

GCC AI Company M&A: Valuing Data Rights, Model Dependency and Compute Obligations

A Transaction Framework for Asset Quality, Contractual Control, Operating Resilience, Valuation and Deal Protection

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Artificial-intelligence acquisitions can transfer a combination of software, data access, trained models, evaluation systems, customer relationships, compute capacity, specialist teams and contractual dependencies. These elements do not carry equal legal control, economic durability or transferability. A target can appear to own a differentiated AI platform while relying on customer data that cannot be reused, third-party models that can change price or terms, cloud commitments that exceed demand, open-source components with obligations, and key employees whose knowledge is only partly documented. The buyer therefore needs a transaction method that converts technical architecture and legal rights into cash-flow, valuation and deal-protection decisions.

This paper develops that method for acquisitions of AI companies with material operations, customers or regulatory exposure in the Gulf Cooperation Council. It uses five linked records: an AI asset map, a data-rights ledger, a model-dependency matrix, a compute-obligation curve and a valuation waterfall. The records support commercial, technology, financial, legal, privacy, cyber-security, accounting, tax, competition and integration diligence. They also provide a common language for management, advisers, investment committees and boards.

The framework starts with the customer outcome and traces every material AI capability through data origin, transformation, model, infrastructure, human review, output, revenue and contractual responsibility. It then tests control, lawful use, exclusivity, reproducibility, portability, change-of-control effects, cost, resilience and exit options. The valuation analysis separates current cash flow, evidenced growth, transferable intangible assets, required remediation, future investment and residual uncertainty. Deal terms are selected according to the risk that remains after diligence rather than through a generic technology warranty package.

All transaction examples, thresholds, financial values and scenarios in this paper are hypothetical modelling inputs. They do not describe an identified company, acquisition, legal conclusion, valuation opinion or investment recommendation. A live transaction requires current jurisdiction-specific legal, regulatory, accounting, tax, cyber-security, technical and commercial advice.

JEL Classification: G34, G32, G38, L22, L86, O32, O33

Keywords: artificial intelligence, mergers and acquisitions, GCC, data rights, model dependency, compute obligations, valuation, due diligence, deal protection, technology transactions

1. Start with the economic outcome, then trace the AI system


An acquisition thesis should state the customer problem, the target's role in solving it and the cash flows expected after completion. The label “AI company” provides limited transaction information. One target may own a narrow prediction model and sell licences. Another may orchestrate several third-party models inside a regulated workflow. A third may provide human services supported by automation. Each business has different assets, margins, liabilities and integration requirements.

The buyer should trace the operating chain from source data to customer outcome. The chain includes collection rights, data preparation, model selection or training, evaluation, deployment, inference, human review, output delivery, feedback and billing. Every link has an owner, a contract, a cost, a failure mode and an alternative. This trace prevents a persuasive product demonstration from becoming a substitute for transaction evidence.

The NIST AI Risk Management Framework organises AI risk through governance, mapping, measurement and management. Its core specifically addresses third-party data, software and supply-chain dependencies.[1] The NIST generative-AI profile adds content provenance, data lineage, testing and upstream dependency analysis.[2] These concepts are useful in M&A because the buyer must understand both the asset being acquired and the operating system required to preserve its value.

The transaction team should build one evidence room around the chain. Technical diagrams, contracts, privacy records, model cards, security reports, cost ledgers, customer cohorts and financial results should reconcile to the same products and legal entities. Unreconciled names, versions or boundaries are diligence findings. They can indicate weak configuration control, incomplete ownership, missing costs or a product claim that does not match the commercial contract.

Figure 1. AI asset map from source rights to realised customer value



The map links each operating layer to its principal transaction question and evidence owner.

Table 1. AI asset map and transaction evidence

Asset or capability	Evidence of control	Economic test	Principal transaction response
training and evaluation data	source, consent, licence, provenance, restrictions	incremental performance and reproducibility	value adjustment, remediation or specific protection
model weights and code	ownership, assignment, licence and repository history	differentiated outcome and replacement cost	asset transfer, escrow, licence or transition covenant
third-party model access	provider terms, price, service level and termination	margin, switching cost and continuity	consent, dual-provider plan or price protection
compute capacity	reservation, purchase, minimum spend and exit rights	utilisation, unit economics and downside exposure	debt-like adjustment, carve-out or renegotiation
customer workflow	contracts, integrations, user roles and acceptance evidence	retention, expansion, pricing and liability	revenue bridge, consent plan and warranty scope
specialist organisation	employment, incentives, documentation and decision rights	continuity, innovation and integration capacity	retention, governance and milestone plan

The evidence request should follow the product architecture and legal-entity perimeter.

2. Define the acquisition perimeter before valuing the platform

AI groups often develop through several legal entities, free-zone registrations, intellectual-property holding companies, offshore subsidiaries, research collaborations and reseller arrangements. The product may be marketed under one brand while customer contracts, employee agreements, cloud accounts and intellectual property sit elsewhere. A share acquisition can transfer hidden liabilities. An asset acquisition can omit licences, permits or contracts that cannot be assigned.

