AI Product M&A Diligence: Testing Data Rights, Model Dependency and Unit Economics before Signing

An Evidence Framework for Valuation, Transaction Protection and Integration Readiness

Authored by Chennakeshav Adya
Independent Researcher
September 2026

Abstract

AI product acquisitions can combine familiar software economics with less familiar dependencies. A target may own customer relationships and application code while relying on third-party foundation models, cloud infrastructure, licensed data, open-source components and specialist employees. Product demonstrations and reported gross margins can obscure restrictions on training data, changing inference costs, weak evaluation evidence, customer concentration, manual service effort and contractual limits on transfer. These issues can affect whether revenue continues, whether the buyer can operate the product after closing and whether an apparent technology premium is supported by transferable economic capability.

This paper develops an evidence-led diligence framework for buyers, boards, investors and transaction teams assessing an AI product before signing. The framework links six questions: what the product does; which data and technology components it depends on; what rights and controls apply; whether performance is valid and repeatable; how customers use and pay for the product; and how revenue converts into contribution margin and cash. It translates those findings into valuation scenarios, transaction protections and a first-100-day integration plan.

A hypothetical acquisition illustrates the method. The author assumes a target with USD 72 million of annual recurring revenue, USD 54 million of reported gross profit and an indicative enterprise value of USD 720 million. The target uses external foundation models for most production inference and relies on a proprietary retrieval layer, customer-configured workflows and domain evaluation sets. All values, scores and outcomes in the example are analytical assumptions. They are not observed company data, forecasts, valuations, transaction terms or probabilities. The example shows how vendor pricing, manual review, data remediation, customer consent and retention risk can change sustainable margin and value. The central conclusion is that an AI premium should follow verified, transferable and economically durable capability rather than product labels or technical novelty.

JEL Classification: G24, G32, G34, L22, L86, O32, O33

Keywords: artificial intelligence, mergers and acquisitions, due diligence, data rights, model dependency, unit economics, valuation, intellectual property, AI governance, transaction protection

1. Define the transaction decision before opening the data room

AI diligence should begin with the buyer's decision and value thesis. The immediate question is whether the buyer should acquire the target on the proposed terms. Supporting questions concern standalone value, strategic value, downside exposure, transaction protection and integration

capacity. A broad technology review without these decision links can produce an impressive catalogue of models and controls while leaving price and signing conditions unresolved.

The investment thesis should state the expected source of value in operational terms. Possibilities include access to customers, lower product-development time, proprietary workflow data, specialist talent, a defensible distribution position, higher conversion, reduced service cost or an entry route into a regulated market. Each claimed source should have a baseline, evidence owner, financial mechanism, dependency and time horizon. The team should also record a counterfactual: build internally, license a product, partner, acquire a different target or continue the current plan.

The target's description as an AI company carries little analytical weight. The product may use a third-party model through an application programming interface, adapt an open model, train a proprietary model or combine several approaches. Economic control can sit in the workflow, data rights, evaluation capability, customer integration, distribution, brand or regulated permission. The diligence perimeter should follow the complete system that creates the customer outcome.

The board should approve materiality thresholds before detailed work starts. A matter may be material because it affects revenue, contribution margin, customer retention, legal use, service continuity, regulatory permission, integration cost or strategic optionality. Materiality should reflect the transaction thesis and downside, not only the target's current financial statements. A small upstream component can be critical when it cannot be replaced within the available time.

The output of this first step is a signed-off diligence charter. It names the decisions, workstreams, evidence standard, escalation thresholds, owners and dates. It also separates confirmatory diligence from thesis testing. Confirmatory work validates representations. Thesis testing asks whether the buyer's expected value creation can actually occur.

2. Map the product as a value chain rather than a single asset

An AI product is a chain of components and operating relationships. A typical chain includes source data, labelling and curation, model training or procurement, fine-tuning or retrieval, orchestration, application code, user interface, monitoring, human review, customer systems, distribution and infrastructure. The buyer needs to know who owns or controls each element, how it changes over time and what happens if one element becomes unavailable.

NIST's Generative AI Profile notes that generative AI value chains can involve many third-party data sources, pretrained models and software libraries, which can make attribution and accountability difficult [1]. The profile recommends risk management across the lifecycle and highlights the importance of component integration, provenance, testing and incident response. These controls support diligence, while commercial value still requires evidence of adoption and economics.

The product map should distinguish features demonstrated in a controlled environment from services used by paying customers. It should identify the production model and version, fallback arrangements, prompt and retrieval logic, data stores, interfaces, security boundaries, human interventions and downstream actions. A feature that generates text for internal review has a different consequence from a system that approves credit, configures industrial equipment or communicates regulated advice.

Ownership and control should be recorded separately. The target may own application code while its ability to operate depends on a revocable model licence, a cloud discount, a founder's undocumented methods or a customer's permission to use its data. The buyer should ask whether each dependency transfers on a change of control, whether consent is needed and whether the supplier can change price, features, geography or usage terms.

The map also identifies the value-bearing layer. A widely available model may support a defensible product when the target has exclusive workflow access, high-quality domain evaluation, embedded distribution and validated customer outcomes. A technically distinctive model may produce weak enterprise value when rights are uncertain, customer adoption is limited or operating cost scales faster than price.

