclusion as a technical score.

13. Analyse change-of-control triggers contract by contract

Change-of-control exposure can appear in customer, supplier, data, cloud, employment, financing, insurance, grant and partnership agreements. Clauses may require consent, notice, repayment, renegotiation or termination. The diligence team should search the full contract set and confirm results through legal review and business-owner interviews.

Materiality should reflect more than contract value. A small data agreement may enable a major product. A low-spend model provider may support most revenue. A customer with a termination right may influence other customers or a regulatory approval. The register should therefore link each clause to technical dependencies and cash flows.

Consent strategy needs an owner, message, sequence, fallback and deadline. Early requests may reveal the transaction or weaken bargaining leverage. Late requests may threaten closing. The parties should align the approach with legal advice, transaction conditions and operating continuity. Interim covenants should prevent avoidable new exposures.

14. Test customer transfer and trust

AI customers may have approved a named model, hosting region, security architecture, data use, subcontractor or human-review process. An acquisition can change one or more of these elements even when the legal supplier remains the same. The buyer should map contractual commitments, assurance reports, questionnaires and side letters to the intended integration plan.

Customer diligence should segment revenue by sensitivity to model, data, ownership and service change. Interviews or permitted confirmation procedures can test reasons for purchase, switching cost, trust and renewal criteria. The buyer should avoid promising future product changes before authority and feasibility are established.

Retention scenarios should reflect customer evidence and contractual rights. A transaction protection could link part of consideration to retained recurring gross profit, customer consent or renewal, subject to careful definitions and seller-control considerations. The operating plan should provide a communication sequence, named account owner and escalation route.

15. Measure critical-person dependence

AI products often concentrate knowledge across a small group that understands data lineage, evaluation design, deployment, incident response or customer configuration. Job titles do not reveal this dependence. The buyer should map decisions, systems, credentials, relationships and tacit knowledge to named roles, then test backup coverage and documentation.

Concentration becomes transaction risk when loss of a person can interrupt service, prevent a release, undermine a claim, delay integration or weaken customer confidence. The analysis should consider notice periods, retention, incentives, immigration, location, restrictive covenants, succession, workload and cultural fit. Employment law and local enforceability require specialist advice.

Retention payments alone do not transfer knowledge. A transition plan should assign artefacts, shadowing, dual approval, access transfer, customer handover and acceptance tests. The buyer can connect completion to escrow release, earn-out operations or integration gates where appropriate and legally supportable.

Figure 4. Critical-person dependency heat map

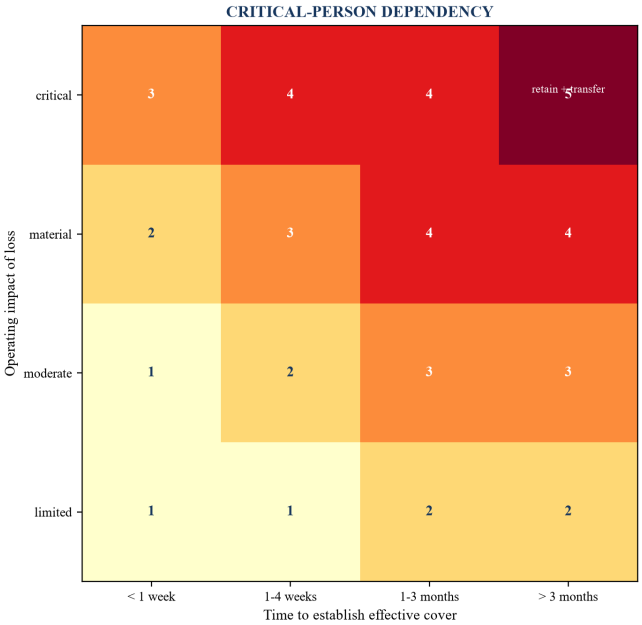

Concentration combines operating impact with the time required to establish credible cover.

Table 4. Critical-person evidence and protection plan

Capability	Evidence	Failure consequence	Protection and test
model release	approvals, repository history and runbook	inability to deploy or roll back safely	deputy, dual control and supervised release
data lineage	catalogue, transformation record and interviews	rights or quality cannot be reconstructed	documented lineage and independent walkthrough
evaluation design	test sets, scoring code and threshold history	claims and changes cannot be validated	versioned suite and reproducibility test
customer configuration	account records, support history and access	service disruption or renewal risk	named handover and customer acceptance
incident response	incident record, contacts and exercises	delayed containment and notification	joint exercise and tested escalation
supplier relationship	contract, correspondence and technical dependency	loss of access or weaker terms	executive transition and documented fallback

Retention should be paired with controlled knowledge transfer and operating acceptance.

