

Data Algorithms and Market Power in AI Enabled Acquisitions

A competition diligence framework for data compute models distribution and ecosystem control

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

An acquisition involving artificial intelligence can change competitive conditions through assets and relationships that ordinary market-share analysis does not fully capture. Control over proprietary data, computing capacity, specialised talent, model access, application programming interfaces, distribution channels and product defaults can reinforce one another. A transaction may also alter whether a target can supply several ecosystems, whether customers can switch or combine models, and whether an incumbent can use information or ranking control to favour its own services.

This paper develops a transaction-specific competition diligence framework for AI-enabled acquisitions. It maps the AI value chain, identifies the decision-relevant units of control, tests alternative theories of harm and connects regulatory evidence to valuation, financing, transaction terms and remedy design. The framework examines data quality and rights, compute and cloud dependencies, model access, interoperability, defaults, self-preferencing, network effects, algorithmic conduct, talent and intellectual property. It also distinguishes observed evidence from management estimates and scenario assumptions. Current materials from competition authorities in the United States, European Union, United Kingdom, France and Canada, together with OECD research, provide the public-policy foundation. Applicable law and market facts require current transaction-specific advice.

The worked case is wholly hypothetical. A global business-software ecosystem proposes to acquire an AI workflow platform for an assumed enterprise value of USD 3.60 billion. The target has USD 420 million of annual recurring revenue, three model-provider relationships, a USD 310 million committed cloud-and-compute obligation and 54 competition diligence questions. Management estimates USD 480 million of gross synergy present value. A downside scenario assumes USD 42 million of incremental interoperability, data-separation and multi-cloud investment. Every company, amount, market position, probability and outcome is a scenario assumption.

JEL Classification: G34, K21, L13, L40, O33

Keywords: artificial intelligence, merger control, competition diligence, data advantage, compute access, interoperability, self-preferencing, ecosystems, algorithms, market power

1. Make competition diligence part of the investment decision

Competition diligence for an AI-enabled acquisition should determine whether the transaction thesis depends on access, control or conduct that may be difficult to preserve after regulatory review. The board needs more than a filing timetable. It needs to understand how the target competes, which inputs make that competition possible, which counterparties can constrain the combined firm and which transaction terms could change those conditions.

The central decision is whether the acquirer can own the target and still deliver the value case under credible competition constraints. That requires a linked assessment of market definition, competitive significance,

transaction structure, data and compute access, distribution, interoperability, customer switching, innovation and remedies. The diligence output should influence price, financing headroom, long-stop dates, covenants, risk allocation and the design of integration plans.

The work should begin before exclusivity where feasible. Early evidence can reveal whether a headline acquisition is economically closer to a vertical integration, platform acquisition, ecosystem extension, input foreclosure, nascent-competitor acquisition or transfer of critical talent and intellectual property. A transaction can contain several of these features at once.

The board should receive four connected outputs. First, an AI-stack map should show the target's dependencies and sources of leverage. Second, a competition evidence book should separate observed facts, management estimates and transaction scenarios. Third, a theory-of-harm matrix should link each concern to evidence, valuation exposure and possible mitigations. Fourth, a decision roadmap should show when the board must price, restructure, remedy or abandon the transaction.

2. Use current authority concerns as diligence hypotheses

The US Department of Justice and Federal Trade Commission state in the 2023 Merger Guidelines that platform mergers can affect competition between platforms, on a platform and to displace a platform. Their guidance identifies network effects, conflicts of interest and the acquisition of data used for matching, sorting or prediction as relevant mechanisms. [1] The same guidelines address transactions that entrench or extend a dominant position by denying rivals access to customers, scale or network effects. [2]

The FTC's January 2025 staff report on cloud-provider investments and partnerships with AI developers describes contractual and economic features that can affect competition. These include equity and revenue-sharing rights, consultation and exclusivity provisions, commitments to spend investment proceeds on the partner's cloud services, access to compute and the sharing of technical, financial and customer information. The report identifies possible switching costs and constraints on access to key inputs. [3]

The European Commission's 2024 competition-policy brief maps generative AI across chips, compute, data, models and downstream services. It considers access to key inputs, vertical integration, partnerships and the possibility that firms can leverage positions across the AI value chain. [4] A joint statement by the European Commission, UK Competition and Markets Authority, US Department of Justice and FTC identifies concentrated control of chips, compute, data and expertise; the extension of market power; and arrangements among key players as areas requiring vigilance. [5]

The CMA's foundation-model update identifies risks involving control of critical inputs, powerful incumbents shaping distribution and partnerships reinforcing existing positions. [6] France's Autorité de la concurrence has separately examined the competitive functioning of the generative-AI sector. [7] Canada's Competition Bureau describes data, compute, algorithms, pricing and merger effects within the AI production chain. [8] These public materials are hypotheses for diligence, not conclusions about a specific transaction.

Jurisdictional assessment should run alongside substantive analysis. Notification thresholds, control tests, call-in powers and theories differ. Minority rights, employee transfers, licences and commercial arrangements can receive different treatment across regimes. The transaction team should document the basis for each filing decision, monitor changes to the deal package and reconcile legal control with the practical economic relationship.