The perimeter schedule should identify every legal entity, branch, beneficial owner, licence, bank account, tax registration, employee, contractor, repository, domain, data store, cloud tenant, model account and customer agreement. It should also show which entity pays compute and which entity invoices each customer. Intercompany agreements need pricing, ownership and service terms that match operating practice.

Change of control may affect more than customer contracts. Model providers, cloud providers, data licensors, research institutions, government programmes, accelerators and strategic investors may hold consent, termination, information or most-favoured rights. A buyer should distinguish automatic transfer, consent-based transfer, prohibited assignment and contractual silence. Silence is not evidence of unrestricted transfer.

The perimeter also determines competition, foreign-investment, sector and data approvals. The UAE Ministry of Economy and Tourism applies Federal Decree-Law No. 36 of 2023 to economic concentration and publishes the notification process.[3] Cabinet Decision No. 3 of 2025 sets notification thresholds based on UAE relevant-market sales or market share.[4] A technology target can have modest global revenue while holding an important position in a narrow UAE market. The competition analysis should therefore define products, customers, substitutes,

geographic reach and data or compute inputs early enough to protect the timetable.

3. Build a data-rights ledger, not a dataset inventory

A dataset inventory states what information exists. A data-rights ledger explains why the target can collect, transform, combine, train on, evaluate with, retain, transfer and commercialise it. Those rights can differ by source, field, person, territory, purpose, model version, customer and time period. The ledger is the legal and economic foundation of an AI valuation.

For each material dataset, the buyer should record origin, data subject or owner, collection channel, contractual right, consent where relevant, stated purpose, permitted processing, permitted model use, derivative rights, output rights, confidentiality, retention, localisation, cross-border transfer, deletion, audit, exclusivity, sublicensing and change-of-control treatment. The record should point to executed documents and system evidence. A policy statement without contract or consent evidence is an incomplete right.

The UAE federal Personal Data Protection Law provides a governance framework for personal-data processing and addresses data-subject rights and cross-border transfers.[5] DIFC Data Protection Regulation 10 addresses personal-data processing through autonomous and semi-autonomous systems.[6] A target operating across mainland UAE, DIFC, ADGM and other markets may face several regimes. The buyer should map each processing activity to the relevant entity, establishment, customer and data flow rather than apply one group-wide conclusion.

Rights to use information for service delivery do not automatically create rights to train a general model, combine customer datasets or reuse outputs for another purpose. A customer may own inputs and outputs while granting the provider a limited processing right. Another customer may permit aggregated learning subject to anonymisation. The ledger should separate operational processing, product improvement, evaluation, research, marketing and resale.

The ledger also needs technical evidence. Data catalogues, access logs, lineage records, pipeline code, deletion workflows and model-training manifests should support the contractual position. If the company cannot identify which data trained a material model, it may be unable to comply with deletion, restriction or provenance obligations. That limitation can affect valuation even if no breach has been established because remediation, retraining and customer assurance have costs.

Figure 2. Illustrative data-rights transferability matrix

TECHNICAL ACCESS DOES NOT ESTABLISH TRANSFERABLE DATA RIGHTS

Owned and documented	Evidenced	Evidenced	Evidenced	Evidenced	Evidenced	Evidenced
Customer licensed	Evidenced	Conditional	Evidenced	Conditional	Restricted	Conditional
Public or open source	Evidenced	Conditional	Evidenced	Conditional	Conditional	Evidenced
Partner supplied	Conditional	Restricted	Conditional	Restricted	Restricted	Restricted
Personal or regulated	Conditional	Restricted	Conditional	Restricted	Restricted	Restricted
Derived and synthetic	Evidenced	Conditional	Evidenced	Conditional	Conditional	Evidenced
	Collect	Train	Evaluate	Commercialise	Transfer	Retain

The matrix distinguishes technical possession from legal and commercial control.

4. Test whether data creates measurable advantage

Lawful control does not establish economic value by itself. The buyer should test whether data improves model performance, customer outcomes, speed, cost, compliance or distribution. A large volume of common or low-quality data can have limited incremental value. A smaller longitudinal dataset tied to verified outcomes may create a stronger advantage.

The analysis should compare model performance with and without the target dataset, where technically and legally possible. It should examine coverage, representativeness, freshness, labelling accuracy, missingness, leakage, bias, duplication and drift. It should also measure the cost and elapsed time required to recreate the dataset through lawful alternatives. Reproduction cost is relevant evidence, but the valuation should remain anchored to expected economic benefit.

Exclusivity requires careful definition. A customer may grant the target access while retaining the same right for competitors. Public sources can be accessible to all, yet a proprietary cleaning, labelling or feedback process can create value. The buyer should identify the scarce element: access, permission, history, ontology, labels, workflow integration, feedback, trust or distribution.

WIPO describes training data, model weights, code, evaluation systems and related records as potential elements of an AI intellectual-property and trade-secret strategy, subject to the rights and protections available in each jurisdiction.[7][8] Transaction diligence should test whether the target treated claimed secrets as confidential through contracts, access control, encryption, repository permissions and employee practices. Value attributed to secrecy needs evidence that secrecy was maintained.