Figure 1. The AI product dependency map follows the customer outcome through every material component

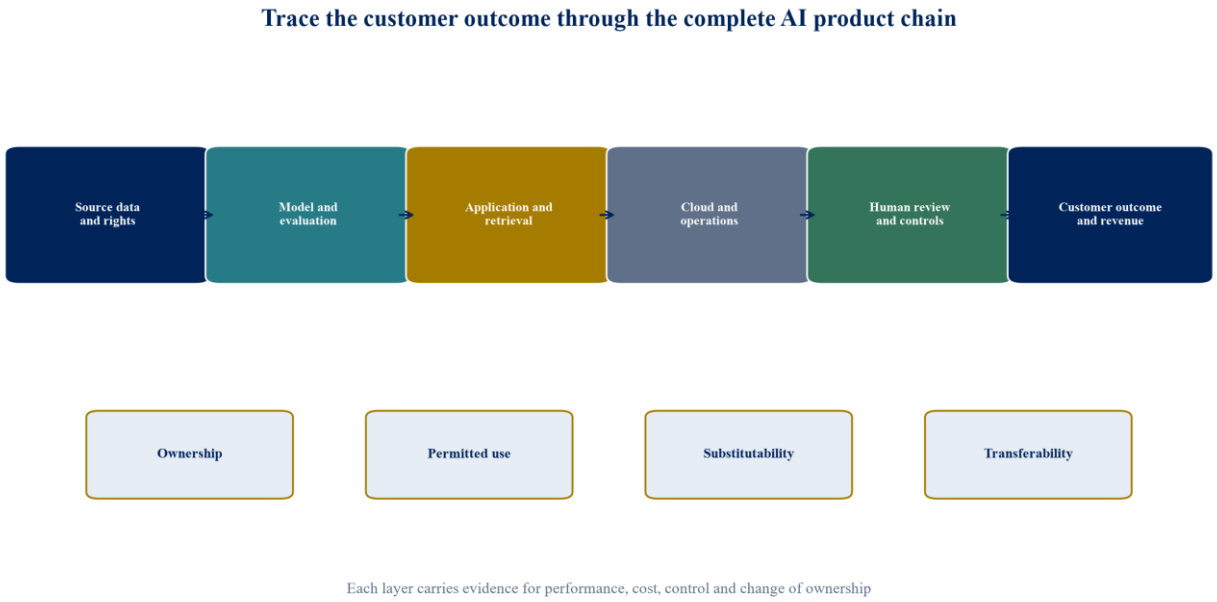

Diligence records ownership, permitted use, substitutability, performance, cost and transferability at each layer.

3. Build an evidence ledger with chain of custody

The buyer should create an evidence ledger before assigning scores or conclusions. Each material assertion should link to its source, date, system, owner and reviewer. Evidence can include executed contracts, data lineage, technical architecture, code repositories, model documentation, evaluation results, production logs, customer usage, invoices, support records, security tests, financial ledgers and board approvals. Management presentations can direct the inquiry, but they do not replace underlying evidence.

Evidence strength depends on proximity to the claimed result. Production outcomes across representative customers carry more weight than a benchmark chosen by the target. Executed licences carry more weight than a summary of rights. Reconciled cloud invoices, employee time and support costs carry more weight than an adjusted gross-margin schedule. Renewal and expansion behaviour carry more weight than customer logos.

Versioning is essential. Models, data sets, prompts, evaluation sets and vendor terms can change quickly. The ledger should identify which version supported each result and whether the production system still uses it. A target may present a strong evaluation from an earlier model while current costs or quality differ. A dated record allows the buyer to distinguish current capability from a historic test.

The ledger should preserve adverse results and exceptions. Failed releases, privacy incidents, model regressions, customer complaints and manual overrides show how the system behaves under pressure. Deleting or excluding these records can create a selection bias in the diligence package. The buyer should reconcile issue registers with support tickets, security logs, customer communications and board papers.

The evidence standard should be risk-based. A low-consequence drafting assistant may require a lighter test than an AI feature that changes a financial, medical, safety or employment decision. The diligence team should document the consequence of error, the available human control and the speed with which harm can propagate. This enables attention and transaction protections to follow material exposure.

Table 1. Core evidence ledger for AI product diligence

Diligence question	Stronger evidence	Warning sign	Transaction use
What operates in production?	Versioned architecture, deployment logs and customer usage	Demonstration environment presented as production	Define acquired capability and closing test
What may be used and transferred?	Executed rights, lineage and change-of-control analysis	Verbal assurance or missing source records	Condition, consent, indemnity or price
Does it perform?	Representative, reproducible evaluation and incidents	Selected benchmark without baseline	Value case, warranty and integration gate
Do customers retain and expand?	Cohort usage, renewals, price and support evidence	Logos without contracted scope or use	Revenue scenario and retention plan
What does delivery cost?	Reconciled compute, people, tooling and support by cohort	Gross margin excludes material delivery work	Margin normalisation and valuation
Can the buyer operate it?	Runbooks, access, succession and tested fallback	Key-person or supplier dependence	Retention, TSA, covenant or walk-away

Each conclusion should be traceable to current underlying evidence and an accountable reviewer.

4. Establish data origin, permitted use and transferability

Data diligence begins with origin. The target should identify first-party customer data, purchased or licensed data, public sources, scraped content, partner data, synthetic data, employee-created material and model-generated outputs. The record should describe collection method, contractual purpose, privacy basis, intellectual-property position, jurisdiction, retention, security and permitted uses. A generic statement that the company owns its data is inadequate.

Rights should be tested against the actual AI lifecycle. Permission to store data for service delivery may not permit training, fine-tuning, retrieval, evaluation, product improvement or resale. A licence may restrict commercial use, geography, model type, sublicensing or transfer on acquisition. Customer terms may differ by cohort and may have changed over time. The buyer should reconcile contract versions to the data used in production and development.

The European Data Protection Board's Opinion 28/2024 considers when an AI model may be treated as anonymous, when legitimate interests may support processing and how unlawful processing during development can affect deployment [2]. It states that anonymity and legal basis require case-specific assessment. The UK Information Commissioner's Office provides an AI and data-protection risk toolkit covering governance, transparency, lawfulness, accuracy, fairness, security and individual rights [3]. The toolkit is under review following UK legal changes, so current advice and applicable law should be checked for the target's facts.

Copyright requires a separate workstream. The US Copyright Office's 2025 report analyses the use of copyrighted works in generative-AI training and licensing considerations [4]. The EU AI Act requires general-purpose AI model providers to maintain a copyright policy and publish a sufficiently detailed training-content summary [5]. These sources do not establish that a particular target has sufficient rights. They show why training and output rights should be evidenced rather than assumed.

