AI Governance for Family-Owned Groups: A Board Policy before the First Incident

A Global Control Framework for AI Use across Holding Companies, Operating Businesses and Family Offices

Authored by Chennakeshav Adya
Independent Researcher
August 2026

Abstract

Artificial intelligence can enter a family-owned group through employee subscriptions, embedded software, customer systems, portfolio companies, advisers and investments before the board has approved a policy. The exposure can cross legal entities and generations. Confidential information can leave controlled environments. Outputs can influence hiring, credit, pricing, investment, health, safety and reputation. Agents can execute actions through connected tools. A single incident can affect several businesses and the family name.

This paper develops a global board-policy framework for governing AI across holding companies, operating businesses, investment entities and the family office. It treats governance as a chain of accountable decisions: define the perimeter, inventory systems and use cases, classify risk, establish permitted and prohibited uses, assign legal and human decision rights, test vendors and controls, monitor performance, and respond to incidents. The framework uses six figures, six tables, a hypothetical group example and a 100-day implementation sequence.

The approach draws on the OECD AI Principles, the United States National Institute of Standards and Technology AI Risk Management Framework and Generative AI Profile, the European Union AI Act, European data-protection guidance, the United Kingdom Information Commissioner's Office and Singapore's Model AI Governance Frameworks. Requirements vary by jurisdiction, role, system and use case. The board should obtain current legal, technical, employment, cyber-security and sector advice before relying on any regulatory interpretation.

JEL Classification: D23, G32, G34, K24, L21, M14, O33

Keywords: artificial intelligence, family business, family office, board governance, generative AI, agentic AI, data protection, cyber risk, incident response

1. Establish policy before the first incident

AI adoption often begins as a productivity choice. An employee uses a public assistant to summarise a contract. A sales team enables an embedded writing tool. A portfolio company deploys a customer bot. A family-office analyst tests an investment memorandum. A vendor adds an AI feature during a software update. Each action can create data, decision, intellectual-property, cyber-security and accountability consequences.

Family-owned groups add structural complexity. The same family name can sit above operating companies, real estate, investment vehicles, philanthropy and personal affairs. Some entities share people, systems and advisers. Others have minority investors, regulated activities or independent boards. A policy issued by one company may have no legal effect in another. An incident can nevertheless transmit through guarantees, shared services, reputation, succession arrangements

and common ownership.

The board policy should define how AI may be acquired, built, configured, connected and used. It should set decision rights before a use case becomes operational. It should also distinguish governance from technical assurance. A policy can allocate authority and minimum controls; qualified teams must still test data, systems, security and impact in context.

NIST structures AI risk management around Govern, Map, Measure and Manage. The OECD principles emphasise human rights, transparency, robustness, security and accountability. Singapore's frameworks add practical guidance for human involvement and agentic systems. These sources support a common operating idea: the organisation should know what an AI system does, who owns the outcome, what evidence supports its use, and when a person must intervene.

Figure 1. AI risk can transmit across a family-owned group

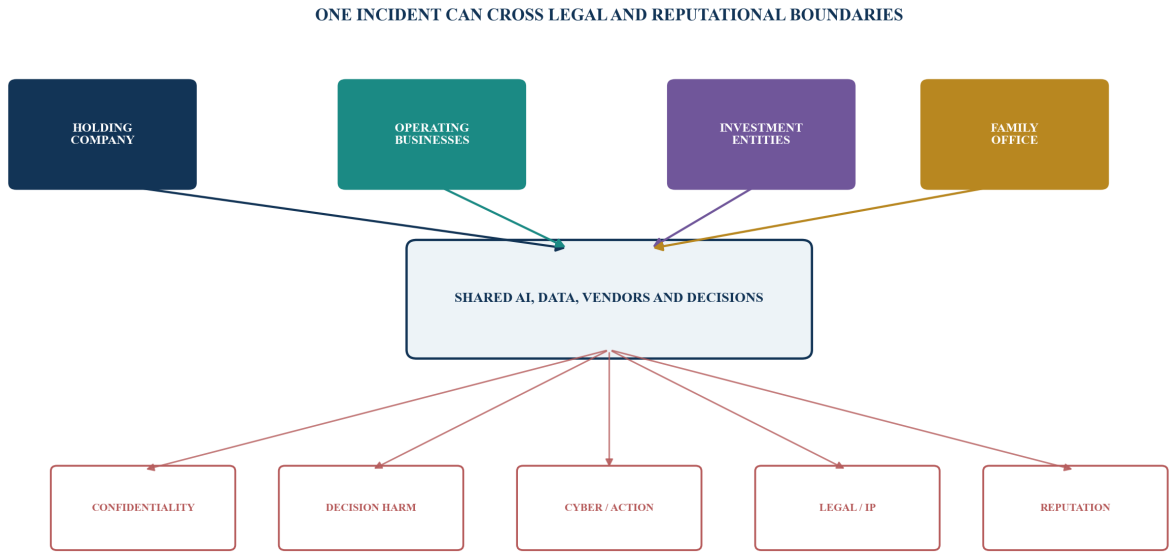

The policy perimeter follows legal entities, shared services, data, decisions and the family name.

