Volume 7 – Issue 80
Issued on: 20 June 2026
International Journal of Research and Studies Publishing
المجلة الدولية لنشر البحوث والدراسات
ISSN: 2709-7064
IJRSP

Electronic Governance Controls for AI-Enabled Agentic Large Language Models in Digital Government (Compliance-by-Design and Evidence-Linked Evaluation)

Dr. Maram Abdulrahman Almaghrabi

Assistant professor, Data Science Department, College of Computing, Umm Al-Qura University, Makkah, Saudi Arabia

Email address: mammaghrabi@uqu.edu.sa

https://orcid.org/0000-0002-8106-078X

Abstract:

Artificial intelligence (AI), particularly large language models (LLMs) and agentic AI systems, is increasingly used to modernize digital government, yet public-sector deployment must satisfy accountability, privacy, and auditability requirements. This paper aims to propose an electronic-governance framework that enables governed use of LLM-enabled agentic systems under binding constraints. Through qualitative documentary analysis, the study examines policy and governance evidence across digital-government strategies, data-governance and privacy instruments, and public-sector AI initiatives. Empirical grounding is strengthened through triangulation with official digital-government maturity indicators, a flagship platform vignette illustrating operational scalability, and a reproducible open-data key performance indicator (KPI) baseline derived from a government complaint-escalation series. The study concludes that scalable public value depends on bounded autonomy implemented through allowlisted tools, controlled state and memory, verification and feedback loops, role-based access control, audit logging, and human-in-the-loop gating. The paper contributes an evaluation-ready input–process–output framework to assess productivity, decision quality, and auditable compliance aligned with Vision 2030 and SDAIA-led transformation, The paper further strengthens feasibility claims through official maturity indicators and flagship implementation evidence, showing advanced service-experience maturity and institutional transformation readiness that support scaling LLM-enabled service augmentation and governance-constrained agentic execution under regulated constraints.

Keywords: Electronic Governance; Digital Government; Artificial Intelligence Governance; Large Language Models; Agentic AI; Accountability; Auditability; Compliance

1. Introduction:

Artificial intelligence is recognized as a general-purpose technology with economy-wide effects on productivity, organization, and innovation (Filippucci et al., 2024). There are two routes currently driving the applied AI revolution. First, generative artificial intelligence, especially Large Language Models (LLMs), has reduced the unit cost of certain tasks like drafting, summarization, classification, translation, and coding. This facilitates scaling and standardization in these text-centric tasks (Brynjolfsson et al.,2024). Second, agentic AI goes beyond just making things by combining foundation models with planning, controlled tool use, state/memory, verification loops, and actions that may be carried out (Abou Ali & Dornaika, 2025). In critical public services, the most justified trajectory is not autonomy but governance-constrained or bounded agency, where execution is restricted by permissions, auditability, and human approval gates.

Saudi Arabia's Vision 2030 elevates data and AI as national priorities, reflected in the establishment of SDAIA and the National Strategy for Data and Artificial Intelligence (NSDAI), positioning the Kingdom to scale AI adoption within a governed public-sector ecosystem (SDAIA, 2023). Accordingly, the central research problem is not whether LLMs and agentic systems matter, but how they can reshape digital-government services under binding regulatory constraints, and which enforceable governance conditions (data governance, privacy controls, and accountability mechanisms) enable safe, scalable modernization.

1.1. Research Questions:

- **RQ1:** How do LLMs and agentic AI reshape digital governance and public-service delivery in Saudi Arabia?
- **RQ2:** Which productivity, decision-quality, and innovation mechanisms are most salient in Saudi digital government modernization?
- **RQ3:** What governance, privacy, and ethical risks arise, and which enforceable controls best mitigate them in the Saudi context?

1.2. Research Contributions

This paper makes three contributions which are:

• Governance-first transformation framing: It conceptualizes regulated agentic transformation as a public sector modernization pathway where AI-enabled productivity gains are conditioned by enforceable governance, accountability, and trust preservation under binding constraints.

• Operational IPO governance model: It proposes an Input–Process–Output (IPO) framework linking national drivers and regulatory guardrails to bounded agentic mechanisms and measurable modernization outcomes put supporting governance-by-design decision making.

• Evaluation readiness: It strengthens feasibility claims through maturity-indicator triangulation (DXMI/Qiyas and GTMI), a flagship vignette, and a reproducible open-data KPI baseline that operationalizes output-layer signals for future empirical evaluation of bounded agent deployments.

The paper is organized accordingly: following the introduction, the literature review synthesis prior work on generative and agentic AI in public-sector contexts, emphasizing governance and compliance as core scaling conditions. Section 3 develops governance-ready foundations of Saudi digital government (SDAIA/NSDAI orchestration, interoperability and data governance, PDPL compliance-by-design, and sovereign Arabic LLM capacity via ALLaM). Section 4 presents the proposed Input–Process–Output (IPO) framework for regulated agentic transformation, Section 5 outlines the methodology, and Section 6 reports findings structured by the IPO layers and supported by indicator-based triangulation and a flagship case vignette. The paper then provides a technical open-data KPI add-on (Saudi Water Authority complaint escalation) and a practical bounded-agent prototype for citizen case triage and routing with risk-based autonomy constraints, before concluding with contributions, limitations, and directions for future research.