16. Reconcile access, credentials and operating authority

The buyer should know who can change production, access sensitive data, approve releases, alter billing, communicate with customers and declare incidents. Credentials held by individuals or suppliers create continuity and security exposure. The diligence team should test access through approved observation and evidence without interfering with production.

The access map should reconcile identity systems, cloud consoles, repositories, model registries, data stores, monitoring, support and third-party portals. It should identify shared accounts, inactive administrators, missing logs and break-glass procedures. Remediation may need to wait until closing to avoid disrupting the target, so the Day-One plan should prioritise sequence and rollback.

Transaction documents can support delivery of access evidence and cooperation, while the integration plan establishes new authority. The buyer should preserve forensic records and respect applicable privacy and employment obligations. Account transfer should never rely on informal password sharing.

17. Inspect incident history as valuation evidence

Incidents reveal how the product behaves under stress and how the organisation responds. The buyer should examine security, privacy, availability, harmful-output, data-quality, customer and regulatory incidents. It should reconcile tickets, status pages, logs, insurance notices, customer credits and board reporting to identify omitted or differently classified events.

The analysis should measure detection time, containment, recovery, recurrence, root-cause quality and completion of corrective actions. Repeated low-severity events can expose weak control design. A single severe event may be less predictive if the cause was removed and the fix was independently tested. The transaction team should preserve context rather than scoring incident count alone.

Unresolved actions can become a closing condition, escrow item, specific indemnity or funded remediation plan. Historical events may also affect warranties, disclosure, insurance and customer retention. Qualified advisers should determine the appropriate legal and accounting treatment.

18. Evaluate monitoring and observability

An AI product needs visibility across input quality, model and prompt versions, retrieval, tool calls, output quality, latency, cost, safety, human review and customer outcomes. The buyer should test whether monitoring covers material segments and whether alerts lead to accountable action. A dashboard is useful only when definitions, data lineage and response ownership are reliable.

Monitoring should distinguish leading and lagging indicators. Input shift, provider notices, latency and review escalation may warn before churn or revenue loss. Customer complaints and failed renewals show realised impact. The transaction model can connect both types to thresholds and response costs.

Post-close covenants or earn-outs that depend on technical metrics require stable definitions and anti-manipulation controls. Where metrics can be influenced by routing, sample selection or threshold changes, the parties need agreed measurement procedures and access to evidence. Simpler commercial measures may be more enforceable when technical instrumentation is immature.

19. Verify security across the AI supply chain

AI supply-chain risk includes model artefacts, training pipelines, code, data, APIs, plugins, cloud infrastructure and third-party tools. The buyer should examine access control, secrets, provenance, vulnerability management, logging, environment separation, incident response and supplier assurance. It should also test for risks specific to the product, including prompt injection, data leakage, unsafe tool use and model extraction.

The UK National Cyber Security Centre's secure-AI guidance organises security across design, development, deployment and operation. This lifecycle view is useful in transactions because the buyer acquires both a current system and the process used to change it. A secure snapshot cannot compensate for an uncontrolled release pipeline.

Security findings should be translated into exploitability, affected assets, likely business impact, remediation and residual exposure. The transaction response may include pre-close cure, price, escrow, insurance review, specific covenant or an integration gate. Security specialists should validate severe findings and remediation.

20. Test resilience and fallback operation

Fallback claims should be demonstrated. The buyer should observe or reproduce provider failover, model substitution, degraded mode, queue management, rollback, backup restoration and incident communication. Tests should reflect realistic customer volume and data constraints. A fallback that preserves uptime but materially degrades quality or breaches a contract is incomplete.

Resilience analysis should identify recovery time, recovery point, capacity, regional dependency, manual effort and customer impact. It should include third-party incidents and internal deployment errors. The cost of maintaining parallel capability belongs in the valuation model when it is required for the underwritten service level.