Transferability is a closing issue. The team should identify change-of-control provisions, consent requirements, data localisation, cross-border transfer mechanisms, customer deletion rights and obligations to return or erase data. It should test whether the buyer can continue using data after integration, combine it with other data or move the product to a different infrastructure environment.

Figure 2. Data rights must remain intact from collection through transaction transfer

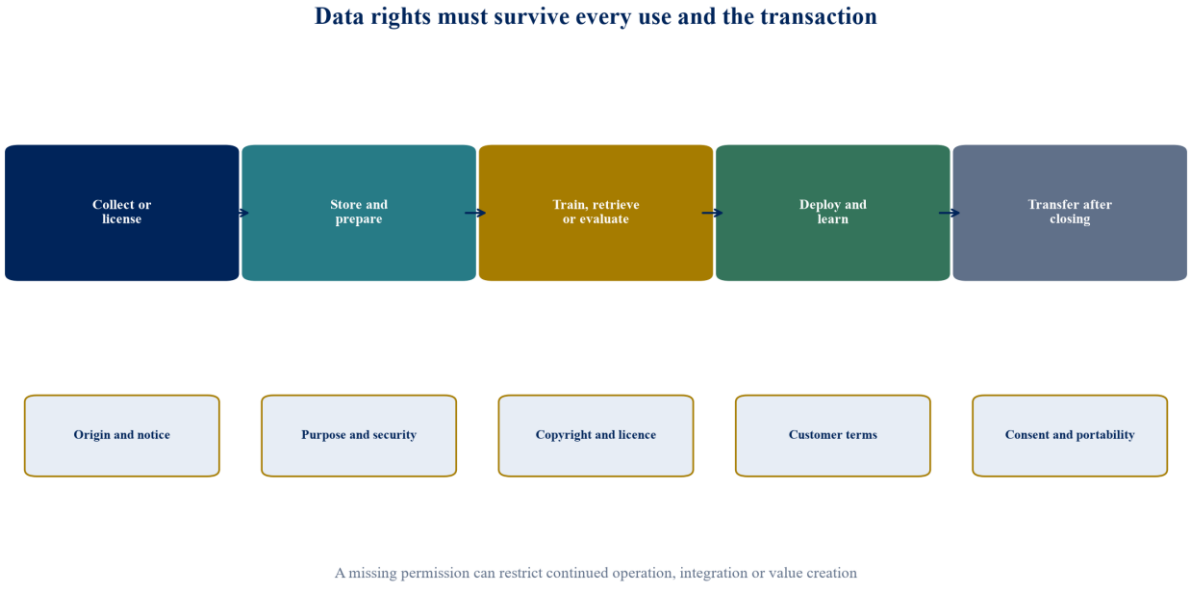

A break in purpose, licence, privacy basis or change-of-control consent can constrain continued operation or value creation.

Table 2. Data-rights diligence matrix

Data class	Questions before signing	Evidence	Possible response
Customer operational data	Does service permission cover the current workflow and model use?	Contract, notice, processing record and system lineage	Consent, product boundary or covenant
Training and evaluation data	May the target train, tune, benchmark and improve commercially?	Licence, source record and model documentation	Remediation, exclusion, indemnity or price
Public or scraped content	What terms, copyright, privacy and technical restrictions apply?	Source inventory, collection method and legal analysis	Remove, replace, license or ring-fence
Third-party licensed data	Does use, sublicensing and change of control continue?	Executed agreement and consent analysis	Closing condition or replacement plan
Synthetic and generated data	What source data, generation method and validation apply?	Provenance, tests and version record	Limit use and require validation
Customer-derived learning	Who owns improvements, labels and usage insights?	Customer terms, product records and IP provisions	Clarify ownership and permitted combination

Legal conclusions require advice for the relevant contracts and jurisdictions; the matrix organises the evidence.

5. Test intellectual property beyond code ownership

AI intellectual-property diligence covers code, model weights, prompts, retrieval logic, evaluation sets, data transformations, documentation, trade secrets, patents, brands and generated outputs. The buyer should verify creation, assignment, licence, open-source obligations, employee and contractor terms, joint-development rights and third-party claims. Repository access and a conventional code scan are necessary parts of this work, while the inquiry extends to the data and model chain.

The target should identify what it created and what it configured. Fine-tuning a third-party model may create useful capability without giving the target ownership or unrestricted portability. Prompt libraries and retrieval pipelines may be protected as confidential know-how when they are documented, controlled and difficult to reproduce. Evaluation sets can be commercially important, but their value depends on lawful source data, representative coverage and continued access.

Open-source components require licence-by-licence analysis. The target should maintain a software bill of materials and a model bill of materials, including libraries, model weights, data sets and deployment tools. The buyer should identify reciprocal obligations, attribution, notice, source-disclosure requirements, use restrictions and licence compatibility. An open licence can support rapid development and portability. It can also limit exclusivity or create compliance obligations.

Output rights depend on jurisdiction, human contribution, contractual terms and the source material involved. The US Copyright Office has stated that copyrightability of AI-assisted outputs turns on human authorship and that prompts alone may not establish protectable authorship [6]. A product's value should therefore avoid assuming exclusive rights in every output. Customer contracts should define permitted use, confidentiality, responsibility and ownership in a manner consistent with applicable law.

Trade-secret protection requires evidence of reasonable control. Access restrictions, confidentiality terms, repository governance, logging, segregation and exit processes matter. If key logic is widely shared, undocumented or stored in personal accounts, a claimed trade-secret moat may be weak. The transaction plan should secure repositories, credentials, assignments and key-person cooperation at signing and closing.

6. Measure model dependency and supplier concentration

Model dependency can affect continuity, cost, performance, compliance and bargaining power. The target should disclose every material production model, provider, version, hosting arrangement and contractual path. The buyer should identify minimum commitments, discounts, revenue sharing, exclusivity, consultation rights, data-use terms, service levels, termination, liability, audit rights, geography, change of control and migration support.

The US Federal Trade Commission's 2025 staff report on partnerships between cloud providers and AI developers describes arrangements involving equity and revenue-sharing rights, cloud-spend commitments, access to compute and information, and potential switching costs [7]. The UK Competition and Markets Authority has examined an interconnected foundation-model value chain involving compute, data, expertise, distribution and partnerships [8]. These findings concern specific market studies. They highlight dependency questions that a buyer should test in the target's own contracts.