2. Background and Related Work:

Large language models (LLMs) are envisioned as a sort of cognitive automation, their integration is seen to reduce the marginal cost of producing, transforming, and standardizing text-based artefacts at a large scale, which makes language-intensive workflows more productive (Filippucci et al., 2024). Empirical studies report measurable improvements in tasks such as drafting, summarization, and communication, and further suggest institutional learning effects when workers repeatedly engage with structured outputs and best-practice patterns (Brynjolfsson et al., 2024). In public-sector contexts, where service delivery depends on high volumes of documentation, case handling, and inter-agency communication LLMs can therefore be interpreted as a language operating layer that increases throughput and consistency, provided that deployment is constrained by governance, safety, and accountability requirements (Filippucci et al., 2024).

However, beyond mere generation, agentic AI enhances the value proposition of foundation models with regard to goal-directed autonomy, which is achieved with system-level architectures that provide planning, task decomposition, controlled tool use, verification, feedback, and governance, including

logging, access control, approval, etc. (Abou Ali & Dornaika, 2025) From this perspective, it is argued that agentic AI is not merely a chat-based feature but rather an architectural pattern that is used to orchestrate end-to-end work with bounded autonomy. As systems move from providing advisory outputs to actionable tasks, governance issues become increasingly complex with concerns related to misaligned actions, unclear operational flows, unsafe execution, and increased exposure to sensitive information, etc. Consequently, prior studies emphasize that, in the case of agentic AI in public service delivery, it is essential to design bounded, auditable socio-technical systems with explicit accountability mechanisms, especially in high-stakes decision-making processes (NIST, 2023).

From this perspective, it is argued that governance and trust are essential prerequisites to the scalable adoption of agentic AI systems. Trustworthy AI increasingly depends on enforceable compliance instruments rather than voluntary principles alone, especially in high-impact environments such as government services (NIST, 2023). In Saudi Arabia, the Personal Data Protection Law (PDPL) outlines essential principles of personal data lawfulness, limiting purposes, and security measures, which become particularly relevant in the context of agentic design with tool execution and memory capabilities, which may increase the privacy risk surfaces if not controlled (SDAIA, 2023). In parallel, national data governance instruments—such as the National Data Management Office (NDMO) governance regulations—shape what systems can access and how they can operate at scale, particularly when interoperability and cross-entity data sharing are prerequisites for agentic action (NDMO, 2020). As a result, governance preparedness becomes a key factor in agentic scalability. Operational autonomy is not only based on model capabilities, but also on controlled access, auditability, and enforced rules limitations that maintain trust, compliance, and responsible execution (NIST, 2023).

Following the literature review, the paper proceeds with a contextual foundations section (Section 3) that draws on official strategies, regulations, and governance instruments to characterize Saudi Arabia's governance-ready digital-government environment. This section establishes the enabling conditions for bounded agentic execution and informs the subsequent IPO framework and empirical triangulation.

3. Governance-Ready Foundations of Saudi Digital Government:

3.1. Centralized Orchestration as a Scaling Mechanism: SDAIA and NSDAI

Saudi Arabia's digital-government trajectory is characterized by an institutional model that privileges coordinated, centrally guided scaling over fragmented or agency-by-agency

experimentation. Within this model, the Saudi Data and AI Authority (SDAIA) serve not only as an innovation enabler but also as a national orchestration actor that aligns entities around shared priorities, standardized governance expectations, and platform-level capabilities (SDAIA, 2020).

In parallel, the National Strategy for Data and Artificial Intelligence (NSDAI) provide strategic direction and an implementation logic that links national ambitions, such as service modernization and economic diversification under Vision 2030, to measurable transformation pathways across the public sector (SDAIA, 2020).

From a governance perspective, this orchestration architecture matters because the public sector's capacity to scale AI is not determined solely by model capability, but by whether institutions can enforce consistent controls, reduce duplication, and maintain accountability across heterogeneous agencies. In practice, direction at the NSDAI level and coordination at the SDAIA level can function as a bridge between policy and implementation. This is especially relevant to generative and agentic AI systems whose value proposition lies in their ability to penetrate all spheres of high-volume workflows while their risk profile is linked to the level of entity integration (NIST, 2023).

3.2. Interoperability and Data Governance as Preconditions for Agentic Execution

Agentic AI systems depend on reliable access to high-quality, interoperable data to support planning, tool execution, and iterative feedback loops (Abou Ali & Dornaika, 2025). However, in public-sector environments, data access and interoperability are governance decisions before they are technical capabilities. In the Saudi Arabia context, the instruments of governance established by the National Data Management Office (NDMO) include requirements for standardization, classification, sharing, and accountability. These requirements determine what agentic systems can access, how they can combine data from different sources, and when they can take action (NDMO, 2020).

This governance layer is critical because agentic systems shift the operational boundary from advice (text generation) to execution (tool-based action). The moment an artificial intelligence system starts to manage the workflow between different administrative entities, such as routing, records, notifications, and eligibility, interoperability is simultaneously enabled and poses risks.

