Responsible AI as Diligence

Scoring Governance before Funding the Model

Authored by Chennakeshav Adya
Independent Researcher
August 2026

Abstract

Artificial intelligence can strengthen product differentiation, reduce service cost and create new revenue. It can also introduce hidden liabilities through uncertain data rights, unreliable outputs, opaque third-party dependencies, cyber-security exposure, weak human oversight and rapidly developing regulation. Investors therefore need evidence that an AI company can operate its models and AI-enabled products responsibly as it scales.

This paper treats responsible AI as an investment-diligence discipline. It translates established international frameworks and official regulatory evidence into a system-level underwriting method. The NIST Artificial Intelligence Risk Management Framework organises risk work around Govern, Map, Measure and Manage. ISO/IEC 42001 specifies an AI management-system approach. The OECD AI Principles frame trustworthy AI around human rights, transparency, robustness, accountability and inclusive growth. Regulation (EU) 2024/1689 establishes risk-based obligations for providers and deployers in scope, including requirements for high-risk systems and general-purpose AI models. United Kingdom government guidance describes assurance techniques such as impact assessment, bias audit, compliance audit, conformity assessment, performance testing and formal verification. The Bank of England and Financial Conduct Authority's 2024 survey found extensive AI adoption in financial services and material reliance on third parties.

The framework begins with an inventory of systems, use cases, users, decisions, data, models, vendors and jurisdictions. It then tests responsibility, legal classification, data provenance, intellectual-property rights, model lineage, evaluation, human oversight, security, monitoring, incident response, customer disclosure and financial impact. Six figures describe the diligence funnel, governance stack, model and data lineage, risk heat map, 180-day remediation programme and investment gate. Six tables provide an inventory, evidence map, data-rights review, hypothetical governance score, post-investment dashboard and diligence-room register.

The worked score and all example thresholds, costs, probabilities and timetables are hypothetical management assumptions created solely to demonstrate the method. They are not benchmarks and should not be represented as market data. Actual obligations and outcomes depend on the product, use case, model, data, customer, contract, jurisdiction and deployment context. This paper provides general information for professional audiences and does not provide legal, regulatory, technical, cyber-security, data-protection, intellectual-property, accounting, investment or securities advice.

JEL Classification: G24, G32, G34, G38, L86, O33

Keywords: responsible AI, investment diligence, AI governance, model risk, venture capital, technology diligence, AI assurance, post-investment monitoring

1. Underwrite the governed system

An investor finances an operating system, not an abstract model. The relevant asset combines data, software, prompts, model weights or application programming interfaces, infrastructure, people, policies, customer workflows and commercial contracts. Governance quality determines whether that system can be sold, monitored, changed and defended over time.

Responsible-AI diligence should therefore ask for evidence at the level where value and harm arise. A board policy can establish intent. It does not reveal which AI systems are live, which decisions they influence, whether training and evaluation data can lawfully be used, how performance is measured, who can stop deployment or what happens when a vendor changes a model.

The underwriting objective is practical. Management should be able to identify every material AI system, show its intended use and limitation, name an accountable owner, produce current evaluation evidence, trace data and model dependencies, explain human oversight, monitor live behaviour and respond to incidents. The investment committee can then connect those controls to revenue quality, margin, regulatory access, customer confidence and exit readiness.

