The AI Operating Model: Ownership, Data, Controls and Benefit Tracking

A Board Framework for Use-Case Ownership, Lifecycle Controls and Auditable Value Realisation

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Artificial intelligence programmes often begin as a collection of tools, pilots and vendor demonstrations. A board needs a different object: an operating model that assigns accountable owners, controls data and model dependencies, governs the lifecycle of each use case and connects expenditure to measurable business outcomes. Without that structure, adoption can expand faster than evidence, risk ownership and financial control.

This paper develops an AI operating model for investor-owned and growth businesses. It starts with an enterprise inventory and a use-case charter, then connects strategic intent, business ownership, data rights, model and vendor controls, lifecycle gates, human oversight, monitoring, incident response and benefit realisation. The framework distinguishes technical performance from operating adoption and financial impact. It also separates forecast benefits from validated outcomes and links each claimed benefit to a controlled baseline, counterfactual, full-cost view and accountable P&L or cash owner.

The framework is informed by international standards and current public guidance from authorities in the European Union, United Kingdom, United States, United Arab Emirates, Saudi Arabia, India, Singapore, Australia and Canada. Some sources apply only to regulated financial institutions, public bodies, providers of specified systems or other defined entities. They provide reference architecture for a proportionate private-company system and do not create duties for another organisation.

All financial values, adoption rates, risk tiers, control thresholds and operating outcomes in this paper are hypothetical management assumptions used to demonstrate the framework. They do not describe a client, technology product, investment, valuation or regulatory conclusion. The paper does not provide legal, regulatory, accounting, tax, cybersecurity, data-protection or investment advice. Organisations should obtain advice appropriate to their activities, systems, users and jurisdictions.

JEL Classification: G34, L21, M15, O31, O32, C45, K24

Keywords: artificial intelligence, operating model, AI governance, use-case portfolio, data governance, model risk, benefit tracking, portfolio companies, private equity, value creation

1. Govern an enterprise system rather than a collection of AI tools

The board's first AI question should concern the operating outcome the company intends to change. A licence, model, assistant or pilot is an input. Value arises when the company changes a decision, process, product or customer proposition and can sustain that change within approved risk boundaries. The operating model connects technology capability to accountable business action.

This distinction matters because a portfolio can look active while remaining economically weak. Teams may count demonstrations, users, prompts or generated documents. Those measures say little about revenue, cost, cash, service, quality or risk. A use case can perform well in a controlled test and fail in production because data access, workflow integration, employee adoption, vendor reliability or customer acceptance is inadequate.

The OECD AI Principles call for accountability, traceability and systematic risk management across the AI lifecycle [1]. NIST organises AI risk management around Govern, Map, Measure and Manage [3][4]. ISO/IEC 42001 frames AI management as a system of policies, objectives and continually improved processes [7]. These sources support a practical corporate conclusion: AI governance belongs inside the operating model rather than beside it as a compliance appendix.

A complete operating chain contains five linked objects. Strategy identifies the decision or workflow worth changing. A use-case charter defines scope, owner and intended outcome. Data and model controls establish whether the capability can operate responsibly. Lifecycle gates determine whether it can progress. A benefit ledger shows whether the change reached the P&L, cash flow, balance sheet, service outcome or risk position. The board receives one integrated view across all five.

Figure 1. The AI operating chain connects strategic intent to controlled business value

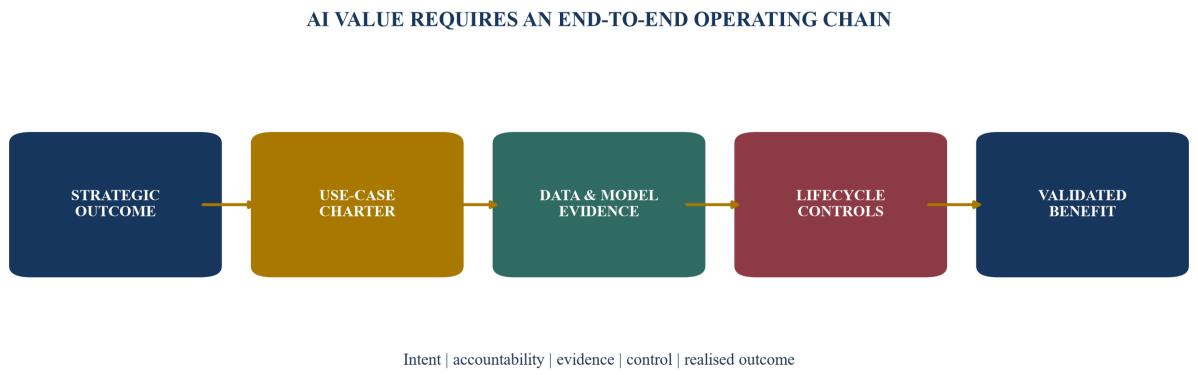

Each layer has a named owner, evidence set, decision right and review frequency.

2. Establish the enterprise inventory before prioritising the portfolio

Management cannot govern what it has not identified. The inventory should cover internally developed models, externally hosted models, embedded AI in enterprise software, employee assistants, automated decision components, analytical models, agents, model-enabled products and material experiments. It should also record shadow use discovered through surveys, access logs, procurement data, network controls and business review.

The unit of account should be the use case rather than the vendor name. One general-purpose model can support customer service, coding, marketing and contract review, with different data, users, decisions and consequences. Conversely, one use case can depend on several models,

retrieval services, data sources and workflow applications. The inventory should represent both relationships.

Each record needs a business owner, technical owner, risk tier, intended users, affected stakeholders, decision role, data categories, model and vendor dependencies, jurisdiction, lifecycle status, approval date, monitoring requirements, benefit owner and retirement route. The record should identify whether the system recommends, drafts, predicts, ranks, decides, acts or controls another system. Autonomy and reversibility often matter more than model sophistication.

The inventory also supports portfolio economics. Duplicate licences, overlapping pilots and incompatible data pipelines can be identified. A common capability may serve several use cases, while a high-cost specialist model may have one narrow owner. Management can then allocate shared cost, negotiate vendors and decide which platforms deserve scale.

Table 1. Minimum enterprise AI inventory

control object	minimum evidence	accountable owner	review trigger	portfolio use	common failure
use case	purpose, workflow and affected decision	business executive	scope or user change	prioritisation	tool name used as the use case
data	source, rights, classification and lineage	data owner	new source or purpose	reuse and control	production data absent from design review
model	provider, version, limits and tests	technical owner	version or configuration change	dependency management	underlying model change unnoticed
risk tier	impact, autonomy, exposure and reversibility	risk owner	material incident or redesign	proportional governance	every use case treated alike
benefit	baseline, counterfactual, cost and owner	finance and business owner	forecast or operating change	capital allocation	activity counted as value
lifecycle	stage, approvals, exceptions and retirement	product owner	gate decision	portfolio throughput	pilots remain open indefinitely

The inventory joins operational, technical, risk and financial evidence in one controlled record.