In other words, data governance, as well as interoperability readiness, acts as a binding constraint to determine whether bounded agentic design can be scaled up without undermining traceability, control, and institutional accountability (NIST, 2023). In other words, agentic feasibility in digital government is a function of (governed integration capacity) not merely algorithmic performance.

3.3. Compliance-by-Design Constraints in Public Services: PDPL

One of the important distinguishing features of high-impact digital government is the prevalence of sensitive identity-related or legally relevant data within the service delivery process. Accordingly, privacy law and compliance expectations become design requirements for AI systems, rather than downstream policy considerations. The Personal Data Protection Law (PDPL) of Saudi Arabia, as reflected in the SDAIA regulatory framework and implementation guidelines, sets the basic requirements for lawful processing, purposes of processing, data protection measures, and accountability (SDAIA, 2023). This is especially relevant for agentic architectures because of the elevated risk of privacy violations arising from the use of autonomy features such as memory/state persistence, tool invocation, and cross-case learning if not strictly constrained during the design phase. Thus, the notion of a (governance-ready) AI system within the Saudi digital government context would necessarily involve (compliance by design) principles that include data minimization, access control, auditing, and data retention that are embedded within the overall system design and are executed deterministically. Accordingly, PDPL-aligned constraints function as enablers rather than barriers: they provide the foundation for reliable scaling by ensuring that productivity gains from LLM-enabled workflows are not achieved at the expense of privacy, access control, or accountability (NIST, 2023).

3.4. Sovereign Arabic LLM Capacity as a National Capability Layer: ALLaM

While the digital government instruments and platform readiness are important factors for the adoption of AI/LLM technology, the Saudi government's investment in sovereign and Arabic-language capability represents an important additional layer of strategic capability, that shape the ability of the nation-state to leverage language model technology that meets its own requirements while allowing for the integration of the technology into the broader digital environment (IBM, 2024).From a digital-government standpoint, sovereign or locally anchored model capacity can support operational and governance priorities: improved Arabic language performance for citizen interactions, tighter oversight over deployment contexts, and potential alignment with national compliance expectations. While sovereign models do not automatically guarantee safety or governance maturity, they can strengthen the feasibility of regulated deployment by enabling more controlled integration patterns and contextualization within national service ecosystems. In sum, ALLaM can be regarded as a capacity layer that complements Saudi Arabia's broader orchestration and governance strategy, providing scalable, Arabic-centric service augmentation while preserving institutional oversight requirements (IBM, 2024).

4. IPO Framework for Regulated Agentic Transformation

4.1. Inputs: Drivers, Guardrails, and Deployment Environments

The proposed Input-Process-Output (IPO) framework posits the impact of artificial intelligence (AI) in Saudi Arabia's digital government as mediated by institutional drivers and enforceable constraints, as opposed to technological diffusion. The input level reveals four drivers as highly relevant. First, the strategic direction, grounded in Vision 2030 and the National Strategy for AI (NSDAI), supplies the national-level mandate and performance orientation for digital government modernization (SDAIA,2020). Second, regulatory and ethical boundaries, including the Personal Data Protection Law (PDPL) and relevant national data governance regulations, specify prescriptive requirements that limit the capabilities of AI systems, the data that AI systems can access, and how accountability is managed (NDMO, 2023).Third, platforms and deployment environments, including digital government platforms and enterprise deployment partnerships, influence how well AI systems can be integrated with operational workflows at scale (IBM, 2024). Finally, capability building and development of the Arabic-first sovereign model enhance the feasibility of AI systems' use and controlled deployment in Saudi Arabia. The above elements collectively define the governance and platform readiness envelope within which generative and agentic AI systems may safely be scaled.

4.2. Processes: Bounded Tool-Use, Oversight, and Verification

At the process level, the framework considers the difference between generative capabilities (text production) and agentic capabilities (planning and executing goal-directed actions). The most defensible approach in a governed public sector context is not autonomy but bounded autonomy. This refers to systems with capabilities to plan, retrieve information, produce text, and execute actions with strict bounds. This suggests that some of the fundamental mechanisms include task decomposition, tool usage with allowlists, memory management, verification, and human-in-the-loop gating for critical decision-making processes (Sapkota et al., 2025).

It is important to note that bounded agentic execution should be considered a design pattern with embedded governance, including role-based access control, comprehensive logging, policy evaluation, and workflow escalation. These are not optional add-ons; rather, they are the operating conditions under which agentic value creation becomes feasible in contexts where mistakes and privacy violations can lead to institutional and legal consequences (NIST, 2023). In this view, agentic systems become scalable only when autonomy is operationalized as a controlled execution capability auditable, constrained, and accountable by design.