AI-specific regulation can change the operating baseline without resolving merger questions. Data protection, intellectual-property, sector regulation, the EU AI Act and the Digital Markets Act can affect access, documentation, technical design and cost. Competition diligence should record these interactions because a regulatory obligation can reduce a claimed advantage, create implementation expenditure or shape a workable remedy.

Table 1. Public competition concerns translated into transaction evidence

Public concern	Transaction question	Evidence to obtain	Deal implication
Control of key inputs	Can rivals obtain suitable data, compute, chips, models and expertise on workable terms?	Contracts, capacity records, rights schedules, technical benchmarks and alternative-supply evidence	Valuation, conditions, access commitments and remedy feasibility
Platform entrenchment	Does the deal remove a participant, interoperability tool or route that constrains the platform?	Usage flows, multi-homing, switching records, API dependencies and roadmap documents	Market definition, competitive significance and perimeter
Self-preferencing	Can the combined firm rank, bundle, default or technically favour its own AI service?	Ranking rules, default settings, product telemetry, incentives and governance	Conduct controls, separation design and integration limits
Partnership control	Do investment, governance, commercial and information rights create influence beyond the equity percentage?	Board rights, vetoes, revenue sharing, exclusivity, spend commitments and information access	Jurisdiction, transaction structure and control analysis
Dynamic competition	Would the target have expanded, partnered or displaced an incumbent without the transaction?	Funding plans, product roadmaps, pipeline, hiring, experiments and board materials	Counterfactual, price, risk allocation and abandon threshold

This table summarises selected public materials for management use. It does not state the complete legal test in any jurisdiction.

3. Map the complete AI value chain

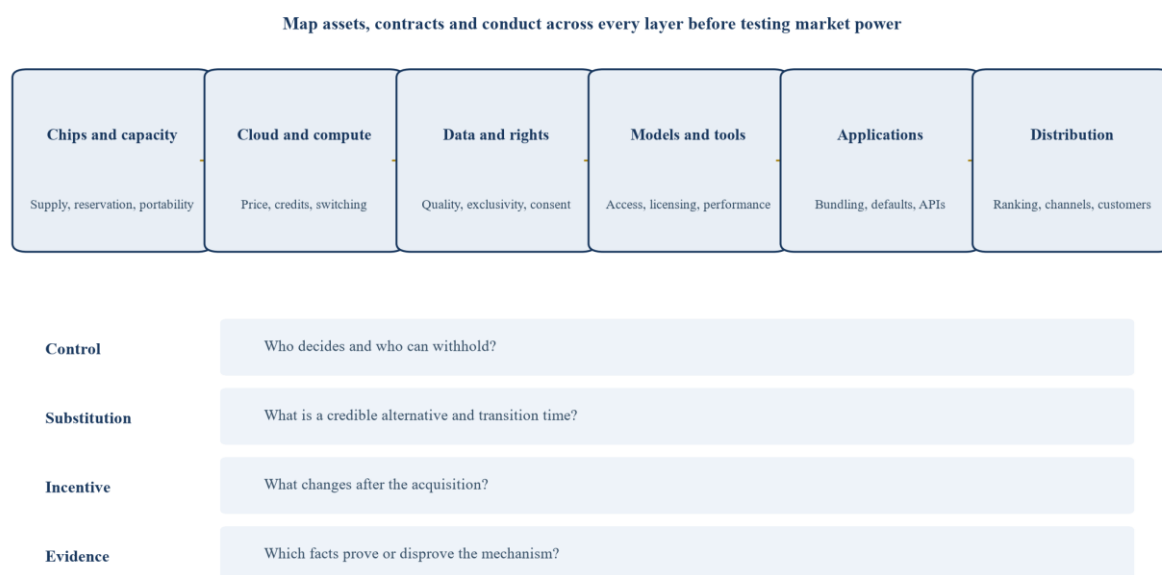
An AI product is usually supported by several layers. Relevant layers may include semiconductor design and fabrication, accelerators, data-centre capacity, cloud infrastructure, training data, data engineering, foundation models, fine-tuning tools, retrieval systems, application orchestration, distribution and customer-facing services. The transaction team should identify where the parties operate and where they depend on third parties.

The map should record control, rights and substitutability at each layer. Ownership is only one form of control. A long-term capacity reservation, preferred access right, exclusive model licence, minimum-spend commitment, technical dependency or default distribution position can shape competitive behaviour. The map should also show which assets can be replicated, how long replication would take and which legal, financial or technical barriers apply.

The value chain should be drawn for current products and credible future products. AI markets can change quickly, while infrastructure, data rights and customer workflows may remain sticky. A target that currently sells an application may be developing a model, agent platform or data network that changes its competitive role. Board materials, engineering plans, hiring requests and customer pilots can be more informative than current revenue classifications.

The team should avoid treating every layer as a market. The map is an evidence tool. Market definition and competitive assessment require legal and economic analysis. The map helps the team identify where substitution evidence, customer views, cost structure and control rights need deeper testing.

Figure 1. AI transaction stack and competition diligence map



The diagram is a proposed evidence map. Transaction-specific analysis determines the relevant markets and theories.