2. Define the perimeter, roles and policy hierarchy

The policy perimeter should begin with an entity map. It should identify the parent, subsidiaries, joint ventures, regulated businesses, investment advisers, family-office entities, trusts, foundations and shared-service companies. For each entity, the board should record ownership, governing body, employees, customers, regulated activities, data locations and technology dependencies.

Roles matter because the same model can sit inside different legal relationships. An entity may develop an AI system, provide it, deploy it, distribute it, import it or use its output. A cloud vendor, system integrator, group technology team and operating company can each carry different obligations. The map should identify controller and processor relationships for personal data and clarify contractual responsibility for training, configuration, monitoring and incidents.

The group policy should set minimum principles and escalation triggers. Entity boards should adopt local schedules covering applicable law, sector rules, works councils, employment practices, customer duties, data residency and regulator expectations. A stricter rule should govern where a shared system cannot reliably separate entities or users.

The hierarchy should connect the AI policy to information security, privacy, records, acceptable use, model risk, procurement, employment, intellectual property, market conduct, financial crime, health and safety, crisis management and delegated authority. Conflicts should be resolved by the accountable legal owner, with advice appropriate to the issue.

Table 1. Board AI-policy perimeter and accountable roles

Policy layer	Board question	Accountable role	Minimum evidence
group principles	which values, prohibitions and escalation thresholds apply across the family group?	parent board	approved policy, entity applicability map and exceptions
entity adoption	which requirements apply to this legal entity and jurisdiction?	entity board	local schedule, legal review and adoption minute
use-case ownership	who owns the business outcome and affected process?	business owner	purpose, users, impact, controls and acceptance criteria
technology	who owns architecture, access, change and resilience?	technology owner	system design, access model, logs, testing and recovery
data	who authorises data sources, purpose, retention and transfer?	data owner	data map, lawful basis, quality, rights and deletion controls
risk and compliance	who challenges classification and control sufficiency?	independent control function	documented review, conditions and residual risk decision
assurance	who tests whether controls operate as represented?	internal audit or qualified assurance	scoped plan, evidence, findings and remediation tracking

Titles differ by group; each duty needs one accountable owner and named support.

3. Build an AI system and use-case register

A group cannot govern systems it has not identified. The inventory should cover internally developed models, vendor applications, embedded AI functions, public tools, robotic process automation with AI components, analytical models, computer vision, recommendation engines, synthetic content, decision systems and autonomous or agentic workflows. It should also capture informal experimentation where group data or decisions are involved.

Discovery should combine declarations with technical evidence. Procurement and accounts-payable records can identify subscriptions. Identity and network logs can identify applications. Browser and endpoint controls can detect unmanaged services where lawful. Software inventories can reveal embedded capabilities. Surveys, interviews and workshops can expose use cases that technology scans miss. The purpose is an auditable perimeter, with employee monitoring designed under applicable law and policy.

The system record and use-case record should be separate. One model may support several uses with different risk. A writing assistant used for public marketing carries different consequences when used for legal advice or employee evaluation. The system record describes the model, provider, version, hosting, access and technical controls. The use-case record describes purpose,

people affected, decision influence, data, owner, deployment and monitoring.

Every record needs a lifecycle state: discovery, sandbox, assessment, approved pilot, production, suspended or retired. Unauthorised systems should enter a remediation path. The board should see material unknowns and overdue reviews, while management owns the working register.

Figure 2. The AI register joins system facts to use-case decisions

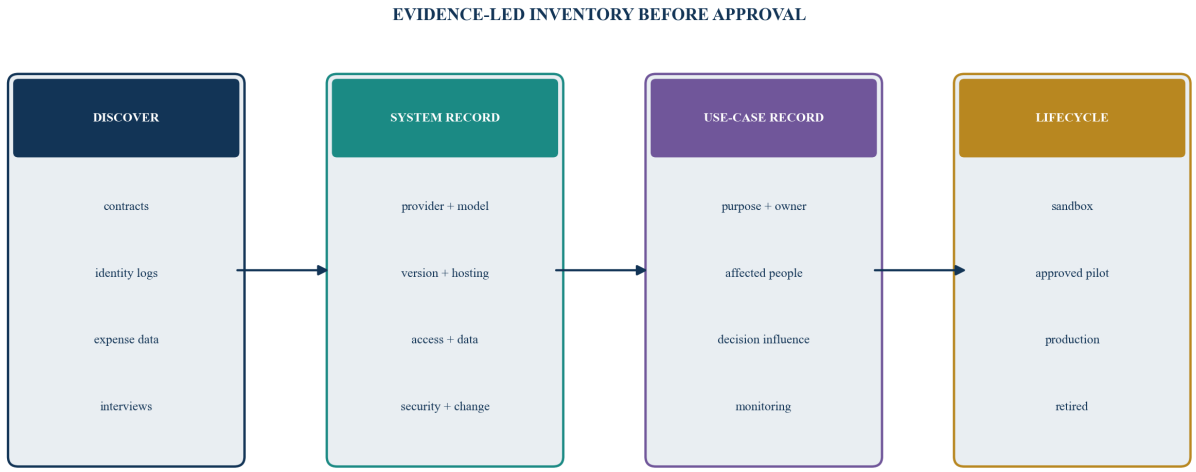

One system can support several use cases; each use case receives its own classification and owner.