Data value can deteriorate after completion. Customer churn, contract renegotiation, deletion requests, platform changes or regulatory restrictions can reduce future access. The forecast should include the operating cost of governance, quality, storage, security, licences and human review. A gross data-asset value without these cash outflows can overstate the acquisition benefit.

5. Separate owned models from orchestrated capability

An AI product can combine proprietary code, fine-tuned open models, third-party APIs, retrieval systems, rules, human review and workflow integrations. The customer buys an outcome rather than a model taxonomy. The buyer still needs to know which components transfer and which remain dependent on external providers.

The model register should identify model name, version, purpose, owner, licence, training method, data lineage, evaluation results, deployment environment, inference cost, hardware requirement, safety controls, known limitations and replacement path. It should reconcile with source repositories, cloud accounts, invoices, architecture diagrams and customer documentation.

Proprietary weights can be valuable where they deliver a differentiated and reproducible outcome. Value may be limited if results depend on undisclosed third-party components, manual corrections or benchmarks that do not represent live performance. The buyer should reproduce material evaluations in a controlled environment and trace exceptions through production logs. Customer-specific acceptance evidence is more informative than a headline benchmark alone.

Open-source components require version and licence review. The diligence team should identify attribution, disclosure, distribution, patent, network-use and modification obligations. A model's licence can differ from the licence for its code, weights, training data or hosted service. The target's commercial use should match each applicable term.

Third-party APIs can accelerate product development and reduce capital requirements. They also introduce price, availability, model-change, data-use, geography and termination risk. The buyer should assess whether the target can switch provider without changing customer outcomes, regulatory approvals or gross margin. A theoretical alternative is insufficient if migration requires months of engineering and revalidation.

Table 2. Model-dependency diligence and value treatment

Dependency	Evidence required	Valuation question	Deal response
proprietary trained model	assignments, lineage, evaluation and reproducibility	does it create durable outcome or cost advantage?	recognise evidenced benefit; fund retention and maintenance
open model	exact licence, version, modifications and deployment	can obligations restrict commercialisation or transfer?	remediation, disclosure, replacement or indemnity
hosted third-party model	contract, price, SLA, data use and termination	how sensitive are margin and continuity to provider action?	consent, transition right, dual sourcing or price protection
customer-specific model	customer rights, training scope and exit obligations	does value remain after customer loss or change of control?	customer-level value, consent and retention adjustment
human-assisted model	workflow, review rate, labour cost and quality	is reported automation economically real and scalable?	normalise margin and include operating investment
experimental model	test record, approved use and production controls	is forecast adoption evidenced?	scenario value or milestone consideration

The classification should be completed for every model that affects a material product, customer obligation or cost line.

6. Convert model dependency into a quantified transaction risk

Dependency becomes financially material through cost, interruption, performance, compliance, customer consent and strategic control. The buyer should calculate exposure for each material model provider under ordinary use, volume growth, price increase, forced migration, service failure and termination. The output belongs in the integrated financial model rather than a standalone technology report.

The US Federal Trade Commission's study of major cloud and generative-AI partnerships describes equity rights, revenue-sharing, consultation rights, cloud-spend commitments, switching costs and access to sensitive technical and business information.[9] The UK Competition and Markets Authority has also examined access to compute, data, models and distribution across the foundation-model value chain.[10] These sources show why a target's provider relationship can affect control and competition as well as cost.

A dependency score should avoid false precision. It can combine concentration, contractual security, technical portability, validation effort, data portability, alternative capacity and time to migrate. The transaction team should retain the underlying evidence and show which factor changes the valuation or deal term. A red score with no decision is only decoration.

The buyer should test the provider's right to use prompts, inputs, outputs, metadata and feedback. Enterprise terms may differ from public terms, and historic use may predate the current contract. The target should demonstrate how provider terms were approved, monitored and reflected in customer disclosures. Any mismatch should be assessed by customer, data type and period.

Model depreciation is another cash-flow issue. Capabilities and prices can change quickly. A proprietary model that required substantial historic development cost may be replaceable by a lower-cost external model. This can reduce standalone asset value while improving future product economics. The buyer should value the best post-close operating architecture, including transition cost, rather than preserve sunk cost in the forecast.

7. Treat compute obligations as a separate deal workstream

Compute can be purchased on demand, reserved, prepaid, committed through minimum spend, financed, hosted in dedicated capacity or bundled with strategic investment. Each structure has different cash, accounting, tax, control and exit effects. The diligence team should obtain complete contracts, amendments, invoices, usage records, credits, rebates, side letters and correspondence for cloud, colocation, accelerators, network and storage.


The contract schedule should show provider, service, region, term, capacity, minimum spend, unit price, escalation, currency, credit, prepayment, renewal, termination, assignment, change of control, data egress, portability, security, service level and remedy. It should also identify whether a related investor, lender or commercial partner influences the commitment. A provider-linked investment can contain economics that do not appear in one agreement.

The buyer should compare contracted capacity with forecast utilisation by model, product, customer and geography. Unused capacity can represent growth optionality, but it can also create unavoidable cash burn. High utilisation can signal efficient demand or insufficient resilience. The forecast should include peak requirements, training cycles, inference variability, redundancy and committed headroom.