4.3. Outputs: Measurable Productivity, Decision Quality, and Auditable Compliance

At the output layer, the IPO model anticipates measurable outcomes in four categories. First, productivity includes improved processing efficiency, including reduced time and increased throughput, for language-intensive operational workflows, including triage, drafting, summarization, and multi-modal communication, which is consistent with evidence that suggests large language models improve productivity, especially for writing- and communication-intensive activities (Brynjolfsson et al., 2023). Second, innovation includes improved Arabic digital products enabled by sovereign capability building and integration capacity, including faster prototyping, which is enabled by Arabic language models and relevant digital platforms (IBM, 2024).Third, decision quality outcomes include improvements in timeliness and consistency for evidence-based governance, which can be facilitated by standardized outputs that reduce variability and promote learning. Fourth, governance modernization outcomes include auditable compliance, traceable actions, and standardized services under enforceable constraints, which align with digital transformation, accountability, and compliance with the PDPL and NDMO (NDMO, 2020; SDAIA, 2023; NIST, 2023). In aggregate, the IPO model captures the core argument of regulated agentic transformation, which holds that digital government scaling in Saudi can be sustained only if strategic direction, capability development, and enforceable governance advance in tandem, ultimately achieving measurable modernization in digital services while preserving privacy, accountability, and trust.

5. Methodology:

This study applies qualitative documentary analysis with structured coding and comparative synthesis to examine how LLMs and agentic AI can modernize Saudi digital government under enforceable governance constraints. The corpus comprises SDAIA/NSDAI materials, NDMO governance instruments, PDPL and implementing guidance, ALLaM announcements, and selected global evidence on generative-AI productivity and diffusion. Analytically, documents were processed through an IPO-aligned workflow: (i) coding sources into driver categories; (ii) extracting and normalizing agentic mechanisms relevant to public-service modernization (planning, controlled tool use and state, verification loops, RBAC, audit logging, and HITL gating); (iii) mapping drivers → mechanisms → expected outcomes via the IPO model; and (iv) conducting a PDPL/NDMO-aligned risk–control assessment to specify enforceable controls for high-stakes services.

Empirical anchoring was strengthened through triangulation using DXMI/Qiyas and GTMI, complemented by a flagship vignette and an open-data KPI baseline (SWA complaint escalation) to operationalize output-layer signals and define evaluation-ready measures for governance-constrained.

5.1. Data and Materials:

This study uses a qualitative, desk-based documentary design supported by indicator triangulation and a reproducible open-data KPI baseline.

Materials comprised: (i) SDAIA/NSDAI strategy documents and related public communications; (ii) NDMO data-governance and interoperability regulations and guidance (standardization, classification and controlled sharing); (iii) PDPL and publicly available implementing guidance treated as binding design constraints; (iv) official digital-government maturity indicators (DXMI and Qiyas) and World Bank GTMI benchmarking; (v) publicly documented Tawakkalna materials used as a flagship vignette; and (vi) a reproducible open-data KPI baseline from the Saudi Water Authority complaint-escalation series (monthly received, closed and average closure duration).

Inclusion criteria. Sources were included if they were official or authoritative publications directly relevant to the study constructs and contained explicit, citable statements or measurements suitable for IPO-aligned extraction and mapping.

Time period. The indicator-based triangulation focuses on the most recent available reporting cycles for DXMI and Qiyas and the referenced GTMI edition used for macro benchmarking. The SWA open-data KPI baseline covers the published monthly span in the released series.

Anonymization. No individual-level or personally identifiable information (PII) was collected or processed. The SWA dataset is aggregated at the month level. The study does not ingest citizen records, case texts, or internal government logs. Any prototype description is conceptual and governance-ready by design (e.g., masking/minimization), without using real personal data.

5.2. Procedures:

The study followed a structured, IPO-aligned workflow to ensure transparent linkage between evidence, mechanisms, and claims about governance-ready scaling.

Step 1: Corpus construction and cataloguing. Documentary sources were identified, screened against the inclusion criteria, and catalogued by type (strategy, governance regulation, privacy/compliance instrument, maturity indicator report, platform vignette material, and open-data series). Each source was mapped to the IPO input layer.

Step 2: Structured extraction and coding. Key statements and reported indicators were extracted and coded using a predefined codebook aligned with the paper's constructs: orchestration /coordination, interoperability and data-governance preconditions, compliance-by-design constraints

(PDPL), sovereign Arabic LLM capability signals (ALLaM), and measurable readiness/impact signals.

Step 3: Mechanism extraction (process layer). From the coded materials, agentic mechanisms relevant to high-stakes public services were identified and normalized into a bounded-autonomy mechanism set: planning and task decomposition, controlled tool use (allowlists), controlled memory/state, verification and feedback loops, role-based access control (RBAC), comprehensive audit logging, and human-in-the-loop (HITL) approval gates for high-risk actions.

Step 4: IPO mapping and synthesis. Evidence from the input layer was mapped to the extracted agentic mechanisms and then to expected output-layer outcomes (productivity/throughput, decision quality/consistency, innovation enablement, and auditable compliance).

Step 5: Indicator triangulation. To strengthen empirical grounding without claiming causal impact, readiness evidence was triangulated using official Saudi digital-government maturity indicators (DXMI and Qiyas) and external benchmarking (GTMI). Indicators were interpreted as outcome-layer readiness signals and mapped to the IPO framework to support the governance-ready scaling argument.

Step 6: Case vignette integration. A flagship platform vignette (Tawakkalna) was incorporated to illustrate how governance-aligned, large-scale service orchestration can provide an execution environment for LLM-enabled augmentation and bounded agentic functions under auditable interfaces and standardized workflows.