Figure 1. Responsible-AI investment diligence funnel

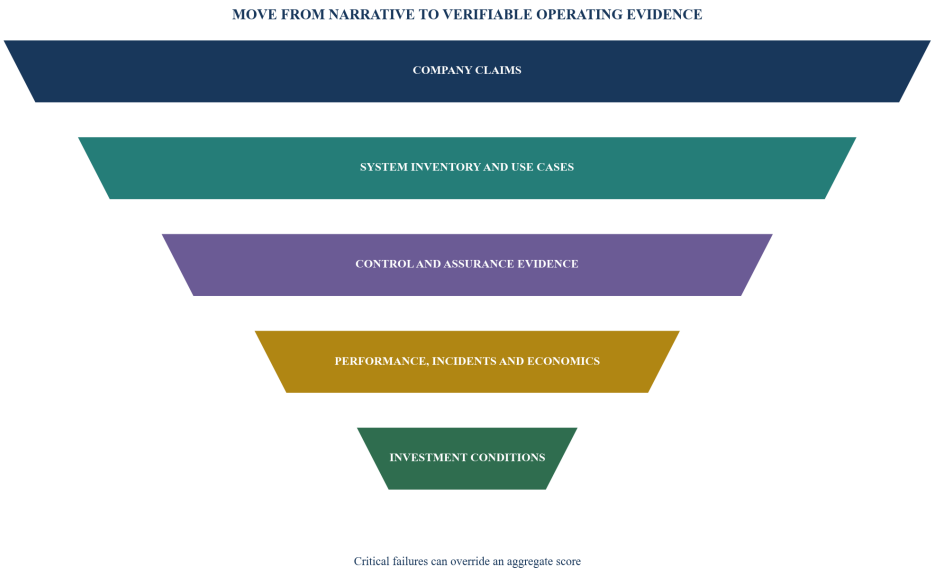

Author framework. Evidence narrows from the company narrative to systems, controls, performance and investment conditions.

2. Inventory systems before scoring governance

The first diligence deliverable is a complete AI-system inventory reconciled to products, infrastructure, vendor contracts and revenue. It should include internally developed models, third-party foundation models, embedded vendor features, decision engines, recommendation tools, automated agents and employee-facing systems that handle sensitive information.

Each record should identify the use case, intended users, affected people, decision or action, autonomy, data categories, model or service, vendor, deployment location, customer, jurisdiction, owner, release version and current status. Materiality should reflect financial, legal, safety, privacy, security and reputational impact.

Reconciliation matters because inventories can omit experimentation, acquired products and features activated by a software vendor. Procurement records, cloud bills, code repositories, data-processing registers, security tools, product roadmaps and customer contracts provide independent completeness tests. A company that cannot identify its AI systems cannot consistently classify, evaluate or monitor them.

Table 1. System-level AI inventory

Field	Diligence question	Evidence	Escalation trigger
Use case and outcome	what does the system do and for whom?	product specification, journey and customer terms	vague purpose or material secondary use
Decision and autonomy	what decision, recommendation or action occurs?	workflow, approval rights and override logs	consequential decision without defined human authority
Data	what training, retrieval, input and feedback data are used?	lineage, licence, consent, retention and access records	missing rights or uncontrolled sensitive data
Model and vendor	which model, version, provider and subcontractor apply?	architecture, model card, contract and change log	undisclosed dependency or unilateral material change
Evaluation	how is fitness for purpose demonstrated?	test plan, datasets, metrics, thresholds and results	no use-case test or failed critical threshold
Live control	how are behaviour, incidents and changes managed?	dashboards, alerts, incident log and release gate	no monitoring, owner or stop mechanism

The inventory should reconcile to product, vendor, cloud, security and revenue records.

3. Map international frameworks to evidence

International frameworks use different legal forms and terminology. Diligence should use them as evidence maps rather than interchangeable certificates. NIST describes the AI Risk Management Framework as voluntary, practical and use-case agnostic. Its Govern, Map, Measure and Manage functions support an operating cycle for identifying context, testing risk and acting on evidence.[1]

ISO/IEC 42001 specifies an AI management system and applies a continual-improvement approach to policies, objectives, processes and controls.[2] An investor should confirm the certified entity, sites, activities, standard edition, certification body and exclusions before relying on a certification claim. A management-system certificate does not independently establish the performance, legality or safety of every product.

The OECD AI Principles were adopted in 2019 and updated in 2024. They address inclusive growth, human rights and democratic values, transparency, robustness, security, safety and accountability.[3] The principles support cross-border consistency and help an investor test whether governance covers affected people as well as the company.

The evidence map should also include applicable sector rules, privacy law, consumer protection, intellectual property, cyber security, product safety and contracts. Responsible-AI diligence is strongest when management can connect a requirement to a control, an owner, a test and current operating evidence.

Table 2. Framework-to-evidence map

Source	Relevant discipline	Evidence expected in diligence	Limitation
NIST AI RMF	Govern, Map, Measure and Manage	context map, risk register, evaluation, monitoring and response	voluntary framework; implementation quality must be tested
NIST generative-AI profile	risks specific to generative AI	content, misuse, privacy, security and model-evaluation evidence	cross-sector profile; use-case controls remain necessary
ISO/IEC 42001	AI management system	policy, roles, objectives, risk process, audit and improvement	certificate scope may exclude a product, entity or location
OECD AI Principles	trustworthy and human-centred AI	accountability, transparency, robustness and affected-person analysis	principles do not replace binding law
EU AI Act	risk-based legal obligations	classification, technical documentation, records and conformity evidence where applicable	application depends on role, system, geography and timing
UK AI assurance guidance	portfolio of assurance techniques	impact assessment, audit, testing and independent evidence	technique selection must fit the system and claim

Applicability should be confirmed for the actual entity, system, use case and jurisdiction.

