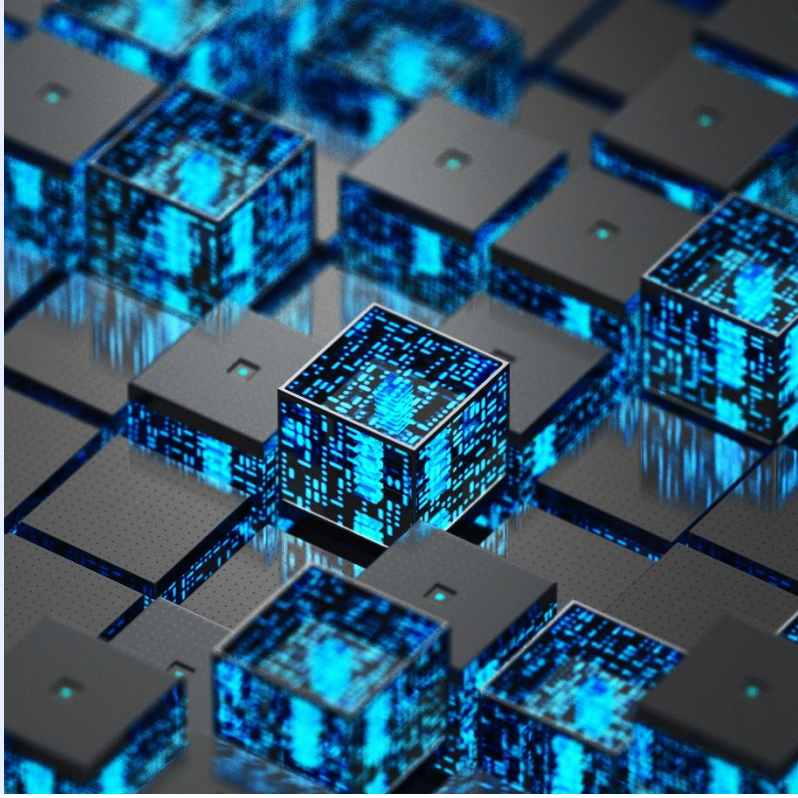
AS&H Clifford Chance

BRIEFING

Responsible AI Policy
A Shift from a Guiding Regime to a Supervisory One

4 November 2026



Key Issues covered

- **Introduction of the RAI Policy**
- **Key Provisions of the RAI Policy**
  - Scope and applicability
  - Fairness, non-discrimination, and ethical AI principles
  - Risk-based regulatory approach
  - Registration and Badge Classifications
  - Governance and accountability obligations
  - Human oversight and intervention
  - Safeguards and Safety Reports
  - Data governance and privacy considerations
  - Roles and Responsibilities
  - AI Regulatory Sandbox
- **Takeaways (How to prepare?)**
  - Public Sector entities
  - Private & Non-Profit Sector Entities
- **AI-First Policy**
- **Conclusion**

Until now, AI governance in the Kingdom of Saudi Arabia (the "**Kingdom**") has operated largely through voluntary guidelines and non-binding principles. The draft Responsible AI Policy (the "**RAI Policy**"), published for public consultation by the Saudi Data & AI Authority ("**SDAIA**"), will mark a fundamental shift. For the first time, the Kingdom will put in place a formal framework for artificial intelligence, one that introduces binding compliance obligations, structured accountability rules, and explicit supervisory and enforcement powers for SDAIA.

For businesses operating in the Kingdom, this is a significant development. AI compliance in the Kingdom will move from a best-practice exercise into a formal regulatory requirement, with real implications for any entity that develops, operates, deploys, or uses AI systems.

While the RAI Policy remains in draft form, businesses should begin assessing the potential impact on their operations and preparing for the anticipated compliance requirements.

