

The AI Capability Balance Sheet: Identifying Which Data, Workflows and Talent Create Defensible Enterprise Value

A Board Framework for Evidence, Control, Economics and Strategic Investment

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Companies are investing in models, data platforms, software, cloud infrastructure, training and workflow redesign. The resulting capabilities often sit across accounting categories and organisational boundaries. Some expenditure is recognised as an asset, much is expensed, and important sources of enterprise value such as workforce knowledge, operating routines and customer trust may not qualify for recognition on the financial balance sheet. A board therefore needs a management framework that makes these capabilities visible without describing them as accounting assets or assigning unsupported values.

This paper develops an AI capability balance sheet for corporate boards, investors and management teams. It records six connected capability classes: data and rights; embedded workflows; specialist and distributed talent; models, software and infrastructure; customer access and permission; and governance, assurance and learning. Each entry carries an owner, evidence date, control position, dependency map, economic mechanism and investment requirement. The framework distinguishes capability existence from legal control, technical performance from commercial adoption, gross productivity from retained cash value, and accounting recognition from strategic relevance.

A hypothetical business-services company illustrates the method. The author assumes annual revenue of USD 480 million, EBITDA of USD 62 million and a three-year AI investment ceiling of USD 45 million. The values, scores and outcomes are analytical assumptions. They are not observed company data, forecasts, valuations, transaction terms or probabilities. The worked example shows how an apparently strong collection of models and data can remain economically weak when customer permission, workflow adoption, portability or key-person depth is missing. The central conclusion is that defensible enterprise value arises from controlled combinations of complementary capabilities, supported by evidence that they improve customer outcomes and cash economics over time.

JEL Classification: G24, G32, L21, M15, O32, O33

Keywords: artificial intelligence, intangible assets, proprietary data, workflow advantage, talent, enterprise value, capital allocation, technology strategy, AI governance, due diligence

1. Define the capability balance sheet as a management instrument

The AI capability balance sheet is a governed inventory of resources and operating relationships that may affect future cash flows, risk and strategic choice. It is not a statement prepared under IFRS Accounting Standards, a valuation report or a register of every technology expense. Its purpose is to help a board decide where an organisation has an advantage, where it depends on

suppliers or individuals, which gaps deserve investment, and what evidence a buyer, lender or partner could verify.

The distinction matters because economic capability and accounting recognition serve different purposes. IAS 38 requires an intangible asset to be identifiable, controlled and capable of producing future economic benefits. Internally generated research expenditure is expensed, while development expenditure is recognised only when specified conditions are demonstrated [1]. A skilled team normally does not meet the control test because employees can leave and legal rights may not give the entity sufficient control over the expected benefits [1]. These accounting conclusions do not imply that research, organisational knowledge or talent are commercially irrelevant. They show why a management inventory must not be labelled as a second set of financial statements.

The inventory should connect each capability to a decision. A data set may support a pricing engine, but the relevant questions concern lawful use, quality, coverage, refresh, exclusivity and measurable performance. A model may achieve a benchmark score, but the board needs to know whether the production workflow improves, customers accept the result, unit cost remains attractive and the company can change providers. A team may contain respected specialists, but enterprise value depends on documented methods, succession, incentives and the ability of other employees to operate the system.

The board should approve the perimeter, definitions and evidence threshold. Materiality can be expressed through revenue, gross profit, customer retention, cost, working capital, regulatory exposure, capital need or strategic option value. An entry that cannot change a decision belongs in an operational asset register rather than the board-level capability balance sheet.

2. Start with the economic mechanism

An AI capability can affect value through revenue, cost, capital intensity, risk or strategic options. Revenue effects include improved conversion, retention, price, coverage and new products. Cost effects include labour productivity, lower error, reduced loss, faster cycle time and avoided hiring. Capital effects include compute commitments, data remediation, integration, working capital and acquisition expenditure. Risk effects include privacy, intellectual property, safety, security, discrimination, service continuity and regulatory exposure. Strategic option value can arise from licensing, partnership, acquisition or entry into an adjacent market.

The mechanism should be stated before a score is assigned. For example, a proprietary claims data set may improve fraud detection. Its economic case depends on whether the model reduces paid losses without rejecting valid claims, whether the insurer has rights to use the data for that purpose, whether regulators and customers accept the process, and whether competitors can obtain comparable data. The data set is only one part of the capability.