3. Design the operating model around six accountabilities

An AI programme needs more than a steering committee. Six accountabilities should be visible. The board approves appetite, material investment and oversight. The executive sponsor owns the portfolio thesis. The business owner owns workflow adoption and the operating outcome. The technical owner owns architecture, implementation and service performance. Independent control functions challenge risk, privacy, security, legal and compliance evidence. Finance controls the benefit baseline and validates realised value.

The roles can be combined in a smaller organisation, but the accountabilities should remain distinct. The person proposing a use case can prepare evidence. A material deployment still benefits from challenge by someone who did not build the model or own the forecast benefit. Independence can be proportionate to risk, cost and organisational scale.

Decision rights should be attached to lifecycle events. A product owner may approve a low-impact configuration change. A risk committee may approve a high-impact deployment. The board may reserve decisions involving material capital, workforce consequences, customer eligibility, safety, regulated outcomes or reputational exposure. Emergency suspension authority should be explicit and immediately available.

Committees should decide rather than receive presentations. Every agenda item should state the decision requested, options, evidence, owner, unresolved limitations and next review. The portfolio office maintains standards and records. It should not absorb accountability from the executive whose function uses the system.

Table 2. Decision rights across the AI operating model

decision	board	executive sponsor	business owner	technical owner	control functions	finance
portfolio appetite and material capital	A	R	C	C	C	C
use-case charter and operating outcome	I	A	R	C	C	C
data and model design	I	C	C	A/R	C	I
independent challenge and risk acceptance	I	C	C	C	A/R	C
production deployment	I or A	A	R	R	C	C
benefit baseline and validated outcome	I	C	R	C	C	A/R
suspension and retirement	I	A	R	R	C	I

R means responsible, A accountable, C consulted and I informed; the exact allocation should reflect organisational size and applicable requirements.

4. Convert strategy into a controlled use-case charter

The charter is the smallest useful contract between strategy, delivery, control and finance. It defines the business problem, current workflow, intended users, affected parties, decision role, proposed intervention, expected outcome, constraints, dependencies, owner, cost envelope, evidence plan and exit conditions. It should be short enough to govern and precise enough to test.

The current workflow deserves equal attention with the proposed model. Management needs the baseline cycle time, error rate, conversion rate, cost, control burden and customer outcome before automation. Weak processes can become faster weak processes. A use case should therefore identify which step changes, which decisions remain human and which downstream process absorbs the output.

The charter should state the model's intended purpose and prohibited uses. A system designed to summarise documents should not drift into customer eligibility decisions without a new assessment. A sales assistant should not ingest confidential deal materials because its interface

makes that convenient. Clear boundaries support access design, monitoring and employee training.

Success criteria should cover four dimensions: technical performance, operational adoption, control effectiveness and financial or service outcome. A use case advances only when evidence exists across the dimensions relevant to its tier. This prevents impressive model scores from carrying a weak operating case through the gate.

The charter should also define a decision hypothesis. It states who will make a different decision, what evidence the system contributes, how quickly the decision must occur and how management will know that the change improved the outcome. This forces the sponsor to identify the human or automated action between model output and value. A forecast that stops at output quality remains incomplete.

Exit criteria are equally important. Management can stop a use case when the problem has disappeared, the intervention cannot meet its threshold, the vendor economics have changed, the control burden exceeds value or a simpler process improvement performs better. The charter should identify how access, data, records, contracts and downstream dependencies will be closed. A planned exit route strengthens capital discipline at the beginning of the investment.

5. Tier use cases by consequence, autonomy, exposure and reversibility

Risk tiering should determine the intensity of governance. The classification can consider the consequence of error, degree of automation, affected stakeholder, scale of exposure, data sensitivity, legal context, model novelty, explainability, external dependence and reversibility. A drafting assistant used by trained employees differs materially from an autonomous system that changes customer access or moves money.

The tier should attach to the use case in context. The same model may be low impact for internal brainstorming and high impact when its output influences recruitment, credit, pricing, safety or health. This use-case view aligns with the risk-based approaches found in the EU AI Act [9][10], NIST AI RMF [3], ASIC's review of licensee governance [27] and OSFI's model-risk framework [29].

A practical structure can contain four internal tiers. Tier one covers bounded assistance with low consequence and clear human review. Tier two covers operational recommendations or content that can affect customers or financial outcomes. Tier three covers material decisions, sensitive data, high autonomy or regulated workflows. Tier four covers prohibited or outside-appetite activity. Internal tiers should never be presented as legal classifications without jurisdiction-specific analysis.

Tiering is dynamic. Growth in volume, new data, a vendor update, wider autonomy, new user groups or a material incident can change the classification. The inventory should therefore record the basis, approver and next review date.

Figure 2. Illustrative internal tiering by consequence and autonomy

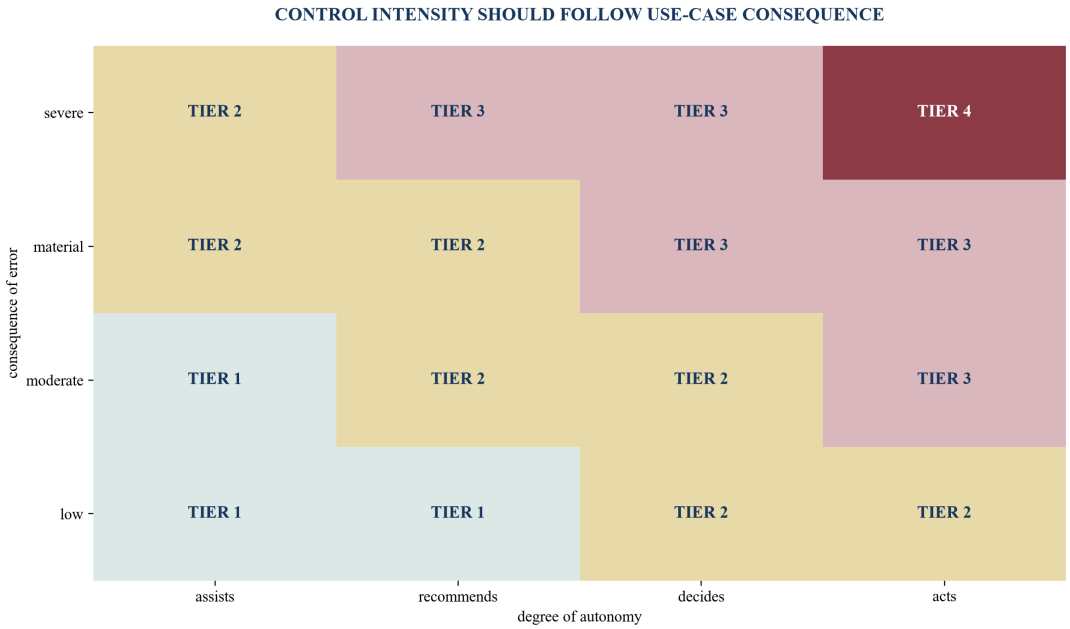

The matrix is a management tool; legal classifications require separate analysis.