Introduction of RAI Policy

SDAIA published the RAI Policy for public consultation. The Kingdom's existing AI framework rests primarily on guidance-based, principle-led instruments which are useful for setting direction, but without binding force or enforcement mechanisms. The RAI Policy will fundamentally change this. It will introduce mandatory rules, formal compliance obligations, and regulatory powers for SDAIA, bringing the Kingdom in step with the new direction of AI regulation globally.

Key Provisions of the RAI Policy

The following presents an overview of the main components of the RAI Policy:

a) Scope and applicability

The RAI Policy will apply to all entities in the public, private and non-profit sectors, and individuals involved in any aspect of AI within the Kingdom, whether as developers, operators, deployers, or users. Obligations extend across the full lifecycle of an AI system, from initial development and

testing through to deployment, ongoing operation, and eventual decommissioning. The RAI Policy will be introduced in phased stages, with mandatory compliance in the first stage applying to entities in high-risk sectors and large companies, as determined by SDAIA.

b) Fairness, non-discrimination, and ethical AI principles

The RAI Policy will formally embed the seven ethical AI principles previously set out in SDAIA's AI Ethics Principles. These include fairness, privacy and security, humanity, social and environmental benefit, reliability and safety, transparency and explainability, and accountability. Developers, operators and users of AI systems will be expected to assess and address risks of bias, unfairness, or discriminatory outcomes in their AI systems throughout the lifecycle.

c) Risk-based regulatory approach

A central feature of the RAI Policy is its risk-based framework which classifies AI systems into four tiers. These are Critical Risk, High Risk, Limited Risk, and Minimal or No Risk. Higher-risk systems will attract more stringent governance, monitoring, and compliance requirements. Notably, the deployment of any AI system classified as Critical Risk is prohibited outright under the RAI Policy.

SDAIA will establish and maintain a general AI risk framework, updated periodically in line with technological developments and international standards.

d) Registration and Badge Classifications

Under the RAI Policy, entities operating across various sectors may undertake several AI governance and compliance measures, including the following:

- **Registration of AI activities:**

  Private sector entities developing, operating, or deploying AI services or products may register with SDAIA and obtain an annually renewable registration certificate. Separately, all AI service providers must register their products on the AI Governance Platform, and government entities' AI offices must register AI products with SDAIA

- **AI Ethics Badges:**

  Entities in the public, private, and non-profit sectors may apply for AI ethics badges for their products, awarded based on compliance maturity and risk profile. Badges are granted across five graduated levels: Aware, Adoptive, Confirmative, Assured, and Visionary. Importantly, high-risk AI systems are required to hold at minimum the "Assured" badge, while limited-risk systems must hold at least the "Confirmative" badge. These badges serve as indicators of an AI system's level of commitment to the AI Ethics Principles issued by SDAIA, with the aim of enhancing trust among users and other stakeholders.

- **High-risk AI systems:**

  Entities in the public, private and non-profit sectors developing, operating, or deploying AI systems classified as high-risk fall within the AI Regulatory Sandbox registration framework, with the detailed arrangements to be established by SDAIA.

e) Governance and accountability obligations

The RAI Policy will place considerable emphasis on internal governance. Entities across the public, private and non-profit sectors involved in the development or operation of AI systems will need to put in place formal governance structures including defined roles and responsibilities, oversight mechanisms, and internal controls. Documentation and recordkeeping throughout the AI lifecycle are also required to ensure traceability and auditability of AI-related decisions.

f) Human oversight and intervention

The RAI Policy will require meaningful human oversight over AI systems, particularly for automated decision-making and high-risk applications. No AI system should operate without human supervision where it may affect individuals or critical functions. Entities across the public, private and non-profit

sectors must appoint one or more responsible AI assessors who are individuals or entities responsible for monitoring systems, implementing safety guardrails, conducting safety assessments, and managing access to control records and safety reports.

g) Safeguards and Safety Reports

Entities across the public, private and non-profit sectors involved in the development, operation, or deployment of AI systems (in particular high-risk systems) are expected to implement appropriate safeguards, monitoring mechanisms, and risk mitigation measures in accordance with the applicable AI risk level.