The CMA's cloud-services investigation assessed technical barriers, egress fees, committed-spend agreements, interoperability and switching.[11] These factors matter in an acquisition even when the target operates outside the UK because global provider terms and architecture can influence its GCC business. The transaction model should estimate the time, direct cost and customer impact of migration rather than assume immediate portability.

Compute commitments may require a debt-like or working-capital adjustment in the equity bridge where payment is unavoidable and the corresponding capacity does not generate forecast value. Prepaid capacity may be an asset only to the extent it is transferable, usable and economically beneficial. A loss-making contract may need a provision or purchase-price adjustment subject to the applicable accounting and transaction terms.

Figure 3. Illustrative compute-obligation curve



Contracted spend is compared with base and downside economic usage to identify surplus commitment and capacity pressure.

8. Reconcile compute unit economics to revenue and gross margin

AI gross margin should be rebuilt from usage and invoices. The model should connect customer volume to tokens, calls, accelerator time, storage, network, human review, monitoring, support and credits. Contracted revenue and compute cost may use different currencies, billing periods and volume tiers. A monthly bridge makes these effects visible.

The buyer should distinguish training, evaluation and inference. Training may be episodic and capital-intensive. Evaluation is a recurring control cost. Inference grows with customer use and can vary according to context length, modality, latency and architecture. Shared platform costs need a rational allocation to products and customers. A target can report strong blended margin while a rapidly growing AI feature loses money.

Customer pricing should be tested against consumption. Per-seat pricing can create risk if usage expands faster than seats. Usage pricing can transfer cost but may discourage adoption. Outcome pricing can support premium value while increasing measurement and liability requirements. The acquisition case should state how pricing can change and which contracts permit adjustment.

Optimisation initiatives need evidence. Smaller models, caching, quantisation, batching, retrieval design, routing and hardware choices can reduce cost. The buyer should review production experiments, quality trade-offs and implementation dates. Forecast savings without tested performance should remain a scenario rather than base-case value.

Compute geography can affect latency, data transfer, customer assurance and regulatory compliance. The target should show where data is stored and processed, which subprocessors participate and how failover works. A migration to lower-cost capacity may require customer notice, consent, revalidation or local infrastructure. Those requirements belong in the integration budget.

9. Test customer revenue for transfer, retention and AI liability

Customer diligence should connect contracts, invoices, collections, usage, outcomes, support, renewals and claims. The buyer should identify which customers purchase a defined software service, a managed outcome, professional services or access to experimental capability. Revenue quality and liability differ across these models.

The contract review should cover change of control, assignment, data processing, model use, output rights, service levels, credits, warranties, indemnities, limitation of liability, audit, security, regulatory cooperation, subcontracting and termination. AI-specific schedules may describe human oversight, prohibited use, performance limitations and incident response. Sales presentations should be compared with signed obligations.

The revenue bridge should separate recurring contract value, usage, implementation, services and one-off work. Customer concentration needs group and end-user views where a reseller or government contractor intermediates the relationship. The buyer should also identify revenue supported by founder involvement or unpriced customisation.

Retention analysis should consider whether the acquired platform remains valuable after changes to model, cloud, brand and ownership. Customers may accept the target because of local relationships, data location, domain expertise or independence from a larger group. Buyer ownership can strengthen distribution or create perceived conflict. Customer interviews should test both effects without disclosing the transaction beyond the approved process.

Claims and near misses are part of valuation. Incorrect outputs, bias, privacy incidents, intellectual-property allegations, missed service levels or regulatory inquiries can create remediation cost and reputational loss. The target should provide incident registers, complaints, credits, reserves, insurance notifications and board reports. An empty register is credible only when the reporting process and source systems support it.

Table 3. Customer cash-flow bridge for an AI acquisition

Bridge item	Evidence	Valuation treatment	Protection option
contracted recurring revenue	executed terms, invoice and collection	starting revenue base	completeness warranty
change-of-control exposure	consent and termination clauses	probability-weighted retention	condition, covenant or holdback
unpriced compute growth	product usage and unit cost	margin normalisation	pricing plan and integration funding
services disguised as software	time records and delivery process	lower scalable margin	forecast reset and earn-out metric design
customer-specific data restriction	data schedule and processing record	limits model reuse and synergy	consent or ring-fenced architecture
output or performance liability	claims, credits and contractual caps	expected cost and downside scenario	specific indemnity or escrow

Each adjustment should reconcile to customer-level evidence and the transaction forecast.

10. Examine cyber-security and operational resilience as value preservation

An AI system expands the attack surface through data pipelines, model endpoints, prompt interfaces, plugins, third-party components and high-value intellectual property. Diligence should cover identity, privileged access, development security, model and data protection, vulnerability management, logging, monitoring, incident response, backups, recovery and supplier risk. The review should reproduce evidence rather than rely on policy documents alone.

The architecture map should identify production and non-production environments, secrets, service accounts, model registries and data stores. Access should be reconciled to active staff and contractors. Repository and cloud histories can reveal credentials, unauthorised data, shadow environments or dependence on personal accounts. These issues can block transfer even when no external incident is known.

