AI Governance before Series A: The Policies Investors Now Expect

A Board Framework for Accountable AI, Data Rights, Model Evidence, Product Claims, Security, Procurement and Investor-Ready Scale

Authored by Chennakeshav Adya, Independent Researcher

August 2026

Abstract

Artificial intelligence can strengthen a startup's product, operating leverage and strategic position. It also creates questions about data rights, reliability, security, customer outcomes, regulatory exposure and the accuracy of commercial claims. These questions become consequential before a Series A process because the company is asking investors to fund a repeatable business, enterprise customers to depend on its systems and employees to scale decisions that may previously have remained with the founders.

This paper develops a board-ready governance framework for an AI-enabled startup preparing for institutional capital. It treats governance as operating infrastructure connecting product design, data provenance, model development, vendor selection, testing, release, monitoring, human oversight, incident response, customer communication and board accountability. It does not assume that every investor applies the same checklist. Actual diligence requirements vary by fund, stage, jurisdiction, sector, customer base and system consequence. The framework identifies evidence commonly relevant to institutional investors, enterprise customers, regulated partners and public authorities.

The analysis uses current official sources across several jurisdictions. The National Institute of Standards and Technology describes its AI Risk Management Framework as a voluntary resource for incorporating trustworthiness considerations into the design, development, use and evaluation of AI systems. NIST's generative-AI profile adds risk-management actions for generative systems. ISO/IEC 42001 specifies requirements for an artificial-intelligence management system. The European Union's AI Act applies progressively, with general-purpose-AI governance already in force and further transparency and high-risk obligations following the official implementation timetable. The UAE Charter for the Development and Use of Artificial Intelligence sets national principles for responsible development and use, while the UAE Personal Data Protection Law governs personal-data processing. Saudi Arabia's AI Ethics Principles apply a risk-based lifecycle approach. Singapore's model frameworks address traditional, generative and agentic AI. The UK Information Commissioner's Office provides an AI and data-protection risk toolkit, and the Financial Conduct Authority applies existing governance, accountability and customer-outcome frameworks to AI in financial services. OECD principles connect trustworthy AI to human rights, transparency, robustness, safety and accountability. United States enforcement materials demonstrate that unsupported AI claims can create securities and consumer-protection exposure.

The recommended programme begins with an inventory of systems, data, models, vendors and decisions. It then assigns risk tiers, owners, permissible uses and release evidence. Six original figures and six implementation tables translate the framework into an investibility architecture, lifecycle gates, an agentic-authority ladder, a unit-economics bridge, a funding and cap-table waterfall and a 120-day execution plan. Financial values, funnel metrics, weights, timelines and valuation effects shown in the paper are illustrative management assumptions. Legal, regulatory,

privacy, cybersecurity, employment, intellectual-property, accounting, tax, competition, valuation and investment conclusions require confirmation by qualified advisers for the relevant company, product and jurisdiction.

JEL Classification: G24, G32, G34, L26, M13, M15, O32, O33

Keywords: AI governance, Series A, venture capital, model risk, data rights, responsible AI, artificial intelligence, generative AI, agentic AI, due diligence, startup governance, enterprise procurement

1. Treat AI governance as investibility infrastructure

An AI startup preparing for Series A needs evidence that its product can scale without losing control of data, models, customer outcomes or commercial claims. Governance provides the operating system for that evidence. It assigns decision rights, defines approved use, requires tests before release, records exceptions and gives the board a way to connect product velocity with risk, cash and customer value.

The starting point is the company's actual AI role. A business may train a model, fine-tune a third-party model, retrieve information for a foundation model, deploy an agent, embed an external service, provide infrastructure or use AI only inside operations. Each role creates different dependencies and duties. The board should document the product outcome, the system boundary, the affected people and organisations, the material decisions, the jurisdictions and the parties that supply data, models, tools and distribution.

Governance should support execution. Product teams need a clear route from experiment to controlled release. Commercial teams need accurate language about capabilities, limitations and customer responsibilities. Enterprise buyers need evidence covering security, privacy, availability, model performance, subcontractors, incidents and exit. Investors need to understand whether growth relies on rights, capabilities and economics the company can retain. Employees need escalation routes when the system behaves outside its approved purpose.

NIST organises AI risk management through Govern, Map, Measure and Manage. ISO/IEC 42001 uses a management-system approach for establishing, implementing, maintaining and continually improving AI governance. These frameworks support a practical board structure: define accountability, understand context, produce evidence, manage residual risk and revisit decisions as the system changes.

The board should approve five linked proofs. Product proof connects the system to a consequential customer outcome. Rights proof covers data, software, models and content. Performance proof covers accuracy, robustness, safety and monitoring. Commercial proof covers procurement, distribution, unit economics and claims. Governance proof covers accountability, incidents, regulation and change. A material weakness in one proof should return the release or investment case for remediation.

Figure 1. AI governance as investibility infrastructure

CONNECT AI EXECUTION TO INVESTOR-READY EVIDENCE

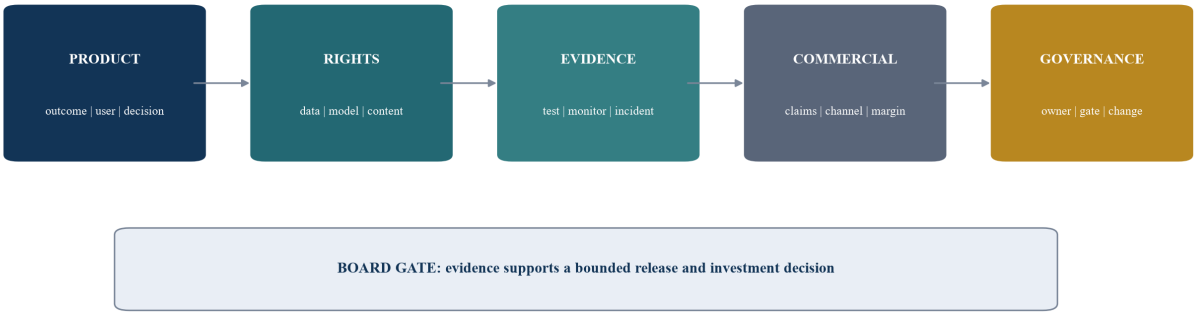

Five linked proofs connect product execution to a board-controlled scale decision.

