Diligencing an AI Target: Revenue Quality, Data Rights, Models and Compute Obligations

A Board Framework for Converting AI Technical Evidence into Price, Protection and Integration Decisions

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

An artificial-intelligence target can report attractive growth while carrying obligations that do not appear in conventional software diligence. Revenue may combine subscriptions, implementation, usage, pass-through model fees and non-recurring pilots. Customer economics may depend on promotional cloud credits or manual expert work. Data used for training, fine-tuning, retrieval and evaluation may have different consent, privacy, confidentiality, copyright and portability conditions. The product may rely on a foundation model that the target does not own, a model licence that restricts use, or an application architecture that cannot move without material re-engineering. Reserved compute, minimum cloud spend, accelerator availability, inference latency and vendor concentration can change both gross margin and the ability to serve contracted demand.

This paper develops a board framework for diligencing an AI target and converting technical evidence into transaction decisions. It separates reported revenue from contracted and economically recurring revenue; builds a customer and cohort evidence ledger; maps data from collection to deletion; establishes model provenance and reproducibility; measures the unit economics of training and inference; identifies cloud, chip and capacity obligations; tests privacy, copyright, cyber and AI-governance exposure; and links exceptions to valuation, purchase-price mechanics, warranties, indemnities, escrow, closing conditions and the integration plan.

The framework draws on current primary and authoritative materials from the United Arab Emirates, European Union, United Kingdom, United States, Singapore, India and Australia. IFRS 15 provides the core revenue-recognition discipline. The European Union Artificial Intelligence Act and General-Purpose AI Code of Practice increase the importance of technical documentation, training-content transparency, copyright policy and systemic-risk controls. NIST materials provide a practical structure for risk management, provenance, testing and secure development. The United States Federal Trade Commission's study of large AI partnerships illustrates why cloud commitments, information access and switching constraints deserve transaction-level attention. Privacy authorities in the UAE, UK and Australia reinforce the need to establish a lawful, purpose-specific basis for personal-data use across the AI lifecycle.

All scores, percentages, prices, customer metrics, model costs, capacity levels, time periods and transaction outcomes in this paper are hypothetical management assumptions used to demonstrate the framework. They do not describe a client, current quotation, forecast, legal conclusion, transaction recommendation or assured outcome. This paper does not provide legal, tax, accounting, regulatory, cybersecurity, technical or investment advice. Qualified advisers should assess the target, transaction documents, models, data, infrastructure and applicable jurisdictions.

JEL Classification: G24, G32, G34, L22, L86, M15, O32

Keywords: artificial intelligence, AI due diligence, revenue quality, data rights, model provenance, cloud commitments, compute obligations, intellectual property, software M&A, transaction risk

1. Treat the AI target as five linked evidence systems

An AI acquisition should begin with a single evidence model that connects revenue, customers, data, models and compute. Conventional workstreams often review these domains separately. Finance tests revenue, legal reviews contracts and intellectual property, technology evaluates architecture, security reviews controls and commercial diligence assesses the market. The separation creates gaps when one fact has consequences across several workstreams.

A customer contract can promise a service level that depends on a third-party foundation model. That model may have a usage policy, regional availability constraint or price schedule that the target cannot control. The service may process personal or confidential data that the customer agreement does not clearly authorise for improvement or training. The reported gross margin may exclude reserved capacity, expert review or implementation effort. The same issue therefore affects revenue quality, liability, gross margin, product continuity and valuation.

The diligence system should assign each material claim an evidence owner, source, date, reconciliation status and transaction consequence. A board should be able to move from a headline such as recurring revenue to the supporting customer contracts, invoices, usage logs, collections, delivery costs, model calls, cloud charges and renewal evidence. It should also be able to move in the opposite direction: from a data-right exception or compute commitment to the customers, products, cash flows and valuation components affected.

The purpose is decision conversion. An exception can be remediated before signing, made a closing condition, priced through the valuation, allocated through an indemnity, protected with escrow, governed by a covenant or accepted within the integration plan. An exception with no documented consequence remains an observation rather than completed diligence.

Figure 1. The linked evidence architecture for an AI acquisition

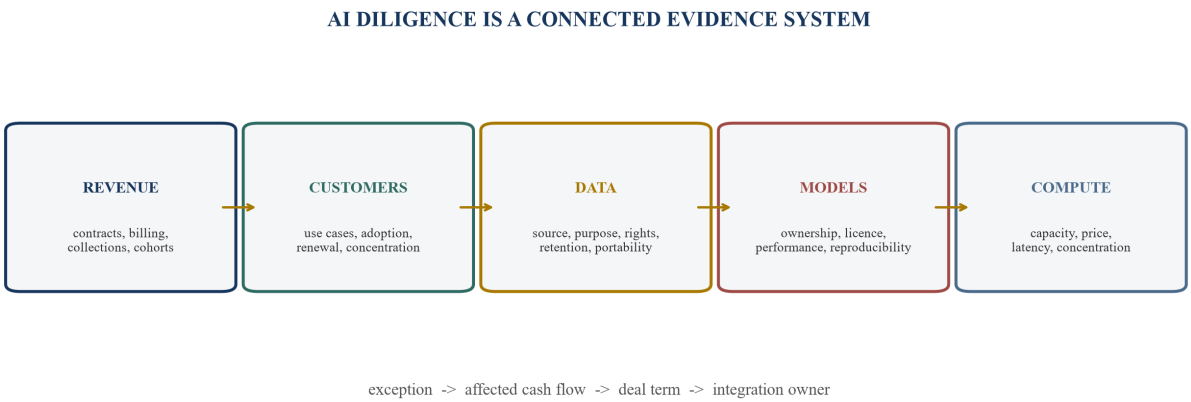

Every material exception should connect to economics, rights, obligations and transaction treatment.

2. Define the product and economic perimeter before testing growth

The diligence team should define what the target actually sells. An AI company can be a model developer, application provider, workflow integrator, data platform, managed service, infrastructure supplier or combination. Each archetype has a different asset perimeter, cost structure and dependency profile. A company described as software may rely on extensive human review. A product described as proprietary AI may orchestrate third-party models with limited internally owned technology. A data platform may derive advantage from contracts and operating processes rather than ownership of the underlying information.

The economic perimeter should identify the legal entities, products, customer contracts, data stores, model artefacts, cloud accounts, accelerator reservations, employees, contractors, licences and regulated permissions required to deliver the acquired proposition. The share purchase perimeter may not capture assets held by founders, affiliates, universities, development partners or customers. A carve-out may require new rights, accounts and service agreements before the buyer can operate independently.