OECD evidence on firms that actively use AI reports substantial differences in objectives, data maturity, skills and organisational practices. In the surveyed group, 78 percent used at least one data-management solution, and many used external as well as internal data [2]. The study describes AI adoption as a company-level transformation that can require integration across functions rather than the purchase of a standalone tool. These results concern surveyed adopters and should not be projected mechanically to a specific company.

The economic mechanism also defines the counterfactual. Management should compare the proposed capability with continued current operations, a vendor product, a partnership, an acquisition and a smaller process redesign. A capability that works technically may destroy value if a simpler alternative delivers the same customer outcome with less cost and risk.

Figure 1. The AI capability balance sheet connects controlled resources to cash and strategic decisions



Each capability needs evidence of existence, control, use, economics and durability before it supports a capital decision.

3. Build an evidence architecture before assigning value

Every entry should carry an evidence record. The minimum fields are the capability description, business owner, technical owner, legal or rights owner, source systems, evidence date, use cases, affected customers, financial mechanism, principal dependencies, control status, investment required, decision rights and next review date. The record should link to underlying documents rather than repeat management assertions.

Evidence should be ranked by proximity to realised economics. Customer contracts, renewal behaviour, usage and price provide stronger commercial evidence than a demonstration. Controlled production tests with measured quality, cycle time, adoption and total cost provide stronger operating evidence than a laboratory benchmark. Executed licences, employment terms, data-processing records and architecture documentation provide stronger control evidence than an informal description. Forecast benefits remain scenarios until they reconcile to realised transactions and cash.

NIST's AI Risk Management Framework organises risk work around Govern, Map, Measure and Manage [3]. Its Generative AI Profile identifies risks involving confabulation, privacy, information integrity, intellectual property, security, bias and third-party components [4]. A capability record can use that structure without turning a voluntary risk framework into a

valuation method. The governance evidence shows whether the company understands and controls the system; it does not prove that customers will pay.

The evidence architecture should preserve failed tests and superseded versions. Model performance, data distributions, vendor terms and customer behaviour can change. A dated evidence chain allows management to identify which conclusions remain current. It also supports transaction diligence because a buyer can trace the capability from contract and architecture through operating results and financial effect.

Table 1. Minimum evidence record for an AI capability

Evidence field	Question	Primary support	Decision use
Existence	What operates today and at what scale?	Production logs, architecture and workflow observation	Distinguish capability from proposal
Rights and control	What may the company use, modify, restrict or transfer?	Contracts, licences, consent records and IP register	Identify enforceability and transaction constraints
Performance	Does the capability improve the required outcome?	Controlled tests, quality measures and incidents	Set operating and assurance thresholds
Adoption	Do employees and customers use the changed workflow?	Usage, cohort, renewal and training records	Separate deployment from embedded use
Economics	How does the capability change price, volume, cost, cash or capital?	Unit economics, ledgers, contracts and reconciliations	Support capital and valuation scenarios
Durability	How quickly can a competitor copy or a supplier weaken it?	Substitution tests, dependency map and switching exercise	Estimate advantage period and downside

The record supports management and diligence; it does not determine accounting recognition or fair value.

4. Test data as a controlled and productive resource

Data creates a potential advantage when the company can use it lawfully and repeatedly to improve an economically important outcome. Volume alone is insufficient. The board should examine provenance, contractual purpose, privacy status, intellectual property, jurisdiction, quality, coverage, labels, refresh frequency, bias, security, retention and transferability. It should distinguish operational access from rights to train, fine-tune, evaluate, improve or commercialise a model.

The European Data Protection Board's Opinion 28/2024 considers when an AI model may be anonymous, when legitimate interests may support processing and how unlawful processing during development can affect later deployment [5]. It emphasises case-specific assessment. The US Copyright Office's 2025 report on generative AI training analyses copyright questions surrounding the use of protected works [6]. These sources address different legal systems and do not decide the rights position of a particular data set. They illustrate why origin and permitted use must be recorded rather than assumed.

Data quality should be tested against the decision it supports. Accuracy, completeness and timeliness are relevant, but representativeness and outcome labels can be more important. A historical data set may reproduce a process that management intends to change. A customer-service corpus may contain sensitive information, undocumented corrections or language patterns

that do not represent new markets. Synthetic data can fill some gaps while creating its own assumptions and validation needs.