2. Define the system, outcome and risk tier

Governance begins with a precise system definition. The company should identify the user, intended purpose, prohibited purpose, inputs, model, tools, retrieval sources, output, action, affected party and accountable human. A generic label such as AI assistant or intelligent platform hides the difference between summarising information, recommending a financial action, approving credit, selecting a candidate, diagnosing a condition or authorising a payment.

The definition should show whether the system only produces content or can also call tools, change records, send communications, move money, alter access, purchase services or affect a person's rights. Agentic systems require a map of available tools, credentials, memory, external services, decision checkpoints and termination controls. Singapore's Model AI Governance Framework for Agentic AI emphasises bounding use cases and powers, meaningful human accountability, lifecycle controls and end-user responsibility.

Risk tiering should follow consequence and exposure. The assessment can consider safety, financial loss, legal rights, vulnerable users, personal or confidential data, scale, reversibility, autonomy, model opacity, cybersecurity, regulatory status and availability of human review. The company should also assess the risk created when several low-consequence components combine into a higher-consequence workflow.

The board should define which tier can be approved by a product leader, which requires security, privacy or legal review, and which requires executive or board approval. High-consequence systems should have stronger evidence, narrower permissions, independent testing and explicit stop rules. Prohibited uses should be written, communicated and enforced through technical and contractual controls.

The classification record should be versioned. A model change, new data source, new customer group, additional jurisdiction, wider autonomy or changed commercial claim can alter the risk tier. A system previously used to draft internal text may need new governance when it begins

responding to customers or executing transactions.

Table 1. AI system definition and risk-tier register

Register field	Evidence question	Example evidence	Reassessment trigger
intended outcome	which user decision or operating result does the system support?	approved use-case statement, user journey and outcome measure	new customer segment, workflow or product promise
system boundary	which models, tools, data stores, vendors and people are inside the service?	architecture, dependency map and responsibility matrix	new foundation model, agent tool, retrieval source or processor
affected parties	whose rights, money, safety, opportunity, privacy or access may change?	impact map, complaints analysis and customer terms	new affected population or higher scale
consequence	what credible harm follows from error, misuse, attack, drift or unavailability?	scenario tests, loss analysis and recovery design	material incident, new autonomy or reduced human review
risk tier	what approval, testing, monitoring and oversight follows from consequence?	approved scoring record and control profile	material model, data, jurisdiction or purpose change
prohibited use	which uses exceed rights, evidence, product scope or risk appetite?	policy, technical restriction, contract and employee training	request for a new use or evidence of misuse

The tier is a management decision supported by qualified review where law or regulated activity is engaged.

3. Establish board accountability and decision rights

The board remains accountable for the company even when a model, vendor or customer performs part of the workflow. It should approve the AI risk appetite, material use cases, accountability model, reporting cadence and escalation thresholds. Management should convert that direction into named owners and operating procedures.

Accountability should follow the full outcome. A product owner can own the intended purpose and customer result. A model owner can own performance, evaluation and change. A data owner can own provenance, quality, permission and retention. Security can own threat controls and incident response. Legal and privacy advisers can assess applicable duties. Commercial leadership can own claims and partner promises. An executive sponsor can resolve conflicts and accept residual risk within delegated authority.

The board should avoid fragmented approval. A system can pass a technical evaluation while failing the rights, customer, procurement or economics case. One release memo should reconcile all material evidence and exceptions. The memo should state the use, tier, version, affected markets, tests, limitations, monitoring, incidents, open issues, accepted residual risk, accountable owners and next review date.

Committees should be proportionate to stage. A pre-Series A company may use a weekly cross-functional release forum supported by a monthly executive review and quarterly board reporting. The forum needs authority to hold a release, narrow its scope or require remediation. Material incidents and changes should bypass the normal cadence.

Board reporting should connect risk and business performance. Useful measures include active systems by tier, overdue reviews, evaluation results, security findings, data-rights exceptions, customer complaints, human overrides, incidents, vendor concentration, model cost, gross margin and cash runway. Reporting should distinguish observed results from forecasts and management scenarios.

4. Build an AI system and dependency inventory

An investor-ready company should know every material AI system it develops, buys, embeds or permits employees to use. The inventory should include production products, customer-specific configurations, internal copilots, experiments connected to real data, shadow tools, models used by contractors and AI features inside enterprise software.

Each entry should identify the owner, purpose, users, risk tier, status, model and version, provider, hosting, data categories, jurisdictions, tools, retrieval sources, downstream decisions, human control, evaluations, monitoring, incidents, contracts and review date. The inventory should link to the underlying evidence rather than becoming a static spreadsheet detached from operations.

Dependency analysis should reach beyond the named model. A service may depend on a cloud region, vector database, orchestration framework, content filter, observability tool, identity provider, external API and customer data store. The company should identify which dependency can change behaviour, terms, price, location, access or availability. It should record substitution time and the effect of termination.

Open-source components require provenance, licence and security review. Model weights, training code, datasets and libraries may carry different licences and obligations. The company should record the version, source, modification, permitted use, attribution, distribution terms and known vulnerabilities. Intellectual-property conclusions should be confirmed by qualified counsel.

Inventory controls should connect to engineering. New model credentials, packages, endpoints and production deployments can trigger registration or update. Procurement can block an unapproved vendor. Access control can limit sensitive data to authorised systems. Monitoring can reconcile live endpoints with approved records. The objective is current operational visibility.

Figure 2. AI lifecycle and evidence gates

GOVERN THE SYSTEM THROUGH ITS FULL LIFECYCLE

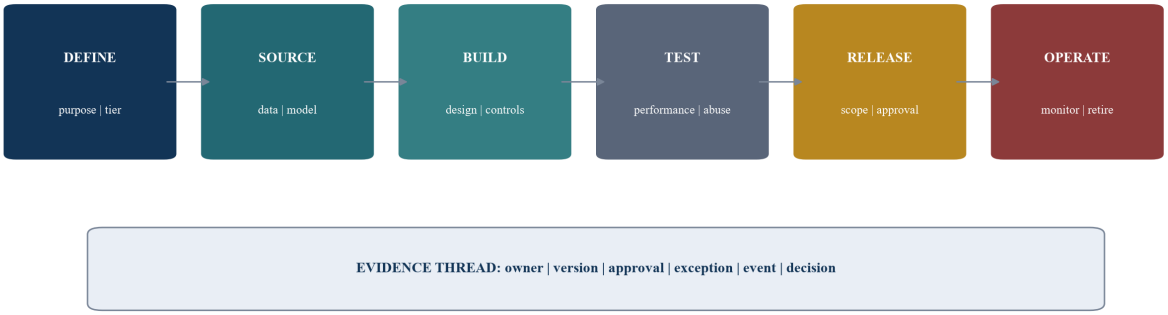

Evidence accumulates from system definition through monitoring and controlled retirement.