4. Define the transaction perimeter through control and dependence

AI transactions can transfer competitive capability without a conventional purchase of all shares. The perimeter may include minority investments, board or observer rights, commercial agreements, licences, capacity commitments, preferred distribution, revenue sharing, employee transfers and rights over future technology. The team should analyse the package as a whole.

The CMA's Microsoft and Inflection decision treated the hiring of almost all of a specialist team together with related arrangements as a reviewable transaction, although it cleared the case on competition grounds. [9] The European Commission's policy brief similarly discusses partnerships and transfers of talent and assets as potential structural changes requiring scrutiny. [4] These examples show why legal form alone cannot define the diligence perimeter.

The rights schedule should capture affirmative rights, vetoes, consultation rights, information rights, budget influence, commercial dependencies and termination consequences. The team should identify whether rights change on milestones, financing events, performance failures or future funding rounds. A right that appears protective in isolation can contribute to influence when combined with other arrangements.

Commercial dependence requires equal attention. If an AI developer must spend most of an investment on one cloud provider, uses provider-specific tooling and faces material data-egress or model-migration costs, the relationship may be difficult to unwind. Diligence should quantify the cost, time, service risk and contractual consequences of moving to a second provider.

5. Form theories before collecting large data sets

A theory of harm is a testable account of how the transaction could weaken competition. The transaction team should state the mechanism, affected customers, relevant assets, expected conduct, timing and evidence that would confirm or reject the concern. Broad labels such as data advantage or ecosystem power are insufficient.

A horizontal theory may ask whether the target is a current or future constraint on the acquirer's AI product. A vertical theory may ask whether control of a model, data source, compute capacity or distribution channel enables foreclosure. An ecosystem theory may ask whether the acquisition increases switching costs or

weakens multi-homing. A conglomerate theory may examine bundling, defaults or privileged integration across complementary products. A labour or innovation theory may ask whether specialist talent, research capability or product optionality is removed from independent competition.

Each theory should identify a counterfactual. The relevant question is what competitive conditions would probably exist without the transaction. Alternatives may include continued independent growth, a new funding round, a partnership with several providers, acquisition by another buyer or commercial failure. Evidence should include contemporaneous plans and constraints rather than a narrative created solely for the filing.

Customer evidence should match the mechanism. A foreclosure theory may require evidence about alternatives, switching cost, required features and expected response to degraded access. A nascent-competition theory may require evidence about likely adoption, product differentiation and the target's ability to scale. Surveys, interviews, bidding records and product telemetry should be designed or interpreted with attention to selection, timing and incentives.

The theory matrix should drive requests. It reduces collection burden and improves evidential discipline. Every major request should connect to a board decision, regulatory question or valuation sensitivity. The team should record what would disprove its initial concern and search for that evidence deliberately.

6. Test data advantage through quality rights and replicability

Data creates competitive value when it improves a product or decision in a way that rivals cannot readily reproduce. Volume alone is an incomplete measure. Diligence should examine relevance, quality, freshness, labelling, uniqueness, legal rights, feedback loops, coverage and the marginal performance gained from additional data.

The data inventory should identify source, owner, permitted use, geography, retention, consent, exclusivity, portability and deletion requirements. It should distinguish data used for training, fine-tuning, retrieval, evaluation, safety, personalisation and product analytics. The team should reconcile contractual rights with actual pipelines and model documentation.

Replicability should be tested with time and cost. Public or licensed data may be available while cleaning, annotation, domain expertise and feedback histories remain difficult to reproduce. Synthetic data may solve some gaps and create different quality or legal risks. Customer-specific data can improve service without being transferable across customers. The team should avoid assuming that access automatically produces durable market power.

The transaction model should show whether combining datasets changes quality, cost or scope. It should also examine whether the combined firm could restrict rivals' access to a shared source, degrade portability or use competitively sensitive platform data to favour its own services. EU Digital Markets Act provisions on data portability and access provide a current regulatory context for these questions. [10]

Model-performance tests should be reproducible. The diligence team should preserve the evaluation dataset, task definition, baseline, model version, parameters and confidence intervals where appropriate. Improvements that disappear on a different benchmark or customer segment may have limited competitive significance. Technical specialists and economists should agree how performance evidence connects to customer choice and economic value.

Table 2. Data advantage diligence register

Diligence field	Core test	Quantitative evidence	Contractual or technical evidence
Source and rights	Can the target lawfully use and transfer the data for the stated purpose?	Records, jurisdictions, consent coverage and expiry	Licences, notices, customer terms and data maps

Diligence field	Core test	Quantitative evidence	Contractual or technical evidence
Quality and relevance	Does the data improve the model or service for the relevant task?	Error rates, lift, coverage, freshness and benchmark results	Evaluation method, lineage and model cards
Replicability	Can a rival assemble a suitable alternative within a competitive period?	Acquisition cost, labelling cost, time and performance gap	Supplier availability, exclusivity and switching plan
Feedback loop	Does product use create data that improves future performance?	User events, learning cadence, marginal improvement and attrition	Telemetry design, consent and retraining controls
Portability	Can customers export data and continue service elsewhere?	Export success, time, cost and lost functionality	API documentation, formats and migration terms
Foreclosure ability	Could the combined firm deny, delay or degrade rival access?	Rival dependence, capacity and revenue exposure	Access rights, governance, ranking and release controls

Evidence ratings should be assigned from verified transaction materials. The examples below describe the required fields.