The strongest data assets often arise from an embedded learning loop. The company performs a workflow, observes an outcome, records a high-quality label, improves the system and delivers a better customer result. Defensibility depends on whether others can reproduce that loop through public sources, partnerships, purchased data or customer access. OECD competition work notes that proprietary data can be important while some training sources may be substitutable [7]. The assessment should therefore test uniqueness rather than attach a premium to the word proprietary.

Table 2. Data defensibility assessment

Dimension	Weak position	Stronger position	Evidence required
Rights	Purpose unclear or reuse restricted	Intended use and transfer rights documented	Contracts, notices, licences and legal analysis
Quality	Incomplete, stale or weakly labelled	Current, representative and outcome-linked	Data profile, lineage and validation results
Scarcity	Public or readily purchasable	Difficult to reproduce within the decision horizon	Market scan and substitution test
Learning loop	Data is collected without measured improvement	Use creates labels and verified performance gains	Versioned evaluations and production outcomes
Workflow access	Data is separated from customer action	Data is generated inside a retained customer workflow	Process map, usage and retention evidence
Portability	Locked to one supplier or environment	Export, documentation and migration tested	Technical export and contractual rights

Scores require documented evidence and should be refreshed when rights, use or market alternatives change.

5. Treat embedded workflows as the operating asset

A model produces enterprise value through a workflow. The workflow combines tasks, data, software, decision rights, human review, controls and customer interaction. It converts technical capability into an output that can be priced, delivered and trusted. A company that owns a strong model but cannot integrate it into sales, service or operations may have less defensibility than a company using a widely available model inside a deeply embedded process.


The workflow map should begin with the customer job and the accountable outcome. It should identify each decision, hand-off, data input, exception, control and priced unit. Management can then compare the baseline with the AI-enabled process. Relevant measures include end-to-end cycle time, error, rework, service level, conversion, retention, cost per completed outcome and human-review load. Time saved in one task is not a cash benefit when delays or corrections appear elsewhere.

Embedding has several levels. An experiment is used by a small group under observation. A deployed tool is available in production. An adopted workflow is used by the intended population with measured compliance. An institutionalised capability is supported by process ownership, controls, training, monitoring, customer terms and budget. A transferable capability also has documentation, interfaces and succession that allow it to survive a transaction or personnel change.

Switching costs can support durability when they arise from customer-specific integration, accumulated outcomes, trusted controls and process improvement. They can also indicate

customer harm or supplier lock-in. The UK Competition and Markets Authority identifies compute, data, expertise and routes to market as important inputs in the foundation-model value chain and has examined the effect of partnerships and integrated firms on competition [8]. A board should understand whether its workflow strengthens its own customer relationship or simply deepens dependence on an upstream provider.

Figure 2. Workflow maturity moves from experiment to transferable operating capability



Progress requires commercial, operating and control evidence; technical deployment alone does not establish maturity.

6. Separate specialist talent from organisational capability

AI capability depends on people who can frame business problems, develop or configure systems, manage data, evaluate performance, operate controls and redesign work. Scarce specialists may be essential. A concentration of knowledge in a few individuals also creates key-person risk. The capability balance sheet should record both expertise and the mechanisms that allow the organisation to retain and reproduce it.

The talent record should cover role, critical decisions, replacement time, succession, incentives, employment and IP terms, documentation, teaching capacity, external dependencies and workload. It should distinguish technical credentials from demonstrated production experience. A model researcher, data engineer, product manager, domain expert, security lead and process owner contribute different capabilities. The absence of one can constrain the whole system.

ILO's 2025 global exposure index estimates that one in four workers is in an occupation with some degree of generative AI exposure and concludes that transformation is more likely than wholesale replacement for most affected jobs [9]. This result is an occupational exposure analysis, not a company workforce forecast. It supports task-level planning and cautions against treating headcount reduction as the only source of value.

Organisational capability becomes stronger when critical knowledge is expressed in evaluation sets, design records, runbooks, controls, code review, model cards, data dictionaries, customer playbooks and decision logs. Documentation is useful only when another qualified person can operate or change the system. A succession test should ask a replacement team to explain the architecture, reproduce a release, diagnose an incident and reconnect performance to the business metric.

Talent retention should be connected to the capital case. Equity, retention awards, career paths, training and acquisitions can secure skills, but each has a cost and time horizon. The board should identify which roles need internal control, which can be sourced through partners and which should be shared across business units.