The buyer should test recovery for critical products. Backup existence does not demonstrate recovery. A timed restoration, model redeployment and customer-communication exercise provides stronger evidence. Dependencies on one employee, provider region, key or undocumented process should be recorded in the integration plan.

AI-specific testing can include prompt injection, data extraction, model abuse, adversarial input, output filtering, rate limits and monitoring. NIST's AI guidance links trustworthiness to security, resilience, transparency, accountability and lifecycle management.[1][2] The diligence report should translate each material weakness into probable cost, customer effect, regulatory exposure and time to remediate.

Insurance policies require scope review. Cyber, technology errors and omissions, professional indemnity, directors' liability and crime cover can contain AI, contractual-liability, intellectual-property, territorial or prior-knowledge limitations. Coverage does not replace risk analysis. It can affect the expected cash cost of identified scenarios.

11. Assess people, knowledge and governance

AI value can depend on a small team that understands data, models, infrastructure, customers and regulatory context. The buyer should map critical roles, decision rights, replacement time, compensation, equity, notice, restrictive covenants, immigration status, inventions, external appointments and retention risk. The map should include technical and commercial knowledge, not only senior titles.

Employee and contractor intellectual-property assignments should cover the jurisdictions, entities and development periods involved. The target should explain contributions from founders, universities, incubators, customers, open-source communities and previous employers. Repository history can help reconcile authorship, but legal review determines rights.

Knowledge transfer should be measurable. Architecture, data definitions, evaluation methods, deployment, incident response, customer configuration and vendor management need current documentation. A model that only one person can retrain or validate carries a continuity cost. The buyer can require documentation and shadow operation before completion or during a funded transition.

Governance should show who approves models, data, product claims, providers, security exceptions and customer use cases. Board minutes, risk committees, release records and incident decisions provide evidence. The UAE's published AI ethics guidance emphasises fairness, accountability, transparency, explainability, robustness, safety, human-centred values, sustainability and privacy.[12] The transaction question is whether the target operationalises relevant principles through roles, controls and evidence.

Retention consideration should reward continuity and value creation without obscuring purchase price or transferring ordinary employment cost into an earn-out. Objectives can include product stability, customer retention, migration, documentation and approved development milestones. Metrics should remain within the participants' reasonable influence and avoid incentives to compromise safety or revenue quality.

12. Convert product evidence into a durable revenue forecast

The forecast should begin with customer cohorts and product economics. For each material customer or cohort, the buyer should model contracted revenue, renewal date, usage, price, compute, support, implementation, service credits and collection. It should then apply product, provider, data-rights and change-of-control risks identified in diligence.


New-business forecasts need pipeline evidence, sales cycle, conversion, implementation capacity and cash collection. A strategic buyer may contribute distribution, but synergy should be separated from standalone value and assigned an integration cost. The seller should not receive full value for buyer-created synergy unless the negotiated structure deliberately shares it.

The model should include data and model maintenance. Drift monitoring, evaluation, red teaming, data acquisition, licences, human review, compliance and security are recurring costs. Capitalising development expenditure for accounting purposes does not remove the cash requirement. The forecast should show the investment needed to preserve each claimed capability.

Downside cases should combine risks that can occur together. A model-provider price increase may coincide with customer pricing pressure. A data-rights remediation can reduce training volume and delay releases. A security incident can trigger customer review and higher insurance cost. Correlated stress is more informative than independent sensitivity tables.

The output should identify the earliest cash constraint, lowest margin, renewal concentration, compute surplus, remediation spend and valuation effect. Management actions need timing, cost and authority. A forecast that assumes immediate price increases, cost cuts and provider migration without contractual or operational support is not an executable case.

Figure 4. Illustrative valuation bridge from reported enterprise value to evidence-backed value



Adjustments represent hypothetical transaction analysis rather than observed market values.

13. Value identifiable assets and residual capability carefully

The income approach can value customer relationships, technology, trade names and other identifiable assets where forecast cash flows and useful lives can be supported. Cost and market evidence can provide cross-checks. The transaction valuation should distinguish asset value, operating-business value and buyer-specific synergy.

IFRS 3 requires an acquirer to recognise identifiable assets acquired and liabilities assumed in a business combination, subject to the standard's requirements.[13] IAS 38 describes an identifiable intangible asset through separability or contractual or legal rights and addresses recognition of acquired intangible assets.[14] The purchase-price allocation is an accounting exercise after the transaction price is determined, yet its concepts help the deal team test what is actually being acquired.

A data-related intangible may derive value from contractual access, curated structure, labels or customer relationships. The underlying facts may be unavailable for ownership. A model-related asset may include code, weights, documentation and evaluation, while a continuing team and future development sit elsewhere in the business value. The valuation analysis should avoid assigning the same cash flow to several assets.

Useful life should reflect legal rights, customer renewal, technology change, data refresh, provider dependence and replacement. Historic development cost is weak evidence when a newer model can replicate performance cheaply. Conversely, a deeply integrated system with long validation and regulatory acceptance may have a longer economic life than its code version suggests.

Goodwill can reflect the assembled organisation, expected synergy and future capabilities that do not qualify as separate assets. The investment committee should still understand the operational assumptions embedded in residual value. A large goodwill balance can be commercially rational, but it should not conceal dependence on uncommitted people, customer consents or future buyer investment.