A closing test can confirm access to critical providers and recovery procedures. A post-close gate can delay platform consolidation until resilience is proven. The buyer should avoid removing redundancy solely to capture an early synergy if that redundancy protects material revenue.

21. Quantify remediation investment

Remediation should be costed as a programme rather than a list of findings. The estimate should include people, specialist advisers, licences, infrastructure, parallel run, customer work, validation, documentation, contingency and management attention. It should identify which work is required to operate, comply, integrate or deliver the acquisition thesis.

Timing matters. A cost incurred before revenue growth has a different value effect from one funded by later cash flow. Delays can also defer synergies or product releases. The valuation model should therefore place cost and benefit in periods and state dependencies.

The buyer should distinguish seller cure, purchase-price adjustment, debt-like item, escrow-funded work and buyer investment. Accounting treatment depends on the facts and applicable standards. The commercial model can show the cash consequence while qualified advisers determine classification.

22. Translate risk into valuation scenarios

Valuation should begin with the observed operating baseline and explicit assumptions. The algorithmic risk register then changes revenue, gross margin, operating cost, investment, working capital or timing. The team can calculate enterprise-value sensitivity under coherent scenarios and identify which assumptions drive the range.

Multiple compression should be used carefully. If lower growth, margin and durability are already in cash flows, an additional arbitrary reduction can double count risk. A lower terminal multiple may be justified when dependency or rights uncertainty affects long-term defensibility beyond the forecast period. The rationale should be documented.

The investment committee should see both expected value and downside liquidity. A transaction may retain expected value while requiring more cash to remediate or withstand disruption. Lenders and sponsors need to assess covenant headroom, debt service and funding availability under the same scenarios.

23. Choose between price and contingent consideration

A fixed price adjustment can address an exposure whose expected cash effect is measurable and largely outside seller control after closing. Contingent consideration can address uncertainty that resolves over time, such as customer retention, consent, performance or margin. The metric, period, control rights and dispute process determine whether the mechanism reflects economics or creates conflict.

Technical metrics can be precise in code and ambiguous in commerce. Model performance may change with sample, routing, threshold and customer mix. Commercial metrics such as recurring gross profit may capture value more directly but can be influenced by integration decisions. The parties need definitions that align control and measurement.

The buyer should model behavioural incentives. An earn-out tied to revenue may encourage costly usage or weak customer selection. A margin measure may discourage necessary investment. A balanced scorecard can become too complex to administer. Transaction specialists should select the smallest measurable set that reflects the unresolved risk.

24. Design escrow and holdback around verifiable release events

Escrow can support recovery for warranty breaches or secure performance of specific obligations, subject to applicable law and negotiated terms. A release event should be objective, time-bound and supported by evidence accessible to both parties. Examples may include receipt of a material consent, completion of a rights cure, transfer of critical access or an independently verified remediation.

The escrow amount should relate to plausible exposure and enforcement, not a generic percentage. The team can model replacement cost, affected gross profit, customer liability, remediation and delay. Caps, baskets, exclusions and insurance interact with the practical recovery.

Technical acceptance tests should be defined before signing when possible. If a test depends on a future environment controlled by the buyer, the agreement should address cooperation and change. Qualified counsel should draft all escrow, holdback and claim provisions.

25. Draft representations around evidence, not adjectives

A representation that a model is accurate, robust or compliant can be difficult to interpret. More useful provisions may address disclosed systems, ownership and rights, material contracts, incidents, claims, documentation, testing, change history and operation in accordance with specified policies or obligations. The appropriate language depends on the transaction and jurisdiction.

Disclosure should connect exceptions to the affected asset, customer and cash flow. A large data room does not create clarity if the buyer cannot identify which document governs a production dependency. Schedules should reconcile to the algorithmic risk register and materiality analysis.

Representations do not replace diligence or operating control. Their value depends on knowledge qualifiers, survival, caps, exclusions, insurance and collectability. Counsel should advise on drafting and enforceability. The transaction team should maintain evidence supporting the commercial rationale.

26. Use interim covenants to preserve the underwritten system

The period between signing and closing can be material for an AI business. Models, prompts, datasets, providers, pricing, staff and customer commitments may change through ordinary operations. Interim covenants should preserve agreed elements of the business while allowing necessary development and compliance.