7. Inventory models, software and infrastructure by dependency

The technology layer can include proprietary models, third-party foundation models, open models, prompts, fine-tuning, retrieval systems, data platforms, orchestration, evaluation tools, security controls, application software, cloud services, chips and network capacity. A useful inventory shows how these components combine to deliver each material workflow. A list of vendor names does not reveal operational dependency.

For every component, management should record ownership, licence, version, purpose, service level, data use, cost basis, minimum commitment, capacity, portability, replacement time, concentration, incident history and exit provisions. The company should identify components whose failure stops revenue, breaches a customer obligation or invalidates an evaluation. It should also identify where a vendor can change price, access, model behaviour or terms.

The economics should be measured per completed customer outcome. Token, query or compute costs are inputs. Total cost includes retrieval, storage, networking, integration, testing, human review, monitoring, security, support, vendor management and failed calls. Contribution margin should be tested under usage, price and quality scenarios. Growth can reduce margin when the commercial price is fixed while inference and review scale with usage.

Portability should be tested rather than described. The company can run a documented migration exercise using representative traffic, data and controls. The test should measure time, performance loss, engineering effort, contract constraints and customer impact. Full portability may be uneconomic. The decision record should show the value received for accepting dependence and the contingency if the supplier position changes.

8. Record customer access, permission and trust

Customer relationships connect capability to revenue. The balance sheet should identify which customers use the AI-enabled process, what they were told, whether contracts permit the use, how price is set, what service levels apply and whether customers can audit, restrict or terminate the arrangement. A capability has weak commercial evidence when customers use the underlying service but have not accepted the changed process or paid for the outcome.

Permission can arise through contract, regulation, procurement approval, user choice or established service terms. It can be narrow. A customer may permit automation for internal assistance but prohibit training on its information. A regulated customer may require

explainability, human review, location controls or named subcontractors. Customer trust becomes operational when these requirements are documented, monitored and reflected in renewals.

Commercial evidence should be cohort based. Management should compare adoption, price, conversion, retention, service quality and support cost for customers using the capability against a relevant baseline. Selection effects should be disclosed. Early adopters may be more tolerant or technically capable than the wider customer base. A paid pilot provides stronger evidence than an unpaid trial, but it may still fail to establish scalable demand.

The SEC has acted against firms that made false and misleading claims about AI use [10]. Its public remarks have also stressed that material claims should have a reasonable basis and company-specific risk disclosure [11]. The direct legal application depends on the entity and jurisdiction. The governance lesson is broader: marketing, investor materials, board packs and transaction documents should describe the same capability and limitations.

Table 3. Customer and market evidence ladder

Stage	Evidence	Remaining question	Capital implication
Hypothesis	Interviews and problem definition	Will customers change behaviour?	Fund discovery within a small limit
Trial	Controlled use with defined success measures	Does the workflow perform in context?	Fund testing and required controls
Paid pilot	Contracted customer and measured unit economics	Does value persist outside the pilot?	Release a bounded commercial tranche
Cohort adoption	Representative usage, renewal and support evidence	Can the company scale without margin or risk deterioration?	Fund capacity and operating integration
Retained advantage	Repeated renewals, price support and learning effects	How durable is the advantage against substitution?	Consider scale, partnership or acquisition

Evidence strength increases when use, payment, retention and operating results can be reconciled for representative customers.

9. Make governance and assurance part of the capability

Governance affects whether a company can deploy AI safely, win customers and sustain an advantage. The capability includes policies, inventory, accountability, risk classification, testing, change control, incident response, audit evidence and board oversight. Governance becomes economically relevant when it shortens approval, reduces failure, supports regulated customers or protects continuity.

ISO/IEC 42001 specifies requirements for establishing, implementing, maintaining and improving an AI management system [12]. NIST provides a voluntary risk framework and playbook [3]. The European Union's General-Purpose AI Code of Practice addresses transparency, copyright, safety and security for providers within its scope [13]. These instruments differ in legal status and application. A company should map them to its roles, jurisdictions and customer obligations rather than claim compliance through a checklist.

Assurance should be use-case specific. The evaluation set should represent material inputs, edge cases and affected populations. Measures can include accuracy, false-positive and false-negative costs, calibration, robustness, security, privacy, bias, hallucination, human override and service

continuity. The approval threshold should reflect the consequence of failure. A marketing draft and a credit decision require different evidence and accountability.