14. Structure the valuation around scenarios and option value

AI markets contain technology, price and adoption uncertainty. A single discounted cash-flow case can conceal which assumptions create value. The buyer should construct a small set of internally consistent scenarios: durable proprietary advantage, successful orchestration, commoditised model layer, constrained data rights, compute-cost pressure and strategic acceleration under buyer ownership.

Each scenario should specify customer retention, new sales, price, usage, gross margin, development cost, provider cost, governance investment, working capital and terminal economics. Probability weights require governance and evidence. They should be approved by the decision-makers who own the acquisition case rather than introduced as a mechanical averaging exercise.

Option value can arise from exclusive data access, adjacent customer distribution, a validated platform or scarce talent. The buyer should state the decision, timing, required investment and probability that converts the option into cash flow. A broad statement that the target opens “multiple AI opportunities” is not a valuation model.

Earn-outs and contingent value rights can bridge uncertainty where measurement is reliable. Revenue milestones may reward low-margin or non-transferable sales. EBITDA can be affected by buyer allocations and investment choices. Product milestones need objective acceptance and governance. Customer retention, gross profit, approved deployment and documented transfer can be combined where they reflect the thesis.

The buyer should retain control of safety, compliance and capital allocation. Contingent consideration should not create pressure to deploy an unvalidated model, overuse restricted data or defer necessary remediation. The transaction documents can expressly preserve these decision rights while defining fair measurement and dispute resolution.

15. Link diligence findings to the equity bridge

The bridge from enterprise value to equity value should consider cash, debt, debt-like items, working capital, tax, transaction leakage and specific obligations. AI transactions can add prepaid compute, minimum-spend shortfalls, customer credits, deferred implementation, model licences, data-remediation costs, retention payments and decommissioning obligations.

Classification should follow substance and the agreed purchase-price mechanism. An unavoidable pre-close compute commitment with no forecast benefit may be debt-like. Ordinary recurring cloud usage may remain in working capital or the forecast. A prepaid amount may be an asset if transferable and economically usable. The parties should avoid resolving accounting questions through labels alone.

Deferred revenue requires a fulfilment-cost analysis. Cash collected before completion may fund post-close compute, support, implementation and liability. The working-capital target should reflect seasonality, billing patterns and the cost to serve. Customer credits or refunds arising from pre-close performance need specific treatment.

Remediation can be reflected through price, specific indemnity, escrow, retention, condition precedent, seller action or buyer-funded integration. The selected mechanism should match measurability, timing, control and claim risk. A known cost is usually easier to price directly than to convert into a disputed warranty claim.

Table 4. Finding-to-value and finding-to-protection matrix

Finding	Cash-flow effect	Equity-bridge question	Deal or integration response
restricted training rights	retraining cost, delay or lower performance	known liability or forecast adjustment?	consent, retraining plan, price or indemnity
provider minimum spend	surplus cash outflow	debt-like, working capital or operating item?	renegotiate, assign, carve out or adjust price
non-transferable customer	revenue and margin at risk	included in enterprise value?	consent condition, holdback or earn-out
undocumented model process	retention and rebuilding cost	ordinary investment or value gap?	transition covenant and retention plan
open-source licence gap	remediation, disclosure or replacement	quantified cost or contingent exposure?	pre-close cure and specific protection
unpriced inference growth	future gross-margin compression	forecast normalisation	customer repricing and architecture plan

The same finding may require several coordinated responses.

16. Design representations around the actual AI stack

Representations should cover title, authority, contracts, intellectual property, data protection, cyber-security, compliance, customers, employees, accounts and litigation. AI-specific drafting can extend these areas through schedules that identify models, datasets, licences, providers, incidents, evaluations and product claims. Generic language should not replace a complete disclosure process.

The seller can represent that scheduled proprietary code and models are owned or validly licensed, subject to disclosed exceptions. Data representations should address collection and processing rights, notices, consents, instructions, transfers, retention, deletion, security and complaints according to applicable law and contract. The drafting should reflect the target's actual role as controller, processor, provider or customer.

Provider representations can cover completeness of material contracts, defaults, disputes, notices, commitments and change-of-control rights. Compute schedules should disclose minimum spend,

prepayment, credits, capacity, liens, financing and related arrangements. Customer schedules should identify AI-specific obligations, performance commitments, credits and restrictions on data or model use.

Knowledge qualifications, materiality, time periods and disclosure standards influence claim value. The buyer should prioritise risks that can be stated and evidenced. A representation that an AI system is “accurate” has little meaning without defined use, metric, dataset, threshold and period. The transaction document should avoid technical promises broader than the customer product itself.

Specific indemnities are suitable for identified contingent exposures where responsibility, scope, duration and loss can be defined. Insurance may support general warranty risk, subject to exclusions and underwriting. Known privacy, intellectual-property, cyber or regulatory matters may require separate treatment. Counsel should align the protection package with governing law and enforcement practicalities.