The parties can define notice or consent for material model substitution, new data use, provider termination, pricing change, critical-person departure, security incident, regulatory contact and customer commitment. The threshold should focus on economic and operating materiality. An excessively rigid covenant can prevent responsible fixes or product delivery.

The buyer should operate through agreed governance and respect applicable competition rules before closing. Clean-team arrangements and information controls may be necessary. Legal advisers should determine the permissible process and drafting.

27. Set closing conditions for existential dependencies

A closing condition may be appropriate when a missing consent, right, access or approval prevents the buyer from operating a material part of the business. The condition should be specific and capable of objective determination. The parties should understand who controls satisfaction and the consequences of delay.

Conditions can also address delivery of artefacts, key agreements or regulatory approvals. A condition should not become a substitute for completing ordinary diligence. The buyer should evaluate termination rights, long-stop dates, financing and customer effects.

Where a full condition is disproportionate, alternatives include waiver with price protection, escrow, a transition service, licence, carve-out or funded replacement. The decision depends on materiality, control, time and enforceability. Specialist advice is essential.

28. Build the transaction-protection term sheet

The protection term sheet should list each material exposure, evidence, cash-flow pathway, control owner, pre-close action, contractual response, post-close test and residual risk. This creates a shared working document across commercial, technical, legal, financial, tax, regulatory, security and integration teams.

The term sheet should preserve the distinction between a proposed commercial response and approved legal drafting. It can identify options, valuation ranges and dependencies before the parties settle language. Decisions and rejected alternatives should be recorded to support governance.

The final package should remain internally coherent. A purchase-price reduction assumes the buyer accepts and funds the risk. An escrow assumes a recoverable event can be defined. A covenant assumes the obligated party can perform. An indemnity assumes the loss and claim process are workable. Each mechanism needs an operating plan.

Figure 5. Algorithmic transaction-protection architecture

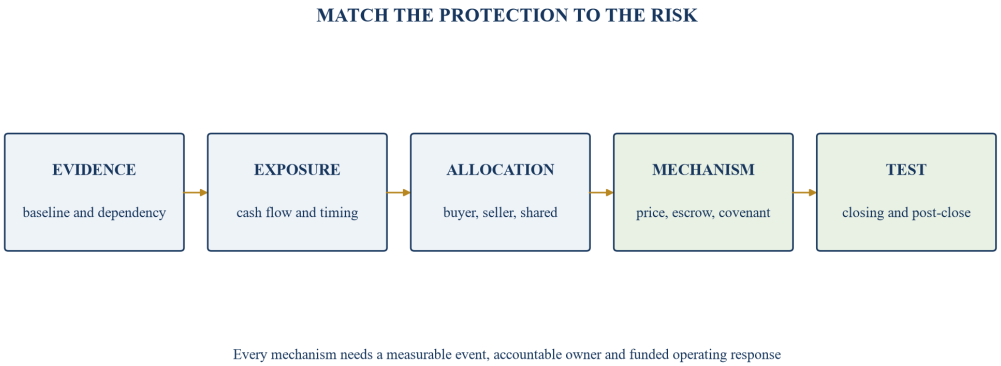

Protections are selected according to timing, control, measurability and severity.

Table 5. Illustrative algorithmic transaction-protection term sheet

Exposure	Evidence and valuation effect	Possible mechanism	Release or monitoring test
model drift outside approved range	reproduced cohort decline and cash-flow scenario	price range, earn-out or remediation reserve	agreed evaluation suite and threshold
upstream provider change	concentration, change rights and replacement lead time	interim covenant, fallback plan and escrow	continued access and tested alternative
restricted data or licence	incomplete right affecting a material feature	condition, warranty, indemnity or carve-out	consent, cure or verified isolation
critical-person concentration	uncovered authority or knowledge affecting continuity	retention, transition covenant and holdback	documented transfer and supervised operation
change-of-control consent	termination or repricing right on material dependency	condition, seller covenant or price protection	executed consent or accepted alternative
regulatory remediation	identified control and investment requirement	funded plan, covenant and integration gate	specialist acceptance and operating evidence

Terms require transaction-specific negotiation and qualified legal advice.

29. Align protections with control after closing

Transaction terms should reflect who controls the result. A seller may have limited ability to deliver a post-close product metric after the buyer changes models, staff, pricing or customers. A buyer may have limited ability to secure a pre-close supplier consent. Misaligned control creates disputes and weakens incentive design.