Step 7: Open-data KPI baseline and prototype specification. The Saudi Water Authority (SWA) open-data series was processed to construct a reproducible KPI baseline (e.g., intake, closures, timeliness, pressure proxies) to operationalize the IPO output layer and define evaluation-ready measures for future bounded-agent deployment.

5.3. Measures and Evaluation:

Evaluation focuses on output-layer readiness and governance-safe measurement. Readiness is triangulated using official Saudi digital-government indicators (DXMI and Qiyas) and external benchmarking (GTMI), mapped explicitly to the IPO output layer. For the technical add-on and prototype, an evaluation-ready KPI set is specified to support future pre/post (or quasi-experimental) assessment, covering: (i) operational performance (throughput/closure rate, timeliness, and pressure proxies), (ii) decision quality (routing accuracy and consistency), and (iii) governance and safety. Detailed metric definitions and computations are provided in the technical add-on and the bounded-agent prototype sections.

5.4. Ethical and Governance Considerations:

This study adopts a governance-first stance suitable for high-stakes public services. It relies solely on publicly available documentary sources, official indicators, and aggregated open data; therefore, no individual-level personal data are collected or processed and no consent procedures are required. Privacy and compliance are treated as binding deployment constraints, aligned with PDPL, NDMO data-governance and interoperability requirements, and relevant SDAIA governance direction. Risk management is operationalized through compliance-by-design controls for bounded agentic execution (data minimization and purpose limitation, RBAC, allowlisted tool use, audit logging, and mandatory HITL gating for high-risk or rights-affecting actions). Given the desk-based design and public aggregated metrics, findings are presented as readiness evidence and evaluation baselines rather than causal estimates of deployed agent impact.

6. Results:

6.1. Input-Layer Findings: Evidence of Regulated, Orchestrated Scaling Beyond Isolated Pilots

The documentary evidence supports the claim that the progress of Saudi Arabia in digital governance is influenced by central coordination (SDAIA), strategic alignment (NSDAI), and regulatory governance instruments such as NDMO governance and PDPL-related requirements. This combination reduces fragmentation across entities and supports diffusion under consistent compliance expectations. Rather than relying on scattered pilots, the national approach favors scaling LLM-enabled service augmentation within a governance-ready ecosystem where platform standards, data governance, and accountability requirements are institutionally embedded.

6.2. Process-Layer Findings: Governance-Constrained (Bounded) Autonomy Enables Agentic Value Under High-Stakes Constraints

Across high-stakes public-sector settings, the strongest and most defensible agentic value arises when autonomy is engineered as governance-constrained (bounded) execution. In this configuration, agentic systems support operational workflows (such as intake triage, summarization, drafting, routing, and low-risk tool actions) through controlled tool-use (allowlists) while preserving human authority over high-impact decisions. As autonomy shifts from advice to execution, governance, privacy, and accountability risks expand; therefore, process-level feasibility depends on enforceable controls such as RBAC, audit logging, verification/feedback loops, and HITL gating for sensitive or rights-affecting cases. The process-layer implication is clear: scalability is driven not only by model capability but by the maturity of controlled access, traceability, and approval workflows that constrain agent behavior deterministically.

6.3. Output-Layer Evidence: Multi-Indicator Triangulation Supports Governance-Ready Scaling of LLM-Enabled Services

To strengthen empirical grounding for the IPO output layer, the study triangulates DXMI and Qiyas with the World Bank GTMI, capturing complementary readiness dimensions: platform service-experience maturity (DXMI), institutional transformation compliance (Qiyas), and macro-level GovTech capability (GTMI). Collectively, these indicators support the viability of scaling LLM-enabled workflows and governance-constrained agentic execution in regulated public services.

Table I. Indicator-Based Readiness Evidence for Regulated Agentic Transformation in Saudi Digital Government.

Indicator (Source)	Year	Reported Value	What it Measures	Relevance to Agentic AI
Digital Experience Maturity Index (DXMI) – Digital Government Authority (DGA)	2025	86.71% (Advanced)	Overall maturity of digital-government platforms and end-to-end service experience (citizen satisfaction, UX, complaint handling, technology readiness).	Indicates readiness for LLM-enabled service workflows in citizen-facing platforms; agentic execution would still require enforceable governance controls (DGA).
DXMI (DGA) – Previous Cycle Comparison	2024	85.04%	Overall DXMI score in the previous cycle (baseline for longitudinal comparison).	Provides a baseline to track changes in service-experience maturity over time in relation to future agent augmentation (DGA).
Digital Transformation Measurement (Qiyas) – Digital Government Authority (DGA)	2024	87.14%	Agency-level compliance with core digital transformation standards and operational controls.	Signals institutional readiness for governed (bounded) agentic execution, including standards compliance and control maturity (DGA).

Qiyas (DGA) – Previous Cycle Comparison	2023	85.53%	Prior cycle Qiyas score (baseline for longitudinal comparison).	Establishes a baseline to track institutional readiness trends relevant to scaling governed agentic capabilities (DGA).
GovTech Maturity Index (GTMI) – World Bank	2025	99.64% GTMI 2025 – World Bank: Saudi Arabia is in Group A (Extensive GovTech Maturity).	National GovTech maturity across core systems, service delivery, citizen engagement, and enablers.	Provides external macro-level benchmarking that corroborates national readiness for scaling AI-enabled public services under governance constraints (World Bank).