Table 2. Minimum AI register fields

Domain	System record	Use-case record
identity	provider, model, version, system owner and hosting	purpose, entity, process, business owner and users
data	training or reference data, prompts, outputs, retention and transfer	data categories, people affected, sensitivity, purpose and lawful basis
decision	model capability and known limitations	decision influenced, consequence, human authority and appeal route
access	identities, privileges, interfaces and connected tools	permitted users, locations, segregation and approval conditions
assurance	provider tests, internal tests, security review and incidents	acceptance tests, impact review, monitoring thresholds and review date
lifecycle	acquisition, version, change process and retirement	discovery, sandbox, pilot, production, suspension or retirement

The register should retain evidence dates, sources and unresolved gaps.

4. Classify the use case before approving the system

Classification should follow the use and its consequences. The board policy can use four tiers: prohibited, restricted, controlled and routine. The labels are internal governance tools; legal categories must be assessed separately. A system can move between tiers when data, autonomy, users, affected people or connected actions change.

Prohibited uses should include activities the group is unwilling or unable to govern, plus any use prohibited by applicable law. Examples may include covert manipulation, unlawful

discrimination, fabricated evidence, unapproved impersonation, circumvention of controls, and entering protected or confidential information into an unauthorised public service. The policy should avoid broad wording that employees cannot apply.

Restricted uses can include material decisions affecting employment, credit, insurance, health, safety, legal rights, customers, investments or access to essential services. They require enhanced legal review, impact assessment, data governance, testing, human authority, documentation and monitoring. Autonomous actions involving money, code, external communications or production systems should also enter a restricted path.

Controlled uses include lower-impact support functions with approved data, tools and human review. Routine uses may cover low-consequence tasks such as formatting public text within approved systems. Classification should consider severity, scale, reversibility, vulnerability of affected people, data sensitivity, automation, opacity, external reliance and ability to detect error.

Figure 3. Risk-tier decision ladder for AI use cases

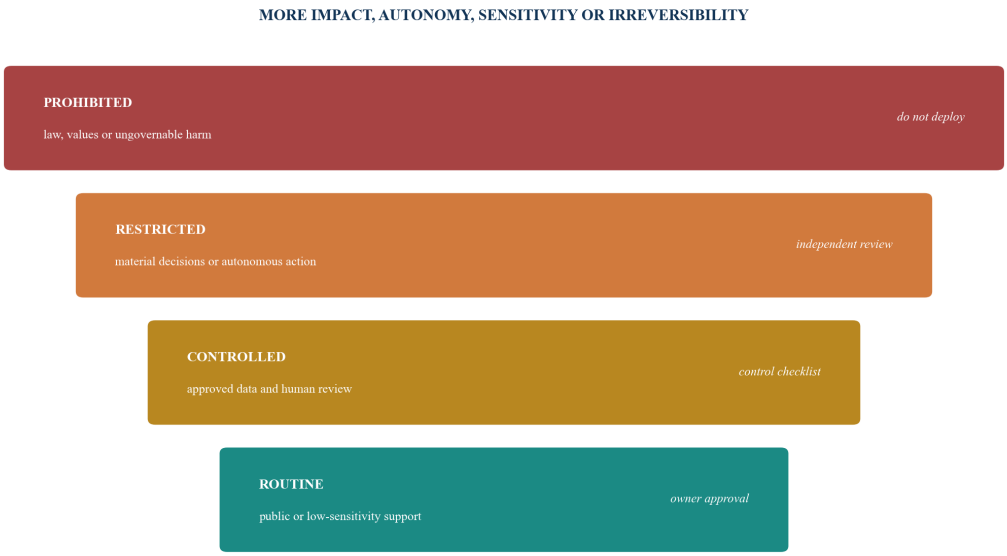

Escalation follows impact, autonomy, sensitivity and reversibility; legal classification remains separate.

Table 3. Board policy for permitted, conditional and prohibited uses

Policy state	Illustrative use	Required action
permitted within approved tool	formatting public material, translation with human review, meeting logistics	use approved account, verify output and retain required record
controlled	research synthesis, internal drafting, coding support, customer-service assistance	authorised data, named owner, source checking, access control and monitoring
restricted	hiring, credit, pricing, legal, investment, health, safety or material customer decisions	impact assessment, legal and control review, testing, human authority, appeal and board threshold
restricted autonomous action	sending communications, changing systems, executing code, moving funds or placing orders	bounded permissions, dual control, sandbox, logs, transaction limits, kill switch and incident response
prohibited by policy	fabricated evidence, unapproved impersonation, covert manipulation, bypassing controls or unsafe data disclosure	block, report, preserve evidence and assess incident obligations
regulatory prohibition	any activity prohibited for the entity, jurisdiction or role	do not deploy; document legal determination and monitor change

Examples require entity-specific legal and sector review before adoption.