The protection should define permitted and required actions, access to evidence, cooperation, change procedures and dispute resolution. If the buyer needs freedom to integrate, a fixed price or escrow may be cleaner than a long technical earn-out. If the seller continues to operate the business, carefully defined performance measures may remain useful.

Governance should include named decision makers and escalation. Technical disagreements can require an independent expert with access to artefacts and agreed methods. Legal advisers should determine how expert determination interacts with contractual claims.

30. Integrate without destroying the evidence

Post-close integration can overwrite logs, repositories, accounts, model versions and customer configurations needed to validate warranties or earn-outs. The buyer should preserve relevant evidence before changing systems. Retention must respect privacy, security, legal hold and contractual requirements.

The integration plan should sequence access transfer, key rotation, monitoring, incident response, rights remediation, people transition and platform change. High-risk migrations should use acceptance gates and rollback. Product improvement can begin early, while changes affecting measurement or claims need controlled baselines.

The deal-risk register becomes an integration control book. Each protection should map to an owner, deadline, evidence and release decision. Finance should connect spending and outcomes to the valuation case. This allows the board to see whether risk was removed, transferred, accepted

or realised.

31. Govern model and provider changes after acquisition

The buyer should establish a change process covering model versions, prompts, retrieval, tools, datasets, thresholds, providers and infrastructure. Changes should state purpose, affected users, evaluation evidence, security assessment, cost, regulatory impact, approval, monitoring and rollback. The level of control should scale with materiality.

Provider notices should feed into the same process. Procurement, engineering, product, risk, legal and finance need a common record of changes that can affect customer outcomes and economics. The organisation should maintain tested fallbacks for dependencies whose failure exceeds risk appetite.

Governance should support speed through clear thresholds and reusable evidence. A slow committee that reviews every prompt change can create shadow practices. Defined delegation, standard tests and exception routes can preserve accountability while allowing responsible development.

32. Connect operating metrics to the investment case

The post-close dashboard should bridge technical, customer and financial measures. It can include performance by material cohort, availability, review rate, incident severity, provider concentration, unit cost, customer retention, gross profit, remediation spend and milestone completion. Each measure needs an owner, definition, source and decision threshold.

Leading indicators should connect to board action. A rise in fallback use may precede cost pressure. A decline in retrieval quality may precede support tickets. A supplier notice may trigger evaluation and customer communication. The dashboard should record the decision and outcome rather than displaying data without governance.

Finance should reconcile realised cash effects to the acquisition model. Benefits and losses should be attributed carefully because integration, market and algorithmic factors interact. The objective is disciplined learning and capital allocation, not artificial precision.

33. Apply accounting and impairment discipline

Acquired technology, customer relationships, goodwill and other assets require accounting under the applicable framework. IFRS 3 governs recognition and measurement in business combinations, while IAS 36 addresses impairment of relevant assets. Algorithmic dependency, rights restrictions, customer loss or weaker forecasts may inform cash-flow assumptions and impairment indicators, subject to accounting judgment.

The diligence model should preserve evidence used in forecasts and valuations. Post-close monitoring can identify changes in performance, cost, useful life or economic benefit. Management and auditors need transparent assumptions and reconciliation.

This paper does not determine accounting classification, fair value or impairment. Qualified accountants and valuation specialists should apply the standards to the transaction facts. The operating framework helps them locate evidence and cash-flow pathways.

34. Provide lenders with an algorithmic downside case

Lenders need visibility into recurring revenue quality, gross margin, concentration, capital needs and downside liquidity. An AI business may show strong growth while depending on a provider that can change price or terms. The financing model should therefore include provider, drift, customer and remediation scenarios.

Covenant headroom should be tested after realistic cash costs and delays. EBITDA adjustments for remediation or one-off integration require scrutiny because some technology and assurance spending may recur. Lenders may seek reporting, consent, liquidity or insurance protections depending on the risk and structure.

The borrower should avoid technical detail that obscures material economics. A concise dependency and mitigation schedule can support credit analysis. Sensitive information should be shared through appropriate controls and legal advice.