The team should map each product claim to a repeatable customer outcome. Automation, prediction, generation and decision support are different economic propositions. The relevant evidence can include time saved, error reduction, conversion, loss avoidance, throughput, service quality or revenue created. A target that cannot connect model performance to a customer outcome may have a technically interesting product with weak commercial defensibility.

The perimeter should also distinguish current production capability from road-map claims. Model demos, pilots, benchmarks and signed contracts should not be blended. A road-map feature may require new training data, regulatory clearance, scarce talent or compute that is not funded. The acquisition case should value present assets and separately assess the probability, cost and ownership of future development.

Table 1. AI business archetypes and primary diligence implications

archetype	economic product	critical asset	principal diligence question	common valuation risk
foundation-model developer	model access or licence	weights, training process and talent	can the model be reproduced, governed and distributed?	training cost and uncertain monetisation
AI application	workflow outcome	product integration and customer adoption	does value persist if the underlying model changes?	model dependency and feature commoditisation
managed AI service	delivered result	people, process, models and data	how much delivery remains manual or customer-specific?	software multiple applied to service economics
data and retrieval platform	trusted information layer	rights, quality and connectors	can the buyer lawfully retain and reuse the data?	rights loss or costly re-permissioning
AI infrastructure	capacity or optimisation	hardware access and orchestration	are capacity, utilisation and customer demand aligned?	long commitments against volatile demand
embedded AI feature	improved host product	distribution and product fit	can contribution be measured separately?	double counting within broader revenue

Many targets combine archetypes; the analysis should allocate revenue, cost and dependency by product.

3. Reconcile reported revenue to contractual and economic reality

The revenue workstream should begin with the accounting ledger and finish with an economic bridge. IFRS 15 requires identification of the customer contract, performance obligations, transaction price, allocation and recognition as obligations are satisfied. AI businesses can create additional judgement because contracts combine platform access, usage, implementation, support, model training, custom development, data preparation and outcome-based fees.

Reported revenue should be reconciled by customer, contract, invoice and cash receipt. The team should test contract start dates, acceptance conditions, cancellation rights, credits, service-level penalties, variable consideration, usage estimates, related-party arrangements and principal-versus-agent treatment. When third-party model or cloud services are resold, gross presentation requires careful control analysis. A large top line can therefore represent a smaller economic fee when the target primarily arranges another provider's service.

The diligence bridge should separate recognised revenue, contracted backlog, annualised recurring measures, usage run rate and management pipeline. These measures answer different questions. Annual recurring revenue may be useful operationally, while its definition can include annualised recent consumption, future ramps, minimum commitments or inactive capacity. Public software filings demonstrate that issuers define annualised measures differently and warn that those measures are not forecasts or substitutes for recognised revenue.

Revenue quality also depends on collectibility and delivery. Signed contracts with disputed implementation, customer credits, low adoption or negative unit economics should not be treated as equivalent to deployed, accepted and collected subscription revenue. The evidence file should contain the original agreement, amendments, order forms, invoices, credit notes, usage,

acceptance, support history and cash receipts. Reconciliation differences should be quantified and assigned.

4. Build a customer and cohort ledger that survives management definitions

Customer analysis should follow legal customers and economic end users. Resellers, channel partners, cloud marketplaces and enterprise groups can obscure concentration. The ledger should connect parent, contracting entity, paying entity, deployment, users, product, region and ultimate economic sponsor. It should identify whether several reported customers are controlled by one group or depend on the same implementation partner.

Cohorts should be measured from an operational start date that reflects the customer's ability to receive value. Bookings, signatures, invoice dates and production activation can occur in different periods. The team should calculate gross retention, net retention, contraction, expansion, reactivation and churn under documented definitions, then reconcile them to contracts and invoices. Usage-based models need both revenue and activity cohorts because price changes or minimum commitments can hide declining adoption.

The review should distinguish product adoption from expert intervention. Implementation engineers, prompt specialists, data scientists or customer-success staff may manually resolve failures. Their cost may sit outside cost of revenue. The team should sample customer journeys and measure deployment time, exception rates, support hours, human review, model escalation and change requests. A product can show strong renewal because the target absorbs substantial unpriced service work.

Customer concentration should include dependency concentration. Several customers may rely on the same data source, foundation model, cloud region or regulated use case. A single change in model policy, data licence or sector rule can affect diversified-looking revenue. The cohort ledger should therefore include the technical and rights dependencies behind each revenue line.

Figure 2. Illustrative revenue-quality bridge from reported run rate to diligence case

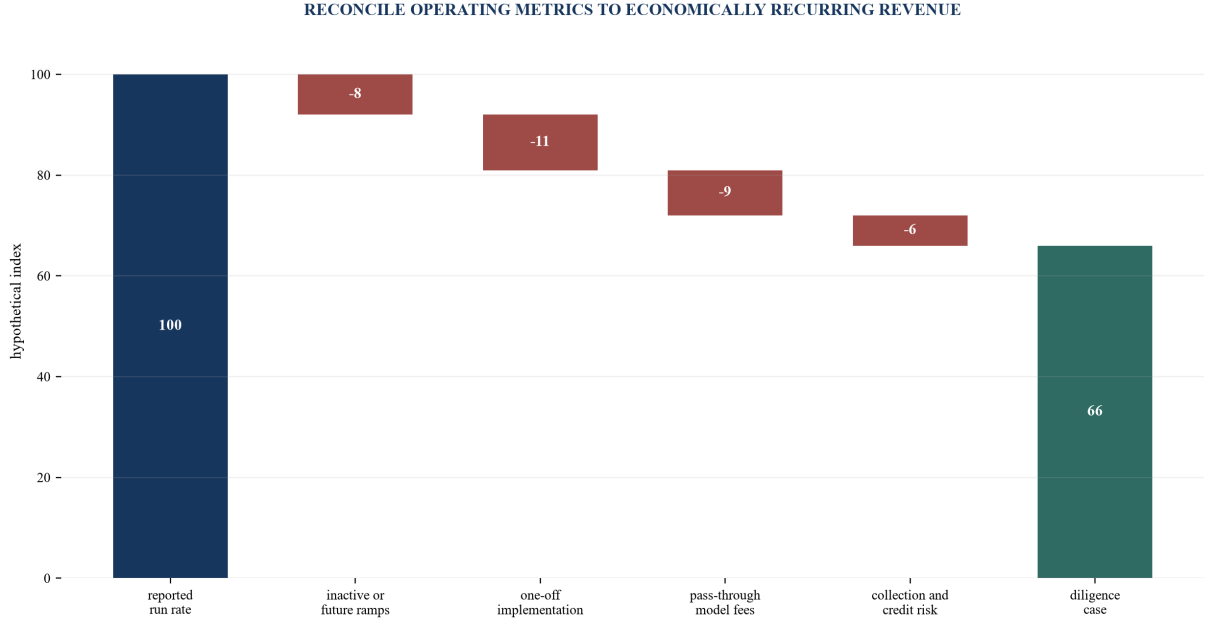

Values are hypothetical management assumptions and do not represent an actual company or valuation.