Technical substitution should be demonstrated. The target should rerun a representative workload on an alternative model or architecture and measure quality, latency, throughput, safety, engineering effort and cost. A fallback model listed in a diagram is not a tested substitute. Switching may require new prompts, retrieval settings, evaluation thresholds, customer approvals, security review and regulatory documentation.

Dependency also exists inside open models. A target may rely on a particular model family, training framework, accelerator, library, data format or maintainer community. The diligence team should identify the operational owner and replacement time for each critical component. It should test model export, deployment automation, observability and the ability to reproduce a release from controlled artefacts.

Supplier concentration can support favourable economics during a promotional period and weaken them at renewal. The financial model should separate contractual prices from list prices and model the end of credits, discounts or committed-use benefits. A buyer should also identify whether the target's preferred access or pricing derives from a relationship that will survive the transaction.

Table 3. Model and infrastructure dependency assessment

Dependency	Evidence to test	Downside question	Transaction response
Foundation model API	Contract, versions, usage and evaluation	Can price, access or quality change materially?	Consent, covenant, price case and fallback
Cloud and accelerator capacity	Commitments, credits, allocation and architecture	What happens when discounts end or capacity tightens?	Margin normalisation and capacity condition
Open model or library	Licence, maintainers, security and reproducibility	Can support, licence or compatibility change?	Controlled fork, support plan or replacement

Dependency	Evidence to test	Downside question	Transaction response
Retrieval and vector layer	Data rights, quality, interfaces and portability	Can the buyer reproduce search quality elsewhere?	Export test and migration budget
Human review operations	Staffing, location, training and quality	Does automation economics depend on hidden labour?	Cost normalisation and retention plan
Specialist talent	Assignments, runbooks and succession	Can the system operate after key departures?	Retention, handover and closing deliverables

A written alternative becomes credible only after a representative migration or fallback test.

7. Reproduce technical performance under representative conditions

The buyer should define the required product outcome before choosing metrics. Accuracy can be useful for classification, while generation may require task completion, factuality, groundedness, safety, latency, consistency and human-review measures. The evaluation should reflect customer populations, languages, devices, data conditions and failure consequences. A single aggregate score can hide weak performance in important segments.

NIST's AI Risk Management Framework organises work around Govern, Map, Measure and Manage [9]. Its resources support test, evaluation, verification and validation across the lifecycle. The buyer can use this structure to examine whether objectives, limitations, risk tolerances, testing methods, monitoring and accountability are documented. Applying the framework does not itself establish product quality or legal compliance.

Reproducibility requires access to the model version, prompts, parameters, retrieval corpus, tools, evaluation data and scoring logic. The buyer should rerun reported tests and compare them with production outcomes. It should inspect contamination, benchmark selection, excluded cases, manual corrections and changes made after the test. Evaluation sets should be versioned and protected from routine product-development leakage.

The diligence team should include adversarial and boundary testing. It should examine prompt injection, data leakage, insecure tool use, hallucination, harmful bias, over-reliance, unsafe automation, denial of service and failure of upstream components where relevant. The test depth should follow the use case and consequence. Results should record severity, frequency, detectability, containment and remediation time.

Production monitoring is as important as pre-release evaluation. The team should inspect drift, incidents, overrides, escalation, customer complaints, latency, cost and model changes. It should verify that the target can connect an incident to the affected model, data, customer and release. A strong product organisation knows where its system fails and how decisions are made when thresholds are breached.

8. Examine security, resilience and operational control

AI systems add attack surfaces to familiar application and infrastructure risks. Sensitive prompts and retrieval content may be exposed. Adversarial inputs may influence model output or connected tools. Training and evaluation artefacts can be poisoned. Credentials, model weights and proprietary data can be stolen. A connected agent can take actions beyond the intended authority if permissions and approval gates are weak.

Security diligence should review architecture, identity and access, secrets, encryption, network boundaries, logging, secure development, vulnerability management, incident response, supplier

assurance, business continuity and recovery. It should identify which controls cover models, data pipelines, orchestration and tool access. Standard certifications can provide evidence about the management system, but the buyer should test whether the certified scope includes the acquired product and current production environment.

Resilience should be demonstrated through failure exercises. The target should show how the product behaves when the primary model is unavailable, a model response changes, a retrieval store is corrupted, a cloud region fails, usage spikes, a safety threshold is breached or a human-review team is inaccessible. Recovery-time and recovery-point objectives should match customer commitments and the consequence of failure.

Operational authority needs clear limits. A product that recommends an action differs from one that executes it. The buyer should map which decisions the system may make, which require human approval and which must stop when evidence is incomplete. Tool permissions should follow least privilege. High-consequence actions should have authentication, segregation, approval, monitoring and rollback.

Security findings should translate into financial and transaction effects. Remediation can require engineering time, infrastructure, external testing, customer notification, insurance or delayed integration. The buyer should include these costs in the value bridge and define which matters must be completed before signing, before closing or during the first 100 days.

9. Classify regulatory roles and obligations by market

Regulatory exposure depends on the product's role, use case, customer, geography and position in the AI value chain. The target may be a model provider, downstream system provider, deployer, importer, distributor, processor, controller or regulated-sector supplier. The same product can have different roles across markets and customer configurations. Diligence should map obligations at the level of each material use case.

The EU AI Act places obligations on general-purpose AI model providers, including technical documentation, information for downstream providers, a copyright policy and a public summary of training content [5]. Providers of models with systemic risk face additional evaluation, risk-management, incident-reporting and cybersecurity obligations [10]. A downstream application company may depend on receiving sufficient model information to fulfil its own obligations. The buyer should verify the target's classification, transition plan and contractual access to necessary documentation.

Sector regulation can matter more than a horizontal AI rule. Financial services, healthcare, employment, consumer protection, product safety, critical infrastructure and professional advice can impose additional standards. The team should identify licensing, approval, recordkeeping, explainability, outsourcing, model risk, customer communication and human-oversight requirements. Legal advice should address the exact product and jurisdictions.