5. Set decision rights and meaningful human authority

Human review should be designed around the decision. A person cannot provide meaningful oversight when they lack time, information, authority, competence or an alternative. The policy should identify who can approve, reject, override, pause and appeal an AI-influenced outcome. It should also state which decisions remain exclusively human.

The board sets risk appetite and material thresholds. Management approves use cases within delegated authority. The business owner owns the outcome, including errors created by overreliance. Technology owns system operation. Data owners govern sources and rights. Legal, privacy, risk, compliance, security and human resources review their domains. Internal audit or qualified assurance evaluates control design and operation without assuming management responsibility.

The review design should match consequence and reversibility. A low-impact draft may require a competent user to check accuracy and tone. A material employment or credit decision may require independent evidence, reason-giving, bias assessment, record retention and an accessible reconsideration path. An agent that can execute actions may require pre-approved boundaries, dual authorisation and immediate suspension capability.

The policy should address automation bias. Users need training to understand that fluent outputs can be wrong, incomplete or unsupported. Review interfaces should expose sources, confidence limitations, exceptions and changes. Management should test whether reviewers actually challenge outputs rather than merely confirm them.

6. Govern data, confidentiality and intellectual property

AI data governance begins before a prompt is entered. The organisation should classify the information, identify the legal owner, confirm purpose and authority, and understand what the provider stores, uses, transfers or exposes. Client, employee, transaction, health, financial, legal, family, security and unpublished investment information can require heightened controls.

The policy should distinguish approved enterprise environments from public or consumer accounts. Contract and configuration should address whether prompts and outputs train models, retention, deletion, sub-processors, locations, access, encryption, audit logs, incident notice and exit. Technical controls should reinforce the rule through identity, data-loss prevention, endpoint management and access restrictions where lawful and proportionate.

Personal-data obligations depend on jurisdiction and role. The UK Information Commissioner's Office identifies accountability, lawful basis, fairness, transparency, security, data minimisation, individual rights and data-protection impact assessment as relevant considerations. The European Data Protection Board's December 2024 opinion addresses anonymity, legitimate interest and the consequences of unlawfully processed training data. Each use requires case-specific assessment.

Intellectual-property review should cover input rights, output rights, licences, open-source obligations, confidential know-how, third-party content, brand assets, training terms and indemnities. Generated content should be checked for provenance and clearance before external use. Legal privilege and professional confidentiality require explicit handling approved by counsel.

7. Control third-party models and embedded AI

Procurement should evaluate the provider and the proposed use. A well-known vendor does not establish that a specific configuration, integration or use case is appropriate. Due diligence should cover corporate identity, service description, model dependencies, training and evaluation, security, data terms, sub-processors, intellectual property, availability, monitoring, incidents, audit rights, liability and termination.

Model and feature changes create special risk. Providers can change capabilities, safeguards, data practices or interfaces. The contract should define notice, material change, version control, testing, rollback and termination rights. The system owner should monitor release notes and retest material workflows before accepting changes.

Embedded AI needs discovery. Customer relationship management, productivity, finance, human-resources, security and industrial software may activate AI functions under existing contracts. The owner should assess the feature before enabling it and confirm whether data flows or contractual roles change.

Concentration also matters. Several group companies may depend on the same foundation model, cloud platform or integration provider. A common failure, policy change or outage can affect multiple entities. The register should expose shared dependencies and alternative arrangements.

Table 4. Third-party AI due-diligence and contracting matrix

Domain	Diligence question	Contract or control evidence
model and service	which models, versions, components and sub-providers deliver the service?	service description, dependency map and change notice
data	what enters the service, where is it processed, how long is it retained and is it used for training?	data terms, configuration, deletion, transfer and audit evidence
testing	how are accuracy, robustness, bias, security and misuse tested for the intended context?	evaluation reports, limitations, customer tests and acceptance thresholds
security	how are identities, interfaces, secrets, logs and incidents protected?	security architecture, certifications, vulnerability process and notice terms
intellectual property	what rights govern inputs, outputs, training material and claims?	licences, warranties, indemnities and infringement process
resilience	what happens during outage, model withdrawal or material change?	service levels, continuity, export, rollback and exit plan
accountability	who owns errors, complaints, regulatory support and remediation?	named contacts, cooperation, records, audit rights and liability allocation

Depth should be proportionate to the use case, data and dependency.

8. Secure the AI lifecycle and connected tools

Security should cover data poisoning, prompt injection, insecure output handling, model extraction, credential exposure, excessive permissions, dependency compromise, denial of service and misuse. The threat model should reflect whether the system is internal, customer-facing, connected to tools, fine-tuned on group data or exposed to untrusted content.

Development and configuration should use separate environments, approved components, secret management, code review, testing and change control. Retrieval systems should govern source collections and permissions. Outputs passed to code, databases, email, finance or industrial systems should be validated before execution. Logs should support detection and investigation while respecting privacy and confidentiality.

Red teaming and evaluation should match the use. Test sets need expected answers, edge cases, adversarial inputs and affected-user perspectives. Measures can cover accuracy, unsupported claims, harmful content, refusal, bias, security, latency, resilience and override. The owner should set acceptance thresholds and record limitations before deployment.