17. Use conditions and covenants to preserve value before completion

Signing and completion can be separated by regulatory review, customer consent, financing or reorganisation. AI assets can change quickly during this period. Interim covenants should preserve ordinary development while restricting actions that materially alter the acquired risk, such as changing critical models, granting data rights, entering large compute commitments, modifying key customer terms or suffering unmanaged personnel departures.

The target should continue security, backup, evaluation, compliance and incident reporting. The buyer may require notice of provider changes, material model releases, data incidents, customer claims and regulatory contacts. Access rights must respect competition law, confidentiality, personal data and clean-team controls.

Conditions should be objective and proportionate. Material customer or provider consent, transfer of intellectual property, release of liens, reorganisation of cloud accounts, completion of data remediation and regulatory clearance can be closing conditions where the transaction depends on them. Operational improvements that can continue after completion may be covenants supported by holdback or milestone consideration.

The long-stop date should reflect competition review, sector approvals and consent timing. UAE economic-concentration review can include a statutory decision period after a complete filing under the applicable framework.[3][4] The transaction timetable should reserve time for information requests and should not assume that a filing is unnecessary before the relevant-market analysis is completed.

18. Plan integration before agreeing the price

Integration determines whether acquired rights and capabilities produce the forecast cash flow. The buyer should create a day-one, first-100-day and first-year plan across customers, people, data, models, compute, security, finance, compliance, governance and brand. Each workstream needs an accountable owner, dependency, budget, metric and decision date.


Day one should preserve service, access, billing, payroll, security and customer communication. The buyer should avoid immediate technical consolidation without understanding data location,

customer commitments and validation. Critical accounts, keys, domains, repositories and provider contacts should transfer through an approved runbook.

The first 100 days should close priority rights gaps, stabilise key staff, validate the architecture, reconcile model costs, confirm customer retention and establish group governance. Provider migration or model replacement should proceed through testing and customer obligations. Integration savings should be measured net of duplicated operation, revalidation and transition support.

The first year should deliver the acquisition thesis. Product and distribution milestones should connect to gross profit and cash, while data and model governance remain auditable. The board should receive a value-creation dashboard that distinguishes acquired performance, buyer synergy, remediation and new investment.

The integration plan also needs an exit architecture. If a provider fails, a model becomes restricted or the group later divests the business, can data, code, models, customer configurations and operations be separated? Designing portability at acquisition protects both resilience and future strategic options.



Greater residual uncertainty shifts protection from ordinary covenants toward conditional or contingent structures.

Table 5. Integration value-control dashboard

Control area	Day-one evidence	100-day measure	First-year value measure
customers	service continuity and communication	consent and retention closure	cohort gross profit and expansion
data rights	access and processing continuity	priority gaps remediated	governed reuse and measurable product benefit
models	production stability and ownership access	evaluations reproduced and dependencies tested	outcome, cost and portability improvement
compute	accounts, capacity and billing controlled	commitments reconciled and optimisation tested	utilisation, unit cost and resilience
people	critical-role continuity	knowledge transfer and incentives active	retained capability and succession depth
governance	incident and approval routes active	group controls implemented	audit-ready evidence and reduced residual risk

The dashboard should reconcile to the approved acquisition model and risk register.

19. Establish an investment-committee decision record

The investment paper should state the acquisition thesis, perimeter, consideration, funding, approvals, customer economics, data rights, model dependencies, compute obligations, people risk, valuation, downside, deal protection and integration plan. Each material conclusion should point to evidence and an accountable owner.

The committee should see the five linked records. The asset map explains the operating chain. The data-rights ledger establishes lawful and transferable control. The model-dependency matrix identifies replaceability and concentration. The compute curve shows committed capacity against economic demand. The valuation waterfall connects findings to price and synergy.

The paper should separate confirmed facts, open items and management assumptions. Open items need conditions, owners and dates. Assumptions need scenarios and monitoring. Accepted risks need a reason, value at risk and response. This discipline allows the board to determine whether price, protection and execution capacity are aligned.

The decision should also identify the point at which the acquisition thesis fails. Examples include loss of a specified customer cohort, inability to transfer material data rights, a compute shortfall that cannot be renegotiated, departure of critical staff or required regulatory remedies. A pre-agreed failure boundary improves negotiation and prevents sunk process cost from replacing investment judgement.

Post-close reporting should use the same definitions approved at signing. Reclassifying customers, compute or development spend can make the acquisition appear successful without creating value. Finance, product, technology and legal owners should reconcile the dashboard to the original model and explain every material variance.

20. Make transaction evidence the basis of AI value

AI M&A value is created where rights, technology, operations and customers remain connected after completion. Data possession without transferable rights can create remediation. Proprietary claims without reproducible performance can create overvaluation. Growth without compute economics can consume cash. Strong technology without customer transfer and retained people can fail during integration.

The five-record method turns these risks into transaction decisions. It enables a buyer to identify which capability is owned, which is licensed, which is contingent and which must be rebuilt. It also connects diligence to forecast cash flow, identifiable assets, consideration, covenants, indemnities, transition work and post-close investment.

The framework supports both strategic and financial buyers. A strategic buyer may have models, compute or distribution that change the target's best operating architecture. A financial buyer may preserve independence and fund provider diversification. Each buyer should value the target under its executable ownership case and retain a standalone reference case to distinguish acquired value from synergy.