Table 3. Illustrative tier triggers and minimum controls

tier	typical context	minimum pre-production evidence	approval	monitoring	stop trigger
one	bounded internal assistance	purpose, approved tool, data rule and human review	business and technology owner	usage, service and exceptions	confidentiality or repeated quality issue
two	operational recommendation or external content	structured tests, workflow control, owner training and benefit baseline	business owner plus control review	quality, adoption, complaints and drift	threshold breach or material customer impact
three	material decision, sensitive data or high autonomy	independent validation, legal analysis, security test, fallback and incident plan	reserved committee or board authority	continuous or high-frequency control set	safety, rights, regulatory or financial boundary
four	prohibited or outside appetite	documented refusal and access prevention	board-approved appetite	attempted-use surveillance	any attempted deployment

Management should tailor the trigger and control set to applicable law, sector and risk appetite.

6. Treat data rights and lineage as operating prerequisites

AI performance and accountability depend on data that the organisation can identify, use, protect and reproduce. The data control file should record source, owner, permitted purpose, lawful or contractual basis, classification, geography, retention, quality tests, transformations, lineage and downstream uses. It should also identify personal, confidential, privileged, copyrighted, licensed and sensitive data.

Data rights can be as important as data quality. A technically useful dataset may be unavailable for model training, retrieval or vendor processing. Customer agreements, employment terms, database rights, intellectual property, privacy rules and sector obligations can constrain use. Legal approval should attach to a specific purpose and data flow rather than to the broad label of AI.

Lineage should follow data from source through preparation, retrieval, prompt context, model interaction, output storage and downstream decision. This supports reproduction, incident analysis, deletion, access review and benefit attribution. The OECD links AI accountability with traceability across datasets, processes and decisions [1][2]. OSFI's final E-23 guideline similarly emphasises suitability, provenance and enterprise data governance for models [29].

Production data should be compared with development and test data. Population change, missing fields, new categories, seasonality and operational workarounds can weaken a model without changing its code. The monitoring design needs data-quality thresholds and an owner who can correct the source process.

7. Control the full model and vendor dependency chain

The relevant system usually extends beyond a single model. It can include a foundation model, embeddings, retrieval database, orchestration layer, safety filter, business rules, application, cloud service, human review and downstream system. An operating model should record the chain and identify which party controls each component.

Vendor diligence should cover intended use, service architecture, data handling, sub-processors, model versions, security, testing evidence, intellectual-property terms, confidentiality, retention, incident notification, service levels, audit rights, change notice, business continuity, termination and data return. A prominent provider name does not replace use-case diligence.

Configuration can materially change behaviour. System prompts, retrieval sources, temperature, tools, permissions, memory, action limits and fallback rules belong in controlled change management. A vendor may update the underlying model while the application name remains unchanged. The company needs a defined trigger for re-testing and re-approval.

Concentration should be visible at portfolio level. Several critical workflows may depend on one model provider, cloud platform or data service. The board should understand outage, pricing, contractual and strategic dependency, then decide which alternatives, fallbacks or manual procedures are proportionate.

Table 4. Evidence across the AI dependency chain

component	evidence	owner	control question	change trigger	fallback
source data	rights, quality, lineage and access	data owner	may this data serve this purpose?	source or purpose change	approved reduced dataset
model	version, provider, limits and evaluation	technical owner	is the model fit for intended use?	version or parameter change	prior approved model or manual route
retrieval	indexed sources, refresh and permissions	product owner	does retrieved context remain authoritative?	corpus or access change	source-only workflow
orchestration	prompts, tools, action limits and logs	engineering owner	can the system act outside scope?	prompt, tool or permission change	read-only mode
vendor	contract, security, continuity and notice	procurement and risk	can dependency be governed through the lifecycle?	material vendor change	alternate provider or exit plan
human process	review, escalation and authority	business owner	can people detect and correct material error?	role or volume change	manual control procedure

The evidence set should be proportional to the use-case tier and the company's ability to control each component.

8. Make build, buy and configure decisions on total operating economics

The build-or-buy decision should include strategic differentiation, data advantage, integration complexity, control requirements, time to value, talent, infrastructure, vendor dependence and full lifecycle cost. The cheapest pilot can become an expensive operating dependency when integration, assurance, monitoring and change are included.

Buying a model does not transfer accountability for the use case. The company still owns purpose, users, data, workflow, decisions, monitoring and customer outcome. Building internally provides greater control over some components and creates obligations for engineering, security, documentation, maintenance and talent continuity.

Configuration is often the dominant corporate path. The company combines external models with proprietary data, prompts, rules and workflow. This can create differentiated value without training a foundation model. It also creates a system whose performance depends on many changeable components. The operating model should govern the configured whole.

The investment case should distinguish reusable platform capability from use-case-specific spend. Identity, logging, evaluation, secure data access and orchestration may support multiple cases. Finance can allocate shared costs transparently while management avoids charging every early use case with the entire platform.

9. Use lifecycle gates that require evidence rather than enthusiasm

The lifecycle can contain six gates: discover, design, validate, deploy, scale and retire. Each gate should have entry criteria, required evidence, decision authority, conditional approvals and a time-bound exception route. A gate closes when the evidence supports the next exposure, not when the project team has completed a slide deck.

Discovery confirms the problem, owner and baseline. Design defines workflow, data, system and control architecture. Validation tests technical performance, human process, security and relevant rights. Deployment introduces controlled production exposure. Scale requires evidence that adoption, controls and benefits remain sound at greater volume. Retirement removes access, dependencies and retained data in a controlled way.

The NIST playbook provides actions across governance, mapping, measurement and management [4]. The NIST generative AI profile adds attention to risks such as confabulation, information integrity, privacy, security and intellectual property [5]. The lifecycle gate should convert such considerations into use-case evidence.

Conditional approval can be valuable when limitations are known and contained. The record should specify scope, mitigant, owner, expiry and evidence required for renewal. Open-ended exceptions allow pilot weaknesses to become permanent production design.