35. Use a worked transaction example

Consider a hypothetical buyer evaluating an AI workflow company with annual recurring revenue of USD 30 million and a stated contribution margin of 58 per cent. The product relies on one external model family for 72 per cent of production requests, a proprietary dataset licensed from three sources, and six senior staff who hold release or customer-specific knowledge. These figures are illustrative assumptions.

Reproduction shows that one high-value customer cohort performs below the aggregate claim. A provider price scenario reduces contribution margin by four percentage points. One data licence requires consent for assignment, and two release capabilities lack effective deputies. The buyer estimates a twelve-month remediation programme and models revenue, margin, investment and delay without treating any assumed probability as observed fact.

The buyer's working response includes a valuation range reflecting the moderate drift case, a closing condition for the material data consent, an interim covenant over provider and model changes, a measured escrow for rights cure, and retention plus knowledge-transfer milestones. Counsel, accountants, security specialists and relevant regulators would need to advise on the actual transaction.

36. Run the investment-committee decision gate

The investment committee should receive the signed-state baseline, dependency map, drift scenarios, rights waterfall, critical-person assessment, protection term sheet and funded integration plan. The paper should separate observed evidence, analytical assumptions, specialist opinions and unresolved questions. It should identify which exposures are accepted, reduced, transferred or capable of terminating the transaction.

Decision criteria can include valuation range, downside liquidity, closing certainty, legal operating perimeter, customer continuity, integration capability and strategic fit. The committee should know which assumptions require verification before signing, before closing and after control transfers.

Approval should record conditions, owners and evidence. A risk accepted implicitly tends to reappear as an integration surprise. A documented acceptance supports accountability and later

review.

The committee should also receive a dependency-adjusted value bridge. This bridge starts with the standalone valuation and shows the effect of customer concentration, provider exposure, rights remediation, critical-person cover, resilience investment and delayed synergies. Each adjustment should point to the relevant diligence evidence and state whether it changes expected cash flow, timing, liquidity or terminal durability. The bridge helps reviewers challenge overlaps and identify where a contractual protection supports the valuation conclusion.

A reverse stress test can identify the combination of drift, customer loss, cost increase and remediation delay that exhausts liquidity or breaches a financing threshold. The exercise is an analytical scenario rather than a forecast. It shows which early indicators matter most and which mitigations must be funded at closing. Where the downside depends on several correlated events, the committee should examine their common causes, including a single provider, dataset, person or regulatory classification.

37. Implement the first 100 days

Day One should secure authority, access, continuity and incident response. The first thirty days should preserve evidence, reconcile assets and rights, test monitoring, confirm customer commitments and begin knowledge transfer. Days thirty-one to sixty should complete priority remediation, reproduce material evaluations and test fallbacks. Days sixty-one to one hundred should scale only the changes that have passed customer, technical, security and economic gates.

The timetable should adapt to transaction risk. A regulated product, material data migration or provider replacement may require longer. Calendar speed should never substitute for acceptance evidence. Exceptions should name the temporary control, owner, deadline and residual exposure accepted by an accountable executive.

The transaction-protection register should be reviewed alongside integration milestones. Escrow releases, earn-out calculations, covenants and warranties may depend on operating evidence. Legal and finance teams should remain connected to the technical work.

38. Maintain an evidence room for the acquired system

The buyer should maintain a controlled evidence room containing architecture, model and data inventories, evaluations, licences, contracts, consents, approvals, incidents, monitoring, change records and customer commitments. Access, retention and privilege require careful design. The room should support operations and assurance without becoming an uncontrolled copy of sensitive data.

Evidence should be versioned and linked to decisions. A model card prepared after diligence should not overwrite the signed-state record. A consent should link to the dataset and technical control it authorises. A remediation test should identify the environment and acceptance authority.

Automation can help reconcile records and detect missing links. Human specialists remain responsible for interpretation and approval. The evidence room supports integration, audits, customer diligence, financing and future exit preparation.

39. Conclude with a disciplined allocation of uncertainty

AI transactions contain operating uncertainty that cannot always be removed before signing. A disciplined process makes the uncertainty visible, connects it to cash flow, assigns control and chooses a measurable response. The buyer gains a more defensible valuation and integration plan. The seller gains clearer requests and a route to demonstrate value.

The core sequence is consistent: reproduce the baseline, map dependencies, test drift, evidence rights, assess people and obligations, model cash effects, select transaction protections and preserve post-close tests. Each step should remain proportionate to materiality and supported by qualified specialists.