Regulatory diligence should inspect actual operating evidence: impact assessments, policies, risk registers, testing, model and data documentation, incidents, complaints, audit findings, customer disclosures and regulator correspondence. A policy created for the transaction carries less evidential weight than a record used in product decisions over time.

The transaction timeline should include approvals and remediation. Change of control may require notification or consent. Product modifications after integration may change classification. Combining buyer and target data can create a new processing purpose or risk profile. The integration plan should preserve compliant service while governance and documentation are aligned.

10. Reconstruct fully loaded AI unit economics

Reported gross margin may omit costs needed to deliver the AI outcome. The buyer should reconstruct economics from usage and ledgers. Direct cost can include model inference, cloud compute, storage, data licences, vector services, observability, safety filters, security tooling, human review, customer-specific configuration, implementation, support, incident handling and usage taxes or network charges. Capitalised development and shared engineering may also affect cash economics even when accounting presentation differs.

The analysis should use customer and workload cohorts. Cost can vary by prompt length, output length, modality, model, region, latency requirement, retrieval volume, tool calls, concurrency and error rate. Average cost per request can mislead when a small number of high-intensity customers consume disproportionate resources. The team should reconcile usage records to supplier invoices and contracted customer volume.

Revenue should be analysed against the unit that drives cost. A seat price can be attractive when usage is bounded and adoption expands retention. It can be risky when a flat contract permits unlimited high-cost inference. Usage pricing can protect margin while creating customer budget uncertainty. Outcome pricing can align value and price, subject to measurement and attribution. The buyer should test minimums, overages, renewals, credits and contractual rights to reprice.

Human effort should be measured, including work described as onboarding, customer success or quality assurance. Hidden manual review can make a product appear more automated than it is. The buyer should distinguish temporary implementation work from recurring delivery work and test whether automation can reduce it without harming quality or control.

The financial model should include vendor price changes, model mix, context growth, customer usage, quality thresholds, fallback costs, currency and committed spend. It should show gross margin, contribution margin and cash by cohort. A credible case reconciles technical drivers to invoices and customer contracts rather than applying a flat margin assumption.

Figure 3. Reconstruct the path from contracted revenue to retained contribution margin

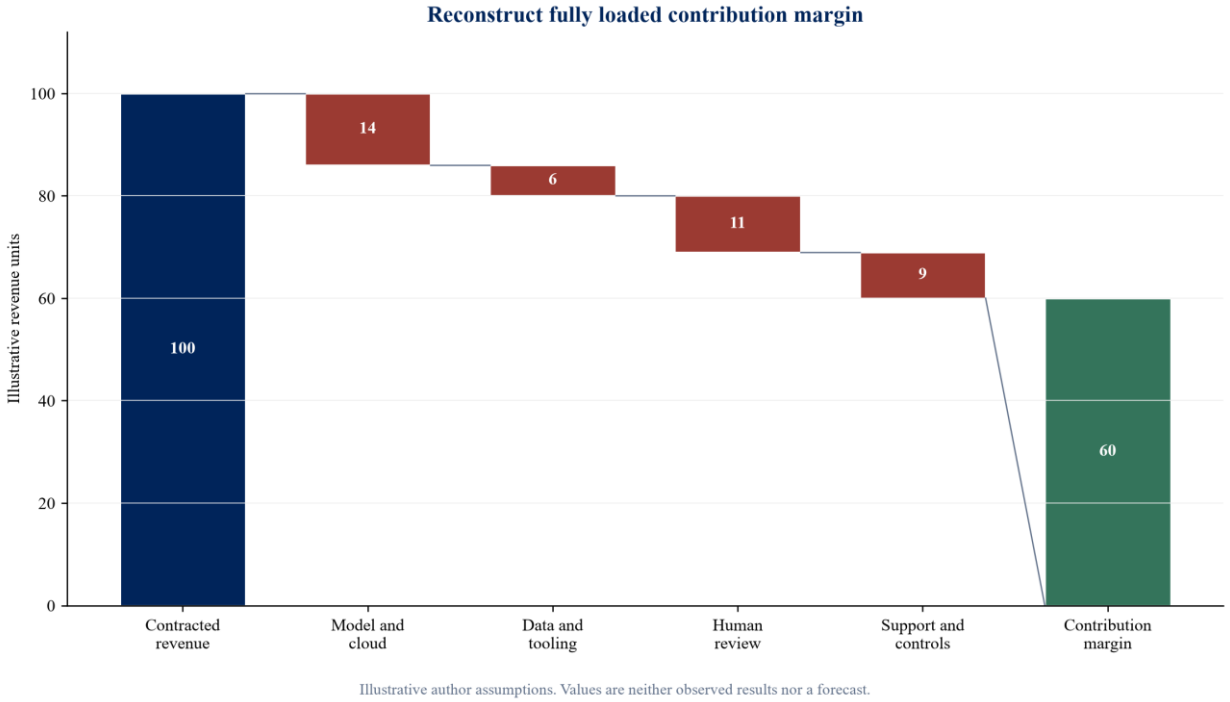

The bridge includes compute, data, human review, support, controls and remediation required to deliver the contracted outcome.

11. Test revenue quality and customer dependence

AI product revenue should be examined through contract, usage, outcome and renewal. Annual recurring revenue can include pilots, cancellable commitments, services, reseller arrangements, minimum usage and discounts. The buyer should reconcile the revenue schedule to executed contracts, invoices, collections and product telemetry. It should identify which features customers actually use and whether those features depend on specific models, data or people.

Cohort analysis should cover retention, expansion, contraction, price, usage, support burden, implementation time and gross margin. Logo retention can hide seat or usage contraction. Revenue growth can be concentrated in a few customers or a reseller. A large customer may also supply critical data, provide a reference, finance development or constrain the product roadmap. These relationships should be valued and risked separately.