Figure 3. Evidence accumulates through six lifecycle gates

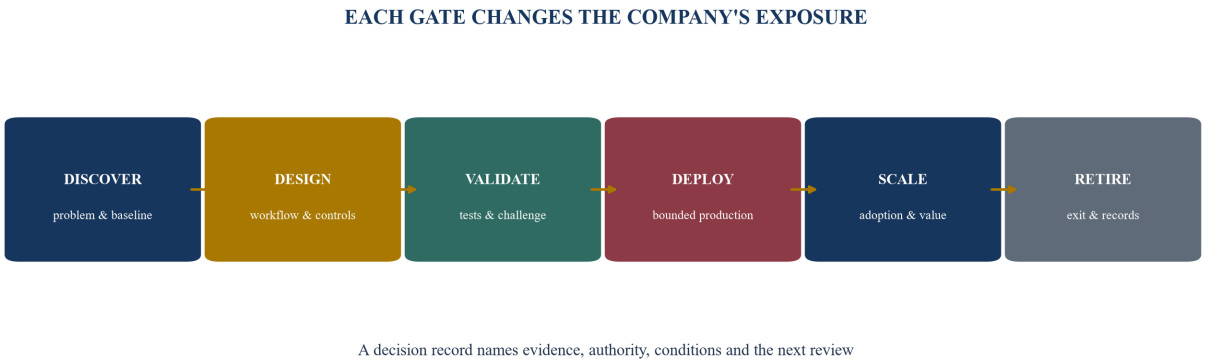

Exposure should increase only after the relevant evidence and authority are present.

10. Validate the system in the context in which it will operate

Validation should test the system, workflow and user behaviour against intended use. Technical evaluation can include accuracy, reliability, robustness, bias, explainability, privacy, security, latency and cost. Operational evaluation should examine escalation, human review, downstream actions, edge cases, exception volume and recovery.

The test set should reflect the production population and material failure modes. Average performance can hide poor outcomes for rare events, specific customer groups or high-value cases. Thresholds should be tied to consequence. A minor drafting error and an incorrect payment instruction require different tolerance.

Independent challenge should be proportionate. The reviewer should understand the use case, data, model, workflow and applicable requirements, then assess whether evidence supports deployment. For higher-impact use cases, the reviewer should be able to reproduce key tests and challenge the choice of metric, population, threshold and mitigant.

Generative systems need qualitative and adversarial evaluation as well as benchmark scores. Testing can cover unsupported statements, prompt injection, confidential-data leakage, harmful content, tool misuse, inconsistent responses and refusal behaviour. The NIST generative AI profile and secure-development guidance from CISA and international partners provide relevant reference points [5][17][18].

Ground truth should be governed. Labels and expected answers may contain expert disagreement, outdated policy or historical bias. The validation record should identify who created the test cases, the source date, review method, uncertainty and population coverage. Where an exact answer is unavailable, management can evaluate whether the system cites authoritative evidence, expresses uncertainty and escalates appropriately.

Evaluation must also reflect frequency and cost. A high-volume assistant may require automated monitoring plus sampled expert review. A rare material decision may warrant full case review. The operating model should budget for continuing evaluation rather than treat the pre-production test as a one-time project cost. Validation evidence loses relevance when the model, data, workflow or affected population changes.

11. Define human oversight as a decision design

Human oversight should specify who reviews what, with which information, under what time constraint and with what authority. A generic instruction to keep a human in the loop can fail when the reviewer has no expertise, too many outputs, weak evidence or incentives to accept the machine's recommendation.

The design can distinguish approval, verification, exception handling and appeal. A reviewer may approve every high-impact outcome, sample low-impact outputs, investigate threshold breaches and provide an escalation path for affected parties. The reviewer needs access to source evidence and known limitations rather than only a polished model answer.

Automation bias should be treated as an operating risk. Training, interface design and quality sampling can test whether employees challenge outputs. Overrides should be recorded with a reason so management can distinguish healthy judgement from inconsistent practice. Repeated override patterns can reveal a model, data or workflow problem.

Authority should match consequence. A user may correct a draft. A functional leader may suspend a workflow. A risk owner may require a control. The board or reserved committee may withdraw a material use case. The action should remain possible during an outage or incident.

12. Deploy through controlled production and change management

Production deployment should use approved environments, access controls, version control, logging, segregation of duties, rollback, capacity planning, service monitoring and release records. The deployment decision should confirm that operating owners, support teams and users are ready alongside the technology.

The release package should identify model and component versions, configuration, data sources, tests, limitations, approvals, monitoring thresholds, incident contacts and fallback. This creates a reproducible baseline. A material change should trigger the appropriate subset of re-testing and approval.

Change classification can distinguish routine, significant and emergency changes. A prompt wording correction may follow a standard route. A new tool permission, model family, data source, autonomous action or affected customer group may require renewed risk assessment. Emergency changes should be reviewed after the event.

Deployment should also protect the non-AI workflow. Manual fallback, service continuity and record access matter when a provider is unavailable or a model is suspended. Resilience should be tested before the system becomes critical.

13. Monitor performance, drift, incidents and control effectiveness

Monitoring should follow the use-case hypothesis. Technical measures can include accuracy, error type, calibration, drift, latency, availability, token or compute cost and security events. Operating measures can include adoption, completion time, rework, overrides, complaints, exceptions and downstream defects. Benefit measures should track the controlled financial or service outcome.

Thresholds need actions. A yellow boundary may require investigation and restricted scale. A red boundary may require suspension, rollback or manual fallback. The owner, response time and escalation route should be documented. A dashboard without action logic is observation rather than control.

Incident management should classify harm, contain exposure, preserve evidence, notify internal and external parties where required, correct the system and confirm recovery. It should include vendor incidents and near misses. Lessons should update tests, thresholds, training and portfolio standards.

Monitoring must remain version-aware. Performance before a model or data change should not be blended with performance after it. The record should support cohort analysis and explain which population, version and workflow produced an outcome.

Figure 4. The control dashboard links signals to management action

TECHNICAL, OPERATING AND CONTROL SIGNALS BELONG TOGETHER

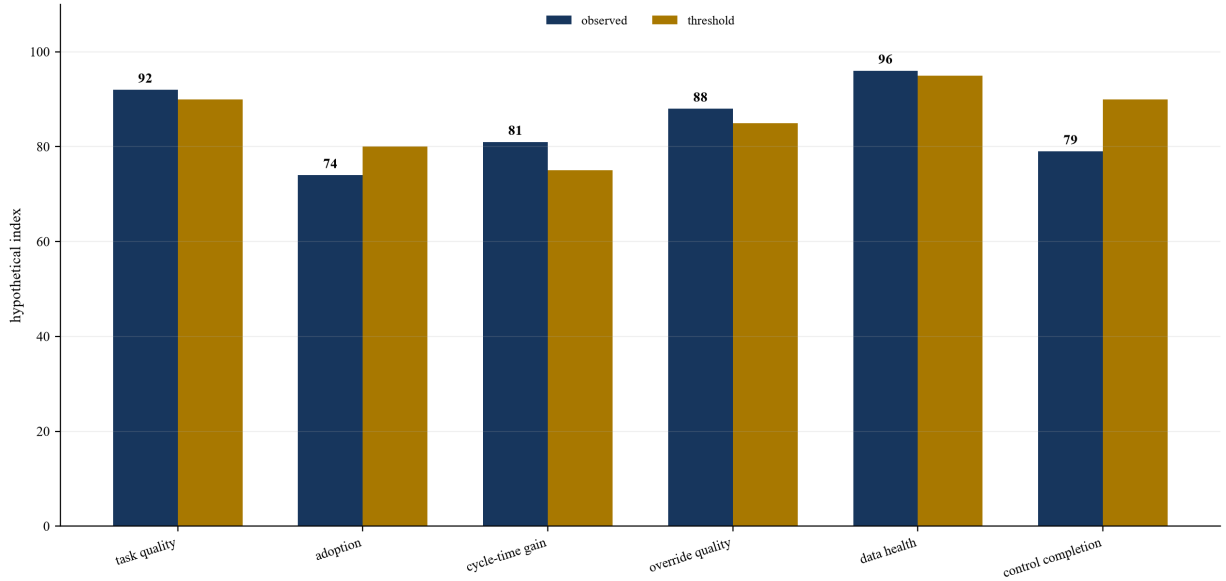

Values are hypothetical management assumptions used to illustrate a balanced monitoring view.