7. Quantify compute cloud and chip dependencies

Compute capacity can be both a cost and a competitive constraint. Diligence should separate model training, fine-tuning, inference, evaluation and peak-demand requirements. It should identify accelerators, locations, providers, reservation terms, pricing, energy constraints, data residency and the portability of workloads.

Cloud credits and strategic investments can obscure economic cost. The team should calculate effective unit cost after credits, committed spend, revenue sharing and minimum usage. It should determine what happens when credits expire or model demand changes. A transaction that increases utilisation may improve economics, while concentration on one provider can increase switching exposure.

Technical portability requires more than contractual freedom. Models may depend on proprietary chips, orchestration, storage, networking, security, monitoring and developer tools. Replatforming can require engineering work, new validation and customer approval. The diligence model should show migration time, duplicate-run cost, performance change and service risk.

Capacity evidence should be forward-looking and bounded. Reserved accelerator capacity may exceed current usage because it supports planned training or redundancy. The team should reconcile reservations, actual consumption, queue times, utilisation, failed jobs and demand forecasts. It should test whether a rival can acquire an economically and technically suitable substitute within the period relevant to competition.

The FTC staff report identifies cloud commitments, input access and switching costs as possible competition issues in AI partnerships. [3] The joint international statement similarly identifies specialised chips and substantial compute as critical inputs. [5] The board should know whether the investment thesis assumes continued preferential access and whether a remedy or supplier change would reduce value.

8. Examine model access and interoperability

Model access determines whether applications and customers can choose, combine or replace AI capabilities. The team should inventory proprietary models, third-party models, open-weight components, fine-tunes, embeddings, retrieval systems and evaluation assets. It should record licences, use restrictions, portability, termination rights and technical dependencies.

Interoperability should be tested at several levels. Data interoperability concerns formats and transfer. Model interoperability concerns interfaces, prompts, context, tools and evaluation. Application interoperability

concerns identity, permissions, workflows and audit trails. Distribution interoperability concerns access to operating-system, device or platform functions on terms comparable with the incumbent's own service.

The European Commission adopted binding specification measures in July 2026 concerning access by competing AI services to relevant Android features and access by third-party search providers to search data. [11] This current action illustrates that interoperability and data access are operational competition questions. Transaction diligence should identify which interfaces are essential, who controls change and how the combined firm could affect quality, timing or access.

The integration plan should preserve credible multi-model and multi-provider options until the competition assessment is complete. Removing adapters, ending licences or redesigning APIs too early can narrow remedy choices and undermine the counterfactual evidence.

Interface governance should be included in diligence. The team should identify version policies, deprecation periods, rate limits, authentication, service levels, developer support and access to testing environments. A formally open interface can remain commercially ineffective if changes are unpredictable, documentation is incomplete or performance differs from the incumbent's internal access.

9. Test distribution defaults and self preferencing

Distribution can convert a capable AI product into a competitive constraint. Relevant channels may include operating systems, cloud marketplaces, productivity suites, enterprise applications, browsers, app stores, devices, resellers and embedded workflows. The team should quantify traffic, activation, conversion, usage and customer acquisition by channel.

Defaults matter when users face search, learning, migration or contractual costs. A default can be explicit, such as a preselected assistant, or technical, such as privileged access to device functions. Bundling can lower customer cost and improve integration; it can also make an independent rival harder to discover or purchase. Diligence should measure both effects.

Self-preferencing tests should examine ranking, recommendation, latency, feature access, data access, pricing, eligibility and commercial incentives. The question is whether the combined firm would have the ability and incentive to favour its own service and whether customers or rivals could respond. The evidence should include experiments, ranking documentation, internal forecasts and historical treatment of third parties.

The EU's 2026 DMA review identifies interoperability, self-preferencing, access to data and cloud dependencies as themes relating to AI services. [12] These obligations do not replace merger analysis. They can affect the factual baseline, remedy options and valuation of distribution advantages.

The distribution analysis should distinguish reach from incremental power. A large installed base can accelerate adoption, while customers may retain meaningful choice if rival services are easy to install, integrate and purchase. The team should examine placement, prompts, defaults, contractual bundles, discount allocation and the technical friction of activating an alternative.

10. Measure ecosystem and network effects

An ecosystem can create value by reducing integration cost and improving the experience across complementary products. It can also increase dependence when identity, data, workflows, developers and purchasing are tied together. Diligence should map the customer journey and determine which elements make switching or multi-homing difficult.

Direct network effects arise when a service becomes more valuable as more similar users join. Indirect effects arise when participation on one side attracts another side, such as developers, sellers or data providers. Learning effects can arise when use improves performance. Scale economies can reduce unit cost. These mechanisms should be measured separately because they create different competitive and valuation implications.

The 2023 US Merger Guidelines specifically address platform participation, network effects and conflicts between a platform operator and participants. [1] The diligence team should ask whether acquiring the target removes a tool that enables multi-homing, a participant important to a rival platform or an emerging route that could displace an incumbent service.

Switching evidence should include customer interviews, churn reasons, migration projects, export success, retraining, integration dependencies and contractual terms. Stated willingness to switch is weaker than observed behaviour. The model should estimate cost and time for defined customer segments rather than one average figure.

