

AI GOVERNANCE · 11 JULY 2026 · 10 MIN READ

UAE AI Governance for Boards: From Pilot to Accountable Deployment

A practical governance model helps UAE boards scale valuable AI while preserving accountability, evidence, human oversight and control through the system lifecycle.

Akash Chetwani, CFA
SENIOR PARTNER · ADVISORY

IRAA Global Advisory

Contents

AI GOVERNANCE · 10 MINUTE READ · JULY 2026

A practical governance model helps UAE boards scale valuable AI while preserving accountability, evidence, human oversight and control through the system lifecycle.

-
- | | |
|----|---------------------------------|
| 01 | Five controls before deployment |
|----|---------------------------------|
-
- | | |
|----|-------------------------------------|
| 02 | Create an inventory before a policy |
|----|-------------------------------------|
-
- | | |
|----|--|
| 03 | Assign accountability across the lifecycle |
|----|--|
-
- | | |
|----|-----------------------------------|
| 04 | Design meaningful human oversight |
|----|-----------------------------------|
-
- | | |
|----|-------------------------------------|
| 05 | Test the system in its real context |
|----|-------------------------------------|
-
- | | |
|----|---------------------------------------|
| 06 | Govern data and intellectual property |
|----|---------------------------------------|
-
- | | |
|----|-----------------------------|
| 07 | Monitor, respond and retire |
|----|-----------------------------|
-
- | | |
|----|---|
| 08 | Turn governance into a recurring management cycle |
|----|---|
-
- | | |
|----|---------------------------------------|
| 09 | Evidence that should survive scrutiny |
|----|---------------------------------------|
-
- | | |
|----|------------------|
| 10 | Official sources |
|----|------------------|
-

Artificial intelligence creates value only when a business can explain what the system does, which data it uses, who is accountable and what happens when it is wrong. Boards should govern AI as an operating system, not a collection of experiments.

The UAE AI Ethics Principles and Guidelines describe eight principles: fairness, transparency, accountability, explainability, human-centred design, privacy preservation, robustness and safety, and sustainability. The guidelines are non-binding, but they provide a useful governance baseline. Binding obligations can also arise from personal-data, consumer, sector, employment, intellectual-property and contract law.

Five controls before deployment

PURPOSE	Define the business decision, user and measurable benefit.
ACCOUNTABILITY	Name an executive owner, technical owner, risk reviewer and human decision-maker.
DATA	Confirm provenance, authority, quality, sensitivity and permitted use.
TESTING	Measure accuracy, bias, security, robustness and failure by relevant user group.
MONITORING	Track outcomes, complaints, drift, incidents, overrides and vendor changes.

Create an inventory before a policy

List AI systems already used by employees, embedded in software, supplied to customers or developed internally. Include generative assistants, scoring, forecasting, fraud detection, recruitment, pricing, customer service and biometric functions. Record the provider, model, data, users, decision, integration and countries involved.

Shadow AI is a material risk. Employees may paste customer, financial or source-code information into public tools because access is easy. Provide approved tools and clear handling rules rather than relying only on prohibition.

Classify each use case by impact. A writing assistant used on public material does not require the same governance as a model recommending credit, employment or healthcare decisions. Factors include effect on individuals, scale, sensitive data, reversibility, safety, financial exposure and regulatory significance.

Assign accountability across the lifecycle

The business owner should remain accountable for the outcome even when technology is supplied by a vendor. Technology teams manage integration and security; data owners manage access and quality; legal and risk teams review obligations; and independent assurance challenges high-impact systems.

Create an AI approval committee or use an existing risk forum. Require a short impact case covering purpose, alternatives, data, model, testing, human oversight, vendor, cost, legal basis and shutdown plan. Approval should expire when the model, purpose or data materially changes.

Board reporting should focus on material systems, not tool counts. Show benefits, incidents, outcome disparities, customer complaints, overrides, unresolved findings and concentrations in a model or cloud provider.

Design meaningful human oversight

“Human in the loop” is not meaningful if the employee lacks time, information or authority to disagree. Define which decisions require review, what evidence the reviewer receives, how overrides are recorded and when a case must be escalated.

Prevent automation bias through training and interface design. Show uncertainty and source limitations. Do not present generated language as confirmed fact or hide that a recommendation comes from a model.

For high-impact decisions, provide a route for affected people to ask questions, correct information and seek reconsideration. UAE personal-data law gives data subjects rights to information about decisions based on automated processing, including profiling, where applicable.

“An AI system is accountable only when a named person can stop it, explain it and correct its consequences.”

Test the system in its real context

Vendor benchmarks do not establish fitness for the company’s users and data. Build a representative test set, including difficult cases, minority groups, incomplete data and adversarial inputs. Define acceptance thresholds before reviewing results.