14. Make security and resilience part of the product design

AI systems create familiar software and data risks alongside model-specific attack paths. Threats can involve prompt injection, data poisoning, insecure tool access, model theft, sensitive-data disclosure, adversarial inputs, unsafe outputs and dependency compromise. Security should therefore begin with the architecture and intended use.

The threat model should identify assets, actors, trust boundaries, interfaces and potential misuse. Controls can include approved environments, least privilege, input and output handling, secret management, network restrictions, retrieval permissions, tool allowlists, content controls, logging, anomaly detection, red teaming and incident response. The correct set depends on consequence and exposure.

Agentic systems deserve particular attention because they can call tools and change external state. Permissions, transaction limits, confirmation steps, separation of duties and reversible actions should be explicit. A model's ability to generate a correct plan does not establish authority to execute it.

Secure-by-design guidance from CISA, the UK NCSC and international partners emphasises security across design, development, deployment and operation [17][18]. The operating model should translate this into gate evidence and an accountable security owner.

15. Integrate privacy, intellectual property and records into the workflow

Privacy analysis should follow the actual data flow and affected people. Management should understand which personal data enters prompts, retrieval systems, logs and outputs; where it travels; who can access it; how long it remains; and how rights can be exercised. Purpose change and automated decision use can require additional analysis.

Intellectual-property controls should cover inputs, training or fine-tuning data, generated outputs, vendor terms, open-source components, employee work and customer material. The company should define when generated content can be published, incorporated into products, relied upon in advice or shared externally.

Records policy should preserve decisions and evidence while respecting retention limits. Material approvals, tests, model versions, prompts or configurations, exceptions, incidents and benefit validation may require controlled retention. Routine conversational data may have a different schedule.

The EU AI Act, related transparency guidance, UK data-protection guidance, UAE and Saudi ethics frameworks and Singapore's governance materials illustrate the growing importance of traceability, transparency and accountable use [9][11][15][20][21][22][25][26]. Applicability and legal effect depend on the specific system, role and jurisdiction.

16. Build the benefit case from a controlled baseline and counterfactual

Benefit tracking begins before deployment. The charter should record the current process, volume, labour, cost, cycle time, quality, conversion, service and risk outcome. It should also define the counterfactual: what management expects without the intervention. A simple before-and-after comparison can misattribute market growth, pricing, staffing or process changes to AI.

Each benefit should have a formula, data source, frequency, owner and validation method. Time saved becomes financial value only when the company changes capacity, cost, throughput or service. A forecast reduction in handling time can be operationally useful. It should not be presented as EBITDA until the bridge to staffing, output or avoided spend is evidenced.

Revenue benefits need similar discipline. An AI-assisted sales process can influence lead quality, conversion, price, retention or representative capacity. The company should identify the mechanism, compare cohorts where feasible, control for other changes and avoid counting the same revenue in several initiatives.

The full-cost view should include licences, compute, implementation, integration, data preparation, assurance, security, training, support, change management and ongoing monitoring. Shared platform cost should be allocated consistently. Exit cost and vendor price sensitivity belong in the case.

Attribution can use several designs. A randomised rollout can provide strong evidence where ethical and operationally practical. Phased implementation, matched cohorts, interrupted time series or process-level decomposition may be more suitable elsewhere. The chosen method should be stated before results are reviewed, with limitations and concurrent changes recorded. Management should resist selecting the method that produces the most favourable number after the event.

Benefits can be financial, service, risk or capability outcomes. They should not be forced into a monetary estimate when the evidence does not support one. A controlled reduction in response variability or improved traceability may matter to the board even when finance cannot validate an immediate P&L effect. The register can show the operational outcome and its decision relevance while preserving the boundary around financial value.

Public-sector benefits guidance provides a useful general principle: benefits should be identified, defined, tracked, realised and sustained through the lifecycle, with accountable owners and assurance [31][32]. The same discipline improves private-company capital allocation.

Table 5. Controlled AI benefit register

benefit	baseline and counterfactual	operating measure	financial bridge	evidence owner	validation decision
cycle-time reduction	comparable task cohort without intervention	minutes per completed task	capacity redeployed, output increased or cost avoided	business owner	finance confirms mechanism and period
quality improvement	historical error and rework rate	accepted output and defect rate	lower rework, loss or service cost	quality owner	control function confirms definition
conversion uplift	comparable customer or seller cohort	conversion, price and retention	incremental contribution after cannibalisation and cost	commercial owner	finance validates attribution
working-capital effect	expected collection or inventory path	days and ageing	cash released after volume and seasonality effects	CFO	treasury confirms cash movement
risk reduction	defined loss or incident baseline	exposure, events and control performance	expected loss or capital effect where supportable	risk owner	methodology independently reviewed
platform reuse	approved standalone delivery cost	use cases and shared services consumed	avoided duplicate cost net of shared platform spend	technology owner	finance approves allocation

Forecast, observed and finance-validated values should remain separate.

17. Translate operational evidence into a P&L, cash and value bridge

Finance should maintain a bridge from gross operational benefit to validated economic outcome. Begin with observed output or performance change. Remove volume, mix, price, market and other concurrent effects. Deduct run cost, implementation amortisation where relevant, incremental control cost and displacement. Identify whether the result reaches revenue, gross margin, operating expense, working capital, capital expenditure or risk.

Accounting treatment should remain separate from management benefit measurement. IAS 38 contains recognition criteria for intangible assets, and the IASB has continued to examine questions involving AI-related data and solutions [33]. Management should obtain accounting advice rather than assuming that an internally described investment or benefit determines financial-statement treatment.

The benefit bridge should use period, cumulative and run-rate views. A use case can create a one-time implementation burden, a recurring cost and a later recurring benefit. Cash timing may differ from accounting expense. The board needs all three dimensions to allocate capital.