5. Test recurrence, consumption and outcome pricing separately

Subscription revenue is not automatically recurring in economic substance. The diligence team should test termination for convenience, renewal history, deployment depth, switching effort, budget ownership, product criticality and customer alternatives. A multi-year contract can remain fragile when service levels are weak, use is low or the customer can reduce committed scope.

Consumption revenue requires a volume-price-cost model. Usage can grow while gross profit declines if queries become more complex, model prices change or customers migrate to lower-priced tiers. The team should measure units that correspond to the product: tokens, documents, images, minutes, agents, transactions, compute hours or business outcomes. It should reconcile metering to invoices and identify free usage, promotional credits, internal traffic, abuse and unbilled consumption.

Outcome-based pricing requires attribution. The contract may link fees to savings, recovered revenue, resolved cases or successful transactions. Diligence should test baseline definitions, measurement rights, disputes, timing, caps and customer control over the outcome. Revenue can be delayed or reversed when evidence is weak. The target may also accept risk that exceeds its ability to influence performance.

Usage and outcome models should be stressed across model substitutions. A cheaper model can improve margin while reducing quality. A stronger model can improve conversion while raising cost or latency. Customer economics should therefore be tested at the task level, with performance, price and compute linked to the same use case. A blended corporate gross margin can conceal unprofitable high-value workloads.

Table 2. Revenue-quality tests by commercial model

model	evidence required	economic test	common exception	transaction response
fixed subscription	term, scope, activation and renewal	contracted gross profit after delivery support	signed but inactive customer	exclude from recurring base until activation
consumption	metering, rates and minimums	contribution per unit at normalised volume	growth driven by free credits	normalise price and cost together
enterprise platform plus usage	platform obligation and variable units	retention of base and expansion quality	minimum masks weak adoption	stress post-minimum renewal
implementation plus licence	distinct obligations and acceptance	repeatable software margin after implementation	services classified as software	separate multiple and earnout measures
outcome fee	baseline, attribution and collection	expected fee after disputes and timing	target cannot control outcome	use conservative probability and protection
reseller or marketplace	principal-agent evidence and fees	net economic take rate	gross billing treated as revenue	adjust revenue and valuation denominator

The transaction team should reconcile every metric to contracts, invoices, usage and cash.

6. Map data rights across the complete lifecycle

Data diligence should begin with a lineage map rather than a list of databases. Each material dataset should be traced from source through collection, transfer, preparation, labelling, training, fine-tuning, retrieval, prompting, evaluation, output, monitoring, retention and deletion. Rights can change at each stage. Permission to process data for a customer's service may not permit use for general model improvement. A public source may still contain personal information, confidential material or protected works.

The ledger should state the data controller or equivalent decision maker, processor relationships, purpose, legal basis, contract rights, confidentiality restrictions, copyright position, geographic location, transfer mechanism, retention, deletion capability and evidence. It should also identify derived datasets, embeddings, synthetic data, labels and model artefacts. Deleting a source record may not remove its influence from a trained model, which changes remediation and warranty design.

The UAE Personal Data Protection Law establishes requirements for processing personal data and cross-border transfers. UK Information Commissioner's Office guidance asks organisations to map machine-learning processes involving personal data and consider data minimisation and procurement diligence. Australian guidance states that public availability alone does not establish lawful use for model training and expects privacy-by-design assessment. India's data-protection regime adds another jurisdictional basis for consent, notice and processing controls.

The buyer should test operational deletion. A policy can state that data is deleted while backups, feature stores, vector databases, evaluation sets, logs or vendor systems retain copies. The target should demonstrate deletion through sampled records and downstream systems. Where complete removal from model weights is impracticable, the diligence record should state the limitation, affected obligations and remediation path.

Table 3. Minimum data-rights ledger

lifecycle stage	evidence	rights question	technical test	deal consequence
collection	notice, consent, contract and source	was the material obtained for this purpose?	sample source to ingestion record	remediation, exclusion or specific protection
preparation and labelling	vendor terms and instructions	can contractors access and create labels?	trace sampled records and access logs	assign ownership and confidentiality gaps
training and fine-tuning	dataset manifest and approvals	does the right cover model improvement?	reproduce dataset version and training run	retraining cost or use restriction
retrieval and prompting	customer terms and isolation design	can data enter context or vendor systems?	test tenancy, logging and leakage controls	covenant, architecture change or warranty
evaluation and monitoring	benchmark and incident records	can examples be retained for assurance?	reproduce evaluation set and results	preserve evidence or obtain permission
output and retention	output terms and deletion records	who may use outputs and for how long?	sampled deletion across downstream stores	liability allocation and integration task

Legal conclusions require qualified advice in every relevant jurisdiction.

7. Establish copyright, confidentiality and database provenance

The diligence team should distinguish legal access from technical access. A crawler, API, customer upload or licensed database can make content available while leaving restrictions on copying, text and data mining, model training, redistribution, attribution, derivative use or commercial output. The ledger should capture licence text, version, territory, purpose, duration, termination and change-of-control provisions.

The European Union Artificial Intelligence Act requires providers of general-purpose AI models to maintain a copyright-compliance policy and publish a sufficiently detailed summary of training content. The Commission's code and guidance provide operational documentation pathways. These requirements make provenance an enterprise capability rather than a retrospective legal memorandum. The buyer should assess whether the target can create the required evidence from actual systems.

The United States Copyright Office's work on generative-AI training demonstrates that legal questions remain fact-specific and commercially significant. Diligence should avoid treating pending litigation or policy debate as a substitute for evidence. The target should identify sources, permissions, technical controls and claims. It should quantify the products and cash flows that depend on contested material.

Confidential information requires separate analysis. Customer data, source code, model outputs and employee know-how can enter prompts, logs or evaluation sets. Contractual confidentiality may prohibit this even when privacy law is not engaged. The target should demonstrate tenancy controls, vendor settings, retention configuration, access restrictions and incident handling. A general statement that enterprise data is not used for training needs contract and configuration evidence.

8. Build a model bill of materials and ownership chain

A model bill of materials should identify every production model, version, provider, licence, weights location, training code, dependencies, datasets, evaluation artefacts, deployment endpoints and responsible owner. It should distinguish internally trained models, fine-tuned models, open-weight models, commercial APIs and customer-supplied models. The product can depend on several models through routing, fallback or specialist tasks.