Production monitoring should detect drift, abnormal access, unsafe outputs, policy violations, complaints, incidents and provider changes. A passing pre-deployment test has a date and scope. It does not provide permanent assurance.

9. Apply stronger boundaries to generative and agentic AI

Generative AI can produce plausible text, images, audio, video and code without establishing truth or rights. NIST's Generative AI Profile identifies risks that can be novel or intensified, including confabulation, dangerous content, privacy, bias, information integrity, intellectual property, cyber security and value-chain integration. The policy should require source checking, content labelling where applicable, provenance controls and human accountability.

Agentic AI adds action. An agent can plan steps, use memory, call tools, communicate and change systems. Singapore's Model AI Governance Framework for Agentic AI, published in January 2026, emphasises bounding autonomy, assigning human accountability, controlling access, testing, transparency and monitoring. These principles can inform internal governance while local law and context determine requirements.

An agent should receive the least authority required. Tools, data, destinations, transaction amounts, time windows and action types should be allowlisted. High-consequence actions should require confirmation or dual approval. Rate limits, budgets, environment segregation, immutable logs, anomaly detection and a tested kill mechanism should support control.

The business owner should document what the agent may decide, what it may recommend and what it may execute. The policy should prohibit an agent from expanding its own permissions or bypassing human gates. Connected workflows should be tested against malicious instructions and compromised data sources.

Figure 4. Approval and control architecture for AI and agents

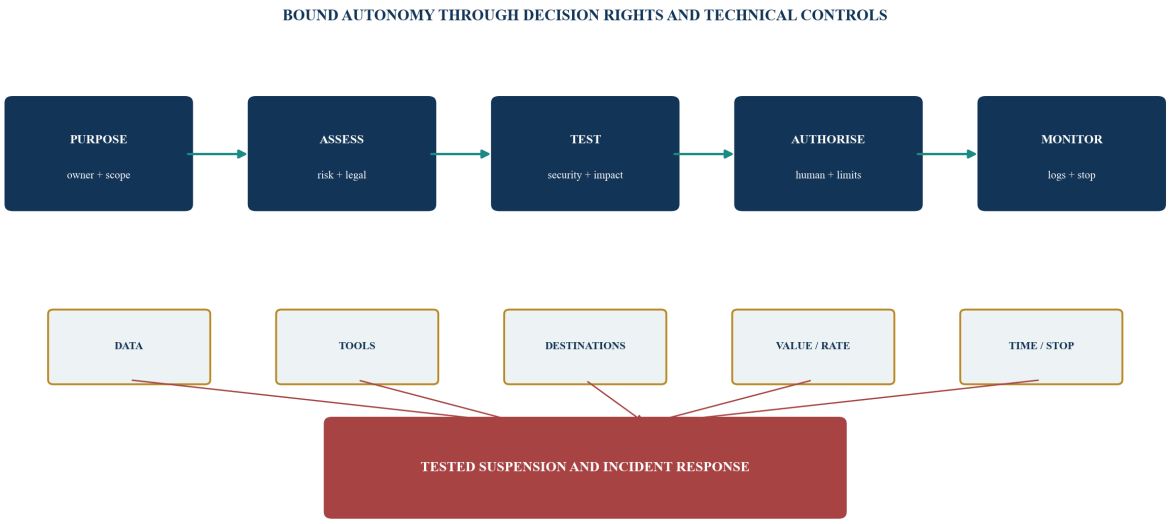

Autonomy expands only inside approved boundaries, with evidence and a tested stop mechanism.

10. Govern people, employment and capability

Employees need clear rules and usable approved tools. A policy that only prohibits behaviour can drive activity into unobserved channels. Training should explain permitted systems, data categories, verification, records, intellectual property, security, escalation and consequences. Role-based modules should address developers, procurement, human resources, investment teams, customer operations, executives and board members.

Employment uses require heightened care. Recruitment, monitoring, performance, scheduling, promotion and termination can affect rights and livelihoods. The review should consider applicable employment, equality, privacy, consultation and automated-decision rules. It should test data quality, proxies, disparate outcomes, explanations, accessibility and meaningful reconsideration.

The group should also plan workforce changes. AI can redesign tasks, controls and skill requirements. Management should assess segregation of duties, job design, supervision, competence and change fatigue. Critical knowledge should not move into an opaque tool without retention and continuity arrangements.

Family members and directors are users too. Personal assistants, messaging tools and devices can contain sensitive family, transaction and governance information. Onboarding and annual declarations should cover board portals, approved accounts, travel, voice and image cloning, impersonation and urgent-payment fraud.

11. Protect information integrity and the family name

Synthetic content can create operational and reputational harm. A false executive message can request money or data. Fabricated audio can affect a negotiation. An inaccurate public answer can be attributed to a family business. The group needs authentication, external publishing controls and an escalation path for impersonation and misinformation.

High-risk communications should use verified channels and secondary confirmation. Payment or account changes should never rely solely on voice, video or email. Public content generated or materially edited by AI should receive factual, legal and brand review appropriate to its consequence. Provenance and labelling requirements should be checked by jurisdiction and channel.