5. Map data rights, provenance, privacy and retention

Data governance should answer whether the company has the right to collect, access, transform, combine, train on, evaluate with, disclose and retain each data category. A customer contract, public webpage, vendor licence, employee upload or open-source repository does not automatically establish every intended right. The company should maintain a purpose and rights record for training data, evaluation data, retrieval content, prompts, outputs, telemetry and human feedback.

The record should identify the source, controller or owner, legal basis or contractual permission, purpose, geographic location, sensitivity, access, processor, retention, deletion and restrictions on model training or reuse. Personal data should be mapped to applicable privacy law. Confidential and regulated data should be mapped to contractual and sector duties. Copyright, database rights, trade secrets and publicity rights may require separate analysis.

The UAE Personal Data Protection Law establishes duties for processing personal data and rights for data subjects. The UK ICO's AI toolkit provides practical risk areas across the AI lifecycle and is under review following the Data (Use and Access) Act; companies should recheck current guidance before relying on it. The EU AI Act, GDPR and sector rules can operate together. Saudi and Singapore frameworks also connect responsible AI to privacy and data governance.

Data minimisation belongs in product design. The company should identify the minimum data required for the approved outcome, limit access, separate environments and set deletion or de-identification rules. Retrieval systems should enforce document-level permissions. Prompt and output logs should be assessed for sensitive data. Model training should be separated from service delivery when rights and customer expectations differ.

Provenance should survive transformation. The company should be able to trace a material output to the model version, prompt or input, retrieval sources, policy configuration and relevant human action. This supports evaluation, correction, customer explanation, incident investigation and

investor diligence. Where complete traceability is technically unavailable, the limitation should be documented and reflected in system scope.

6. Govern foundation models, vendors and open-source components

Third-party models can accelerate product development while creating concentration, pricing, data, performance and continuity risks. The company should assess the provider, service, contract, model card or technical documentation, data-use terms, security, privacy, hosting, subprocessors, service levels, change policy, evaluation access, intellectual-property terms, indemnities, incident notice, termination and data return or deletion.

The vendor assessment should be use-case specific. A model acceptable for drafting internal text may be unsuitable for a customer decision with financial or legal consequences. The company should test the exact model, configuration, prompts, tools, retrieval data and guardrails used in production. Provider benchmarks do not replace company evaluation.

Model change management is material. Providers may update weights, routing, filters, context limits, pricing or terms. The company should know whether a version can be pinned, how advance notice works and what happens when a version is retired. A material change should trigger regression testing and a release decision. If the service uses dynamic model routing, the approval should cover the routing logic and every eligible model.

Commercial resilience requires an exit plan. The board should understand the effort to replace the provider, export data, reproduce evaluations, migrate prompts and tools, maintain customer service and reapprove the system. Multi-model design can reduce some concentration while raising integration, testing and cost complexity. The choice should follow verified economics and service requirements.

Table 2. Model, vendor and data dependency register

Dependency	Diligence evidence	Operating control	Exit evidence
foundation model	model documentation, terms, security, privacy, hosting and change policy	approved version, configuration, evaluation, access and monitoring	substitute model test, migration steps, time and cost
training or tuning data	source, rights, provenance, quality, sensitivity and restrictions	access, purpose limitation, versioning, validation and deletion	retraining or removal method and affected model versions
retrieval content	content authority, freshness, permissions, confidentiality and jurisdiction	document-level access, source citation, update and removal	export, re-index, deletion and customer transition
open-source component	repository, version, licence, vulnerability and maintenance status	software bill of materials, scanning, patching and approved use	replacement component and compatibility test
orchestration or agent tool	permissions, credentials, callable actions, logs and failure behaviour	allowlist, least privilege, approval point, limits and kill control	credential revocation, workflow migration and state recovery
cloud and observability	location, security, data use, retention, subprocessors and availability	encryption, isolation, logging, service monitoring and incident response	data export, deletion evidence, alternative region or provider

Contractual and legal conclusions require review of the executed terms and the intended use.

7. Set development, testing and release controls

AI development needs a reproducible path from experiment to production. The company should record the problem, baseline, data, model, configuration, code, evaluation set, acceptance criteria, reviewer, release scope and monitoring plan. The evidence should allow a qualified person to understand what changed and why the release was authorised.

Testing should cover the approved use and credible failure modes. Performance measures can include accuracy, precision, recall, calibration, retrieval quality, factual support, instruction following, stability, latency and cost. Consequence-based tests can cover bias, harmful content, privacy leakage, prompt injection, data exfiltration, insecure tool use, denial of service, model extraction, adversarial inputs and failure under incomplete or unusual data.

Generative systems require evaluation of outputs in context. A single overall score can hide material failures by language, customer group, workflow or task. The evaluation set should represent the production distribution, important edge cases and known abuse paths. Human review should use defined criteria, trained reviewers and reconciliation where judgement matters. Synthetic tests can expand coverage while remaining clearly identified as synthetic.

Release criteria should be approved before results are seen. The record should state minimum performance, zero-tolerance conditions, acceptable residual risk, unresolved issues, rollout limit and stop rules. A model can be technically strong and still remain outside release because rights, security, procurement, customer communication or support are incomplete.

Controlled rollout reduces exposure while generating evidence. The company can limit customers, data, transactions, tools, autonomy or volume. Monitoring should compare production with evaluation assumptions. Material drift, incident or customer harm should lead to investigation and a documented decision to continue, narrow, roll back or retire.

8. Measure performance, robustness, bias and failure

Model performance should be defined in relation to a business decision and the people affected. A summarisation system can be evaluated for coverage, unsupported statements and citation accuracy. A forecasting system can be evaluated for error, calibration and stability. A ranking system can be evaluated for relevance, exposure and group outcomes. An agent can be evaluated for task success, tool selection, action accuracy, policy compliance and recovery.

Robustness testing examines variation. Inputs may change in length, language, format, noise, ambiguity or intent. Data distributions may shift. A retrieval source may become stale or unavailable. A vendor can update a model. The company should test which changes affect the outcome and set monitoring for observable signals.