Ownership analysis should trace employee and contractor assignments, university or grant conditions, prior-employer material, open-source software, model licences and third-party code. A company can own application code while lacking rights to weights, data or essential libraries. Change-of-control or assignment clauses can interrupt use at closing. The target should also identify contributions made through personal accounts or repositories.

Model licences should be reviewed in the version actually deployed. Terms, acceptable-use policies and distribution rights can change. Open-weight availability does not mean unrestricted use, and a permissive software licence may not govern associated model weights or data. The team should capture attribution, copyleft, field-of-use, hosting, redistribution and derivative-model conditions.

The bill of materials should connect to products and revenue. A dependency that serves a minor internal feature has a different transaction consequence from one that drives the principal

customer outcome. This mapping supports valuation, warranty scope, closing conditions and replacement planning.

Figure 3. Model provenance and reproducibility chain

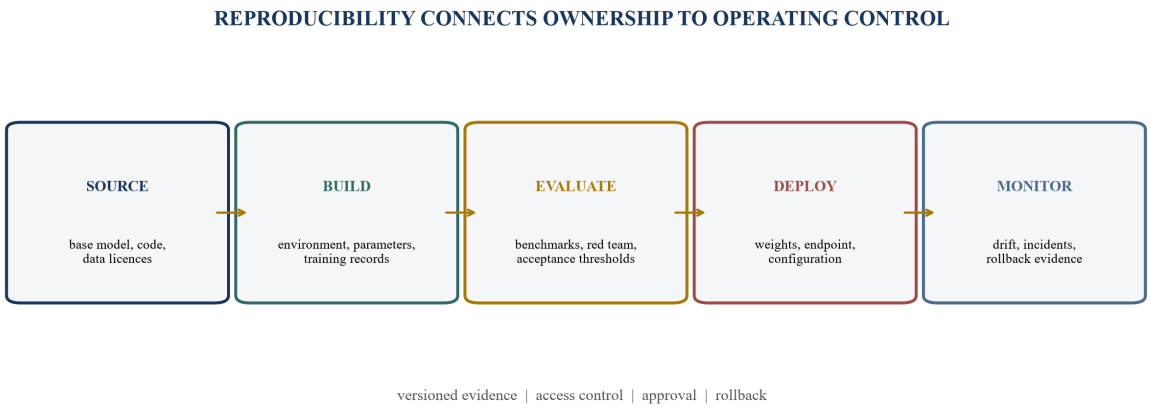

The target should be able to reproduce the deployed artefact from controlled inputs and records.

9. Test whether the deployed model can be reproduced

The buyer should require a controlled reproduction exercise for material internally developed or fine-tuned models. The target should identify the source code commit, data version, environment, dependencies, parameters, random seeds, training logs, checkpoints, evaluation suite and approval record. The exercise need not repeat an economically prohibitive full training run when another procedure can establish control, but the limitation should be explicit.

Reproducibility is an operational question as well as an intellectual-property question. A model may depend on undocumented preprocessing, a departed employee's local files, expiring cloud artefacts or a vendor image that cannot be recreated. The production endpoint may differ from the approved model because configuration or routing changed after evaluation. The team should compare model registry, deployment records and live behaviour.

The test should include rollback. When performance degrades, a provider changes terms or an incident occurs, the company needs a known safe state. Rollback evidence includes preserved artefacts, configuration, compatibility, data migration and an approved decision path. A backup file without a tested restoration process offers weak continuity.

Reproducibility also supports valuation. Proprietary performance that cannot be independently recreated may depend on tacit knowledge rather than controlled assets. The buyer can respond through retention, closing deliverables, documentation, holdback or a lower value assigned to the claimed advantage. The response should follow quantified dependency rather than a general technology discount.

10. Validate performance against the acquisition use case

Model benchmarks should match the product's operating context. Public benchmarks can support comparison, while they may not represent customer data, latency, safety, cost or failure severity. The target should maintain an evaluation suite linked to use cases, user groups, languages, jurisdictions and material harms. Results should include distributions and failure modes rather than one average score.

The team should separate model performance from system performance. Retrieval, prompts, tools, guardrails, user interface, workflow design and human review influence the customer outcome. A foundation-model upgrade can improve a benchmark while breaking structured outputs or increasing latency. Diligence should test the complete production path under representative load.

Evaluation governance should record dataset provenance, version, contamination risk, acceptance thresholds, review and release decisions. Repeated tuning against one benchmark can overstate generalisation. Customer examples used in evaluation require appropriate rights and isolation. High-risk use cases need evidence proportionate to the consequence of error.

The acquisition thesis should specify the performance that creates value. If the buyer plans to deploy the product into a new geography, language, sector or customer base, existing evidence may not transfer. The integration case should budget for new validation, data, controls and approvals. Synergy based on untested extension should remain a contingent management assumption.

11. Measure model drift, incidents and human intervention

Production evidence should show how the company detects drift in data, behaviour, cost and customer outcome. The target should define monitored signals, thresholds, owners, response times and escalation. Model drift can arise from changing inputs, upstream model updates, product configuration, user behaviour or adversarial activity. The team should compare written controls with logs and incident records.

Human intervention needs measurement. Reviewers may correct outputs, approve actions, label cases or handle exceptions. The target should report intervention rate, time, skill level, quality and cost by use case. Unrecorded intervention can inflate automation claims and understate cost of revenue. It can also create labour, privacy and confidentiality issues when reviewers are external or cross-border.

Incident diligence should sample failures from detection through closure. The file should include affected customers, data, model version, root cause, containment, notification, remediation and recurrence prevention. Absence of incidents can mean strong performance, incomplete monitoring or narrow definitions. Customer complaints, support tickets and credit notes can provide independent signals.

The NIST AI Risk Management Framework and generative-AI profile emphasise governance, mapping, measurement and management across the lifecycle. Singapore's governance frameworks add accountability, data, trusted development, incident reporting, testing, security and provenance dimensions. These sources provide a practical basis for testing whether the target operates an evidence-producing control system.

Table 4. Model evidence and transaction consequence matrix

evidence domain	minimum record	failure signal	economic exposure	possible deal treatment
ownership and licence	assignments, licence versions and dependency map	missing right or restricted field of use	product interruption or claim	closing cure, indemnity or exclusion
reproducibility	code, data, environment and run record	deployed model cannot be recreated	loss of proprietary capability	retention, holdback or valuation adjustment
performance	use-case evaluation and thresholds	benchmark does not match production	churn, liability or remediation	warranty, covenant or integration budget
release governance	approval, registry and deployment evidence	unapproved configuration in production	control failure and customer breach	closing condition and enhanced oversight
monitoring	drift, incident and intervention metrics	blind spot or manual dependency	margin and continuity risk	normalisation and operating plan
portability	tested replacement or export path	vendor or architecture lock-in	switching cost and bargaining exposure	price adjustment and migration plan

A technical exception becomes actionable when affected products and cash flows are identified.