Indicator evidence (Source). DXMI (2025) and Qiyas (2024) report advanced maturity/compliance levels, supported by external GTMI benchmarking (DGA, 2025; DGA, 2024; World Bank, 2025)

6.3.1. Longitudinal Readiness Signals (DXMI and Qiyas Trends)

Consecutive measurement cycles are visualized to capture trends in platform maturity and institutional compliance.

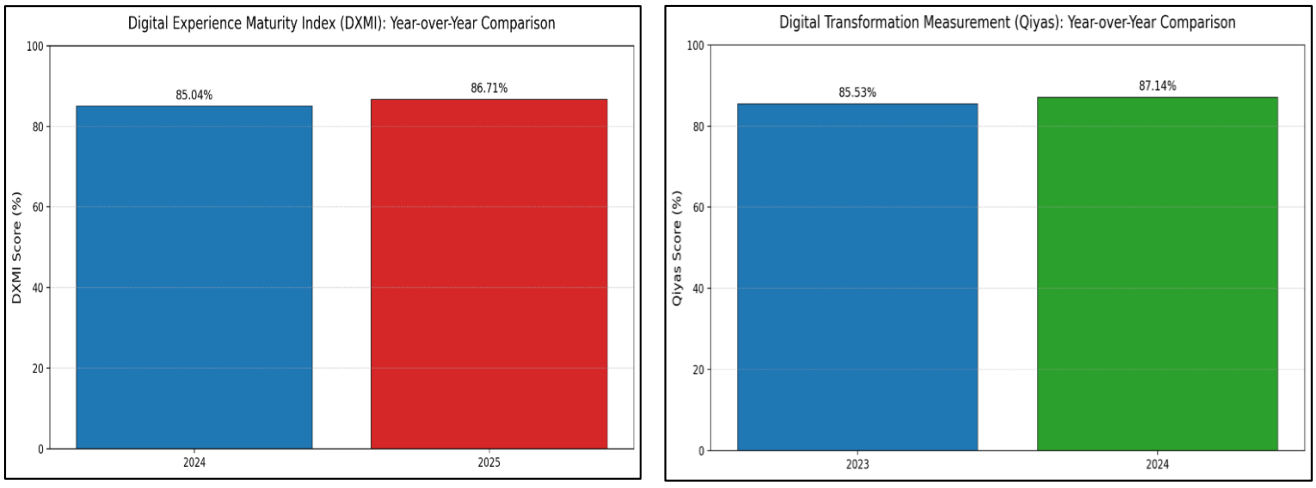

Figure 1. DXMI – Year-over-Year Comparison - Source: Author's own work.

This figure shows year-over-year change in citizen-facing digital service maturity.

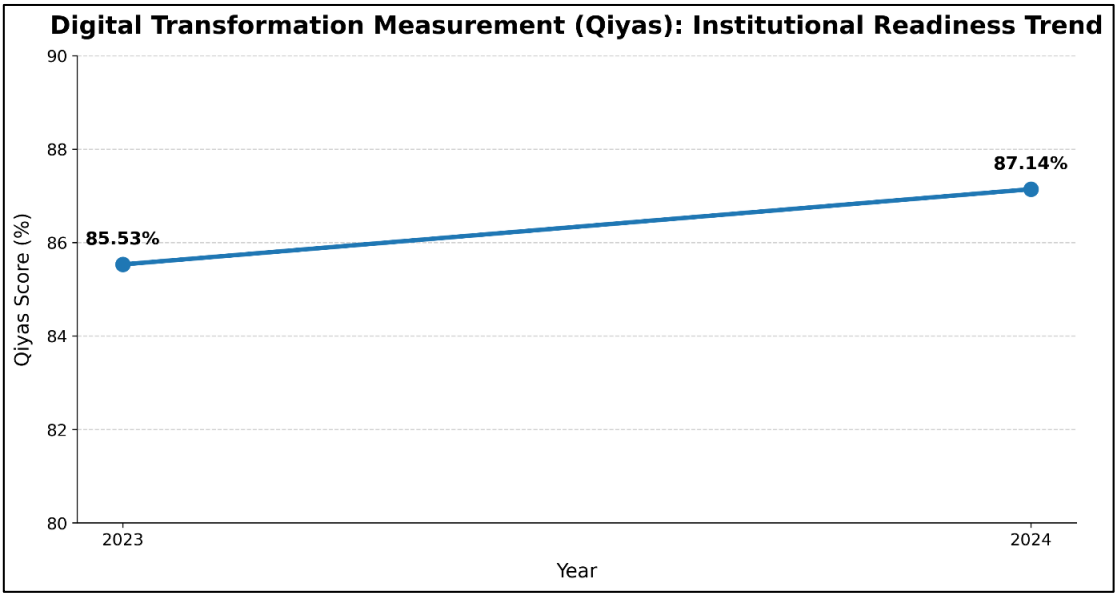

Figure 2. Qiyas – Institutional Readiness Trend. Source: Author's own work.