Fairness analysis should follow the decision and legal context. The team should identify relevant groups, potential allocation or quality harms, data limitations, permissible attributes and mitigation options. Aggregate parity does not by itself establish a fair outcome. Qualified advisers should confirm employment, credit, insurance, consumer and other sector requirements where applicable.

Failure management should be designed before launch. The system needs a safe response when confidence is low, evidence conflicts, a tool fails or the requested action exceeds authority.

Options include asking for clarification, presenting uncertainty, refusing the action, routing to a human, using a deterministic fallback, restricting scope or stopping the workflow.

The monitoring plan should assign an owner to each metric and threshold. It should include performance, latency, cost, security, privacy, complaints, override, incident and customer outcome. Alerts need investigation procedures and decision authority. Investors and enterprise buyers can then review a working control system rather than a policy without operating evidence.

Figure 3. Generative and agentic AI authority ladder

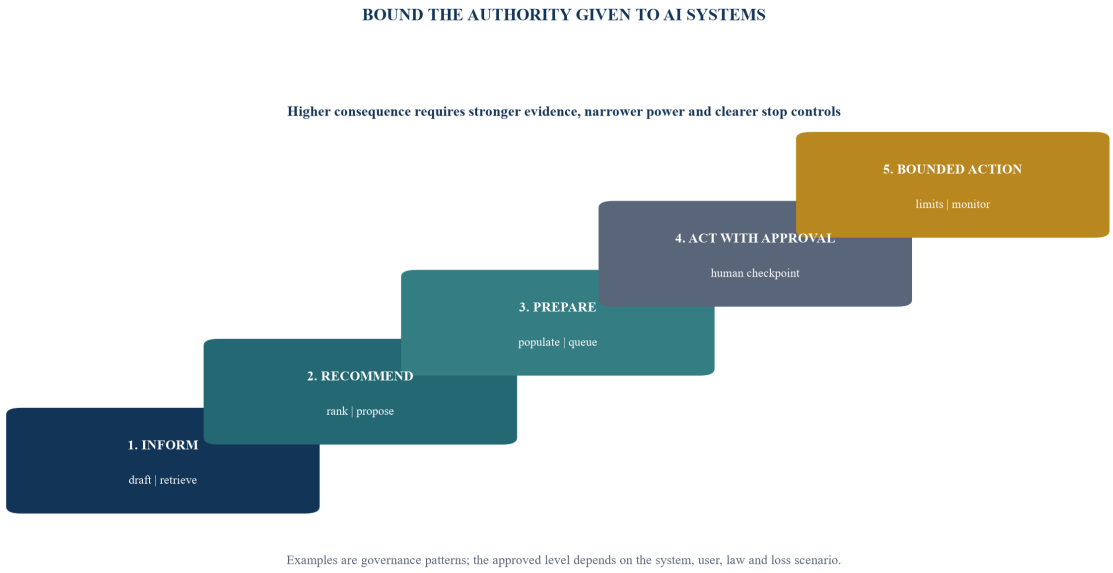

Authority rises only when evidence, permissions and human accountability support the consequence.

9. Control generative AI and agentic action

Generative AI can create fluent outputs that remain incomplete, unsupported or unsuitable for a particular decision. The product should distinguish generated content from verified facts, require source grounding where appropriate and disclose meaningful limitations to the user. High-consequence outputs should reach a qualified human before action.

Agentic AI adds action rights. An agent may call tools, browse systems, modify records, communicate externally, place orders or move funds. Governance should define the allowable goal, tools, data, credentials, transaction limits, time limits, rate limits, approval checkpoints, prohibited actions, logs, monitoring, termination and recovery. Permissions should follow least privilege and be issued for the shortest practical period.

The authority ladder in Figure 3 separates informing, recommending, preparing and acting. A company can gain operating value at lower levels while accumulating evidence for a higher level. The board should resist treating autonomy as a product objective by itself. The commercial case should state which action improves the customer outcome and why the added authority is worth the added risk and control cost.

Human oversight should be meaningful. A person cannot review hundreds of complex actions in seconds or approve a recommendation without relevant information. The process should give the reviewer time, competence, evidence, authority and a usable way to reject or modify the action. Repeated automatic approval is a signal that the control needs redesign.

The company should test compound failure. An agent may receive a malicious instruction from retrieved content, select an over-privileged tool, expose confidential data and then create an external communication. Scenario testing should cross component boundaries. Logs should preserve the goal, plan, tool calls, inputs, outputs, approvals, errors and final action while respecting privacy and security requirements.

10. Secure models, data, prompts and interfaces

AI security should sit inside the company's wider security programme and address model-specific threats. The security team should map assets, identities, data flows, trust boundaries, tools and dependencies. Threat scenarios can include prompt injection, indirect injection through retrieved content, data poisoning, malicious model files, insecure deserialisation, credential theft, model extraction, sensitive-data disclosure, denial of service and supply-chain compromise.

Controls can include input and content isolation, tool allowlists, least privilege, network restrictions, secrets management, environment separation, model and package verification, signed releases, dependency scanning, rate limits, output validation, retrieval permission checks, monitoring and tested incident response. Security claims should reflect evidence from the deployed system and configuration.

The company should treat prompts and policies as code when they materially affect behaviour. They need version control, review, testing and rollback. Retrieval indexes and policy configurations also need change records. A secure base model can become unsafe through an over-privileged tool, weak prompt, exposed secret or incorrect customer-data boundary.

Security tests should include realistic adversarial paths. Red-team exercises can test whether a user, document, web page, integration or compromised account can redirect the model or agent. Findings should be risk-ranked, assigned and retested. The release record should state any accepted limitation and its compensating control.

Incident response should cover detection, containment, evidence preservation, customer and regulator communication, model or tool disablement, credential rotation, data review, restoration and lessons learned. The company should know how to stop a model, agent or integration without taking down unrelated services. Enterprise customers and investors may ask for evidence that the plan has been exercised.