12. Reconstruct training, fine-tuning and inference economics

AI gross margin should be rebuilt from workload evidence. The model should include training, fine-tuning, evaluation, inference, retrieval, storage, networking, observability, safety filters, third-party APIs, support and human review. Costs should be allocated by product and customer where possible. Corporate cloud invoices alone can hide credits, commitments, shared environments and capitalised development.

Training economics require a run ledger. The ledger should capture successful and failed runs, accelerator hours, energy or cloud cost, data preparation, engineering time, checkpoints and evaluation. Failed experiments are part of the cost of capability. A business that capitalises selected development costs or excludes research expense from gross margin needs a transparent bridge to cash economics.

Inference economics should connect request type, model, input size, output size, latency, caching, batching and quality. Average cost per query can be misleading when a small set of long or agentic tasks consumes most compute. The team should analyse percentiles and customer mix. It should test price changes, model substitution and peak demand.

Unit economics should include service obligations. A lower-cost model may create more support or review. An apparently expensive model may increase conversion or reduce human intervention. The relevant metric is contribution to the contracted outcome after all delivery resources. The acquisition model should retain a workload-level bridge rather than use one assumed percentage margin.

Figure 4. Illustrative AI workload contribution bridge

MEASURE THE COMPLETE COST OF DELIVERING THE AI OUTCOME

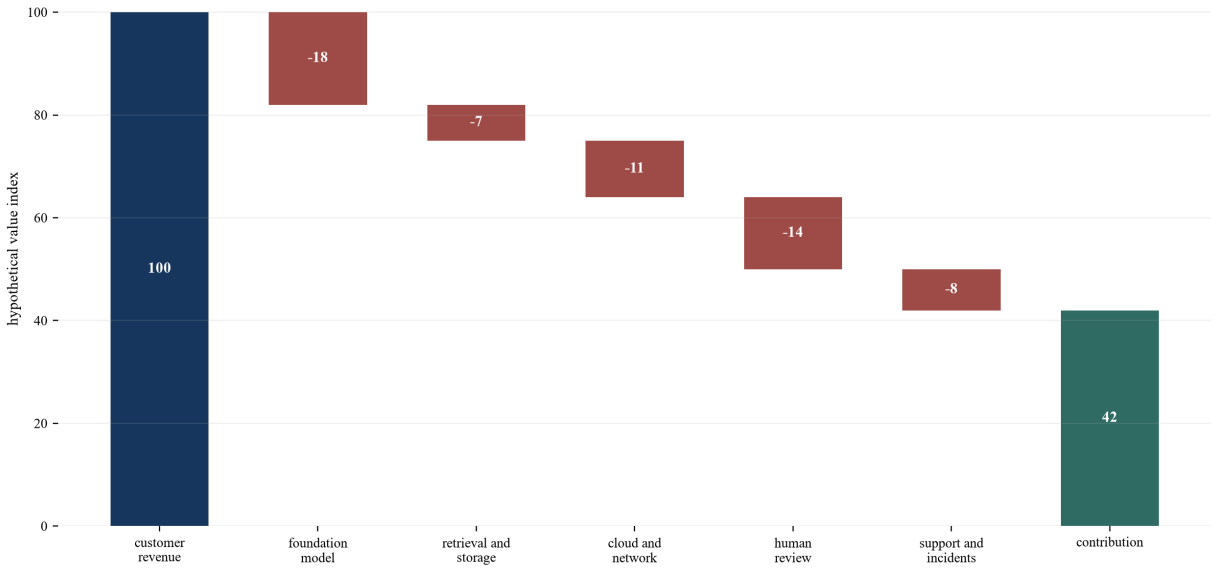

Amounts are hypothetical management assumptions used to demonstrate the diligence method.

13. Audit compute capacity, commitments and switching constraints

Compute diligence should identify accounts, providers, regions, services, accelerator types, reserved instances, minimum spend, prepaid credits, discounts, capacity reservations, take-or-pay obligations and termination terms. The target may have negotiated favourable pricing that does not transfer on change of control. It may also carry a commitment sized for a growth plan that the buyer does not adopt.

The United States Federal Trade Commission's 2025 staff report on large AI partnerships describes billions of dollars of cloud commitments, discounted resources, information sharing and potential switching constraints. The report concerns a specific set of partnerships and does not determine the treatment of another transaction. It demonstrates the diligence importance of linking financing, compute and strategic rights.

Capacity evidence should cover availability as well as price. A contract can state a discount without guaranteeing the required accelerators, region or start date. The team should compare reserved capacity, actual utilisation, queueing, service incidents and workload forecasts. It should identify which customers and model releases depend on scarce capacity.

Switching analysis should be tested through architecture and operations. Proprietary accelerators, model services, data stores, networking, security tools and deployment pipelines can create technical lock-in. Data-egress cost, migration time, retraining and performance regression can be material. A slide claiming multi-cloud capability should be supported by a deployed or tested path.

Table 5. Compute-obligation register

obligation	evidence	value question	downside test	transaction response
minimum cloud spend	master agreement, schedules and invoices	is committed spend supported by demand?	slower growth and lower utilisation	debt-like adjustment or seller cure
prepaid credit	payment and expiry records	is the credit transferable and usable?	expiry before workload deployment	haircut or closing confirmation
capacity reservation	resource, region and period	does capacity match product requirements?	model or customer mix changes	assignment, resizing or migration
volume discount	pricing tiers and eligibility	does pricing survive change of control?	discount removed at closing	price normalisation and consent
model API commitment	units, policy and termination	can usage move to another model?	provider price or policy change	replacement plan and covenant
hardware or colocation	title, lease, energy and maintenance	is the asset fully funded and operational?	delay, failure or lower demand	capex bridge and specific protection

The register should reconcile contracts, invoices, utilisation and the operating forecast.

14. Test cybersecurity across data, models and the AI supply chain

AI cybersecurity diligence should extend conventional software controls. Threats include data poisoning, model extraction, prompt injection, insecure tool use, training-pipeline compromise, malicious model artefacts, dependency vulnerabilities, secret leakage and misuse of autonomous actions. The analysis should identify attack surfaces across data ingestion, development, model registry, deployment, retrieval, APIs and end-user interaction.

NIST's secure-development profile for generative AI extends the Secure Software Development Framework to model development. CISA and international partners have published practices for securing data used to train and operate AI systems. These materials support evidence requests for provenance, integrity, access, logging, testing and incident response.