The above figure Shows year-over-year change in institutional compliance with digital transformation standards (DGA, 2024).

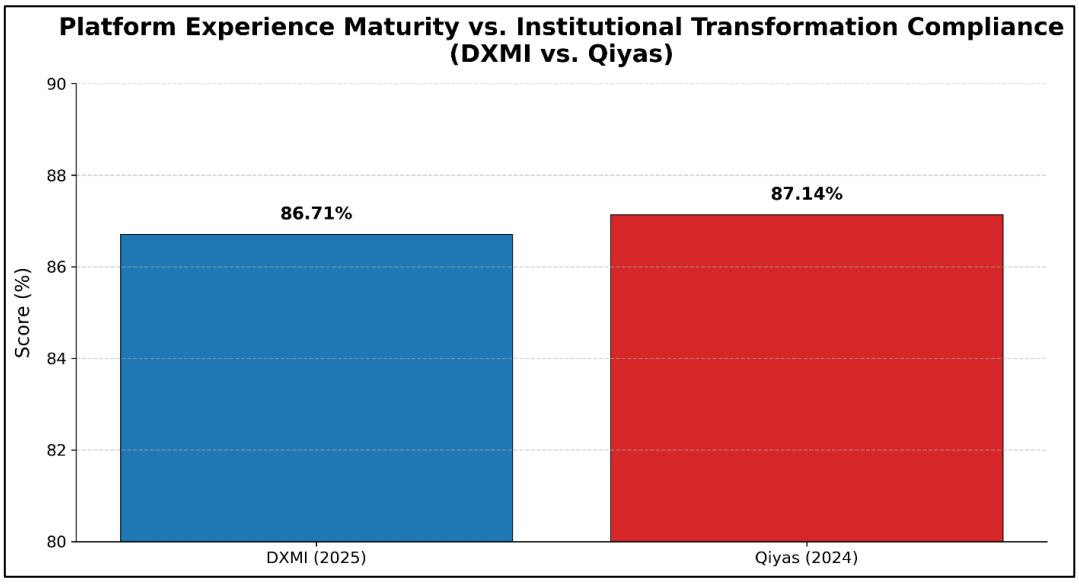

Figure 3. DXMI vs. Qiyas. Source: Author's own work.

This comparative visualization contrasts platform-level experience maturity (DXMI) with agency-level transformation compliance (Qiyas). Together, they provide complementary evidence of readiness for regulated agentic adoption: service quality maturity plus governance capacity (DGA, 2025; DGA, 2024).

6.3.2. IPO Integration: Mapping Indicators to Agentic Readiness

Table II. IPO Mapping of Digital Government Indicators to Agentic AI Readiness

INPUTS (National Drivers)	PROCESSES (Agentic Mechanisms)	OUTPUTS (Readiness Evidence + Outcomes)
• Vision 2030 + NSDAI strategic direction • Regulatory guardrails: PDPL + NDMO data governance • Digital-government platforms & infrastructure • Sovereign capability building (Arabic-first LLM readiness)	• Planning & task decomposition • Tool-use (bounded API actions via allowlists) • Verification & feedback loops • Memory/state (controlled continuity) • Human-in-the-loop (HITL) gating • Audit logging + access control (RBAC)	Outcome-layer readiness indicators: • DXMI (service-experience maturity) • Qiyas (institutional transformation compliance) • GTMI (macro GovTech benchmark) Enables governed scaling of: • LLM-enabled language workflows (summarization, drafting, routing) • Bounded agentic execution (traceable tool-use under governance) Expected measurable outcomes: • Productivity + throughput • Decision quality + consistency • Auditable compliance + standardized delivery

Source: Author's own work.

Table 2 maps DXMI/Qiyas/GTMI evidence to the IPO model, positioning DXMI and Qiyas as outcome-layer readiness signals that support governed agentic mechanisms under enforceable national constraints.