Boards should reject simple assurances that a model is proprietary, accurate or portable when the operating evidence is incomplete. They should also reject technical findings that lack a cash-flow pathway. The transaction succeeds when legal rights, operating control, customer outcomes and funded economics remain aligned after control changes.

The framework should remain proportionate. A small productivity tool and a regulated decision system do not require identical diligence. Materiality can be set through affected revenue, customer harm, regulatory exposure, replacement lead time, cash requirement and strategic importance. The team can deepen testing where several indicators converge and document why lower-risk areas received limited work. This risk-based approach protects decision quality while controlling diligence cost and management burden.

Future buyers should expect the evidence standard to rise as regulations, customer procurement and technical assurance mature. A target that maintains reproducible evaluations, traceable rights, controlled releases, tested fallbacks and transferable knowledge can reduce transaction friction. Those capabilities also improve ordinary operations. Deal readiness therefore begins before a sale process, through governance that makes the system understandable, changeable and accountable.

References

[1] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), 2023. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf

[2] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, 2024. https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

[3] National Institute of Standards and Technology, AI RMF Playbook. https://airc.nist.gov/airmf-resources/playbook/

[4] European Commission, Guidelines on obligations for General-Purpose AI providers, 2025. https://digital-strategy.ec.europa.eu/en/faqs/guidelines-obligations-general-purpose-ai-providers

[5] European Commission, Guidelines on the scope of obligations for providers of general-purpose AI models, 2025. https://digital-strategy.ec.europa.eu/en/library/guidelines-scope-obligations-providers-general-purpose-ai-models-under-ai-act

[6] European Commission, General-Purpose AI Models in the AI Act: Questions and Answers. https://digital-strategy.ec.europa.eu/en/faqs/general-purpose-ai-models-ai-act-questions-answers

[7] European Commission, AI Act regulatory framework. https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[8] European Commission, General-Purpose AI Code of Practice. https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai

[9] Central Bank of the UAE, Guidance Note on Consumer Protection and Responsible Adoption and Use of Artificial Intelligence and Machine Learning, 2026. https://rulebook.centralbank.ae/en/rulebook/guidance-note-consumer-protection-and-responsible-adoption-and-use-artificial-intelligence

[10] UK National Cyber Security Centre, Guidelines for Secure AI System Development, 2023. https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development

[11] UK Competition and Markets Authority, AI Foundation Models: Update Paper, 2024. https://www.gov.uk/government/publications/ai-foundation-models-update-paper

[12] UK Competition and Markets Authority, Merger Assessment Guidelines, 2021. https://www.gov.uk/government/publications/merger-assessment-guidelines

[13] US Department of Justice and Federal Trade Commission, Merger Guidelines, 2023. https://www.justice.gov/atr/2023-merger-guidelines

[14] UK Information Commissioner's Office, Guidance on AI and Data Protection. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/

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

[16] International Organization for Standardization, ISO/IEC 42001 Artificial Intelligence Management System. https://www.iso.org/standard/81230.html

[17] International Organization for Standardization, ISO/IEC 23894 Artificial Intelligence Risk Management. https://www.iso.org/standard/77304.html

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

[19] IFRS Foundation, IAS 36 Impairment of Assets. https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/

[20] UK Takeover Panel, The Takeover Code. https://www.thetakeoverpanel.org.uk/the-code/download-code

[21] European Union Agency for Cybersecurity, Multilayer Framework for Good Cybersecurity Practices for AI, 2023. https://www.enisa.europa.eu/publications/multilayer-framework-for-good-cybersecurity-practices-for-ai

[22] US Cybersecurity and Infrastructure Security Agency, Roadmap for Artificial Intelligence, 2023. https://www.cisa.gov/resources-tools/resources/roadmap-artificial-intelligence

[23] US Securities and Exchange Commission, Cybersecurity Risk Management, Strategy, Governance, and Incident Disclosure, 2023. https://www.sec.gov/files/rules/final/2023/33-11216.pdf

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

[25] Bank for International Settlements, Financial Stability Institute, Regulating AI in the Financial Sector: Recent Developments and Main Challenges, 2024. https://www.bis.org/fsi/fsipapers24.htm

About the Author

Authored by Chennakeshav Adya, Independent Researcher

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.