4. Classify legal exposure by role and use

The EU Artificial Intelligence Act illustrates why a company-level label is insufficient. Regulation (EU) 2024/1689 applies to defined actors including providers and deployers and contains specific regimes for prohibited practices, high-risk systems, transparency and general-purpose AI models.[4] A non-EU company can be in scope where relevant statutory conditions are met, including use of system output in the Union.

Diligence should create a role and classification record for every material system. It should identify the provider, deployer, importer, distributor or product manufacturer as applicable; intended purpose; prohibited-practice assessment; high-risk analysis; transparency duty; general-purpose model role; geographic nexus; and connected sector obligations.

The record should cite legal analysis and operating facts. A contract that calls the company a technology provider does not settle its statutory role. Product changes can alter classification, particularly when a general tool becomes intended for a consequential employment, credit, education, essential-service or safety use.

Management should monitor implementation dates and official amendments. Proposals should be identified as proposals until enacted. The investment model should include the cost of documentation, quality management, human oversight, logging, testing, registration or conformity work that the company's actual role requires.

5. Trace data and intellectual-property rights

Data rights can determine whether an AI product is commercially durable. The company should trace training, fine-tuning, retrieval, prompt, input, output and feedback data from source to use. Evidence should cover ownership or licence, collection basis, consent where relevant, privacy notices, purpose, geographic transfer, retention, deletion, security and restrictions on model training or vendor reuse.

Intellectual-property diligence should test source-code ownership, employee and contractor assignment, open-source obligations, model licences, dataset terms, web-scraping practices, customer-content rights and output allocation. The analysis should distinguish a contractual allocation from the separate question of whether an output qualifies for protection in a particular jurisdiction.

Customer contracts can prohibit use of confidential data to train shared models. Vendor terms can reserve rights to prompts or inputs unless an enterprise configuration is selected. A system can also reproduce protected or confidential material. Management needs controls that connect contractual permissions to product configuration and user behaviour.

Table 3. Data and intellectual-property rights review

Asset or flow	Evidence	Control test	Investment implication
Training dataset	source, licence, consent and provenance	sampled records reconcile to permitted use	remediation, retraining or claim exposure if rights are weak
Retrieval corpus	content inventory, permissions and access	user can retrieve only authorised content	customer confidentiality and data-separation risk
Prompts and inputs	product terms, privacy notice and vendor setting	sensitive data is blocked, minimised or contractually protected	enterprise-sales and privacy exposure
Model and code	ownership, assignment, model licence and open-source register	repository and dependency scan reconcile to records	freedom-to-operate and exit readiness
Output	customer terms, evaluation and similarity controls	known reproduction and attribution risks are tested	customer warranty and indemnity position
Feedback and logs	purpose, retention, access and deletion	operational logs match stated policy and contract	monitoring value versus privacy and secrecy cost

Legal conclusions require current jurisdiction-specific advice.

6. Map the model and vendor supply chain

Modern AI products often combine several external services. A user request may pass through identity, orchestration, retrieval, embedding, foundation model, safety, logging and analytics providers. Each service creates operational, contractual, security and change risk.

The architecture should identify the exact provider, model, version, hosting region, data path, retention setting, subprocessors, service level, usage limits and exit route. The company should preserve evaluation results by version. A vendor upgrade can change output quality, safety, latency and price even when the product interface appears unchanged.

Third-party exposure is commercially material. The Bank of England and FCA's 2024 survey reported that one third of responding firms' AI use cases were third-party implementations, a higher share than in their 2022 survey.[5] An investor should assess concentration, bargaining power, continuity, audit rights, notification, portability and the cost of switching.

Figure 2. AI governance and operating stack

Author framework. Governance connects board accountability to the technical and customer system.