Governance can also expose a weak capability. Missing lineage, undocumented prompts, unclear ownership or unreconciled incidents may show that the company cannot control a system it describes as proprietary. Remediation deserves capital when the workflow is economically material and the company has a credible route to evidence. Otherwise, the board may restrict the use, partner or exit the capability.

10. Translate capability into unit economics and cash

The economic bridge should start with a measurable unit, such as a resolved claim, approved application, designed component, completed service request or retained subscription. Management records baseline price, volume, labour, technology, error, rework, loss, working capital and capital expenditure. It then compares the AI-enabled unit using production evidence and separates recurring benefits from implementation effects.

Gross time saving should not be booked as value. The bridge should state how released capacity changes payroll, hiring, throughput, customer response or revenue. It should deduct model and cloud cost, data preparation, integration, review, assurance, security, support, training, change management and transition duplication. It should also reflect customer price sharing, because buyers may demand part of the productivity benefit at renewal.

Cash timing matters. Upfront integration and data remediation can precede benefits by several periods. Minimum cloud or model commitments can create fixed cost. Revenue benefits may depend on procurement cycles. Workforce changes can require consultation, notice or restructuring cost. The board should review EBITDA, operating cash flow and peak funding rather than rely on a single return percentage.

Scenario analysis should preserve key uncertainties. Adoption, performance, customer price, usage, vendor cost, review intensity and implementation time can be varied separately. Management should identify the threshold at which the investment no longer meets the company's capital standard. The capability balance sheet then records the evidence needed to narrow each uncertainty.

11. Score defensibility through control, combination and time

Defensibility is the ability to sustain economically attractive outcomes against competition, imitation, supplier power and internal degradation. It is not an intrinsic property of a data set, patent, model or employee. The assessment should combine legal control, operational embedding, commercial adoption, learning effects, switching cost, assurance, portability and reinvestment capacity.

A weighted score can support comparison if mandatory gates remain separate. Data rights, customer permission, safety or regulatory requirements should not be offset by high performance. The score should also show confidence. A high score supported by management assertion is weaker than a moderate score supported by contracts and production cohorts. The committee should see both position and evidence quality.

The durability period should be explicit. A model advantage may decay quickly as competitors gain access to similar models. A workflow advantage can last longer when integration and accumulated outcomes are difficult to reproduce. A regulatory permission may create a barrier while also increasing compliance cost. A strong brand can support adoption but can lose value after a material incident. Each mechanism has a different review frequency.

The score should influence action rather than serve as a league table. High capability and strong economics may justify scaling. High capability with weak customer economics may require repricing or a smaller scope. Low capability in a strategically important workflow may justify acquisition or partnership. Low capability and low materiality should preserve capital.

Figure 3. Illustrative defensibility map combines capability strength and evidence confidence



Bubble size represents assumed three-year capital need for six hypothetical capabilities; all scores are author assumptions.

12. Measure supplier and key-person dependency

Dependency can reduce the value of an otherwise successful capability. The matrix should identify single suppliers, minimum commitments, exclusive terms, renewal dates, price-change rights, service concentration, data export, model substitution, security responsibilities and replacement time. It should include specialist contractors and key employees as well as technology providers.

The CMA's foundation-model review describes concentration and vertical relationships across compute, data, model development and distribution [8]. OECD work on AI compute reports substantial concentration of infrastructure and a need to understand access and capacity [14]. These market observations do not determine a company's negotiating position. They support a company-specific map of bottlenecks and alternatives.

Management should quantify the economic consequence of failure. A supplier outage may stop an optional feature or a regulated service. A 30 percent price increase may be immaterial for one workflow and eliminate contribution margin for another. Departure of a specialist may slow experimentation or prevent release approval. The matrix should show time to replace, cost to replace and the first customer or financial obligation at risk.

Mitigation can include dual sourcing, reserved capacity, export rights, modular architecture, escrow, succession, documentation, retention, insurance and customer contingency. Each measure has cost. The board should fund mitigation in proportion to the cash flow and obligation at risk rather than demand theoretical independence from every supplier.

13. Distinguish accounting, management and valuation views

The financial statements, capability balance sheet and valuation model answer different questions. IAS 38 determines recognition and measurement of intangible assets within its scope [1]. IFRS 3 requires an acquirer to recognise identifiable intangible assets separately from goodwill when the applicable conditions are met in a business combination [15]. IAS 36 governs impairment testing for relevant assets and cash-generating units, including annual assessment for goodwill and certain intangibles [16]. The management inventory can include economically important capabilities that are not recognised assets.