Forecast confidence should be explicit. A range with defined drivers is more decision-useful than a single precise number built on weak adoption assumptions. Sensitivities can show user adoption, task volume, provider price, error rate and time-to-scale.

Figure 5. Hypothetical bridge from gross operating benefit to validated annual value

FINANCE SHOULD BRIDGE ACTIVITY TO VALIDATED ECONOMIC VALUE

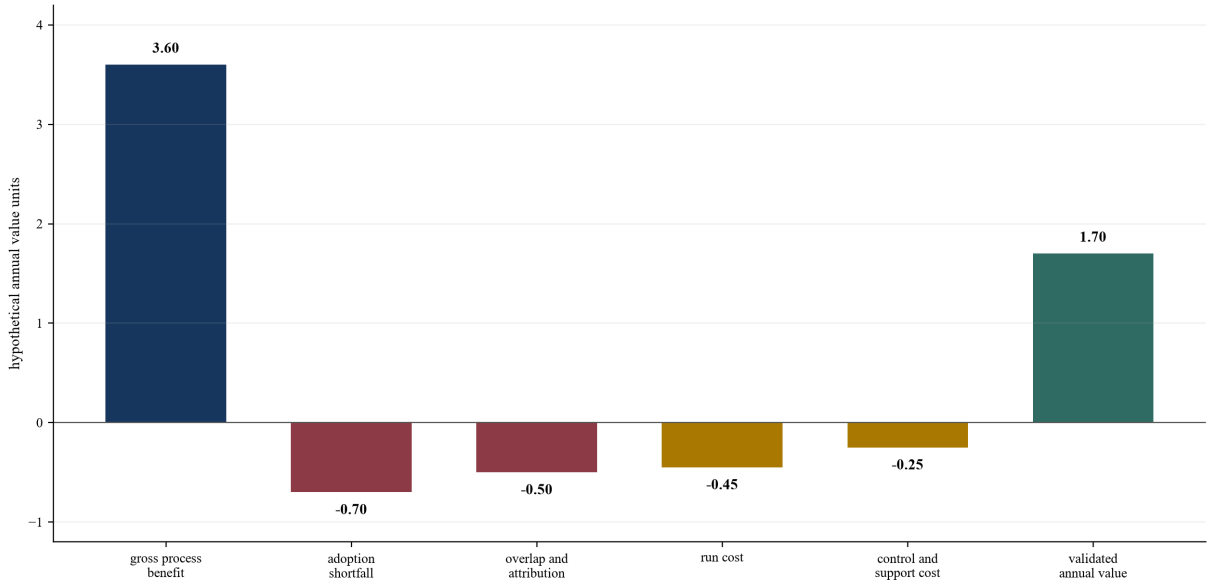

Values are hypothetical management assumptions and do not represent a client outcome or forecast.

18. Treat adoption and workflow redesign as part of the investment

AI benefit depends on work changing. Adoption should be defined as correct use in the intended workflow rather than login or licence activity. The company should know who uses the system, for which tasks, with which review, and whether the output changes the downstream decision.

Workflow redesign can remove handoffs, change role boundaries and create new control steps. It can also expose hidden work previously absorbed by employees. Process mapping should identify where time moves, which queues disappear, which exceptions grow and which skills become more valuable.

Training should be role- and use-case-specific. Users need purpose, boundaries, data rules, review expectations, escalation, examples and known limitations. Managers need to interpret monitoring and benefit evidence. Technical and control teams need lifecycle responsibilities. Board members need the portfolio economics and material exposure.

The workforce plan should address capacity rather than assume immediate cost removal. Time released can support volume growth, service, quality or headcount avoidance. Realised economic effect depends on management action. Employee consultation and local employment requirements may also apply.

19. Allocate capital through portfolio gates

The portfolio should compete for scarce capital, data, engineering, control capacity and management attention. A common scoring framework can compare strategic relevance, addressable value, evidence quality, time to impact, reuse, risk, dependency and implementation burden. Scores inform judgement and should not automate the investment decision.

Funding can be staged. Discovery receives a small evidence budget. Validation receives resources after the baseline and architecture are credible. Production funding follows successful tests and control readiness. Scale funding follows adoption and early benefit evidence. This reduces

sunk-cost escalation.

Portfolio review should identify stop decisions. A use case can be technically successful and economically unattractive. Another may remain valuable but require a different vendor or workflow. Capital should move when evidence changes.

Shared platforms need separate governance. Identity, data access, evaluation, monitoring and model gateways can lower marginal cost across the portfolio. Their value can be measured through enabled use cases, avoided duplication and control quality, alongside service cost and concentration risk.

Capital review should consider the value of learning. A bounded experiment may deserve funding because it resolves a material uncertainty about data, adoption, customer response or technical feasibility. The decision record should state the uncertainty, evidence sought, maximum exposure and date by which management will scale, redesign or stop. This treats experimentation as a purchased option with a defined expiry rather than an indefinite pilot.

The portfolio should also recognise bottlenecks. Ten funded use cases can compete for one data owner, security reviewer or workflow team, delaying all of them. Capacity planning should sequence work around the scarce capability that determines throughput. The board can then decide whether to add capacity, narrow the portfolio or change the architecture.

20. Apply jurisdiction and sector overlays without fragmenting the core

The enterprise can maintain one core operating model and attach overlays for legal entity, geography, sector and use case. The core contains inventory, ownership, tiering, lifecycle, evidence, monitoring and benefits. An overlay adds applicable rights, approvals, records, disclosures, testing or supervisory expectations.

The EU AI Act uses a risk-based framework with obligations that vary by system and actor; enforcement and application timing have continued to develop [9][10][11][12]. UK authorities use cross-sector principles alongside existing regulators, while the PRA applies model-risk expectations to specified firms [14][16]. US sources include voluntary NIST frameworks, sector rules and enforcement against unsupported claims [3][5][19].

The UAE and Saudi Arabia publish ethics and governance guidance emphasising accountability, transparency, privacy, safety and responsible adoption [20][21][22][23]. India's RBI FREE-AI work addresses responsible AI in finance [24]. Singapore's governance frameworks and AI Verify programme support practical testing and responsible deployment [25][26]. ASIC has identified gaps between AI adoption and governance among reviewed licensees [27]. Canada's OSFI has finalised a broad lifecycle model-risk guideline with AI considerations and a 2027 effective date [29][30].

These sources differ in legal status, scope and timing. The operating model should record applicability and advice for each use case. A global policy statement should not substitute for a local legal analysis.

Cross-border architecture requires an operational view of data location, remote access, vendor processing, support teams, incident notification and regulatory access. A contract may identify one service region while telemetry, human support or sub-processors operate elsewhere. The

data-flow record should reflect actual processing and should be refreshed when the service changes.

Sector overlays can add controls without creating a separate enterprise system. A financial use case may require model-risk, conduct and customer-outcome evidence. A healthcare, employment, safety or critical-infrastructure use case can require different professional, rights and resilience analysis. The common inventory and gate structure preserves portfolio visibility while the overlay preserves domain-specific accountability.