Figure 3. Data, model and vendor lineage

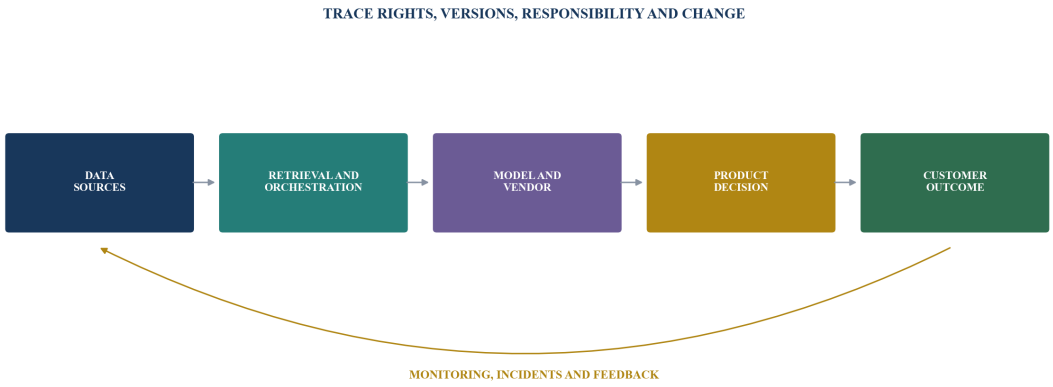

Author framework. Diligence should trace every material dependency from source to customer outcome and back to monitoring evidence.

7. Evaluate fitness for the intended use

Generic model benchmarks rarely prove fitness for a specific customer workflow. Diligence should begin with the intended task, population, operating environment, consequence and acceptable failure. Management should explain which metric represents value and which metrics protect customers, employees, counterparties and the company.

Evaluation evidence should identify the dataset, provenance, representativeness, sample size, version, metric, threshold, confidence interval where appropriate, reviewer and date. It should cover normal performance, edge cases, adversarial behaviour, distribution shift and predictable misuse. Generative systems may require factuality, groundedness, refusal, harmful-content, privacy, security and human-preference tests alongside task success.

The NIST generative-AI profile extends the AI RMF with risks and actions relevant to generative systems.[6] The investor can use it to test whether management's evaluation plan covers the company's actual exposures. The purpose is a reasoned control design, not completion of a checklist.

Claims should reconcile across sales material, model cards, customer contracts and test results. If management advertises a percentage improvement, diligence should establish the baseline, dataset, method and repeatability. A result from a laboratory dataset should not be represented as live customer performance without supporting evidence.

8. Test human oversight as a real control

Human oversight works only when the person has authority, information, competence and time. A nominal approval box can legitimise automation bias while adding little protection.

The workflow should identify what the human sees, what training is required, how uncertainty is communicated, which decisions can be overridden, when escalation is mandatory and how disagreement is recorded. The investor should sample actual cases and logs. High override rates can indicate model weakness; extremely low rates can indicate rubber stamping.

The company should define a stop mechanism for a system that breaches a critical threshold or generates harmful outcomes. Responsibility should remain clear during incidents. A vendor, customer and company cannot each assume that another party is monitoring the same risk.

9. Underwrite security and foreseeable misuse

AI expands the security surface through model endpoints, prompts, retrieval stores, training pipelines, generated code, plugins, agents and vendor integrations. Threat analysis should cover unauthorised access, prompt injection, data exfiltration, model extraction, poisoning, insecure output handling, privilege escalation, supply-chain compromise and abuse of the system's capability.

The diligence room should include architecture, threat models, access design, secrets management, logging, penetration tests, red-team exercises, vulnerability remediation and incident response. Testing should reflect the actual product and permissions. An agent with authority to send messages, move money or change records presents a different exposure from a drafting assistant.

Security findings should flow into valuation and financing. Remediation can delay enterprise revenue, increase infrastructure cost and require specialist leadership. Material unresolved findings can become closing conditions, indemnity subjects or reserved-budget items.

10. Monitor live behaviour and change

Pre-release testing captures a controlled snapshot. Live performance can move when customers, data, prompts, vendors, models, policies or attackers change. The company needs monitoring that detects material drift and a release process that reassesses risk after a change.

Indicators should reflect the use case. They may include task success, error severity, groundedness, override, complaint, harmful-output, latency, cost, data leakage, refusal, security alert and incident metrics. Thresholds should identify a warning, escalation and mandatory stop.

Model and prompt versions should be traceable to each significant output. Vendor change notices should enter the release gate. Emergency changes should be reviewed after deployment. A company that cannot reproduce which configuration served a customer will struggle to investigate complaints and establish contractual performance.