Developer and partner incentives can be leading indicators. An acquisition may make third parties more willing to build because the ecosystem becomes larger, or less willing because they expect discrimination or appropriation. Diligence should review partner churn, API usage, investment decisions, roadmap feedback and the treatment of applications that compete with the platform's own products.

11. Review algorithms as conduct and evidence

Algorithms can affect price, ranking, matching, eligibility, capacity, recommendations and access. Diligence should identify decision systems that interact with competitors, suppliers or customers and determine what data, objectives, constraints and human controls govern them.

The team should examine whether a platform algorithm can favour an affiliated product, whether pricing systems respond to competitors in ways that facilitate coordination and whether access rules can degrade rivals. Canada's Competition Bureau notes that AI-powered pricing algorithms have prompted discussion of tacit algorithmic coordination. [8] The transaction assessment should focus on evidence and mechanism rather than assuming that automation itself creates harm.

Model governance records can also provide important evidence. Training objectives, evaluation criteria, feature importance, experiments, incident logs and change approvals can show how a system actually operates. Technical teams should explain these materials in decision language without disclosing unnecessary proprietary information.

The combined-firm design should include accountable owners, auditability and change controls for competition-sensitive algorithms. If the transaction creates a new conflict of interest, governance may need separation, access restrictions, testing or independent oversight.

12. Treat talent intellectual property and acquires as competitive assets

Specialist researchers, engineers, product leaders and data experts can be scarce inputs. The diligence team should identify critical teams, mobility constraints, retention plans, research agendas and dependencies on individual leaders. It should distinguish ordinary hiring from a coordinated transfer of a competitive capability.

Intellectual property diligence should cover model weights, code, patents, trade secrets, datasets, evaluation methods, licences and rights arising from employment or research collaborations. Ownership gaps can weaken value. Broad licences or retained rights can affect exclusivity and the competitive counterfactual.

The Microsoft and Inflection case demonstrates that employee transfers combined with licences and related arrangements can constitute a reviewable merger situation in the UK. [9] Boards should assess substance, assets and continuing business capability when structuring an acquire or strategic licence.

Retention economics should be connected to competition risk. A transaction may require incentives to preserve the target's capability while regulators assess whether that capability should remain independent. The financing model should therefore include retention duration, failed-deal arrangements and restrictions on integration.

13. Evaluate dynamic competition and innovation

Current revenue can understate an AI target's competitive importance. A small firm may possess technology, talent, customer access or an architecture capable of challenging an incumbent. The team should test the probability, timing and scale of independent expansion using contemporaneous evidence.

Relevant evidence includes product roadmaps, experiment results, customer pilots, pipeline, funding plans, hiring, compute reservations, partnership discussions and board decisions. The team should examine technical milestones and commercial constraints. A roadmap should not be treated as inevitable; failure probabilities, funding needs and execution risks belong in the counterfactual.

Innovation analysis should identify overlapping research paths and alternative approaches. The question can be whether the deal removes independent experimentation, reduces incentives to develop an interoperable product or redirects a technology into the incumbent's ecosystem. It can also be whether combining assets accelerates innovation in a way that is verifiable and transaction-specific.

OECD research emphasises uncertainty in AI market development and the importance of dynamic competition, key inputs and linkages across the value chain. [13] The board should use scenarios and decision ranges, not false precision.

Innovation efficiencies require the same discipline as risk claims. The team should identify the complementary assets, development milestones, resources, timing and probability of delivery. It should test whether the benefit requires ownership or could arise through licensing, supply or distribution agreements. Benefits should be linked to customer outcomes and adjusted for integration cost and execution risk.

14. Build a competition evidence room

The evidence room should be organised around theories and decisions. Core folders should cover transaction documents, strategy, market and customers, data, compute, models, distribution, interoperability, algorithms, talent, intellectual property, partnerships, valuation, synergies and regulatory analysis.


Each material item should have a source, date, owner and evidential status. Observed data should remain distinguishable from management estimates and transaction scenarios. Derived analyses should preserve formulas and lineage. Customer evidence should record methodology and avoid selective quotations.

The team should reconcile ordinary-course documents with transaction narratives. Internal documents can use different terminology and market boundaries. The analysis should explain differences rather than rewriting history. Key metrics should reconcile to source systems, audited accounts where relevant and the transaction model.

Requests should be proportionate and controlled. Competition-sensitive information may require clean-team handling. Personal data, source code, model weights and security information require specialised controls. Evidence governance should support the deal decision while preserving legal, confidentiality and cybersecurity requirements.

The evidence room should include a controlled issue log. Each issue should state the question, evidence received, unresolved contradiction, responsible owner, decision consequence and next action. Superseded analyses should remain traceable. This record helps the board understand how conclusions changed and supports consistency across filings, lender materials and transaction documents.

Figure 2. Competition evidence funnel from source to board decision



The funnel is a proposed diligence-control sequence.

15. Define the hypothetical transaction case

The case concerns a global provider of business software, cloud services and enterprise distribution acquiring an AI workflow platform. The assumed enterprise value is USD 3.60 billion. The target has USD 420 million of annual recurring revenue and serves customers through direct sales, cloud marketplaces and integrations with enterprise applications.