Table 3. Evaluation, security and release evidence

Evidence area	Pre-release record	Production measure	Material response
task performance	representative evaluation set, baseline, acceptance criterion and reviewed result	quality, drift, override and task failure by use	investigate, limit cohort, retrain, reconfigure or roll back
robustness	edge cases, distribution shifts, dependency failure and recovery tests	unusual inputs, fallback use, availability and recovery	qualify output, switch fallback, reduce scope or hold action
fairness and outcome	affected groups, harm analysis, performance slices and review path	outcome differences, complaints and successful challenge	pause affected use, investigate data and decision, remediate
security	threat model, adversarial test, dependency scan and closed critical findings	attack signal, access anomaly, secret exposure and data event	contain, disable tool or model, rotate, notify and review
privacy and rights	data map, purpose, rights, retention, deletion and processor assessment	access, retention exception, data request and deletion completion	stop unsupported use, correct data state and document outcome
release authority	approved version, scope, limitations, owner, monitoring and stop rules	exceptions, overdue reviews, material changes and incidents	convene release authority and record continue, narrow, roll back or retire decision

The required depth should follow the system's consequence, customer commitments and applicable obligations.

11. Build human oversight and customer recourse

Human oversight should be tied to consequence. The company should identify which decisions require human confirmation, which can be sampled, which need specialist review and which can operate automatically within approved limits. The design should state what the reviewer sees, what alternatives are available and how disagreement is recorded.

Customers need clear information about the service. The communication should explain the outcome, material role of AI where relevant, important data, limitations, customer responsibilities and route to support. A person affected by a consequential decision may need an explanation, correction or appeal under applicable law or contract. The product should route that request to someone who can review the underlying evidence and change the outcome where authorised.

Complaint records are valuable governance data. Management should classify the system, model version, customer group, issue, consequence, resolution and root cause. Repeated confusion can reveal a weak disclosure. Repeated corrections can reveal data or model problems. Slow resolution can reveal missing ownership. The board should see material patterns with the commercial dashboard.

Employee governance also matters. Staff should know which tools are approved, which data can be used, how outputs must be checked and when external communication requires review. Training should reflect role and risk. Developers need secure and responsible design practices. Commercial teams need claims discipline. Customer teams need escalation and explanation. Executives need incident and board obligations.

The company should support a speak-up route for AI concerns. Employees and contractors may observe data misuse, unsupported claims, pressure to bypass tests or customer harm. The route should protect evidence, assign investigation and provide escalation beyond the immediate

product owner.

12. Govern claims, sales, demos and investor materials

AI claims should be specific, evidenced and current. Words such as autonomous, proprietary, accurate, unbiased, secure, compliant or real time can create a misleading impression when they omit material dependencies or limitations. The company should maintain a claims register linking each statement to test results, product scope, model version, date, reviewer and approved channel.

United States enforcement illustrates the exposure. The Securities and Exchange Commission charged two investment advisers in 2024 over false and misleading statements about their purported use of AI. The Federal Trade Commission has repeatedly warned businesses to support AI performance claims and to honour privacy and confidentiality commitments. These materials support a general control principle: product, marketing and fundraising statements should match demonstrable capability and actual operation.

Demos need the same discipline. A curated example, human-assisted result or prototype should be described accurately. The company should distinguish production features from roadmap items. If the demonstration uses synthetic data, a staged integration or manual review, the audience should receive that context. Sales teams should not promise a regulatory status, security property or performance level that has not been approved.

Investor materials should reconcile product and financial claims. If the model is third-party, the deck should not imply ownership of the foundation model. If margin assumes a lower future inference cost, the forecast should identify the assumption. If revenue depends on enterprise security approval or regulatory permission, the plan should state the gate. If customer outcomes are based on a pilot, the cohort and method should be clear.

The board should approve material AI language used in fundraising. A data room should contain the claims register, model and vendor map, evaluations, security evidence, customer limitations, rights analysis, incidents and open issues. Accurate disclosure supports diligence and reduces the risk that commercial enthusiasm becomes a misstatement.

Figure 4. AI unit-economics bridge from model cost to contribution

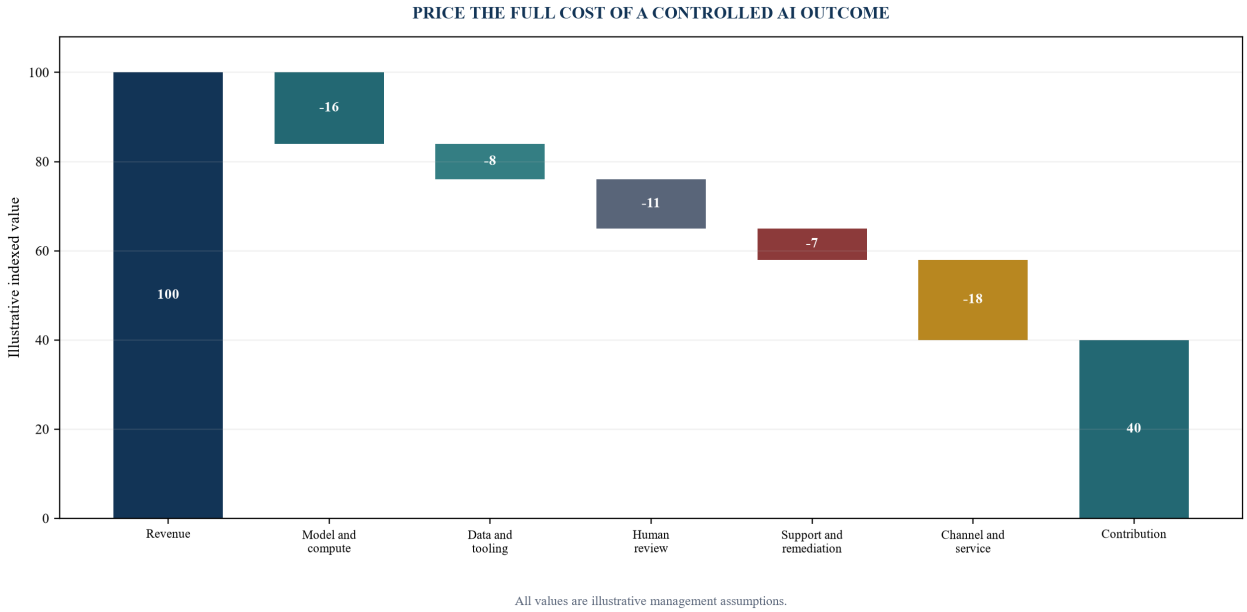

Values are illustrative management assumptions and should be replaced with verified cohorts and contracts.

13. Map regulation and market-entry obligations

AI regulation varies by jurisdiction, sector, activity and role. A startup should map where it develops, supplies and deploys the system; where users and affected people are located; which data is processed; and which regulated products or decisions are involved. The map should distinguish binding law, regulator rules, contractual requirements, voluntary standards and proposals.