11. Examine incidents, complaints and redress

An empty incident register can mean a safe product, a new product or weak detection. Diligence should test the reporting culture and evidence sources rather than treat zero as automatically positive.

The review should reconcile incidents with security tickets, customer support, complaints, service credits, refunds, litigation, regulator contact, model alerts and vendor notices. Each material event should show detection, containment, customer impact, root cause, corrective action, owner and closure evidence.

Customer redress should match the consequence. A user needs a channel to challenge an outcome where appropriate and obtain meaningful review. Contractual limitations should be assessed alongside consumer, employment, data-protection and sector obligations.

12. Score governance with critical-failure gates

A score can make evidence comparable across systems and portfolio companies. It should not conceal a critical failure. The investment committee should define gates that require remediation, funding conditions or rejection irrespective of the aggregate result.

The following worked example assigns weights and scores only to demonstrate the method. Every figure is a hypothetical management assumption. It is not a benchmark. A real assessment should define scoring criteria before reviewing evidence, keep an audit trail and use more than one reviewer for material judgments.

Table 4. Hypothetical responsible-AI governance score

Domain	Hypothetical weight	Illustrative score out of 5	Weighted result	Example evidence gap
Accountability and inventory	12%	3.5	8.4	acquired product not fully inventoried
Legal classification and transparency	12%	3.0	7.2	two jurisdiction analyses need refresh
Data and intellectual-property rights	16%	2.5	8.0	legacy dataset provenance incomplete
Model and vendor management	10%	3.0	6.0	exit test not completed
Evaluation and robustness	18%	3.5	12.6	edge-case coverage inconsistent
Human oversight and redress	10%	3.0	6.0	override-quality review absent
Security and misuse	12%	4.0	9.6	agent red team requires expansion
Monitoring and incidents	10%	3.5	7.0	customer metrics need reconciliation
Total	100%		64.8 out of 100	remediation required before scale capital

All weights, scores and thresholds are hypothetical management assumptions and are not market benchmarks.

Critical gates can include missing rights to a core dataset, an unlawful prohibited practice, a safety-critical system without fit-for-purpose evaluation, an undisclosed material incident, inability to identify models in production or lack of authority to suspend a harmful system. The gate should state the decision and evidence required to clear it.

Figure 4. Illustrative AI diligence risk heat map

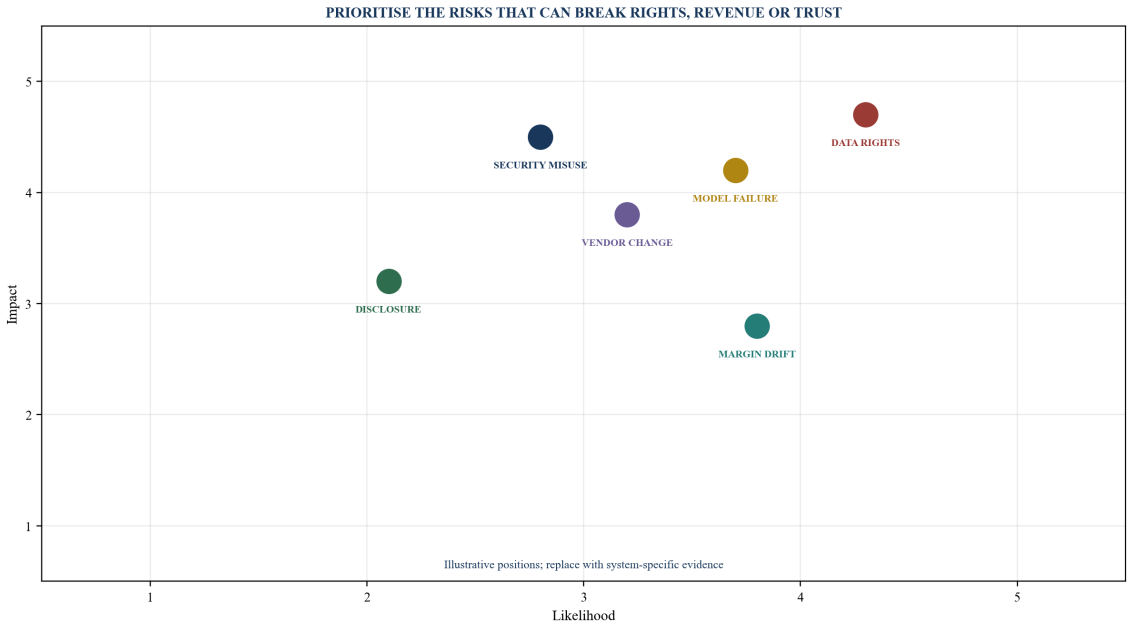

Positions are hypothetical management assumptions used solely to demonstrate prioritisation.