The target currently accesses three foundation-model providers. It uses one strategic cloud provider for most training and inference workloads under a USD 310 million committed obligation extending beyond the proposed closing. Its product connects customer documents, workflow data and feedback to domain-specific agents. Customers can export some records, while configurations, evaluation histories and workflow context require additional migration work.

Management estimates USD 480 million of gross synergy present value from distribution, infrastructure, procurement and product integration. The transaction model assumes a 160-day signing-to-closing period. A competition diligence programme identifies 54 questions across data, compute, models, distribution, interoperability, algorithms, talent and the counterfactual.

The downside scenario assumes USD 42 million of incremental investment over three years for multi-cloud capability, data separation, interoperability, auditability and third-party access. The model does not assume that an authority will require these measures. It tests the board's ability to sustain the investment case if constraints affect exclusivity, bundling, defaults or integration timing.

Table 3. Hypothetical transaction assumptions

Item	Scenario assumption	Decision use	Competition implication
Enterprise value	3,600.0	Price and financing	Tests value under remedy and delay scenarios
Target annual recurring revenue	420.0	Scale and valuation	Requires source reconciliation and segment evidence

Item	Scenario assumption	Decision use	Competition implication
Gross synergy present value	480.0	Strategic value case	Separate distribution, cost, product and data components
Cloud and compute commitment	310.0	Liquidity and dependency	Test portability, concentration and termination economics
Downside investment	42.0	Remedy and integration reserve	Funds multi-cloud, separation, interoperability and controls
Signing-to-closing period	160 days	Timetable and financing	Model regulatory review and delayed integration
Model-provider relationships	3	Substitution and resilience	Test rights, performance and credible alternatives
Competition diligence questions	54	Programme scope	Assign evidence owner and decision consequence

Every value and outcome is a scenario assumption. USD millions unless stated otherwise.

16. Model the counterfactual and transaction scenarios

The base counterfactual assumes that the target remains independent, completes a planned funding round and continues multi-model distribution. A second counterfactual assumes a strategic partnership without control. A third assumes acquisition by a buyer without a competing ecosystem. None is treated as fact. The team assigns probabilities from verified funding, governance, technical and commercial evidence.

The transaction scenarios should test full approval, extended review, behavioural constraints, interoperability obligations, data separation, multi-cloud requirements, divestiture of selected assets and transaction failure. Each scenario should show probability, timing, revenue, cost, capex, customer impact, financing carry and implementation consequences.

Synergies need mechanism-level evidence. Distribution synergy should identify channel, customer segment, conversion rate and cannibalisation. Infrastructure synergy should account for existing commitments and migration cost. Product synergy should identify development dependency and release timing. Data synergy should show rights, technical method and measurable performance improvement.

The board should focus on value at risk by theory. A constraint on default distribution may affect revenue synergy. A multi-cloud requirement may affect cost. Data separation may delay product integration. A prohibition or divestiture can change the transaction perimeter. The decision model should avoid one undifferentiated regulatory discount.

17. Allocate review effort through a risk heat map


The heat map should combine ability, incentive, effect, evidence quality and remedy complexity. High exposure with weak evidence requires urgent work. High exposure with strong evidence still requires board attention because the transaction may need repricing or restructuring.

Data exclusivity can carry high exposure where the dataset is important, non-replicable and transferable. Compute concentration can carry high exposure where capacity alternatives are unavailable within the required period. Self-preferencing can carry high exposure where the acquirer controls a major route to market and can change ranking, defaults or feature access.

The map should include positive evidence. Credible customer switching, interoperable APIs, available capacity, multi-homing, low incremental data value and independent governance can reduce a concern. Evidence should be tested, not asserted.

The heat map is a management tool. It does not predict an authority's decision. Its purpose is to allocate diligence, valuation and negotiation resources before optionality disappears.

Figure 3. Hypothetical AI acquisition competition risk heat map



Scores are scenario assumptions for the worked case and are not regulatory findings.

18. Design remedy readiness before it is requested

Remedy readiness protects transaction optionality. The team should identify which concerns could be addressed through perimeter change, access, interoperability, data separation, licensing, supply commitments, governance or monitoring. It should test whether a measure is technically feasible, commercially sustainable and capable of verification.

Structural options may include divesting a product, dataset, model, team, contract or distribution right. A viable package needs assets, people, rights, customers and transition support sufficient to compete. Behavioural options may address access, ranking, interoperability, portability, licensing, capacity or information flows. Their effectiveness depends on clarity, incentives and monitoring.

The team should calculate cost and value leakage for each option. It should identify tax, financing, intellectual-property, privacy, cybersecurity and operational consequences. A remedy can affect both the target and the acquirer's wider ecosystem.

Integration design should preserve separability. Product roadmaps, data architectures, contracts and teams should retain clear boundaries until risk is resolved. This does not imply that a remedy will be required. It gives the board executable choices and improves negotiation credibility.

Remedy testing should include customers and counterparties where legally and commercially appropriate. A technically elegant access commitment may fail if customers cannot use it, rivals lack complementary assets or monitoring requires information that cannot be verified. The team should define success metrics, reporting sources, dispute routes and the consequences of non-compliance.