The crisis plan should identify the family spokesperson, entity spokesperson, legal lead, technology lead and law-enforcement or platform contacts. Monitoring should focus on credible threats and respect privacy. Evidence should be preserved before content is removed where lawful.

12. Map regulation and standards by entity and role

The regulatory map should be current and entity-specific. The European Union AI Act applies through a staged timetable and assigns obligations by system category and actor role. The European Commission states that most provisions apply from 2 August 2026, while certain provisions began earlier and some high-risk system obligations have later timing. Guidance, standards and enforcement arrangements continue to develop. Applicability and current dates require legal confirmation.

The United States has a sectoral and multi-level environment. NIST AI RMF is voluntary and provides a governance structure rather than a legal determination. Federal, state, sector, contract, consumer-protection, employment, privacy, intellectual-property and securities requirements may apply. The group should maintain a jurisdiction and use-case map rather than rely on one national label.

The United Kingdom has a principles-based regulatory approach across existing regulators, alongside data-protection and sector obligations. The ICO states that its AI and data-protection guidance is under review following the Data (Use and Access) Act 2025. The version and applicability should therefore be checked when a decision is made.

Singapore provides practical model frameworks for traditional, generative and agentic AI. These are useful governance references; entity obligations still depend on applicable law, sector and

role. The OECD principles offer a widely adopted basis for interoperability across jurisdictions.

The policy should record the legal source, effective date, affected entity, role, accountable owner and evidence. Regulatory change should trigger reassessment of classification, documentation, testing, contracts and incident duties.

13. Design incident response around evidence and decisions

An AI incident can involve confidentiality, personal data, cyber security, unsafe action, discrimination, financial loss, intellectual property, misinformation, regulatory breach or service interruption. The incident plan should connect existing crisis, cyber, privacy, legal, risk and communications procedures rather than create an isolated channel.

The first actions are to protect people, contain the system, preserve evidence and stop further harmful action. The team should identify affected entities, systems, versions, data, users, decisions and connected tools. Logs, prompts, outputs, configuration, model and vendor records should be preserved within legal limits.

The incident commander should coordinate technical and business decisions. Legal and privacy teams assess notification and privilege. The business owner assesses affected decisions and customer remediation. Security investigates access and compromise. Communications controls internal, family and external statements. Finance tracks loss and insurance. The vendor should support evidence, containment and remediation under contract.

Lessons should update the register, controls, training and vendor terms. Restart should require defined evidence and authority.

Figure 5. AI incident decision chain

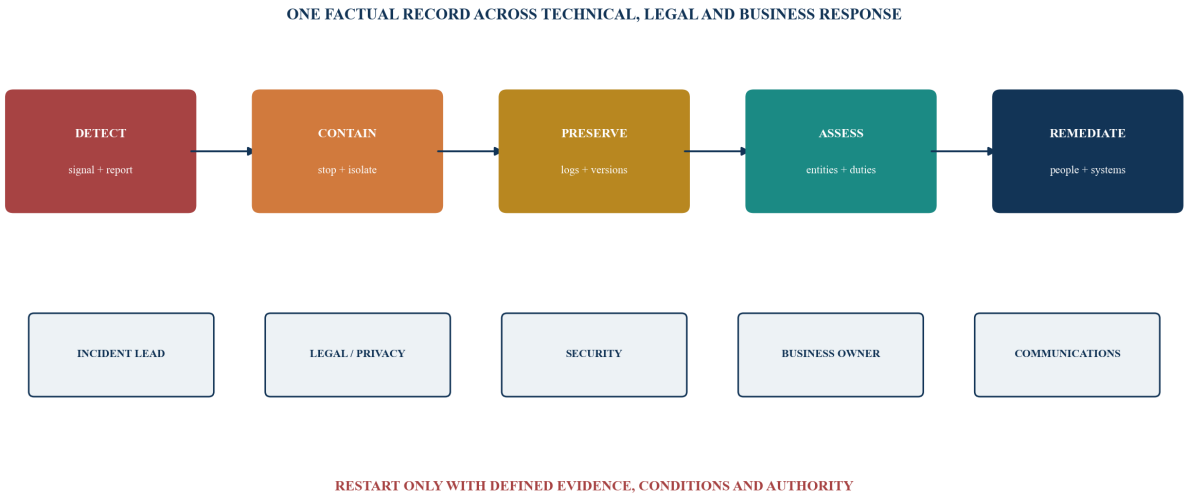

The sequence protects people and evidence while coordinating entity-specific duties.

14. Demonstrate the framework with a hypothetical global group

Consider a hypothetical family-owned group with a holding company in the United Kingdom, consumer and industrial businesses in the European Union and United States, a technology subsidiary in Singapore, and a separate investment office. The group has 4,800 employees and consolidated annual revenue of USD 1.2 billion. These figures are management assumptions created solely to demonstrate the method.