The European Union AI Act applies progressively. The official service-desk timeline states that definitions, AI-literacy requirements and prohibitions applied from February 2025; general-purpose-AI obligations and governance from August 2025; and most remaining rules and Article 50 transparency duties from August 2026, subject to the published timetable and Digital Omnibus amendments. High-risk-system rules have later dates. A company should confirm its role, system category, transition and market status with qualified counsel.

The UAE Charter for the Development and Use of Artificial Intelligence provides national principles for responsible AI. The UAE Personal Data Protection Law governs personal-data processing. Sector rules can add requirements, especially in financial services, health and government work. Saudi SDAIA's AI Ethics Principles use risk categories and lifecycle controls for stakeholders developing, deploying or using AI systems in the Kingdom.

The United Kingdom applies existing privacy, consumer, equality, product, financial-services and other law alongside its cross-regulator AI approach. The ICO toolkit supports AI data-protection risk analysis. The FCA states that existing frameworks, including Consumer Duty, senior-management accountability, governance and controls, apply to AI in financial services. Singapore's IMDA frameworks provide practical governance for traditional, generative and agentic AI. The OECD principles support interoperability around human rights, transparency, robustness, safety and accountability.

Market entry should have a legal and operating gate. The board should require a written role and obligation map, current effective dates, product controls, customer disclosures, regulator or

scheme engagement where needed and an owner for change monitoring. A framework or voluntary standard can support governance without establishing legal compliance by itself.

Table 4. Multi-market AI governance and market-entry map

Market	Official position	Company design implication	Launch evidence
United Arab Emirates	national AI Charter and federal personal-data law; sector rules may also apply	map principles, data processing, system consequence, customer disclosure and sector duties	qualified role analysis, data map, controls, terms and change owner
Saudi Arabia	SDAIA AI Ethics Principles apply a risk-based lifecycle approach; personal-data and sector rules remain relevant	classify risk, embed lifecycle controls, document privacy and assess regulated use	current Saudi advice, risk record, data controls and required approvals
European Union	AI Act applies progressively with role, system and risk-based obligations	determine provider, deployer, importer or other role; map GPAI, transparency and high-risk duties	current classification, transition dates, technical documentation and post-market controls
United Kingdom	existing legal and regulatory frameworks apply; ICO and sector regulators publish AI guidance	map privacy, consumer, equality, security, product and sector outcomes	data-protection assessment, accountable owner, testing, explanation and sector review
Singapore	IMDA model frameworks address traditional, generative and agentic AI	bound agent powers, define human accountability, implement lifecycle controls and inform users	use-case risk assessment, technical controls, testing and training
United States	federal and state laws and sector rules apply; SEC and FTC materials emphasise accurate claims and lawful data use	map product, consumer, employment, financial, privacy, intellectual-property and state requirements	jurisdictional advice, claims evidence, privacy commitments and complaint route

This table summarises official framework positions at publication; each company must reverify current law, role and effective dates before launch.

14. Convert governance into enterprise procurement readiness

Enterprise customers often evaluate an AI supplier through security, privacy, resilience, legal, procurement and business-owner reviews. A founder can shorten that process by maintaining a current evidence room. The room should match actual product scope and avoid generic certifications or policies that do not cover the deployed service.

The evidence pack can include the system description, architecture, data flow, security programme, penetration test, vulnerability process, incident plan, privacy assessment, processor list, data locations, model and vendor register, evaluations, business-continuity plan, service levels, insurance, customer terms, intellectual-property position and exit support. Where a certification exists, management should identify its scope, date and exclusions.

Customer questionnaires should feed the governance system. Repeated questions reveal which evidence buyers require. Repeated exceptions reveal product or contract gaps. The company can create standard responses backed by controlled documents and named owners. Material customer-specific commitments should enter the product and monitoring backlog.

Procurement economics should be measured. Security review, legal negotiation, integration, model evaluation, customer-specific hosting and support consume cash. Revenue forecasts should include approval probability, implementation time, renewal, expansion and service cost. A large contract with broad liability, custom controls and low gross margin can weaken the financing case.

The board should review the relationship between governance maturity and sales execution using observed company data. Useful measures include questionnaire turnaround, time to security approval, material exceptions, contract cycle, implementation time, customer incidents, renewal and contribution. The company should avoid claiming a causal valuation uplift without transaction evidence.

15. Link governance to unit economics and valuation

AI unit economics should include the cost of delivering a controlled outcome. Compute and model fees are one component. Data acquisition, retrieval, observability, security, human review, support, incident remediation, customer-specific controls, compliance, insurance and vendor minimums can materially change contribution.

The finance team should trace a cohort from eligible user to successful and retained outcome. It should record model calls, tokens or compute, retries, tool use, human interventions, support, refunds and loss. Segmenting by workflow, customer and model can reveal where a high headline margin hides expensive failure or review.

Model improvement can affect economics through higher task completion, lower retries, reduced human review, shorter handling time or stronger retention. A larger model may improve quality while increasing cost and latency. A smaller or specialised model may improve contribution for a bounded task. Routing should be governed and evaluated as part of the product.

Valuation should follow durable cash-flow evidence, strategic rights and risk-adjusted growth. Governance can support the case by protecting data rights, strengthening enterprise procurement, reducing dependency surprises, substantiating claims and making operating performance measurable. The actual valuation effect depends on company results, market conditions, investor judgement and transaction terms. It should not be presented as a fixed premium.

The board should use scenarios. A base case can reflect current model prices, observed conversion and approved review. An upside case can reflect evidenced automation and signed distribution. A downside case can reflect vendor price increases, longer procurement, a model incident, additional human review or delayed regulation. Every scenario should show cash runway and the decision available to management.