19. Connect competition findings to valuation financing and terms

Competition risk affects cash flows, timing, financing and optionality. The valuation model should link each material theory to revenue, margin, capex, working capital, tax and terminal value. It should show the value of synergies that depend on exclusivity, default placement, data combination or rapid integration.

Financing models should include review delay, ticking fees, bridge duration, commitment expiry, hedging, retention and duplicated operating cost. Lenders may require information about regulatory conditions and remedy exposure. The team should control sensitive information while answering legitimate underwriting questions.

Transaction terms can allocate risk through conditions, covenants, efforts standards, remedy commitments, termination rights, reverse termination fees, long-stop dates and information obligations. The economic model should inform those provisions. A buyer should understand the maximum remedy burden consistent with its investment case before negotiating a broad commitment.

The board paper should show the walk-away threshold and who can approve changes. A higher probability of approval does not compensate for an unbounded downside if the transaction terms require remedies that destroy the thesis.

20. Implement a ninety day competition diligence roadmap

The first ten days should establish the decision, transaction perimeter, jurisdiction map, AI-stack map, evidence taxonomy and clean-team controls. The team should identify the leading theories and assign owners.

Days 11 to 30 should reconcile data, compute, model, distribution and partnership evidence. Customer and technical work should begin. The valuation team should create theory-specific sensitivities rather than wait for the legal analysis to finish.

Days 31 to 60 should test counterfactuals, switching, interoperability, foreclosure incentives, dynamic competition and remedy feasibility. The team should challenge both adverse and favourable hypotheses. Material gaps should reach the transaction committee with defined decisions.

Days 61 to 90 should finalise the board decision record, filing evidence, financing cases, term positions and remedy-readiness pack. The programme should continue after signing as new authority questions, market developments and transaction changes arise.

Table 4. Ninety day competition diligence roadmap

Period	Required output	Accountable owner	Verification
Days 1 to 10	Transaction perimeter, stack map, theories, evidence taxonomy and controls	Transaction executive and competition lead	Board-approved scope and source register
Days 11 to 30	Data, compute, model, distribution and partnership evidence	Workstream owners and clean team	Reconciled metrics, contracts and technical evidence
Days 31 to 60	Counterfactual, customer, switching, incentive and remedy tests	Economist, technical lead and valuation lead	Challenge sessions and sensitivity model
Days 61 to 90	Board decision, filing evidence, deal terms and remedy-readiness pack	Transaction committee	Signed decision record and indexed evidence room
Monthly thereafter	Market update, evidence refresh and scenario reforecast	Competition lead and finance lead	Committee minutes and variance report

Period	Required output	Accountable owner	Verification
Before closing	Final conditions, continuing obligations and release gates	Legal control owner and integration lead	Closing certificate and controlled implementation plan

Timing should be adapted to the actual transaction process and authority timetable.

21. Give the board a compact decision dashboard


The dashboard should show the leading theories, evidence quality, value at risk, timetable, contractual exposure and next decision. It should avoid a single red-amber-green score that conceals uncertainty. Each theory should have an owner, current conclusion, contrary evidence and deadline.

The board needs separate views of regulatory probability and economic consequence. A low-probability issue can warrant action where the consequence is prohibition or a remedy that destroys the thesis. A likely issue with limited value impact may be manageable through price or implementation design.

The dashboard should identify assumptions nearing expiry. Model performance, customer behaviour, capacity availability, product roadmaps and regulatory positions can change during a long process. Evidence should have refresh dates.

Readiness should also be visible. The board should know whether the filing narrative reconciles to transaction documents, whether the model includes remedy cases, whether financing remains available and whether the integration plan preserves required boundaries.

Figure 4. Hypothetical board decision dashboard



Values are scenario assumptions and do not predict an authority decision.

22. Use AI competition diligence to preserve decision quality

AI-enabled acquisitions require an integrated view of assets, contracts, technical architecture, customer behaviour and market incentives. A conventional market-share screen can miss how data, compute, models, distribution and ecosystem control reinforce one another.

The framework begins with the board's investment decision and maps the full AI stack. It defines testable theories, builds a controlled evidence room and connects findings to valuation, financing, terms, remedies and

integration. It treats rights, dependencies and technical interfaces as economic facts rather than background detail.

The hypothetical case shows how an assumed USD 3.60 billion acquisition can be tested through 54 questions and specific downside investments. The numbers do not forecast a live transaction. They demonstrate how the board can identify which parts of USD 480 million of assumed synergy depend on contested forms of access or integration.

Current public materials show sustained authority attention to AI inputs, partnerships, platform power, interoperability, self-preferencing and dynamic competition. That attention should translate into earlier evidence collection and clearer decision thresholds. Teams should update the analysis as law, technology, market structure and transaction facts change.

This paper provides a management framework. It does not determine the legal or regulatory treatment of any acquisition. Parties should obtain current competition, corporate, data-protection, intellectual-property, employment, financing and sector-specific advice in every relevant jurisdiction.

Frequently asked questions

Why is ordinary merger diligence insufficient for some AI acquisitions?

AI competition can depend on data rights, compute access, model interfaces, distribution, defaults, feedback loops and specialist talent. These mechanisms may not be visible in revenue shares or conventional product categories.

What is the first competition diligence deliverable?