Discovery identifies 64 AI-enabled systems and 117 use cases. Twenty-nine systems are approved enterprise tools, 18 are embedded vendor features, 11 are employee subscriptions and six are internally developed. The initial use-case review identifies 12 restricted uses, 69 controlled uses, 30 routine uses and six uses suspended pending investigation.

The restricted set includes recruitment screening, factory safety alerts, dynamic customer pricing, credit recommendations, investment research, customer complaint triage and two agents connected to code and external communications. None is approved solely because a vendor markets it as compliant. Each receives a local legal map, impact assessment, data review, acceptance test, human-authority design and monitoring plan.

The board approves a group policy and four entity schedules. It prohibits confidential information in unapproved public tools, fabricated evidence, unapproved impersonation and autonomous movement of funds. It requires dual approval for material external communications and production code changes by agents. Six suspended uses are contained; two are retired, three are migrated to approved environments and one proceeds to a controlled pilot after remediation.

Table 5. Hypothetical family-group AI baseline and first decisions

Area	Management assumption	First board or management action
group footprint	UK holding company; EU, US and Singapore operations; separate investment office	adopt group principles and four entity schedules
scale	USD 1.2bn annual revenue; 4,800 employees	set materiality and reporting thresholds
systems	64 AI-enabled systems	assign system owners, versions and dependency records
use cases	117 identified	classify each use and remove duplicate or obsolete records
restricted uses	12	require legal, impact, testing and human-authority gates
suspended uses	6	contain, preserve evidence and decide retire, migrate or remediate
agentic systems	2 connected to code or communications	constrain tools, require dual approval, log actions and test stop controls

All values are management assumptions for method demonstration; they are not observed group data.

The family office has a separate issue. Analysts have used consumer assistants for public research and draft memoranda. The board requires an approved enterprise environment, source retention, confidential-data restrictions, investment-decision accountability and clear separation between tool output and approved investment evidence. The investment committee retains responsibility for every decision.

The 90-day dashboard shows complete ownership for 60 of 64 systems, tested incident contacts for all material entities, closure of the six suspended uses, and overdue vendor evidence for four controlled systems. The exceptions remain visible until resolved. The example demonstrates how policy becomes a portfolio of decisions rather than a statement of principles alone.

15. Execute a 100-day implementation sequence

During days 1 to 20, the board appoints an executive sponsor and policy owner, defines the group perimeter, issues an interim acceptable-use notice, preserves relevant evidence and identifies material systems. Management creates the entity, system and use-case registers and names incident contacts.

During days 21 to 40, the group classifies use cases, identifies prohibited and restricted activity, maps data and legal roles, and contains urgent exceptions. Procurement reviews material vendors and embedded features. Technology establishes approved access and initial detection controls.

During days 41 to 60, owners complete impact, privacy, security, intellectual-property and human-authority assessments for restricted uses. Technical teams build acceptance tests, permissions, logs, monitoring and suspension controls. Entity boards review local schedules.

During days 61 to 80, the group pilots training, incident exercises and control evidence. It tests agent boundaries and approval gates. Management closes or escalates exceptions, updates contracts and defines board metrics. Independent control functions challenge high-impact approvals.

During days 81 to 100, the board adopts the policy and entity schedules, approves material residual risks, and receives the first dashboard. Internal audit sets a risk-based assurance plan. The register becomes a recurring management record with event-driven reassessment.

16. Install recurring governance and event-driven review

The operating cadence should match change. System owners monitor production. Business owners review performance, complaints, overrides and exceptions monthly. Control functions review restricted uses at least quarterly. Boards receive material exposure, incidents and overdue remediation on an agreed schedule.

Events should trigger immediate reassessment. These include a new use, model or provider; material version change; new data category; expansion to another entity or jurisdiction; connection to a tool; increased autonomy; significant complaint; control failure; incident; acquisition; divestment; or regulatory change.

Acquisitions require early discovery. Diligence should identify material AI systems, data rights, contracts, incidents, claims and compliance gaps. Integration should preserve evidence and prevent uncontrolled connection to group systems. Divestments need data separation, licence transfer and access removal.

Succession also matters. A next-generation family member may sponsor innovation or join the board. Governance should use the same decision rights and evidence standards. The policy should remain an institutional asset that survives changes in leadership, vendors and technology.