Customer interviews should test why the product was selected, which outcome it changes, who owns the budget, how quality is measured, what alternatives exist and what would cause switching. The buyer should discuss model changes, data use, security, service levels, pricing and change of control where permitted. Interviews should be reconciled with product and contract evidence.

The team should identify whether the product replaces labour, augments a workflow, reduces risk or creates a new service. Each mechanism has a different budget and adoption path. Productivity claims should show how released capacity creates cash value through reduced cost, avoided hiring, greater volume, faster conversion or differentiated price. Time saved without an operating decision is not a realised benefit.

Revenue durability also depends on the buyer's integration plan. Customers may value the target's neutrality, privacy architecture, model choice or independent roadmap. An acquisition by a larger platform can strengthen distribution and resources while creating concerns about data sharing, price or product priority. The retention plan should address these concerns before public announcement where legally and practically possible.

12. Assess people, knowledge and delivery capacity

AI product capability often concentrates in a small group of technical and domain specialists. Diligence should map roles to critical decisions, code, data, evaluation, customer relationships, security and operations. The buyer should identify replacement time, retention risk, incentives, notice periods, location, immigration, employment assignments, consulting arrangements and competing commitments.

Knowledge transfer should be tested rather than inferred from documentation. A replacement team should be able to explain the architecture, reproduce a release, rerun an evaluation, trace a data source, respond to an incident and reconcile product usage to cost. Runbooks, model cards, data dictionaries, design records and customer playbooks are valuable when they enable another qualified person to operate the system.

The target's operating model may rely on founder approval for model choice, customer exceptions or release decisions. This can support quality at an early stage and create a scaling constraint. The buyer should map decision rights and determine which responsibilities need succession before closing. Retention packages should follow the capabilities required to preserve value, with objectives and time horizons linked to integration.

Delivery capacity affects the revenue case. A growing pipeline may require implementation engineers, domain experts, data integration and customer success. The buyer should compare contracted backlog with available teams, cycle times and hiring plans. Revenue that cannot be implemented within customer deadlines should not receive the same value as deployed recurring revenue.

Culture and control matter during integration. A high-speed product team can lose effectiveness under unclear approvals, while an informal control environment can expose a regulated buyer. The integration design should preserve product accountability, set clear release gates and avoid duplicated authority. The first 100 days should focus on critical rights, continuity, customer trust and financial visibility before wider platform consolidation.

13. Connect accounting evidence to valuation without confusing the two

Accounting and transaction valuation answer different questions. IFRS 3 requires an acquirer to recognise identifiable intangible assets separately from goodwill when the relevant criteria are met [11]. IAS 38 addresses recognition and measurement of intangible assets, including research and development [12]. These standards can affect purchase-price allocation and post-acquisition earnings. They do not set the negotiated enterprise value or validate an AI premium.

The buyer should identify potential technology, customer, contract, brand, data-related and other intangible assets with valuation specialists and auditors. Rights, separability, contractual terms, useful life, obsolescence and economic benefit matter. Some strategic capabilities, including assembled workforce and synergies, may form part of goodwill rather than separately recognised

assets. The transaction model should anticipate amortisation, impairment testing and integration costs.

Valuation should begin with sustainable cash economics. The base case should normalise revenue quality, fully loaded contribution margin, required development, security and compliance cost, working capital, capital commitments and taxes. The downside should reflect customer loss, higher model cost, removal of disputed data, delayed approvals, talent attrition and remediation. Strategic value should be separated from standalone value and assigned an owner, investment need and probability-weighted timing.

Comparable-company multiples require business-model consistency. A company with proprietary-model economics differs from an application company paying variable inference cost. A high-margin subscription differs from a managed service with recurring human delivery. Growth, retention, contribution margin, capital intensity and concentration should be reconciled before applying a multiple.

The buyer should create a value-at-risk ledger. Each unresolved diligence issue should connect to revenue, cost, time, capital or terminal value. The ledger prevents technical findings from becoming isolated observations and provides a basis for price, escrow, earn-out, indemnity or closing conditions.

14. Convert findings into price and transaction protections

Diligence creates value when findings change a decision or transaction term. Each material issue should have an owner, evidence status, financial exposure, remediation plan and proposed response. Responses include price adjustment, holdback, escrow, earn-out, warranty, indemnity, covenant, condition precedent, consent, retention arrangement, transitional service and walk-away.

Representations should address the facts that matter to the product and jurisdiction. Areas may include data origin and permitted use, privacy compliance, intellectual-property ownership, open-source and model licences, material suppliers, security incidents, evaluation records, regulatory classification, customer contracts, product claims and change-of-control consents. Drafting requires legal counsel and should follow the exact evidence and risk allocation.

Specific indemnities can address identified exposures that are not adequately allocated by general warranties. Escrow or holdback can support recovery. An earn-out can bridge a value gap when future revenue or margin depends on uncertain adoption, subject to clear definitions, operating covenants and dispute mechanisms. A completion condition can require critical consent, assignment, remediation or a tested fallback before ownership transfers.

Transaction protection should not substitute for operability. A contractual claim has limited value if the product cannot serve customers after closing or if the seller lacks resources to satisfy it. Critical continuity matters should be resolved through conditions, transition and integration preparation. The buyer should identify the minimum viable operating state at closing.

Table 4. From diligence finding to transaction response

Finding	Value effect	Evidence needed	Possible protection
Training rights are incomplete	Remediation cost, product restriction and claim risk	Source-by-source rights and usage map	Removal, licence, indemnity, escrow or condition
Primary model cannot be substituted quickly	Continuity, margin and bargaining exposure	Contract review and representative fallback test	Consent, covenant, transition support and downside price
Reported margin excludes recurring human review	Lower sustainable contribution margin	Time, usage, payroll and cohort reconciliation	Normalised EBITDA or revenue multiple adjustment
Largest customers can terminate on change of control	Revenue and financing downside	Executed contracts and customer plan	Consent condition, holdback or earn-out
Evaluation does not represent production	Product and liability uncertainty	Reproducible testing across material cohorts	Closing test, warranty and funded remediation
Key capability is concentrated in founders	Execution and continuity exposure	Role map, assignments and succession test	Retention, handover, vesting and transition covenant

The response depends on severity, detectability, control, seller recourse and the buyer's ability to remediate.