A valuation requires a defined purpose, basis, date, methods, assumptions and evidence. Income, market and cost approaches may be relevant to particular intangible assets or businesses. WIPO guidance describes cost, market and income methods for intellectual property and trade secrets while cautioning that value is context dependent [17]. WIPO has also highlighted data, algorithms and trade secrets as emerging assets that present valuation challenges [18]. These sources do not support a standard premium for AI.

The valuation bridge should identify how a capability changes forecast cash flows, risk, reinvestment and the duration of advantage. It should avoid adding a separate intangible value when the same benefit is already reflected in enterprise cash flows. A relief-from-royalty or excess-earnings method may be relevant to identifiable assets in an appropriate engagement, but assumptions about revenue, margins, contributory assets, attrition, useful life and discount rates require evidence.

The capability balance sheet can improve valuation diligence by showing what exists, who controls it, how it is used and what it costs. It can also identify impairment indicators when expected benefits weaken, a key contract is lost, a system becomes obsolete or remediation cost rises. Accounting and valuation conclusions remain the responsibility of qualified professionals applying the relevant standards and facts.

14. Use the balance sheet for build, buy, partner and divest decisions

Build is more credible where the workflow differentiates the company, data can be used for the intended purpose, internal teams can operate the system and customer evidence supports the economics. Buy can accelerate access to a product, team, rights, customers or embedded workflow. Partner can provide speed and flexibility when contract, portability and value capture are acceptable. Divest or discontinue may be appropriate where a capability is immaterial, structurally uneconomic or outside the company's risk appetite.

The strategic route should be tested against the full capability gap. Buying a model does not supply data rights, customer adoption or operating redesign. Acquiring a team can create retention and integration risk. A partnership can leave the company dependent on the partner's model, distribution or pricing. Internal development can consume time while the market changes. The decision needs a combined view of time to evidence, total capital, control, downside, integration and strategic options.

Acquisition diligence should reconcile the target's claims to contracts, repositories, model evaluations, production logs, customer cohorts, cloud bills, employment terms and financial records. The US Merger Guidelines discuss platforms, access to competitively significant inputs and transactions that may entrench a dominant position [19]. Relevant merger control, foreign investment, data, sector and employment requirements should be considered for the actual transaction and jurisdictions.

Table 4. Strategic route matrix for closing an AI capability gap

Route	Appropriate evidence	Value opportunity	Principal risk
Build	Differentiating workflow, usable data and capable team	Control and deep integration	Delay, adoption failure and accumulated cost
Buy	Target controls a scarce complementary capability	Speed, ownership and strategic option	Overpayment, retention, integration and obsolescence
Partner	Capability is available and contract protects use and exit	Fast learning with lower initial capital	Supplier capture, continuity and lock-in
Divest or stop	Weak economics, low strategic fit or unacceptable risk	Releases capital and management attention	Lost option or separation cost

The matrix guides board discussion; each route requires company-specific commercial, technical, financial, legal and regulatory diligence.

15. Apply the method to a hypothetical business-services company

Assume a multinational business-services company with annual revenue of USD 480 million, EBITDA of USD 62 million and a three-year AI investment ceiling of USD 45 million. The company processes customer documents, resolves service requests and prepares regulated reports. It has twelve years of workflow data, 2,400 employees, a central technology team and several third-party model and cloud contracts. These facts are author assumptions and do not describe an actual company.

Management initially lists fifteen AI initiatives. The capability balance sheet consolidates them into six material capability groups: document classification, service-agent assistance, regulatory knowledge retrieval, exception prediction, customer analytics and model assurance. The review finds that the company possesses extensive data, but rights differ by customer and period. Some contracts permit processing to deliver the service while restricting reuse for training. Data lineage is strong in two workflows and incomplete in three.

The service-agent assistant has the strongest adoption evidence. A controlled cohort shows lower handling time with stable customer satisfaction under the author's assumptions. Human review remains necessary, and the company has not converted all released time into cash. The investment case therefore uses avoided hiring and additional capacity rather than assuming an

immediate headcount reduction. Customer contracts permit assistance, subject to security and location controls.

The regulatory knowledge system has strong specialist content and customer need, but it depends on two senior experts who approve changes. Documentation and succession are incomplete. The first capital tranche funds knowledge capture, evaluation sets and an approval workflow. The document-classification capability uses a third-party model under terms that allow price changes and limit some export functions. Management funds a portability test before scaling.