Figure 6. Family-group AI governance calendar

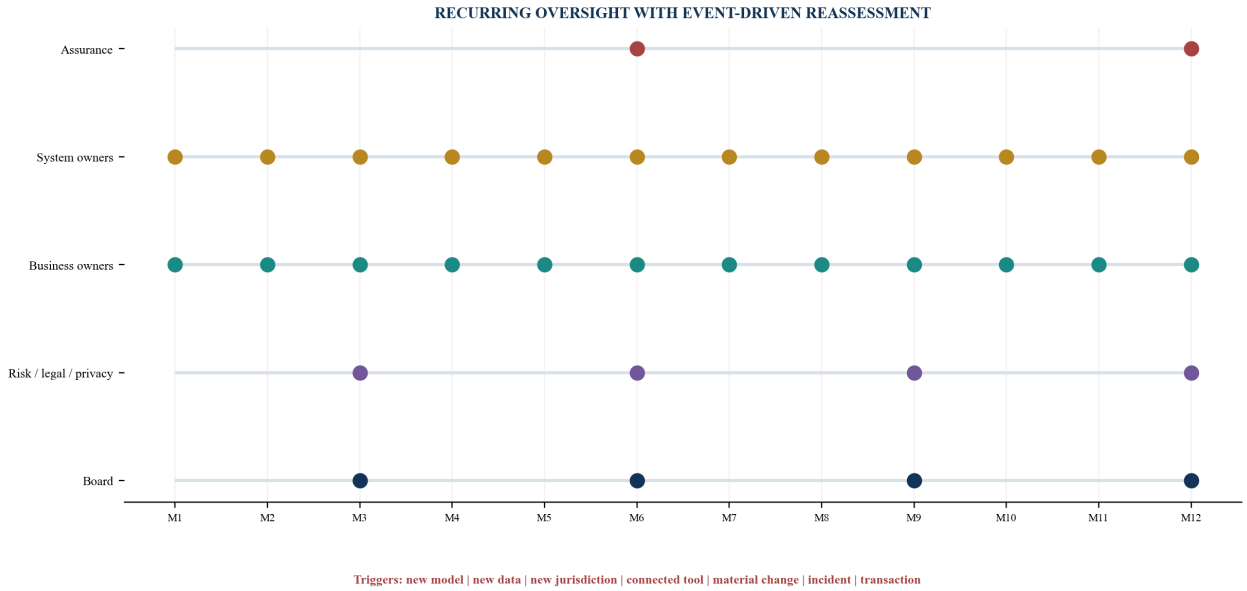

Event-driven reviews sit alongside recurring oversight; cadence can be stricter by entity or use.

17. Give the board a closed-loop scorecard

The board should trace each material AI use to an entity, system, owner, decision, data set, risk tier, authority, control and current evidence. The dashboard should show unknown systems, restricted uses, incidents, overdue reviews and remediation.

Table 6. Board-ready AI-governance implementation scorecard

Domain	Governing question	Acceptance evidence
perimeter	are legal entities, shared systems, vendors and material uses included?	reconciled entity, system, use-case and dependency registers
classification	is each use classified through impact, autonomy, data and reversibility?	approved rationale, legal map and escalation decision
decision rights	can accountable people approve, reject, override, pause and appeal?	authority matrix, user design, training and tested intervention
data and IP	are purpose, rights, confidentiality, retention and transfer controlled?	data map, impact review, contract, configuration and clearance
security and agents	are permissions bounded and connected actions observable and stoppable?	threat model, allowlists, logs, tests, limits and kill mechanism
vendors	are model, data, change, incident, resilience and exit risks governed?	due diligence, contract, monitoring and concentration plan
incidents	can the group contain, preserve, assess, notify and remediate?	exercised plan, contacts, evidence procedure and restart gate
oversight	does reporting drive closure and reassessment?	dashboard, exceptions, remediation dates and assurance findings

Acceptance requires current evidence and a named owner for every material exception.

The scorecard should avoid a false completion percentage. A group can have many registered systems and still lack control over the few that affect people, money, safety or reputation. The board should focus on material exceptions and evidence age.

AI governance becomes durable when it is connected to investment, procurement, employment, cyber security, data, transactions and operating management. The policy provides the rule set. The register provides the factual perimeter. Testing provides evidence. Decision rights provide accountability. Incident response protects people and value when assumptions fail.

References

[1] Organisation for Economic Co-operation and Development, OECD AI Principles, updated May 2024, https://oecd.ai/en/ai-principles

[2] United States National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 1.0, 26 January 2023, https://www.nist.gov/itl/ai-risk-management-framework

[3] United States National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1, 26 July 2024, https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

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

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

[6] European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models, 18 December 2024, https://www.edpb.europa.eu/news/edpb-opinion-on-ai-models-gdpr-principles-support-responsible-ai_en

[7] United Kingdom Information Commissioner's Office, Guidance on AI and Data Protection: Accountability and Governance Implications, under review following the Data (Use and Access) Act 2025, https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/what-are-the-accountability-and-governance-implications-of-ai/

[8] United Kingdom Department for Science, Innovation and Technology, Implementing the UK's AI Regulatory Principles: Initial Guidance for Regulators, February 2024, https://assets.publishing.service.gov.uk/media/65c0b6bd63a23d0013c821a0/implementing_the_uk_ai_regulatory_principles_guidance_for_regulators.pdf

[9] Singapore Infocomm Media Development Authority and Personal Data Protection Commission, Model Artificial Intelligence Governance Framework, Second Edition, https://www.imda.gov.sg/-/media/imda/files/infocomm-media-landscape/sg-digital/tech-pillars/artificial-intelligence/second-edition-of-the-model-ai-governance-framework.pdf

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

[11] Singapore Infocomm Media Development Authority, Model AI Governance Framework for Agentic AI, Version 1.0, 22 January 2026, https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf

[12] United States National Institute of Standards and Technology, AI Risk Management Framework Playbook and Resources, https://www.nist.gov/itl/ai-risk-management-framework/ai-risk-management-framework-resources

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.