21. Test the framework through a hypothetical service-operation case

Consider a hypothetical portfolio company processing 240,000 customer service requests each year. Management proposes an assistant that retrieves approved knowledge, drafts a response and recommends a routing code. Employees remain responsible for review and submission. The company expects shorter handling time and lower rework.

Discovery establishes the current workflow: median handling time, quality review, repeat contacts, complaint rate, queue, labour mix and seasonal volume. The charter identifies customer operations as business owner, technology as system owner, the knowledge team as data owner and finance as benefit validator. Personal and confidential data are mapped before design.

The use case is classified in the middle tier because outputs reach customers and can affect service or rights, while trained employees retain approval. Validation uses representative cases, difficult categories, adversarial prompts and a separate test of retrieval permissions. The company sets thresholds for factual support, routing accuracy, confidential-data handling, review quality and complaint outcomes.

The initial pilot shows a hypothetical 24 per cent reduction in handling time for eligible requests. Only 58 per cent of requests are eligible, and adoption reaches 70 per cent of eligible volume. Management therefore avoids applying 24 per cent to the entire cost base. Finance models realised capacity from actual eligible volume, adoption, rework and incremental run cost.

Scale approval depends on controlled knowledge refresh, user training, threshold stability and a workforce plan for released capacity. The board receives forecast, observed and validated benefits separately. No value is attributed to the model simply because it is deployed.

22. Give the board a decision pack rather than an AI activity report

The board pack should show portfolio exposure, decisions and evidence. A one-page dashboard can include use cases by tier and lifecycle, capital committed, forecast and validated value, adoption, material dependencies, control exceptions, incidents, overdue reviews and decisions required. Detail belongs behind the dashboard.

Every material use case should show owner, intended outcome, stage, risk tier, latest evidence, unresolved limitation, next gate and benefit status. Benefits should be separated into forecast, observed and finance-validated values. Costs should include committed, incurred and run-rate views.

Narrative matters when metrics change. Management should explain why adoption, model performance, control thresholds or value moved; what decision followed; and who owns the action. The pack should preserve version and period consistency.

The board should also see concentration: critical workflows by provider, cloud, data source and key person. This connects AI oversight with operational resilience, cyber risk, procurement and succession.

Assurance should be visible as a coverage map. Internal audit, risk, compliance, security, privacy, finance and external specialists may each test part of the system. The board should see which material assertions have been reviewed, the period and scope, unresolved findings and management response. Multiple reviews do not create complete assurance when they examine the same narrow control and leave ownership, adoption or benefit attribution untouched.

Board information should mature with the portfolio. Early packs can emphasise inventory completeness, appetite and gate design. A scaling portfolio should emphasise production performance, incidents, concentration and validated value. A mature portfolio should also show retirement, control efficiency and whether shared capability is reducing marginal cost. The information architecture should change when the board's decisions change.

Table 6. Board AI operating-model scorecard

board question	primary measure	evidence owner	decision threshold	management action	board decision
where is exposure concentrated?	use cases by tier, autonomy and dependency	risk and technology	appetite or concentration boundary	diversify, restrict or accept	approve appetite response
is lifecycle control effective?	overdue gates, exceptions and incidents	portfolio office	material overdue or recurring breach	remediate, suspend or retire	challenge and reserve action
is adoption changing the workflow?	eligible volume, correct use and rework	business owner	adoption or quality below plan	redesign workflow or training	confirm scale posture
is value reaching finance?	forecast, observed and validated value	CFO	variance or attribution weakness	correct case or reallocate capital	approve funding change
can the service continue?	fallback tests, outage and recovery evidence	technology and operations	tolerance breach	test or strengthen resilience	accept residual exposure
what must be decided now?	gate request, options and limitations	executive sponsor	reserved authority	execute approved option	approve, condition, defer or stop

Each line should lead to a decision, escalation or confirmed operating action.

23. Implement the operating model in one hundred days

The first thirty days establish visibility and authority. Appoint the executive sponsor, approve interim appetite, identify material use cases and tools, define the inventory, assign owners and stop any activity outside clear boundaries. Finance selects a small number of priority benefit baselines.

Days thirty-one to sixty build the control spine. Define risk tiers, lifecycle gates, minimum evidence, change classification, incident route, vendor requirements and benefit register. Select two or three priority use cases that are commercially relevant and sufficiently bounded to test the system.

Days sixty-one to ninety run the first controlled cycle. Complete design, data mapping, validation, security review, user training and production readiness for the selected cases. Record conditional approvals and exceptions. Establish monitoring and a finance-owned benefit bridge.

Days ninety-one to one hundred close the first board cycle. Present the inventory, exposure, evidence, early operating results, benefit confidence, decisions and next-quarter capital allocation. The operating model becomes routine only when governance, delivery and finance use the same records.

Figure 6. A one-hundred-day implementation sequence

THE FIRST CYCLE CONNECTS OWNERSHIP, CONTROL AND VALUE

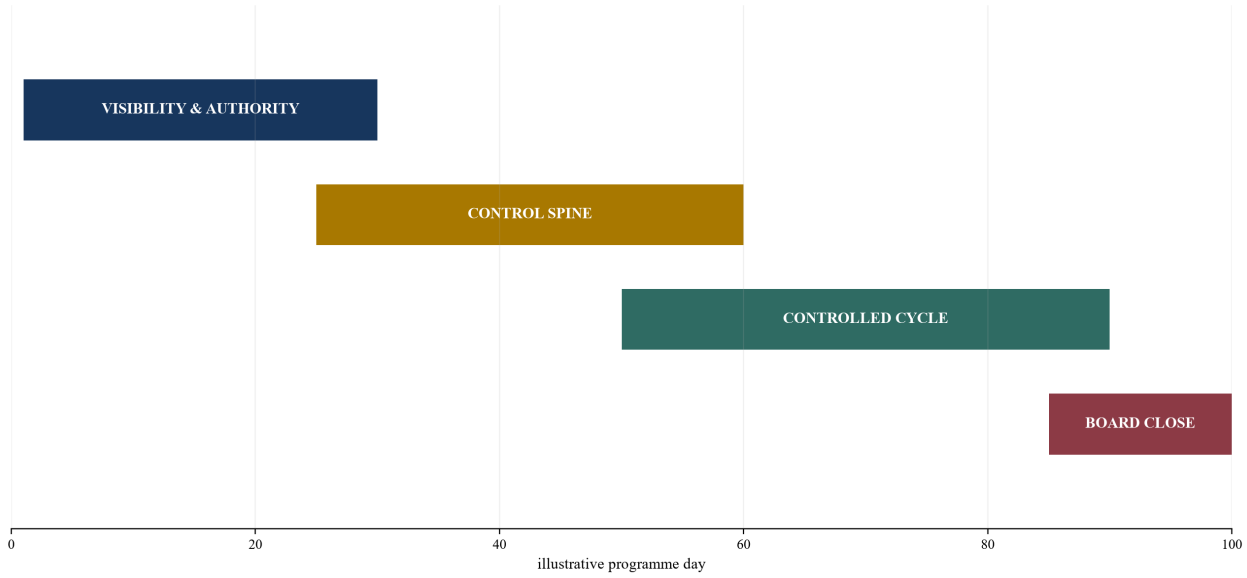

The timing is illustrative and should expand where system consequence, regulation or organisational readiness requires it.