A disciplined acquisition does not require certainty about every future AI capability. It requires evidence about current rights and performance, an explicit view of dependencies, funded operating choices and protection for residual risk. These elements give management and boards a defensible basis for price and a practical route from signing to realised value.

The transaction file should remain usable after completion. Executed contracts, consent evidence, data lineage, model versions, evaluation results, compute schedules, customer bridges and integration decisions should be dated and preserved. This audit trail supports regulatory engagement, purchase-price allocation, warranty claims, future financing and eventual exit diligence. It also allows the combined company to demonstrate that its AI assets are governed as operating capabilities rather than marketing assertions.

References

- [1] National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework (AI RMF 1.0),” January 2023.
<https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10>
- [2] National Institute of Standards and Technology, “Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile,” July 2024. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [3] UAE Ministry of Economy and Tourism, “Economic Concentration,” process and published decisions under the UAE competition framework. <https://www.moet.gov.ae/en/web/guest/economic-concentration>
- [4] UAE Cabinet, “Cabinet Decision No. 3 of 2025 on the Ratios Related to the Implementation of Federal Decree-Law No. 36 of 2023 Regulating Competition.” <https://www.moet.gov.ae/documents/20121/0/Cabinet%2BDecision%2BNo.%2B%283%29%2Bof%2B2025%2BOn%2Bthe%2BRatios%2BRelated%2Bto%2Bthe%2BImplementation%2Bof%2BFederal%2BDecree-Law%2BNo.%2B%2836%29%2Bof%2B2023%2BRegulating%2BC.pdf/c8e816c6-4670-a5cd-906f-b5f0837011e9>
- [5] UAE Government, “Data protection laws,” including Federal Decree-Law No. 45 of 2021 Regarding the Protection of Personal Data. <https://u.ae/en/about-the-uae/digital-uae/data/data-protection-laws>
- [6] Dubai International Financial Centre, “Regulation 10: Processing Personal Data through Autonomous and Semi-autonomous Systems.”
<https://www.difc.com/business/registrars-and-commissioners/commissioner-of-data-protection/regulation-10>

- [7] World Intellectual Property Organization, “Learning Machines: An introduction to AI and IP for small and medium-sized enterprises,” August 2026. <https://www.wipo.int/web-publications/learning-machines-an-introduction-to-ai-and-ip-for-small-and-medium-sized-enterprises/en/index.html>
- [8] World Intellectual Property Organization, “WIPO Guide to Trade Secrets and Innovation: Trade secrets and digital objects.” <https://www.wipo.int/web-publications/wipo-guide-to-trade-secrets-and-innovation/en/part-vii-trade-secrets-and-digital-objects.html>
- [9] US Federal Trade Commission, “FTC Staff Report on AI Partnerships and Investments 6(b) Study,” January 2025. <https://www.ftc.gov/reports/ftc-staff-report-ai-partnerships-investments-6b-study>
- [10] UK Competition and Markets Authority, “AI Foundation Models: Update paper,” April 2024. <https://www.gov.uk/government/publications/ai-foundation-models-update-paper>
- [11] UK Competition and Markets Authority, “Cloud services market investigation: final decision,” July 2025. <https://www.gov.uk/cma-cases/cloud-services-market-investigation>
- [12] UAE Minister of State for Artificial Intelligence, Digital Economy and Remote Work Applications Office, “AI Ethics Principles and Guidelines.” <https://ai.gov.ae/wp-content/uploads/2023/03/MOCAI-AI-Ethics-EN-1.pdf>
- [13] IFRS Foundation, “IFRS 3 Business Combinations.” <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
- [14] IFRS Foundation, “IAS 38 Intangible Assets.” <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
- [15] IFRS Foundation, “IFRS 13 Fair Value Measurement.” <https://www.ifrs.org/issued-standards/list-of-standards/ifrs-13-fair-value-measurement/>
- [16] World Intellectual Property Organization, “Artificial Intelligence and Intellectual Property.” <https://www.wipo.int/en/web/frontier-technologies/artificial-intelligence/index>
- [17] Dubai International Financial Centre, “Data Protection Guidance for Compliance.” <https://www.difc.com/business/registrars-and-commissioners/commissioner-of-data-protection/guidance>
- [18] UAE Government, “The UAE's national strategy for artificial intelligence 2031.” <https://u.ae/en/about-the-uae/strategies-initiatives-and-awards/strategies-plans-and-visions/government-services-and-digital-transformation/uae-strategy-for-artificial-intelligence>
- [19] Bank for International Settlements, “The geography of AI firms,” BIS publication No. 1343, April 2026. <https://www.bis.org/publ/work1343.htm>
- [20] UAE Ministry of Economy and Tourism, “Federal Decree-Law No. 36 of 2023 Regulating Competition.” <https://www.moet.gov.ae/documents/20121/0/Federal%2BDecree-Law%2BNo.%2B%2836%29%2Bof%2B2023%2BRegulating%2BCompetition.pdf/d5f2b12d-7688-5180-e63b-44a8dbe878d6>

About the Author

Chennakeshav (CK) Adya 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.