The author assumes USD 16 million for data and workflow remediation, USD 9 million for assurance and security, USD 8 million for product and integration, USD 5 million for talent and succession, USD 4 million for commercial testing and USD 3 million as an uncommitted reserve. Release occurs in stages. No second tranche is approved without paid customer or verified operating evidence, stable quality and a contribution-margin path.

The worked example changes the strategic conclusion. The company does not lack AI tools. Its principal gaps are rights consistency, workflow evidence, supplier portability and organisational depth. The board allocates capital to those complementary capabilities and considers a small acquisition only if it supplies a documented regulatory workflow and retained customer base rather than a standalone model.

16. Convert the inventory into investment priorities


Priority should reflect materiality, gap, evidence, time and strategic fit. The committee can begin with mandatory remediation where rights, safety, security or customer obligations are not met. It then considers capabilities that protect material revenue, improve cash within the ownership horizon or open a credible strategic option. Low-evidence initiatives receive small discovery budgets and explicit expiry dates.

Each investment case should state the business outcome, baseline, capability gap, alternatives, total cost, peak cash, owner, dependencies, evidence milestones, downside and stop condition. A milestone might require a representative paid cohort, a maximum error rate, completion of a rights review, a portability test, documented succession or a minimum contribution margin. Completion of a software release is not sufficient when the economic or control evidence is missing.

Shared capabilities can reduce duplication. Common data standards, evaluation, security, procurement, legal playbooks and model inventory may serve several workflows. Their costs should be allocated transparently and their service measures defined. A central platform can become a new dependency if business units cannot challenge its economics or obtain support.

The portfolio should preserve an uncommitted reserve. AI markets, regulation and supplier terms can change during the programme. A reserve allows management to respond to customer evidence, acquire a scarce capability or mitigate an unexpected dependency. It also prevents every proposed initiative from being approved merely because a three-year budget exists.

Figure 4. Twelve-month capability investment calendar



The calendar replaces assumptions with evidence before larger capital, partnership or acquisition commitments.

17. Run a twelve-month board cycle

Months one and two establish the perimeter, owners, evidence standard and capability inventory. Management maps material workflows, data rights, technology dependencies, customer permissions and critical roles. Finance defines the unit-economic bridge. Legal, risk, security and sector specialists identify mandatory controls. The board approves immediate restrictions or remediation where required.

Months three and four verify existence and baseline economics. Teams observe workflows, reconcile production logs, contracts, usage, cost and quality, and document current model versions. Customer teams identify renewal and procurement dates. Management removes entries that are immaterial or only proposals. Initial defensibility and confidence scores are assigned with source links.

Months five and six run controlled commercial and operating tests. The company measures paid demand, adoption, end-to-end performance, human review and total cost. It tests supplier portability and succession for critical roles. Finance builds cash and downside scenarios. The investment committee releases limited capital against verified milestones.

Months seven to nine scale the capabilities that meet commercial, operating and control thresholds. Weak initiatives are redesigned, partnered or stopped. Management evaluates acquisitions only where the capability gap, integration route and economics are clear. The board reviews customer claims, regulatory changes, incidents and vendor concentration.

Months ten to twelve refresh valuation, impairment and strategic-option analysis. Evidence is prepared for financing, partnership or transaction diligence where relevant. The board compares the new capability position with the starting inventory, records cash spent and benefits realised, resets capital ceilings and approves the next review calendar.

The cycle should continue quarterly for material capabilities. A capability can weaken through employee departure, rights change, model obsolescence, customer rejection, supplier pricing or control failure. It can strengthen through adoption, labelled outcomes, integration, assurance and documentation. The balance sheet is useful because it records these changes in a decision-ready form.

18. State limitations and decision principles

The framework does not predict AI capability, adoption, regulation, supplier economics or competition. It does not replace accounting, valuation, legal, tax, employment, privacy, cybersecurity, technical or sector advice. The worked example, scores, financial values and investment sequence are author assumptions used to explain the method.

Measurement has limitations. Customer and employee cohorts can contain selection effects. Model performance can change across populations and versions. Benefits may overlap. Workflow redesign can create outcomes that are difficult to attribute to a single asset. Data rights can differ by contract and jurisdiction. A score can conceal these differences if the committee does not examine the supporting record.