15. Design the first 100 days before signing

The integration plan should begin during diligence. It should identify the operating state required on day one, the first 30-day control actions, the 60-day evidence programme and the 100-day investment decisions. Ownership, budgets, dependencies and customer communication should be agreed before signing where possible.

Day-one priorities include access, credentials, supplier continuity, customer support, privacy roles, incident escalation, release authority and key-person retention. The buyer should avoid changing models, data flows or customer terms without assessing their effect. A stable operating perimeter creates time to verify the combined system.

During the first 30 days, the buyer should reconcile model and data inventories, confirm licences and consents, lock critical repositories, validate monitoring and integrate financial reporting. It should establish a joint product and risk committee with authority over releases and material exceptions. Customer teams should identify renewal and change-of-control risks.

By day 60, the combined team should rerun representative evaluations, test fallback, build customer-level unit economics and prioritise remediation. It should compare the signed value case with actual usage, cost, retention and delivery capacity. Differences should be recorded in the value-at-risk ledger.

By day 100, the board should decide which capabilities to scale, migrate, retain independently, combine or stop. Capital should follow verified customer and economic evidence. The buyer should also refresh the purchase accounting plan and integration forecast as evidence improves. The operating cadence should continue after the initial period through monthly product economics and quarterly strategic review.

Figure 4. The diligence-to-integration path converts evidence into signing and capital decisions

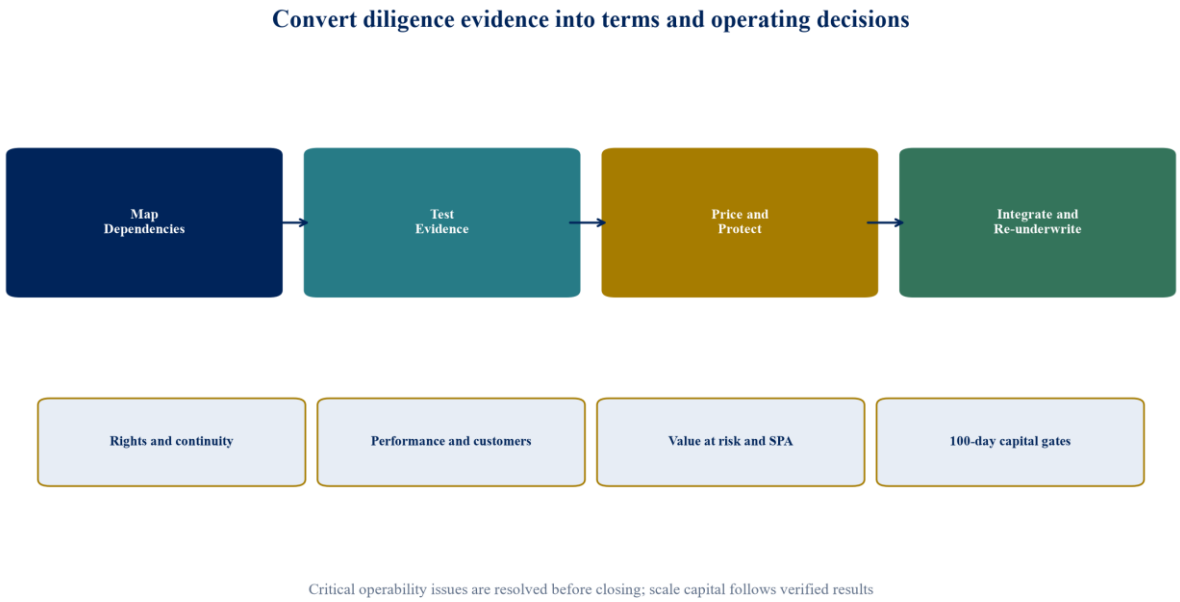

Critical continuity and rights issues are resolved before closing; scale investment follows post-close evidence gates.

16. Apply the framework to a hypothetical acquisition

Consider a hypothetical AI workflow company with USD 72 million of annual recurring revenue, 75 percent reported gross margin and an indicative enterprise value of USD 720 million. The company serves regulated professional-services customers. It uses external foundation models for most inference, combines them with a proprietary retrieval layer and maintains domain evaluation sets. The author assumes that the largest ten customers represent 48 percent of revenue. These values are illustrative assumptions.

The initial thesis attributes value to embedded workflows, customer access, evaluation capability and faster product entry for the buyer. Diligence confirms strong usage and renewal in mature cohorts. It also identifies three issues. First, a legacy training corpus lacks complete source and permitted-use records. Second, the target's cloud discount expires within 14 months and its tested fallback covers only 35 percent of production volume. Third, reported cost of revenue excludes a recurring quality-review team and part of customer-specific retrieval maintenance.

The buyer removes the legacy corpus from the base case until rights and replacement are established. It rebuilds contribution margin using current invoices, full review cost and customer-level usage. Under the author's assumptions, sustainable contribution margin falls by eight percentage points before remediation. The downside case assumes a further six-point effect from vendor repricing and inefficient fallback, plus loss of two concentrated customers at renewal. These are scenarios, not forecasts.

The transaction response combines a lower upfront price, an earn-out tied to retained revenue and contribution margin, a specific data-rights indemnity, closing consents for critical suppliers and customers, and retention arrangements for the evaluation and platform leads. A funded 100-day

plan replaces the disputed corpus, completes fallback tests and introduces customer-level economics.

The example shows why one multiple cannot resolve the decision. Customer evidence supports strategic value, while rights, dependency and cost evidence change the amount, timing and protection of consideration. The buyer proceeds only after the minimum viable operating state is documented and critical consents are obtained.

17. Run a sequenced diligence workplan

The first week should establish the product map, value thesis, materiality thresholds and evidence ledger. The buyer requests architecture, data and model inventories, material contracts, customer and usage schedules, cost ledgers, evaluation records, incidents, security evidence, IP records and organisation maps. Workstream leaders agree common customer, product and model identifiers so evidence can be reconciled.