The team should inspect controls and sample their operation. Access lists should reconcile to current personnel and service accounts. Model artefacts should be signed or otherwise integrity-controlled. Secrets should be separated from code and prompts. External models and packages should pass an approval process. Red-team findings should have owners and closure evidence.

Agentic systems require particular attention because they can call tools, change records or initiate transactions. Singapore's 2026 framework for agentic AI emphasises bounding powers, meaningful human checkpoints, lifecycle controls and user responsibility. Diligence should map each agent's permissions, data, allowed tools, approval points, monitoring and kill mechanism. The transaction team should identify whether the buyer's environment expands those powers after integration.

15. Map regulation by role, product, geography and use case

The regulatory map should identify whether the target acts as developer, provider, deployer, distributor, importer, data controller, processor or regulated service provider in each market. Obligations depend on the role and use case. The same model can support low-risk internal productivity and a higher-risk employment, credit, health or safety decision.

The European Union Artificial Intelligence Act has phased obligations and specific requirements for general-purpose AI models. Commission guidance published through 2026 explains provider scope, code pathways and enforcement timing. The diligence file should state which products are placed on the Union market, whether significant model modification changes the target's role, and what documentation, transparency, copyright, safety or incident obligations apply.

Privacy, consumer, discrimination, sector and cyber rules remain relevant alongside AI-specific regulation. The UK ICO notes that its AI guidance is under review following legal change, so the team should verify the current position at signing and closing. UAE, Indian, Singaporean and Australian requirements should be mapped to actual processing and customers rather than grouped as a generic international compliance item.

The map should distinguish present compliance from a future integration plan. The buyer may intend to deploy the product into additional jurisdictions or regulated workflows. That expansion can require new evidence, notices, testing, human oversight and approvals. Synergy value should reflect the cost and timetable required to make the intended use lawful and operational.

Figure 5. Regulatory classification path for each material AI use case

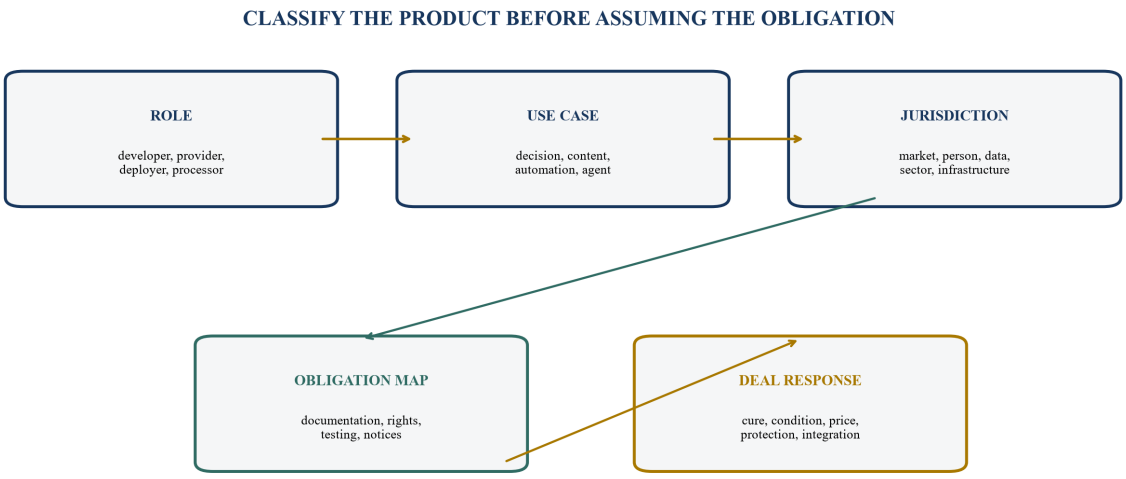

Qualified advisers should determine obligations using current law and transaction facts.

16. Translate exceptions into valuation and purchase-price mechanics

Valuation should reflect the cash-flow path affected by each exception. A data right that can be cured at modest cost differs from one that requires retraining the core model. A non-transferable cloud discount affects future gross margin. A customer contract with weak activation affects recurring revenue and working capital. The adjustment should be linked to probability, timing, cost, affected revenue and alternative.

The team should maintain separate bridges for revenue quality, gross margin, remediation cost, growth investment and synergy. This prevents double counting. Removing pass-through revenue can reduce the revenue multiple denominator while leaving gross profit largely unchanged. Adding a compute normalisation can affect margin and valuation. A retraining programme can be a one-time cash cost and a delay to growth. The model should show each mechanism once.

Purchase-price mechanics can address known balance-sheet or commitment items. Debt-like treatment may be considered for unavoidable prepaid or minimum-spend obligations when appropriate under the agreed definition. Working-capital mechanics may need to address deferred revenue, customer credits, usage accruals and cloud prepayments. Accounting and transaction advisers should define the treatment from the actual documents.

Contingent value should use measures the parties can verify after closing. An earnout based on reported AI revenue can be distorted by pricing, model substitution, integration or cost allocation. Measures can combine activated customers, collected gross profit, verified usage or specified milestones. Governance, access, accounting policy and dispute resolution should be agreed before the metric becomes economically important.

17. Convert identified risk into targeted transaction protection

Representations and warranties should reflect the evidence system. Generic intellectual-property and compliance clauses may not address training-data provenance, model licences, production versions, customer permissions, cloud commitments or AI incidents. The transaction documents should define the covered assets, systems and knowledge standard with qualified legal advice.

Specific indemnities can address identified exposures whose probability or amount is difficult to include in headline value. Escrow or holdback can support recovery when the exposure has a credible claim period. Closing conditions can require assignment, consent, rights remediation, security work, deletion, documentation delivery or capacity confirmation. Covenants can preserve controls between signing and closing.

Protection should remain connected to operational remediation. A warranty does not recreate a missing dataset, secure a cloud region or retain a key model engineer. Each legal term should have an integration owner and evidence date. The buyer should decide which conditions are essential for closing and which can be completed after closing without impairing value or compliance.

Disclosure quality matters. A broad data-room dump can technically disclose material while leaving the buyer unable to assess it. The exception ledger should identify the document, affected product, affected cash flow, owner, proposed treatment and open decision. The transaction committee should review changes as evidence evolves.

18. Prepare Day One continuity and the first hundred days during diligence

Day One planning should protect customer service, model access, data processing, cloud capacity, security monitoring and key talent. The buyer should confirm that accounts, licences, endpoints, certificates, secrets, payment methods and support contacts remain valid through change of control. A technically successful acquisition can still fail operationally if a founder-owned account or non-transferable agreement interrupts service.