The RAI Policy will also impose ongoing monitoring, testing, documentation, and reporting obligations. Before deploying any high-risk AI system entities must conduct comprehensive safety testing, apply the relevant safeguards, and document the process in a formal safety report. Reports must cover the safety guardrails applied, test results across all input and output types, misuse scenarios tested, and post-deployment monitoring plans. Deployment is prohibited unless safety reports and technical testing results are submitted to and reviewed by SDAIA. High-risk systems must also undergo annual audits by an accredited third party

h) Data governance and privacy considerations

The RAI Policy will be closely aligned with the Kingdom's broader data legislative framework, including the Personal Data Protection Law ("**PDPL**"). It will set clear expectations for the lawful and responsible use of data in AI systems, covering data quality, accuracy, fairness, security, and privacy. Entities across the public, private and non-profit sectors involved in the development, operation, or deployment of AI systems will need to ensure that data used in AI development and deployment complies with applicable Saudi data protection requirements.

i) Roles and Responsibilities

The RAI Policy will assign specific responsibilities to each category of participant in the AI lifecycle. This includes developers, operators, AI service providers, content creators, users, and public sector entities based on their level of control and involvement. Responsibility will not be concentrated only on end-users, developers must embed responsible AI requirements from the design stage; operators and deployers must implement and maintain appropriate controls; and all parties must comply with applicable legal and ethical standards throughout the system's lifecycle.

j) AI Regulatory Sandbox

The RAI Policy will establish an AI Regulatory Sandbox which is a supervised testing environment allowing entities to develop and trial AI systems under SDAIA's oversight before market deployment. Companies developing high-risk AI systems and providers of AI system safety certification services must obtain a certificate of successful Sandbox participation before they can offer their services in the market.

Takeaways (How to prepare?)

a) Public Sector Entities

Public sector entities are directly in scope subject to the phased implementation described above. Compliance will be led through each entity's dedicated AI office. In particular, public sector entities (through their AI office) will need to:

- conduct an inventory of existing AI systems;
- establish internal AI governance and oversight frameworks, and assess whether current governance arrangements align with the proposed requirements;
- classify AI systems according to their level of risk;
- implement human oversight mechanisms, particularly for high-risk or sensitive AI applications;

- ensure compliance with applicable data governance, cybersecurity, privacy, and ethical AI requirements;
- maintain monitoring, documentation, and reporting procedures relating to AI systems;
- define and document the roles of all personnel involved in the access and development of AI systems;
- conduct continuous assessments of AI risks and operational impacts; and
- ensure that AI deployment aligns with public interest, transparency, fairness, and non-discrimination principles.

Public sector entities operating high-risk AI systems will be subject to enhanced obligations, including mandatory registration, and safety reporting to SDAIA. The RAI Policy places the public sector at the center of implementing responsible AI governance across the Kingdom.

b) Private & Non-Profit Sector Entities

Private and non-profit sector entities developing, operating, deploying, or using AI systems in the Kingdom will need to build formal governance and compliance programs. In practical terms, this means:

- conduct an inventory of existing AI systems;
- identify and classify AI systems according to risk levels;
- establish internal AI governance, oversight, and accountability structures;
- implement safeguards, monitoring procedures, and risk mitigation controls;
- ensure lawful and responsible use of data in compliance with the PDPL and related frameworks;
- maintain technical documentation, testing records, and safety reports;
- implement transparency and explainability measures for AI-generated outcomes;
- ensure meaningful human oversight for AI systems;
- continuously monitor AI systems after deployment; and
- prepare for possible registration, certification, AI Ethics Badge applications, or AI Sandbox participation for high-risk AI systems.