Figure 5. Funding roadmap and cap-table waterfall by evidence gate

MAKE DILUTION A FUNCTION OF VERIFIED PROGRESS

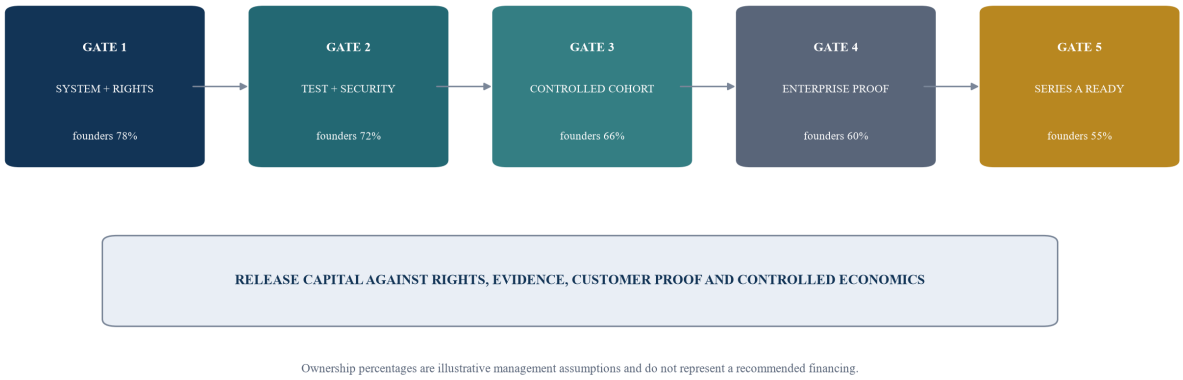

Funding tranches and ownership values are illustrative management assumptions; company decisions require a verified cap table and approved financing terms.

16. Stage funding, cap table and use of proceeds

The funding plan should connect capital to specific evidence. Early funds can establish system definition, data rights, a secure architecture and representative evaluation. The next tranche can support controlled customers, enterprise procurement and operating metrics. Series A capital can then fund distribution, product expansion and organisational capacity against a better-defined risk and economics case.

Use of proceeds should separate product development, model and data cost, security, privacy and legal work, enterprise integrations, sales, customer support, incident capacity and working capital. A budget labelled engineering can hide required governance investment. Investors should be able to see which spending produces a release gate and which spending supports growth after the gate.

The cap table should include issued shares, options, warrants, convertibles, SAFEs or other instruments, liquidation preferences, pro rata rights, information rights, consent rights and any model or data partner equity. Management should model ownership and proceeds across plausible financing terms. The illustrative percentages in Figure 5 demonstrate a method and do not recommend a transaction.

The board should protect operating flexibility. Financing documents, strategic partnerships and customer agreements can affect model choice, data use, exclusivity, pricing, fundraising, intellectual property and change of control. These constraints should be reviewed together. An attractive distribution agreement can weaken optionality if it grants broad exclusivity or ownership of improvements.

Runway planning should include AI cost volatility and procurement timing. The downside case can model a higher inference cost, slower enterprise approval, additional human review or a temporary release hold. The board should approve triggers for hiring, commitments and another financing process. Forecasts remain management estimates until supported by actual collections,

contracts and operating data.

Table 5. Series A data-room and investor evidence checklist

Diligence area	Core evidence	Key reconciliation	Unresolved item treatment
product and market	use cases, roadmap, customer workflow, pilots, contracts and retention	product claims to shipped scope and customer evidence	identify roadmap status, owner, dependency and date
AI system and model	inventory, architecture, model versions, vendors, evaluations and monitoring	live endpoints to approved systems and reported performance	disclose limitation, risk tier, mitigation and release condition
data and intellectual property	provenance, rights, privacy map, licences, employee and contractor assignments	every material dataset and component to an enforceable right	obtain advice, narrow use, replace source or hold affected feature
security and resilience	policies, threat model, test results, incidents, recovery and subprocessors	customer promises to demonstrated scope and current findings	state open finding, severity, compensating control and remediation
commercial and economics	pipeline, contracts, pricing, usage, model cost, service cost and collections	revenue and forecast to customer and cohort evidence	separate signed, deployed, billed, collected and forecast values
governance and regulation	board minutes, policies, accountability, role maps, legal advice and change logs	material use and markets to approvals and obligations	describe status, adviser, consequence and decision gate
financing and cap table	securities, options, rights, scenarios, use of funds and runway	legal records to financial model and fully diluted ownership	resolve inconsistency before transaction documentation

Investor requirements vary; this checklist organises commonly relevant product, commercial, legal and operating evidence.

17. Run a 120-day pre-Series A programme

Days one to twenty establish visibility. Management defines the product outcome, inventories systems and dependencies, maps data and intellectual-property rights, identifies jurisdictions and assigns risk tiers. The board approves the AI risk appetite, accountability and material use cases. The first gate requires a current inventory and an acceptable path for critical rights.

Days twenty-one to forty-five establish lifecycle controls. Product, engineering, security, privacy and commercial teams define evaluation criteria, threat scenarios, vendor requirements, claims approval, incident response and customer recourse. The company tests the production configuration and closes critical findings. The second gate requires reproducible evidence and a bounded release scope.

Days forty-six to seventy establish customer and procurement proof. The company runs a controlled cohort, completes enterprise evidence, reconciles claims with results and measures service cost. It records complaints, overrides, security events and model changes. The third gate requires a demonstrable customer outcome, controlled performance and executable procurement route.

Days seventy-one to ninety-five establish economics and market readiness. Finance builds the unit-economics bridge, customer cohorts, model-cost sensitivity, funding roadmap and cap-table scenarios. Legal and regulatory advisers confirm the applicable market-entry analysis. The fourth gate requires evidence-based contribution and current jurisdictional conditions.

Days ninety-six to one hundred and twenty prepare the transaction. Management builds the data room, reconciles product and financial statements, prepares board materials and closes priority exceptions. The fifth gate authorises a defined fundraising scope, use of proceeds, disclosure position and operating plan. Material gaps remain identified with owners and consequences.