The first hundred days should prioritise control and evidence before optimisation. Workstreams can establish the model bill of materials, consolidate access governance, reconcile data rights, preserve evaluation baselines, confirm incident routes, validate cost allocation and assign regulatory accountability. Architecture migration should follow a tested plan when immediate change would threaten performance or customer commitments.

The buyer should set a model-change protocol. Foundation models, prompts, retrieval, guardrails and routing may change rapidly. A post-close change can alter output quality, cost, privacy, customer obligations or regulatory classification. The protocol should require risk-based testing and approval, with rollback and customer communication where needed.

Talent planning should identify the people who understand data provenance, model behaviour, infrastructure and customer implementation. Retention should be tied to documented transition responsibilities. Concentrated knowledge should become controlled documentation, training and operational ownership rather than remain a permanent dependency.

Figure 6. Illustrative 100-day AI acquisition control programme

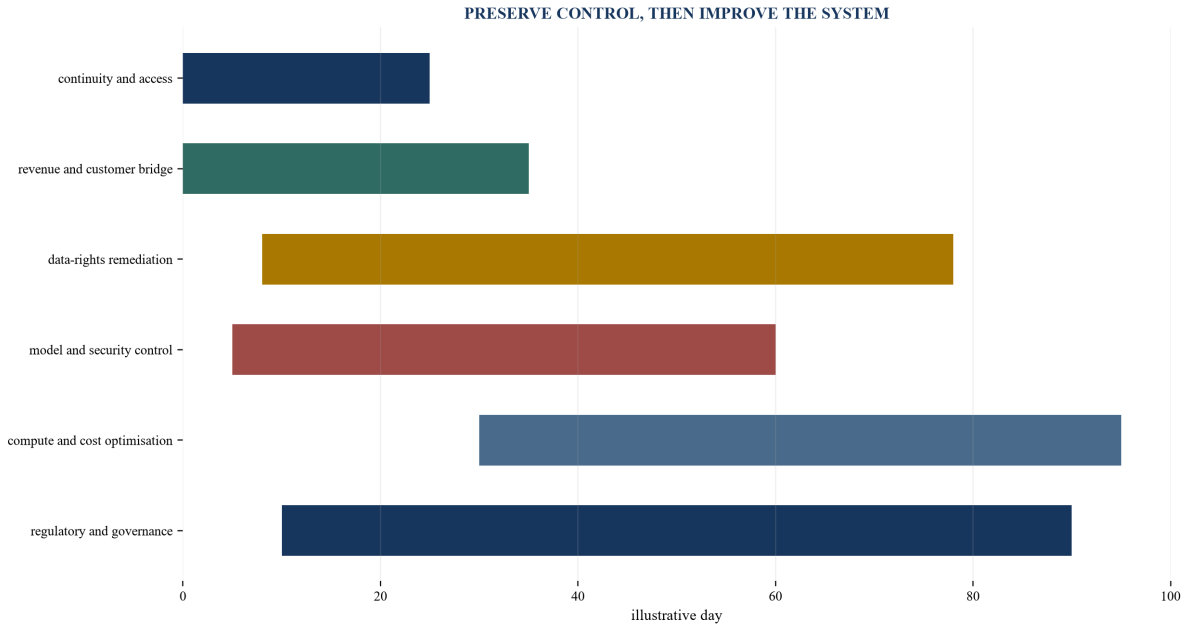

Timing is a hypothetical management assumption; the actual sequence depends on the target and transaction.

19. Apply the framework to a hypothetical acquisition case

Consider a hypothetical AI application company serving regulated enterprises. Management reports USD 24 million of annualised recurring revenue and 68 per cent gross margin. The company uses a commercial foundation-model API, proprietary retrieval and workflow software, customer documents and an external review team. These figures are illustrative management assumptions and do not describe an actual business.

The revenue ledger shows that USD 4 million relates to signed contracts that have not entered production, USD 2 million annualises a recent consumption peak and USD 3 million represents pass-through model charges presented gross. Two customers account for 34 per cent of collected gross profit. Renewal is strong for deployed customers, while implementation averages five months and requires more expert effort than the accounting allocation shows.

The data ledger confirms customer permission for service delivery. Several older contracts do not clearly permit retention of prompts for general product improvement. Evaluation sets include customer examples without a complete permission record. The company can segregate those datasets and proposes a new evaluation process, while retraining is not required because the material was not used in core model weights.

The model bill of materials shows that the proprietary advantage lies in retrieval, workflow integration, evaluation and domain operations. The foundation-model dependency is replaceable in principle, although only one provider meets current latency and quality thresholds. The cloud agreement includes a two-year minimum spend and a discount that requires change-of-control consent. The buyer models margin under current terms, loss of discount and a staged migration.

The transaction committee reduces the recurring-revenue base to deployed and economically active contracts, adjusts gross margin for expert review and normalises model fees. It requires cloud consent as a closing condition, establishes a specific remediation covenant for customer-data permissions, expands relevant warranties and funds a post-close portability test. The buyer preserves synergy from its distribution channel as a separate case rather than embedding it in stand-alone value.

The example shows the discipline of connected evidence. No single exception automatically stops the transaction. Each exception changes a specific cash flow, legal protection, closing requirement or integration workstream. The board can therefore evaluate value and executability using a transparent bridge.

20. Use a board gate that records evidence, assumptions and residual risk

The board paper should present the acquisition thesis, product perimeter, revenue bridge, customer cohorts, data-rights ledger, model bill of materials, performance evidence, compute obligations, regulatory map, cyber findings, valuation bridge, transaction protections and integration plan. It should identify evidence that has been verified, management assumptions used in the model and advice still required.

The board should challenge dependency. Which revenue depends on one customer, model, cloud provider, data right, employee or regulatory classification? How quickly can the dependency be replaced? What cash flow is at risk during replacement? A diversified supplier list is insufficient when the production path depends on one untested component.

Conditions should be measurable. Examples include receipt of a cloud consent, assignment of specified intellectual property, removal of unpermitted datasets, delivery of the model register, closure of named critical findings, activation of a customer contract or confirmation of a regulatory analysis. Broad conditions create ambiguity and weaken accountability.

Residual risk should have an owner and monitoring plan after closing. Acceptance is a board decision when evidence shows the exposure, alternatives and protection. The board should also specify stop conditions. A missing core right, unreproducible essential model, unfinanceable compute obligation or customer economics that fail after normalisation can invalidate the acquisition thesis.