Responsible AI governance is no longer a purely ethical or technical matter. Under the draft RAI Policy, it will be a formal policy requirement. Entities operating AI systems in the Kingdom should integrate AI governance into their broader compliance, risk management, cybersecurity, and data protection frameworks without delay.

AI-First Policy

SDAIA has published another draft policy which is the AI-First Policy (the "**AIF Policy**") for public consultation. The AIF Policy will apply to public sector entities and recommends that state-owned companies and private sector entities benefit from its principles where appropriate. The AIF Policy will seek to establish a government-wide framework encouraging public entities to treat AI as the preferred option when developing services, operations and initiatives, where AI is demonstrated to be suitable and capable of delivering measurable value. The AIF Policy will also identify the sectors and use cases in which the AI First approach should be prioritized, sets out the key principles underpinning AI adoption, and establishes the strategic pillars intended to support the effective and sustainable adoption of AI across public sector entities.

The AIF Policy is intended to serve as a high-level and principle-based policy that will complement existing laws and future guidance to be issued by SDAIA and other competent authorities. Public entities will be expected to consider the policy when developing their services and operations, taking into account their institutional readiness and applicable legal and regulatory requirements.

Together, the two policies will provide a clearer indication of the Kingdom's intended AI policy direction: encouraging AI adoption while simultaneously introducing governance and accountability requirements.

Conclusion

The RAI Policy will potentially be significant development in the Kingdom's AI landscape once approved. It will mark a clear transition from voluntary guidelines and a move toward a binding regulatory framework with mandatory compliance obligations, risk-based oversight, defined roles and responsibilities, and formal supervisory and enforcement powers for SDAIA.

The RAI Policy remains in draft form and may evolve following the public consultation process. However, its approach is clear, and entities across the public, private and non-profit sectors may consider reviewing their AI governance arrangements. Entities across the public, private and non-profit sectors involved in the development, operation, or deployment of AI applications and solutions, particularly those operating in high-risk sectors, should begin assessing their readiness and adoption arrangements.



Dr. Fahad A. Abuhimed
Managing Partner

Email: fahad.abuhimed
@ashcliffordchance.com
Mobile: +966 55 502 2298



Hanan Hassanain
Partner

Email: hanan.hassanain
@ashcliffordchance.com
Mobile: +966 56 296 5481



Muath Alsowayan
Senior Associate, Riyadh

Email: muath.alsowayan
@ashcliffordchance.com
Mobile: +966 59 013 0153



Zyad Alabduljabbar
Junior Associate, Riyadh

Email: Zyad.Alabduljabbar
@ashcliffordchance.com
Mobile: +966 55 517 6340

This publication does not necessarily deal with every important topic or cover every aspect of the topics with which it deals. It is not designed to provide legal or other advice.

ashcliffordchance.com

AS&H Clifford Chance, Building 15, The Business Gate, King Khaled International Airport Road, Cordoba District, Riyadh. P.O. Box: 90239, Riyadh 11613, Kingdom Of Saudi Arabia

© AS&H Clifford Chance 2026

Fahad Abuhimed, Majid Alsheikh, Mansoor Alhagbani and Clifford Chance Company for legal is a professional limited liability company, Unified National no. 7009448064, Chamber of Commerce membership no. 369478, with paid up capital of SAR. 150,000.

Abu Dhabi • Amsterdam • Barcelona • Beijing • Brussels • Bucharest** • Casablanca • Delhi • Dubai • Düsseldorf • Frankfurt • Hong Kong • Houston • Istanbul • London • Luxembourg • Madrid • Milan • Munich • Newcastle • New York • Paris • Perth • Prague** • Riyadh* • Rome • São Paulo • Shanghai • Singapore • Sydney • Tokyo • Warsaw • Washington, D.C.

*AS&H Clifford Chance, a joint venture entered into by Clifford Chance LLP.

**Clifford Chance has entered into association agreements with Clifford Chance Prague Association SRO in Prague and Clifford Chance Badea SPRL in Bucharest.