13. Connect governance to unit economics

Responsible-AI controls affect both cost and revenue. Evaluation, monitoring, data licensing, security, assurance, specialist staff and human review can raise cost to serve. Strong evidence can support enterprise procurement, regulated customers, insurance and lower remediation risk.

The financial model should include model inference, retrieval, monitoring, logging, review, appeal, support, evaluation and compliance costs by customer or transaction. It should test vendor price changes, higher human-review rates, larger context windows, data-licence fees and required geographic hosting.

Gross margin can deteriorate when management prices a software product before measuring the complete governed service. The diligence team should reconcile cloud and vendor invoices to product telemetry and revenue. Customer-specific evaluation and reporting obligations should appear in the contract margin.

14. Test claims in enterprise contracts

AI contracts increasingly allocate responsibility for data, configuration, output, security, documentation, monitoring and regulatory cooperation. Diligence should compare the product's actual capability with warranties, service descriptions and sales claims.

The review should cover input and output rights, vendor training, confidential information, security, data processing, model changes, performance, prohibited uses, customer responsibilities, audit, incident notice, indemnity, liability, termination and portability. A company can accept obligations that its upstream vendor does not support.

Customer due-diligence questionnaires provide commercial evidence. Repeated delays or concessions can reveal missing controls. A strong evidence library can shorten procurement and reduce bespoke responses, creating a measurable return on governance investment.

15. Condition investment on remediation

Diligence findings should become an owned execution plan. Each item needs a risk statement, affected system, required action, evidence of completion, owner, budget, due date and escalation threshold.

Closing conditions are appropriate when the investor cannot accept the exposure at completion. Post-closing covenants can address important work that has a credible owner and funded timetable. Reserved capital can protect delivery where remediation cost is material. A board observer or committee can receive defined indicators without assuming operational responsibility.

The transaction document should avoid vague commitments to follow best practice. A verifiable condition might require a reconciled system inventory, independent security test, documented dataset rights, completion of specified evaluations, adoption of a vendor-change gate or closure of a named incident action.

16. Establish post-investment monitoring

Portfolio monitoring should focus on change and material exposure. The board does not need every technical metric. It needs a reliable view of systems, critical thresholds, incidents, remediation, vendor concentration, regulatory change and commercial impact.

Table 5. Post-investment responsible-AI dashboard

Indicator	Board question	Evidence source	Escalation example
Material systems and changes	what entered, changed or left production?	inventory and release register	unclassified material system deployed
Evaluation status	do current versions meet approved thresholds?	evaluation registry and exceptions	critical metric breach or expired test
Data and vendor exposure	have rights, terms or dependencies changed?	contract, lineage and vendor register	core right disputed or vendor exit impaired
Human oversight	are reviewers effective and sufficiently resourced?	queue, override, quality and staffing data	backlog or quality below approved threshold
Incidents and complaints	what affected customers or obligations?	incident, complaint and support records	material harm, notification or repeat cause
Economics	does governed delivery sustain target margin?	product telemetry, vendor invoices and finance	cost per outcome exceeds approved case

Thresholds should be set for the actual system and approved risk appetite.

17. Build a 180-day remediation programme

During the first thirty days, management should reconcile the system inventory, appoint accountable owners and close any critical data-rights, security or legal-classification question that prevents continued use. The board should approve critical stop thresholds.

During days 31 to 60, the company should standardise system records, evaluation plans, vendor evidence, human-oversight design and incident taxonomy. It should select the highest-risk systems for independent testing or assurance.

During days 61 to 120, the company should complete priority evaluations, deploy monitoring, test escalation and reconcile customer claims to evidence. Finance should measure the complete cost of governed delivery. Commercial teams should receive approved statements about capability and limitation.

During days 121 to 180, internal audit or an independent reviewer should sample operating evidence. Management should close findings, update the board score and integrate governance gates into product, procurement, sales and acquisition processes.