References

[1] OECD, OECD AI Principles, updated 2024, https://www.oecd.org/en/topics/ai-principles.html

[2] OECD, AI, Data Governance and Privacy: Synergies and Areas of International Co-operation, 2024, https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/06/ai-data-governance-and-privacy_2ac13a42/2476b1a4-en.pdf

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

[4] National Institute of Standards and Technology, AI RMF Playbook, updated 2026, https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook

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

[6] National Institute of Standards and Technology, Trustworthy and Responsible AI Resource Center, 2026, https://airc.nist.gov/

[7] International Organization for Standardization, ISO/IEC 42001:2023 Artificial Intelligence Management System, 2023, https://www.iso.org/standard/42001.html

[8] International Organization for Standardization, ISO/IEC 23894:2023 Artificial Intelligence Risk Management, 2023, https://www.iso.org/standard/77304.html

[9] European Commission, AI Act: Regulatory Framework for Artificial Intelligence, updated 2026, https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai

[10] European Commission, The Enforcement Framework of the AI Act, 2026, https://digital-strategy.ec.europa.eu/en/policies/enforcement-ai-act

[11] European Commission, Guidelines on Transparency Obligations for Providers and Deployers of Certain AI Systems, 2026, https://digital-strategy.ec.europa.eu/en/policies/guidelines-ai-transparency-obligations

[12] European Commission, Guidelines on Obligations for General-Purpose AI Providers, 2025, https://digital-strategy.ec.europa.eu/en/faqs/guidelines-obligations-general-purpose-ai-providers

[13] EUR-Lex, Regulation EU 2016/679, General Data Protection Regulation, 2016, https://eur-lex.europa.eu/eli/reg/2016/679/oj

[14] UK Government, Implementing the UK's AI Regulatory Principles: Guidance for Regulators, 2024, https://assets.publishing.service.gov.uk/media/65c0b6bd63a23d0013c821a0/implementing_the_uk_ai_regulatory_principles_guidance_for_regulators.pdf

[15] Information Commissioner's Office, Guidance on AI and Data Protection, updated 2023, https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/

[16] Bank of England, Prudential Regulation Authority, SS1/23 Model Risk Management Principles for Banks, updated 2026, https://www.bankofengland.co.uk/prudential-regulation/publication/2023/may/model-risk-management-principles-for-banks-ss

[17] Cybersecurity and Infrastructure Security Agency and UK National Cyber Security Centre, Guidelines for Secure AI System Development, 2023, https://www.cisa.gov/news-events/alerts/2023/11/26/cisa-and-uk-ncsc-unveil-joint-guidelines-secure-ai-system-development

[18] Cybersecurity and Infrastructure Security Agency, Choosing Secure and Verifiable Technologies, 2024, https://www.cisa.gov/resources-tools/resources/choosing-secure-and-verifiable-technologies

[19] Federal Trade Commission, Artificial Intelligence: Enforcement and Business Guidance, updated 2026, https://www.ftc.gov/industry/technology/artificial-intelligence

[20] UAE Minister of State for Artificial Intelligence, Digital Economy and Remote Work Applications Office, UAE Position on AI Policy, 2024, https://ai.gov.ae/wp-content/uploads/2024/10/UAE-Guiding-on-ai-policy-EN-V3.pdf

[21] UAE Artificial Intelligence Office, AI Ethics Principles and Guidelines, 2022, https://ai.gov.ae/wp-content/uploads/2023/05/MOCAI-AI-Ethics-EN.pdf

[22] Saudi Data and Artificial Intelligence Authority, AI Ethics Principles, https://sdaia.gov.sa/en/SDAIA/about/Documents/ai-principles.pdf

[23] Saudi Data and Artificial Intelligence Authority, Regulations and Policies, updated 2026, https://sdaia.gov.sa/en/SDAIA/about/Pages/RegulationsAndPolicies.aspx

[24] Reserve Bank of India, FREE-AI Committee Report: Framework for Responsible and Ethical Enablement of Artificial Intelligence, 2025, https://www.rbi.org.in/Scripts/PublicationReportDetails.aspx?UrlPage=&ID=1299

[25] Singapore Personal Data Protection Commission, Model Artificial Intelligence Governance Framework, 2020, https://www.pdpc.gov.sg/help-and-resources/2020/01/model-ai-governance-framework

[26] Infocomm Media Development Authority of Singapore, Model AI Governance Framework for Generative AI, 2024, https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/press-releases/2024/model-ai-governance-framework-for-generative-ai

[27] Australian Securities and Investments Commission, REP 798 Beware the Gap: Governance Arrangements in the Face of AI Innovation, 2024, https://asic.gov.au/regulatory-resources/find-a-document/reports/rep-798-beware-the-gap-governance-arrangements-in-the-face-of-ai-innovation/

[28] Australian Prudential Regulation Authority, CPS 230 Operational Risk Management, updated 2026, https://www.apra.gov.au/standards/cps-230

[29] Office of the Superintendent of Financial Institutions Canada, Guideline E-23 Model Risk Management, 2025, https://www.osfi-bsif.gc.ca/en/guidance/guidance-library/guideline-e-23-model-risk-management-2027

[30] Office of the Superintendent of Financial Institutions Canada and Financial Consumer Agency of Canada, AI Uses and Risks at Federally Regulated Financial Institutions, 2024, https://www.osfi-bsif.gc.ca/en/about-osfi/reports-publications/osfi-fcac-risk-report-ai-uses-risks-federally-regulated-financial-institutions

[31] UK Infrastructure and Projects Authority and Cabinet Office, Guide for Effective Benefits Management in Major Projects, 2017, https://www.gov.uk/government/publications/guide-for-effective-benefits-management-in-major-projects

[32] UK Government, The Government Efficiency Framework, 2025, https://www.gov.uk/government/publications/the-government-efficiency-framework/the-government-efficiency-framework--2

[33] IFRS Foundation, Intangible Assets: Other Potential Test Cases, Including Data Resources and Artificial Intelligence-Related Data and Solutions, 2026, https://www.ifrs.org/content/dam/ifrs/meetings/2026/january/iasb/ap17d-other-potential-test-cases-data-ai.pdf

About the Author

Chennakeshav Adya, Independent Researcher