Table 6. Board gate for an AI acquisition

domain	minimum evidence	challenge question	pause signal
product and revenue	reconciled contracts, activation, usage, invoices and cash	what revenue is recurring in economic substance?	material bridge cannot be reconciled
customers	cohorts, concentration and delivery evidence	is adoption repeatable without hidden service effort?	weak use or concentrated collected margin
data rights	lifecycle ledger and sampled provenance	can required data be retained, transferred and used?	core capability depends on missing rights
models	ownership, licences, registry, evaluation and reproduction	can the buyer lawfully operate and recreate the product?	essential model or artefact is unavailable
compute	commitments, capacity, utilisation and portability	can contracted demand be served through downside?	obligation exceeds funded demand or cannot transfer
security and regulation	risk map, incidents, controls and current advice	which role and use case determines each obligation?	material exposure has no viable remediation
value and protection	cash-flow bridge and documented deal response	has each exception been priced or allocated once?	double counting or unbounded residual risk
integration	Day One continuity, owners, budget and sequence	can control transfer without interrupting customers?	critical account, licence or person cannot transition

The board should require evidence across every linked system before authorising control.

21. Build a repeatable AI transaction-diligence capability

Buyers that expect repeated technology transactions should maintain an AI diligence operating model. The model can include standard data requests, evidence schemas, model and data ledgers, workload economics, regulatory classifications, cyber tests, valuation bridges and transaction-term libraries. Standardisation improves comparison while every target still requires fact-specific analysis.

The capability should integrate specialists early. Finance, commercial, legal, privacy, intellectual property, technology, security, cloud economics, regulation, tax and integration teams need a shared exception ledger. Sequential review creates late surprises because a technical fact may alter the accounting, valuation or contract. A short daily cross-workstream decision process can resolve ownership and consequences.

The buyer should retain evidence after closing. Diligence records can seed model governance, vendor management, customer compliance, cost control and incident response. Repeating the same discovery after acquisition wastes time and can lose transaction knowledge. Records should be transferred into accountable operating systems with appropriate confidentiality and access.

The framework should learn from outcomes. Post-close reviews can compare diligence assumptions with actual retention, margin, incidents, compute cost, model changes and remediation. The organisation can refine request lists and thresholds using evidence rather than market fashion.

The central principle is traceable value. An AI target is investable when the buyer can trace customer value to contracted revenue, lawful data, controlled models, available compute and a governed operating system. Diligence should make that trace visible, quantify breaks and convert them into price, protection and execution. The result is a board decision supported by evidence and an integration plan capable of preserving the asset being acquired.

References

[1] IFRS Foundation, IFRS 15 Revenue from Contracts with Customers https://www.ifrs.org/issued-standards/list-of-standards/ifrs-15-revenue-from-contracts-with-customers/

[2] IFRS Interpretations Committee, Principal versus Agent: Software Reseller, May 2022 https://www.ifrs.org/content/dam/ifrs/supporting-implementation/agenda-decisions/2022/principal-versus-agent-software-reseller-may-2022.pdf

[3] United States Securities and Exchange Commission, PTC Form 10-Q for the quarter ended 31 December 2025 https://www.sec.gov/Archives/edgar/data/857005/000119312526039347/ptc-20251231.htm

[4] European Union, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence https://eur-lex.europa.eu/eli/reg/2024/1689/oj?locale=en

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

[6] European Commission, Guidelines for providers of general-purpose AI models https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers

[7] European Commission, Template for general-purpose AI model providers to summarise training content https://digital-strategy.ec.europa.eu/en/faqs/template-general-purpose-ai-model-providers-summarise-their-training-content

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

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

[10] National Institute of Standards and Technology, Secure Software Development Practices for Generative AI and Dual-Use Foundation Models https://www.nist.gov/publications/secure-software-development-practices-generative-ai-and-dual-use-foundation-models-ssdf

[11] Cybersecurity and Infrastructure Security Agency, AI Data Security Best Practices https://www.cisa.gov/resources-tools/resources/ai-data-security-best-practices-securing-data-used-train-operate-ai-systems

[12] United States Federal Trade Commission, Staff Report on AI Partnerships and Investments, January 2025 https://www.ftc.gov/reports/ftc-staff-report-ai-partnerships-investments-6b-study

[13] United States Federal Trade Commission, Partnerships Between Cloud Service Providers and AI Developers https://www.ftc.gov/system/files/ftc_gov/pdf/p246201_aipartnerships6breport_redacted.pdf

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

[15] 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/about-this-guidance/

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

[17] United Arab Emirates, Federal Decree by Law No. 45 of 2021 Concerning the Protection of Personal Data https://www.uaelegislation.gov.ae/en/legislations/1972/download

[18] UAE Ministry of Economy and Tourism, Regulation of Competition https://www.moet.gov.ae/en/regulation-of-competition

[19] Government of India, Digital Personal Data Protection Act 2023 https://www.meity.gov.in/static/uploads/2024/06/2bf1f0e9f04e6fb4f8fef35e82c42aa5.pdf

[20] Government of India, Digital Personal Data Protection Rules 2025 https://www.meity.gov.in/static/uploads/2025/11/53450e6e5dc0bfa85ebd78686cadad39.pdf

[21] Infocomm Media Development Authority, Model AI Governance Framework for Generative AI https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2024/gen-ai-and-digital-foss-ai-governance-playbook

[22] Infocomm Media Development Authority, Model AI Governance Framework for Agentic AI https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/press-releases/2026/new-model-ai-governance-framework-for-agentic-ai

[23] Office of the Australian Information Commissioner, Guidance on privacy and developing and training generative AI models https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/guidance-on-privacy-and-developing-and-training-generative-ai-models

[24] Office of the Australian Information Commissioner, Guidance on privacy and commercially available AI products https://www.oaic.gov.au/privacy/privacy-guidance-for-organisations-and-government-agencies/guidance-on-privacy-and-the-use-of-commercially-available-ai-products

[25] United States Copyright Office, Copyright and Artificial Intelligence https://www.copyright.gov/ai/

[26] United States Copyright Office, Copyright and Artificial Intelligence Part 3: Generative AI Training https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf

[27] World Intellectual Property Organization, How to Prepare for IP Due Diligence https://www.wipo.int/en/web/ip-commercialization/w/blog/how-to-prepare-for-ip-due-diligence-the-ultimate-guide-for-ventures

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

[29] European Data Protection Board, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models https://www.edpb.europa.eu/our-work-tools/our-documents/opinion-board-art-64/opinion-282024-certain-data-protection-aspects_en

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

About the Author

Authored by Chennakeshav Adya, Independent Researcher

Chennakeshav Adya is an independent researcher and Managing Partner of Matchpoint Partners. His work focuses on corporate finance, capital strategy, transactions, operating transformation and the governance systems that convert strategic choices into measurable execution.