Figure 5. Responsible-AI 180-day remediation roadmap

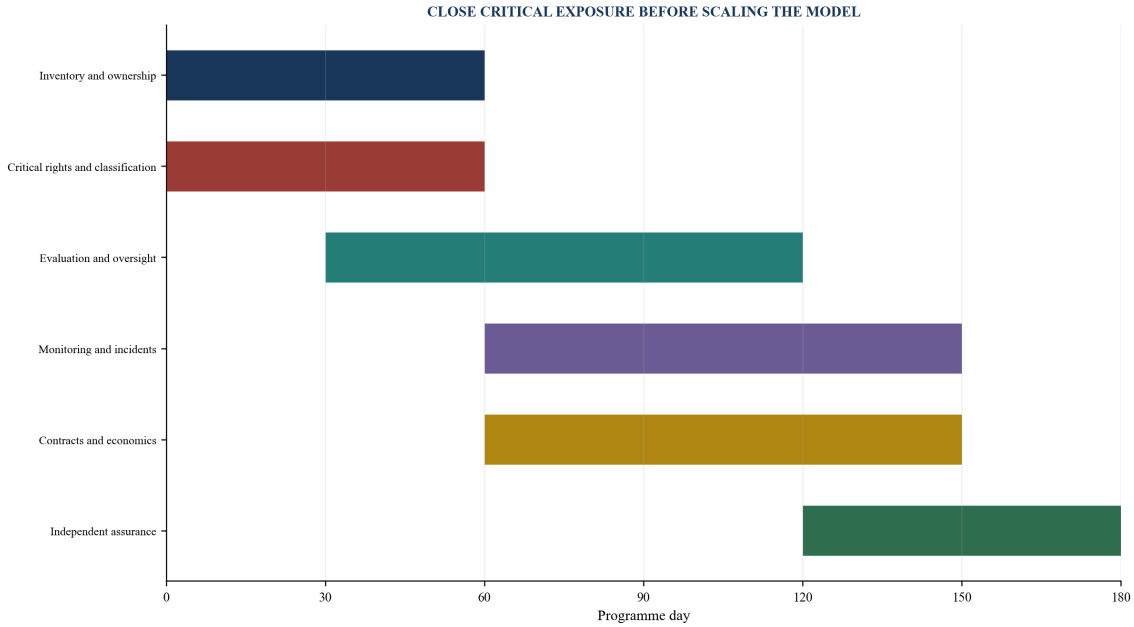

Author framework. Actual priorities and duration depend on the systems and findings.

18. Organise the diligence room

The diligence room should allow a reviewer to move from the company list to a system, its requirements, controls, evidence, results, incidents and remediation. Files should be current, versioned and tied to an owner.

Table 6. Responsible-AI diligence register

Workstream	Core evidence	Principal test	Red flag
Governance	policy, roles, committee papers and risk appetite	can accountable leaders stop or constrain a system?	responsibility diffused across teams
Inventory and classification	reconciled register and legal analysis	are all material systems and roles identified?	production system absent or stale classification
Data and intellectual property	lineage, rights, notices, assignments and licences	can every material asset be used as deployed?	core dataset or code ownership unresolved
Model, vendor and security	architecture, contracts, testing and exit plan	are dependencies controlled and replaceable?	untested critical vendor or material security issue
Evaluation and human oversight	plans, datasets, results, review and overrides	does evidence demonstrate fitness for intended use?	claim unsupported or reviewer lacks authority
Monitoring and incidents	telemetry, thresholds, logs, complaints and actions	can the company detect, contain and learn from failure?	no stop control or undisclosed material event

Evidence should reconcile across legal, product, security, finance, customer and board records.

The investor should sample rather than accept a polished index. Selected systems should be traced end to end. Contract representations, investor materials and customer claims should reconcile to the same evidence.

19. Use an investment gate that can reject the model

The investment committee should decide whether governance supports the proposed valuation and growth plan. Approval can require five elements: lawful and contractually supported inputs; fit-for-purpose evaluation; accountable human and technical controls; live monitoring and incident response; and economics that include the full cost of governed delivery.

The committee should name the critical failures, closing conditions, post-closing actions, reserved budget, reporting thresholds and authority to pause scale. A high composite score should not override a failed critical gate.