Accounting recognition and enterprise value can diverge. An internally developed capability may be strategically important while associated expenditure is expensed. A recognised intangible asset may lose economic relevance. A workforce can create value while failing the accounting definition of an asset. The management inventory should preserve these distinctions and avoid attaching unsupported monetary values to every entry.

The board can use four principles. First, require evidence of a material economic mechanism. Second, assess combinations of capabilities rather than isolated tools. Third, separate legal control, technical performance, adoption and cash economics. Fourth, stage capital and strategic commitments behind dated proof. These principles allow the company to invest decisively while keeping claims and valuation accountable to evidence.

The AI capability balance sheet gives management a common language for data, workflows, talent, technology, customers and governance. Its value lies in the decisions it improves: what to protect, what to build, where to partner, what to acquire, when to stop and which evidence must be produced before more capital is committed.

References

1. IFRS Foundation. IAS 38 Intangible Assets. Issued standards and supporting material. <https://www.ifrs.org/issued-standards/list-of-standards/ias-38-intangible-assets/>
2. OECD, Boston Consulting Group and INSEAD. The Adoption of Artificial Intelligence in Firms: New Evidence for Policymaking. 2 May 2025. <https://doi.org/10.1787/f9ef33c3-en>
3. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework 1.0. January 2023. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
4. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. 26 July 2024; updated 8 April 2026. <https://doi.org/10.6028/NIST.AI.600-1>
5. 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
6. 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>

7. OECD. Competition in Artificial Intelligence Infrastructure. Global Forum on Competition, 2025.
<https://one.oecd.org/document/DAF/COMP/GF%282025%294/en/pdf>
8. UK Competition and Markets Authority. AI Foundation Models: Technical Update Report. 16 April 2024.
https://assets.publishing.service.gov.uk/media/661e5a4c7469198185bd3d62/AI_Foundation_Models_technical_update_report.pdf
9. International Labour Organization. Generative AI and Jobs: A 2025 Update. 20 May 2025.
<https://www.ilo.org/publications/generative-ai-and-jobs-2025-update>
10. 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>
11. United States Securities and Exchange Commission. Office Hours with Gary Gensler: AI Washing. 4 September 2024.
<https://www.sec.gov/newsroom/speeches-statements/gensler-office-hours-ai-washing-090424>
12. International Organization for Standardization. ISO/IEC 42001:2023 Information Technology - Artificial Intelligence - Management System. December 2023. <https://www.iso.org/standard/42001>
13. European Commission. General-Purpose AI Code of Practice. 10 July 2025. <https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai>
14. OECD. AI Compute. Updated 2026. <https://www.oecd.org/en/topics/sub-issues/ai-compute.html>
15. IFRS Foundation. IFRS 3 Business Combinations. Issued standards and supporting material.
<https://www.ifrs.org/issued-standards/list-of-standards/ifrs-3-business-combinations/>
16. IFRS Foundation. IAS 36 Impairment of Assets. <https://www.ifrs.org/issued-standards/list-of-standards/ias-36-impairment-of-assets/>
17. World Intellectual Property Organization. WIPO Guide to Trade Secrets and Innovation, Part IV: Trade Secret Management. <https://www.wipo.int/web-publications/wipo-guide-to-trade-secrets-and-innovation/en/part-iv-trade-secret-management.html>
18. World Intellectual Property Organization. WIPO Supports Aligning Valuation Standards with the Modern Economy. 15 May 2026. <https://www.wipo.int/en/web/ip-financing/w/news/2026/wipo-supports-aligning-valuation-standards-with-the-modern-economy>
19. United States Department of Justice and Federal Trade Commission. Merger Guidelines. 18 December 2023.
https://www.ftc.gov/system/files/ftc_gov/pdf/2023_merger_guidelines_final_12.18.2023.pdf
20. World Intellectual Property Organization and Luiss Business School. World Intangible Investment Highlights 2026. 2026. <https://www.wipo.int/publications/en/details.jsp?id=4860>
21. International Valuation Standards Council. Making Intangibles More Tangible: Series Lessons. 2024.
<https://ivsc.org/ivsc-perspectives-paper-making-intangibles-more-tangible-series-lessons/>
22. 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>
23. OECD. Technology Diffusion. Updated 28 October 2025. <https://www.oecd.org/en/topics/sub-issues/technology-diffusion.html>
24. 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>
25. OECD. OECD AI Principles. Updated 2024. <https://oecd.ai/en/ai-principles>

About the Author

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.