Test more than accuracy. Evaluate false positives and negatives, consistency, hallucination, explainability, latency, security, privacy leakage and resilience when source data changes. For generative AI, test prompt injection, confidential-data disclosure, unsafe content and unsupported citations.

Document limitations and prohibited uses. A model approved for internal summarisation should not silently become a customer eligibility engine. Use access controls and system design to enforce boundaries.

Govern data and intellectual property

Confirm that personal data has a lawful basis and is used for a compatible purpose. Minimise inputs, restrict sensitive information, set retention and understand whether a provider uses prompts or outputs for model training. Review cross-border processing and security.

Identify ownership and licence terms for training data, prompts, model outputs and code. Generated output can reproduce protected or confidential material. Require human review before publication or customer delivery and preserve sources where factual reliance matters.

Monitor, respond and retire

Models change through retraining, provider updates, new data and changing behaviour. Monitor performance and impact against the approved baseline. Set triggers for review, restriction or suspension.

Connect AI incidents to cybersecurity, privacy, legal and business-continuity response. Preserve logs and versions sufficient to reconstruct a material decision. Notify authorities or affected people where applicable obligations are triggered.

Retirement must remove access, revoke integrations, preserve required records and manage data held by vendors. Avoid an unsupported model continuing as an invisible dependency.

01	Inventory: Find internal, embedded and customer-facing AI.
02	Classify: Rank impact, data, reversibility and regulatory risk.
03	Approve: Document purpose, owner, evidence and boundaries.
04	Validate: Test outcomes, security and failure in context.
05	Monitor: Track performance, drift, complaints and changes.
06	Respond: Maintain override, incident and retirement procedures.

Responsible AI governance should make useful systems easier to deploy because evidence and ownership are clear. It should also stop attractive demonstrations from becoming high-impact processes before the business understands their limitations.

Turn governance into a recurring management cycle

A one-time approval is not enough because AI services, underlying models and business data can change after launch. Management should establish a review calendar based on impact. A low-risk drafting assistant may need a light annual check, while a system influencing credit, recruitment, pricing, fraud escalation or customer eligibility should be reviewed more frequently and after every material change. The review should confirm the approved use, current model and vendor version, data sources, validation results, incidents, complaints, overrides and outstanding actions.

The board does not need to approve every use case. It should approve the risk appetite, receive a consolidated view of high-impact systems and challenge whether management has reliable evidence. Useful reporting includes the number of deployed and experimental systems, high-risk approvals, overdue reviews, material vendor changes, incidents, human overrides, customer complaints and models operating outside agreed thresholds. Metrics should explain exposure and response, not simply count pilots.

Evidence that should survive scrutiny

For each material system, preserve a short decision record: the business purpose, accountable executive, affected people, input and output data, chosen model, known limitations, prohibited uses, testing approach, approval conditions and monitoring owner. Link that record to contracts, impact assessments, test evidence, security reviews and change history. If a decision later affects a customer, employee, regulator or court, the organisation should be able to reconstruct what the system did and why people relied on it.

Procurement, legal, privacy, security, risk and the operating team should use one approval path. Separate reviews often leave gaps between a vendor promise, a privacy notice and the actual workflow. A common launch gate makes unresolved issues visible and prevents commercial pressure from converting a pilot into production without ownership. The result is not bureaucracy for its own sake; it is a repeatable way to scale AI while protecting people, reputation and enterprise value.

Official sources

1. **UAE AI Ethics Principles and Guidelines**
2. **UAE position and guiding principles on AI policy**
3. **Federal Decree-Law No. 45 of 2021 on Personal Data Protection**

Online edition: <https://iraaglobal.com/insights/uae-ai-governance-board-accountable-deployment/>



Akash Chetwani, CFA

SENIOR PARTNER · ADVISORY

Cross-border advisory for professionals, founders, investors, and globally connected families.

PHONE

+971 56 921 0222

EMAIL

info@iraaglobal.com

PROFILES

[LinkedIn](#) · [Contributor page and articles](#)

A connected advisory firm for businesses operating across borders.

OUR EXPERTISE

Audit & Assurance

Corporate Tax & VAT

Accounting & Bookkeeping

Finance & CFO Advisory

AML, Risk & Compliance

Business Advisory & Setup

Human Resources Advisory

AI, Technology & Digital Assets

India-UAE Cross-Border Advisory

Private & Family Advisory

Insight beyond numbers.

IRAA Global Advisory

503 Mustafawi Carpet Building, Sharaf DG Metro Exit 1
Dubai, United Arab Emirates

+971 56 921 0222 · +971 50 677 9455
info@iraaglobal.com · www.iraaglobal.com