Weeks two and three focus on rights, reproducibility and economics. Legal and data teams trace material sources and permissions. Technical teams reproduce evaluations and observe production workflows. Finance teams reconcile revenue, usage, supplier invoices and human delivery cost by cohort. Commercial teams interview selected customers and test retention, alternatives and change-of-control issues.

Week four brings the findings together. The team produces a data-rights matrix, dependency map, evaluation report, customer-quality analysis, unit-economics bridge, integration plan and value-at-risk ledger. Each red or amber item has a quantified scenario where possible, evidence gap, owner and proposed transaction response.

Before signing, the investment committee receives a concise decision paper. It states what has been verified, which assumptions remain, the base and downside valuation, critical consents, transaction protections, integration funding and walk-away conditions. Supporting evidence remains linked in the diligence record.

The workplan should allow issues to change the timetable. A missing licence, failed fallback or customer consent may require more evidence than a fixed auction schedule permits. The buyer should decide explicitly whether to extend diligence, change terms, accept a bounded risk or stop. Silence and time pressure should not convert an unresolved assumption into a fact.

18. Conclusion

AI product M&A diligence requires a connected view of product, rights, dependency, performance, customers and economics. A target can possess capable technology and weak transferability. It can report recurring revenue and carry material concentration or cancellation exposure. It can show attractive gross margin while excluding compute, review, support or control costs needed to deliver the contracted outcome.

The buyer should trace every material feature through the data and model chain to a customer result and a financial consequence. Evidence should be current, versioned, reproducible and reconciled. Data and intellectual-property rights should match actual use and survive change of control. Model and cloud alternatives should be tested. Technical performance should represent

production. Customer usage and renewal should support revenue quality. Unit economics should include all delivery and control costs.

Findings then become valuation scenarios, transaction protections and integration actions. Critical continuity and rights issues belong before closing. Scale capital belongs after evidence gates. This discipline gives the investment committee a defensible answer to the central question: which capabilities will the buyer control, operate and monetise after the transaction, and what should it pay for them today?

References

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

2. 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. 18 December 2024. https://www.edpb.europa.eu/documents/opinion-of-the-board-art-64/opinion-282024-on-certain-data-protection-aspects-related-to_en

3. UK Information Commissioner's Office. AI and data protection risk toolkit. Accessed September 2026. 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/

4. United States Copyright Office. Copyright and Artificial Intelligence, Part 3: Generative AI Training, pre-publication version. May 2025. https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-3-Generative-AI-Training-Report-Pre-Publication-Version.pdf

5. European Commission. General-purpose AI obligations under the AI Act. Updated 2026. https://digital-strategy.ec.europa.eu/en/factpages/general-purpose-ai-obligations-under-ai-act

6. United States Copyright Office. Copyright and Artificial Intelligence, Part 2: Copyrightability. January 2025. https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf

7. United States Federal Trade Commission. Partnerships Between Cloud Service Providers and AI Developers. January 2025. https://www.ftc.gov/system/files/ftc_gov/pdf/p246201_aipartnerships6breport_redacted.pdf

8. UK Competition and Markets Authority. AI Foundation Models: Technical Update Report. 11 April 2024. https://assets.publishing.service.gov.uk/media/661e5a4c7469198185bd3d62/AI_Foundation_Models_technical_update_report.pdf

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

10. European Commission. Guidelines on obligations for General-Purpose AI providers. Updated 2026. https://digital-strategy.ec.europa.eu/en/faqs/guidelines-obligations-general-purpose-ai-providers

11. IFRS Foundation. IFRS 3 Business Combinations. 2022 issued standards. https://www.ifrs.org/content/dam/ifrs/publications/pdf-standards/english/2022/issued/part-a/ifrs-3-business-combinations.pdf

12. IFRS Foundation. IAS 38 Intangible Assets. 2022 issued standards. https://www.ifrs.org/content/dam/ifrs/publications/pdf-standards/english/2022/issued/part-a/ias-38-intangible-assets.pdf

13. IFRS Foundation. IAS 36 Impairment of Assets. 2022 issued standards. https://www.ifrs.org/content/dam/ifrs/publications/pdf-standards/english/2022/issued/part-a/ias-36-impairment-of-assets.pdf

14. United States Department of Justice and Federal Trade Commission. 2023 Merger Guidelines. 18 December 2023. https://www.justice.gov/atr/media/1329301/dl

15. United States Federal Trade Commission. FTC Issues Staff Report on AI Partnerships and Investments Study. 17 January 2025. https://www.ftc.gov/news-events/news/press-releases/2025/01/ftc-issues-staff-report-ai-partnerships-investments-study

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

17. European Commission. Guidelines for providers of general-purpose AI models. Updated 28 April 2026. https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers

18. European Commission. General-Purpose AI Code of Practice. 2025. https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai

19. National Institute of Standards and Technology. AI RMF Playbook. Updated June 2026. https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

20. International Organization for Standardization. ISO/IEC 42001:2023 Artificial intelligence management system. 2023. https://www.iso.org/standard/81230.html

21. United States Securities and Exchange Commission. SEC Charges Two Investment Advisers with Making False and Misleading Statements About Their Use of Artificial Intelligence. 18 March 2024. https://www.sec.gov/newsroom/press-releases/2024-36

22. OECD. Artificial Intelligence, Data and Competition. 2024. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/11/artificial-intelligence-data-and-competition_7d437898/e7e88884-en.pdf

23. European Union Agency for Cybersecurity. Securing Machine Learning Algorithms. December 2021. https://www.enisa.europa.eu/publications/securing-machine-learning-algorithms

24. UK National Cyber Security Centre. Guidelines for secure AI system development. November 2023. https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development

25. World Intellectual Property Organization. WIPO Conversation on Intellectual Property and Artificial Intelligence. Accessed September 2026. https://www.wipo.int/about-ip/en/frontier_technologies/ai_and_ip.html

About the Author

Authored by Chennakeshav Adya

Independent Researcher

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

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

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

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