Figure 6. The 120-day AI governance and Series A roadmap

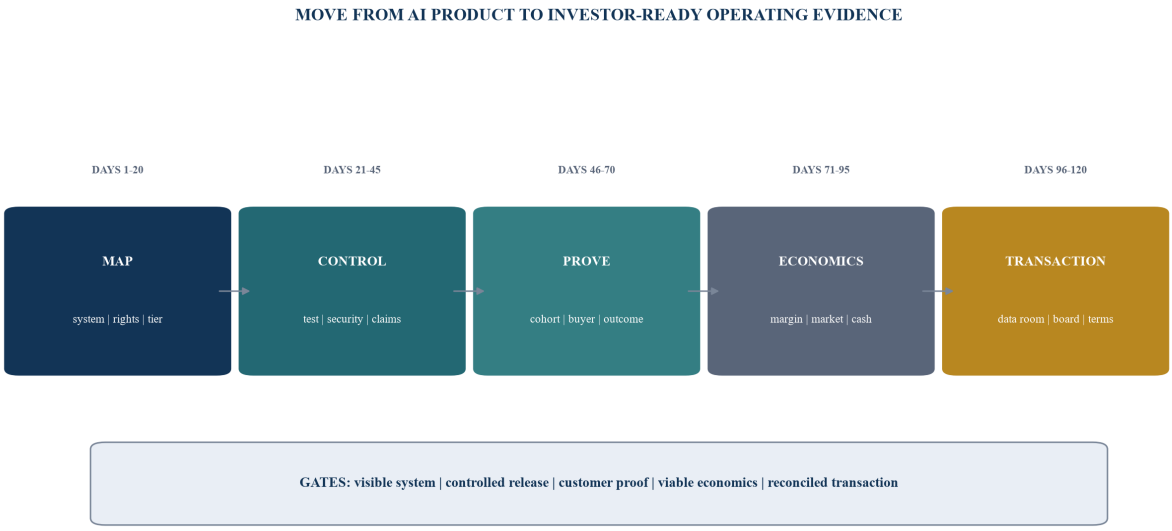

Each phase ends with evidence that the board can approve, qualify or return for remediation.

18. Make the board and investor decision explicit

The final board pack should present the product outcome, system inventory, risk tiers, data and model rights, architecture, vendor dependencies, evaluations, security evidence, human oversight, customer recourse, claims register, regulatory map, procurement readiness, unit economics, funding plan, cap table and data-room status. Each item should identify its source, owner, date and open matter.

The board should approve a bounded decision. It can authorise a defined system version, use, customer cohort, jurisdiction, autonomy level, distribution route and capital budget. Conditions can include closed security findings, current data rights, approved claims, tested monitoring, customer support capacity and a fixed review date. Expansion should require the next evidence gate.

Management should state uncertainty plainly. A pending legal analysis, vendor commitment, customer approval, model improvement, fundraising outcome or forecast conversion should retain its current status. The company should distinguish signed, deployed, billed, collected and forecast revenue. It should distinguish tested capability from roadmap and approved market access from an application or plan.

The investor decision should reflect both growth and controllability. The diligence question is whether the company can convert capital into customer value through systems and rights that remain available, measurable and governable. A strong governance programme gives the board and investor a common evidence base for that judgement.

The operating decision continues after financing. New markets, models, data, agents, acquisitions and customer commitments should re-enter the same system of inventory, tiering, evidence, approval and monitoring. Governance remains part of product and capital execution as the company grows.

Table 6. Board AI governance and Series A decision checklist

Gate	Required evidence	Accountable owner	Decision
system and rights	use-case definition, inventory, risk tier, data rights, model rights and jurisdiction map	product, data and legal owners	approve the system boundary, intended use and conditions
development and release	representative evaluation, security test, privacy review, vendor assessment and stop rules	engineering, model, security and privacy owners	approve a controlled version, scope, cohort and monitoring plan
customer and procurement	outcome evidence, customer communication, recourse, enterprise evidence and executable terms	commercial, product and operating owners	approve the channel, customer commitment and service capacity
claims and disclosure	claims register, demo status, investor statements, limitations and open issues	CEO, commercial and legal owners	approve current external representations and required qualifications
economics and funding	verified cohorts, contribution bridge, model-cost sensitivity, runway and cap-table scenarios	CFO and CEO	approve use of proceeds, financing range and downside actions
market and scale	current role analysis, regulatory obligations, partner readiness, incidents and governance capacity	executive sponsor and board	continue, narrow, remediate, roll back or authorise bounded scale

Each gate requires current evidence, a named owner and a recorded board or delegated decision.

Implementation conclusion

AI governance before Series A begins with visibility. The company defines each system, consequence, right and dependency, then assigns owners and risk-based evidence gates. It connects model performance to data rights, security, customer recourse, claims and economics.

The 120-day programme produces five board decisions: approve the system and rights map; approve controlled development and release; verify customer and procurement proof; validate economics and market entry; and authorise a reconciled financing process. The result is a practical operating record that investors, enterprise customers and regulators can examine.

Investor requirements remain specific to the fund, stage, product, market and transaction. Valuation, funding, cost and timing effects shown in the figures are illustrative management assumptions. Company decisions should use verified contracts, cohorts, technical evidence and qualified professional advice.

References

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

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

[3] - International Organization for Standardization. ISO/IEC 42001:2023, Artificial intelligence management system. https://www.iso.org/standard/42001

[4] - European Commission AI Act Service Desk. Timeline for the Implementation of the EU AI Act. https://ai-act-service-desk.ec.europa.eu/en/ai-act/eu-ai-act-implementation-timeline

[5] - UAE Artificial Intelligence Office. The UAE Charter for the Development and Use of Artificial Intelligence, July 2024. https://ai.gov.ae/wp-content/uploads/2024/07/UAEAI-Methaq-EN2-3.pdf

[6] - United Arab Emirates Legislation. Federal Decree by Law No. 45 of 2021 Concerning the Protection of Personal Data. https://uaelegislation.gov.ae/en/legislations/1972

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

[8] - Infocomm Media Development Authority of Singapore. Artificial Intelligence in Singapore and Model AI Governance Frameworks. https://www.imda.gov.sg/AI

[9] - Information Commissioner's Office. AI and Data Protection Risk Toolkit. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/ai-and-data-protection-risk-toolkit/

[10] - Financial Conduct Authority. AI and the FCA: Our Approach. https://www.fca.org.uk/firms/innovation/ai-approach

[11] - Financial Conduct Authority. AI in Financial Services: Shaping Our Approach through Industry Engagement, 8 June 2026. https://www.fca.org.uk/news/blogs/ai-financial-services-approach

[12] - Organisation for Economic Co-operation and Development. AI Principles. https://www.oecd.org/en/topics/ai-principles.html

[13] - U.S. Securities and Exchange Commission. SEC Charges Two Investment Advisers with Making False and Misleading Statements about Their Use of Artificial Intelligence, 18 March 2024. https://www.sec.gov/newsroom/press-releases/2024-36

[14] - U.S. Federal Trade Commission. AI Companies: Uphold Your Privacy and Confidentiality Commitments, 9 January 2024. https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/01/ai-companies-uphold-your-privacy-confidentiality-commitments

About the Author

Authored by Chennakeshav Adya, Independent Researcher