Figure 6. Responsible-AI investment gate

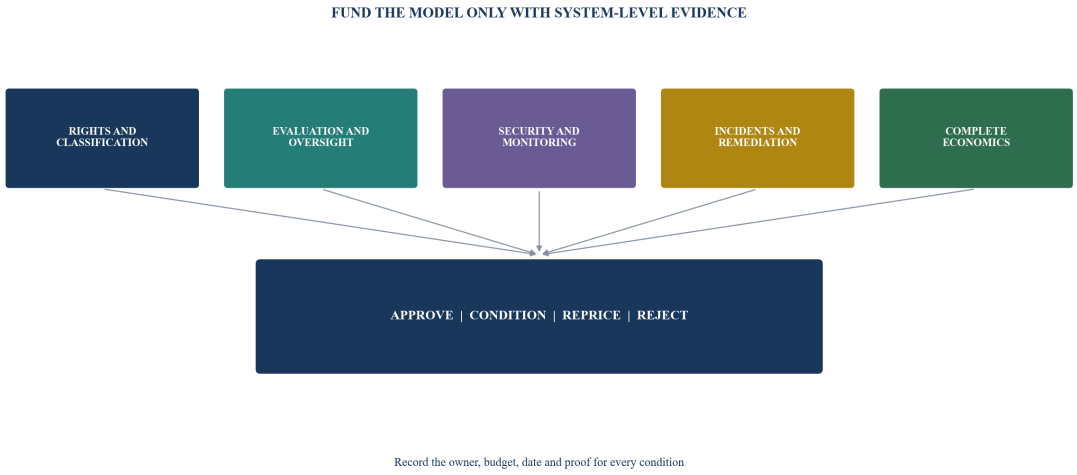

Author framework. The decision reflects evidence, critical failures, economics and the funded remediation plan.

20. Make governance a commercial capability

Responsible-AI governance can become a repeatable commercial capability when it is embedded into product design, procurement, contracting and board reporting. A reconciled inventory reduces discovery time. Standard evaluation records make claims defensible. Vendor and data lineage accelerate customer review. Monitoring and incident processes support regulated and enterprise buyers.

The durable advantage comes from evidence that travels with the product. As models, vendors and laws change, the company can identify affected systems, rerun relevant tests, update documentation and communicate accurately. That capacity supports international expansion and acquisition integration.

Investment diligence should reward credible operating evidence and price the cost of missing controls. The resulting decision is more useful than a generic responsible-AI label: it identifies which systems can scale, which conditions must be met and which exposures can destroy value.

Conclusion

AI governance belongs inside commercial and investment underwriting. The investor should identify every material system, connect it to a use case, trace its data and model dependencies, test legal classification, examine evaluation and human oversight, review security and live monitoring, investigate incidents and include the complete control cost in unit economics.

International frameworks provide a disciplined evidence map. The NIST AI RMF, ISO/IEC 42001, OECD AI Principles, EU AI Act and United Kingdom assurance guidance each contribute useful requirements or techniques. Their value in diligence depends on system-level implementation evidence.

A scorecard can support comparison. Critical-failure gates preserve judgment where rights, legality, safety, security or disclosure are materially deficient. Funding conditions can turn findings into dated execution with accountable owners, reserved capital and board monitoring.

The central investment question is whether the company can demonstrate that its AI systems work as intended, remain within authorised and contractual boundaries, respond to change and sustain attractive economics. That evidence supports responsible scale and a more defensible valuation.

References

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

[2] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10

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

[4] International Organization for Standardization, ISO/IEC 42001 Artificial intelligence management system, https://www.iso.org/standard/42001

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

[6] Organisation for Economic Co-operation and Development, The 2024 OECD AI Principles update, https://oecd.ai/en/wonk/evolving-with-innovation-the-2024-oecd-ai-principles-update

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

[8] European Commission, Navigating the AI Act, https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act

[9] United Kingdom Department for Science, Innovation and Technology, Introduction to AI assurance, https://www.gov.uk/government/publications/introduction-to-ai-assurance/introduction-to-ai-assurance

[10] United Kingdom Department for Science, Innovation and Technology, Portfolio of AI assurance techniques, https://www.gov.uk/guidance/portfolio-of-ai-assurance-techniques

[11] Bank of England and Financial Conduct Authority, Artificial intelligence in UK financial services 2024, https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024

[12] National Institute of Standards and Technology, AI Resource Center, https://airc.nist.gov/

About the Author

Chennakeshav Adya is an independent researcher and Managing Partner of Matchpoint Partners. His work examines strategy, capital formation, valuation, transactions and operating execution across private and public markets.