Create an AI-stack map showing where each party operates, which inputs and interfaces it controls, which third parties constrain it and how the transaction changes rights, incentives or dependencies.

How should a buyer assess a claimed data advantage?

Test relevance, quality, freshness, rights, exclusivity, replicability, feedback effects and measurable performance lift. Data volume without these attributes does not establish a durable advantage.

Why do cloud commitments matter?

Capacity reservations, credits, minimum spend, proprietary tools and migration costs can affect effective economics, switching and access to key inputs. The team should model their cost and portability.

What is self preferencing in an AI transaction?

It can arise when a platform controls ranking, defaults, feature access, data, eligibility or pricing and can use those controls to favour an affiliated AI service. Ability, incentive and customer response require evidence.

How should the counterfactual be built?

Use contemporaneous funding, product, customer, hiring, partnership and technical evidence to test credible paths without the transaction. Assign uncertainty rather than presenting one future as certain.

When should remedy readiness begin?

Begin when a credible theory could affect perimeter, access, interoperability, data, distribution or integration. Early testing preserves executable choices and informs price and transaction terms.

What should the board retain?

Retain the stack map, theory matrix, source register, verified metrics, counterfactual evidence, customer and technical analyses, valuation sensitivities, term decisions, remedy tests, minutes and final decision record.

References

- [1] US Department of Justice and Federal Trade Commission, 2023 Merger Guidelines, Guideline 9 on Multi-Sided Platforms, 18 December 2023. <https://www.justice.gov/atr/merger-guidelines/applying-merger-guidelines/guideline-9>
- [2] US Department of Justice and Federal Trade Commission, 2023 Merger Guidelines, Guideline 6 on Entrenching or Extending a Dominant Position, 18 December 2023. <https://www.justice.gov/atr/merger-guidelines/applying-merger-guidelines/guideline-6>
- [3] US Federal Trade Commission, FTC Staff Report on AI Partnerships and Investments 6(b) Study, January 2025. <https://www.ftc.gov/reports/ftc-staff-report-ai-partnerships-investments-6b-study>
- [4] European Commission Directorate-General for Competition, Competition in Generative AI and Virtual Worlds, Competition Policy Brief 3/2024, 19 September 2024. https://competition-policy.ec.europa.eu/document/download/c86d461f-062e-4dde-a662-15228d6ca385_en
- [5] European Commission, UK Competition and Markets Authority, US Department of Justice and US Federal Trade Commission, Joint Statement on Competition in Generative AI Foundation Models and AI Products, 23 July 2024. https://competition-policy.ec.europa.eu/document/download/79948846-4605-4c3a-94a6-044e344acc33_en
- [6] UK Competition and Markets Authority, AI Foundation Models Update Paper and Technical Update Report, April 2024. <https://www.gov.uk/government/publications/ai-foundation-models-update-paper>
- [7] Autorité de la concurrence, Opinion on the Competitive Functioning of the Generative Artificial Intelligence Sector, 28 June 2024. <https://www.autoritedelaconcurrence.fr/en/publications/opinion-competitive-functioning-generative-artificial-intelligence-sector>
- [8] Competition Bureau Canada, Artificial Intelligence and Competition, 20 March 2024. <https://competition-bureau.canada.ca/en/how-we-foster-competition/education-and-outreach/artificial-intelligence-and-competition>
- [9] UK Competition and Markets Authority, Microsoft and Inflection Inquiry, Full Text Decision, 24 October 2024. <https://www.gov.uk/cma-cases/microsoft-slash-inflection-ai-inquiry>
- [10] European Commission, Digital Markets Act Developer Portal, End User Data Portability. https://digital-markets-act.ec.europa.eu/developer-portal/end-user-data-portability_en
- [11] European Commission, Guidance to Google for AI Interoperability on Android and Sharing of Google Search Data under the Digital Markets Act, 16 July 2026. https://digital-markets-act.ec.europa.eu/commission-provides-guidance-google-ai-interoperability-android-and-sharing-google-search-data-under-2026-07-16_en
- [12] European Commission, Digital Markets Act Review Questions and Answers, 2026. https://digital-markets-act.ec.europa.eu/about-dma/dma-review-qa_en
- [13] OECD, Artificial Intelligence Data and Competition, OECD Artificial Intelligence Papers No. 18, 24 May 2024. <https://doi.org/10.1787/e7e88884-en>
- [14] European Commission, Interoperability under the Digital Markets Act. https://digital-markets-act.ec.europa.eu/developer-portal/interoperability_en
- [15] European Commission, Virtual Worlds and Generative AI Competition Consultation and Workshop. https://competition-policy.ec.europa.eu/about/reaching-out/virtual-worlds-and-generative-ai_en
- [16] UK Competition and Markets Authority, CMA AI Strategic Update, 29 April 2024. <https://www.gov.uk/government/publications/cma-ai-strategic-update/cma-ai-strategic-update>
- [17] US Department of Justice and Federal Trade Commission, 2023 Merger Guidelines, complete official publication. <https://www.justice.gov/atr/merger-guidelines>
- [18] European Commission, Merger Control Procedures. https://competition-policy.ec.europa.eu/mergers/procedures_en

About the Author

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

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 and leads 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.

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

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.