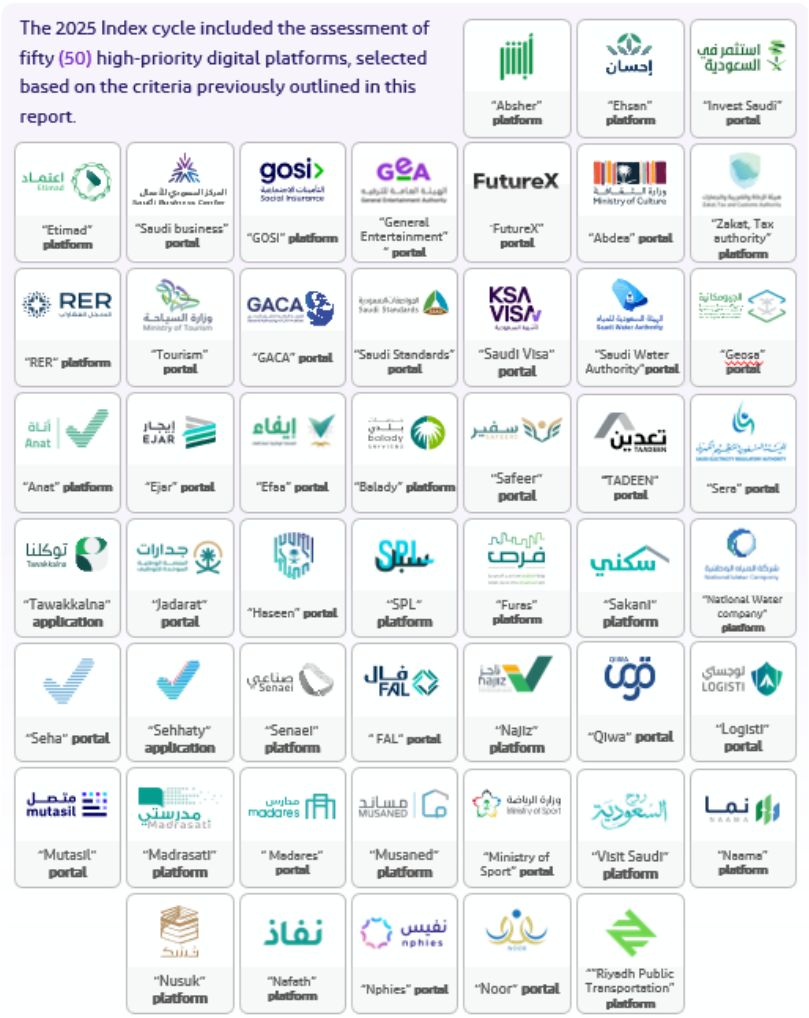

Figure 4. DXMI 2025 Portfolio Scope (50 High-Priority Digital Platforms)

Source: Digital Government Authority (DGA), Digital Experience Maturity Index for Government Services (2025) – General Report.

This figure highlights the breadth of the DXMI 2025 assessment and strengthens interpretation of DXMI as a national indicator of digital service maturity. The portfolio-level coverage implies the presence of scalable operational ecosystems with standardized workflows and controlled integration surfaces, as conditions supportive of governance-constrained agentic functions such as triage, routing, and policy-aligned response generation under auditable public-sector constraints (DGA, 2025).

6.4. Case Vignette: Tawakkalna as a Flagship Platform Demonstrating Governance-Aligned Readiness for Scaling LLM-Enabled Public Services and Bounded Agentic.

To complement indicator-based readiness evidence, the study introduces Tawakkalna as a flagship vignette illustrating Saudi Arabia's demonstrated capacity to deliver large-scale, citizen-centric digital

services under high-trust operating conditions. Although Tawakkalna is not an agentic AI system per se, it represents a governance-ready execution environment in which LLM-enabled service augmentation and governance-constrained agentic functions can be integrated through standardized workflows, controlled integration points, and auditable interfaces. As a unified channel for cross-agency service delivery, Tawakkalna exemplifies the shift from service digitization toward service orchestration at scale, reinforcing dimensions captured by maturity indicators (experience consistency, responsiveness, standardization).

We introduce a Saudi open-data KPI baseline drawn from a complaint-escalation workflow. This add-on operationalizes the IPO model's output layer using measurable service-performance signals (demand, throughput, backlog pressure, and timeliness) and illustrates how such signals can be used to target and evaluate governance-constrained (bounded) agent augmentation under enforceable governance controls.

7. Technical Add-on: Open-Data KPI Analysis for Complaint Escalation (Saudi Water Authority)

7.1. Purpose and Positioning

This technical add-on provides a lightweight, evaluation-ready illustration of the IPO model's output layer using Saudi open government data on a complaint-escalation workflow. We quantify monthly service-pressure and timeliness signals and demonstrate how these KPIs can (i) identify pressure regimes in which a bounded AI agent is most likely to yield operational value under enforceable governance constraints, and (ii) define measurable outcome metrics for future impact evaluation. Importantly, this analysis is descriptive and does not estimate causal effects of deployed agents. Instead, it establishes a reproducible measurement baseline and a KPI set that can be reused in pre/post evaluations (or quasi-experimental designs) once bounded-agent support is implemented.

7.2. Data Source and Scope

We retrieved the dataset from the official Saudi open-data ecosystem (Saudi Water Authority open-data library / open-data center), titled Complaint Escalation Platform (CSV), in CSV format under portal's open-data license terms. The dataset is published as monthly aggregates spanning January–September (9 months) and contains no personally identifiable information (PII) in the released series. (SWA, 2025).

The dataset comprises four operational fields:

- Month

- Received complaints, R_t
- Closed complaints, C_t
- Average closure duration (days), D_t

This structure supports monitoring of (i) demand intake (R_t), (ii) throughput (C_t), and (iii) responsiveness (D_t) in an escalation setting. A key limitation is that monthly aggregation does not allow case-level queue reconstruction; therefore, backlog-like measures are treated strictly as pressure proxies (Section 7.5), not as true backlog estimates.

7.4. Data Preparation and Quality Checks

Given the pre-existing organization of the data at a monthly aggregate level, no further temporal aggregation was required. The following data preprocessing steps were taken:

1. **Chronological normalization:** standardized month labels to a numeric month index to ensure correct time ordering.
2. **Numeric parsing:** extracted numeric values for closure duration from strings where needed (e.g., 8 days → 8).
3. **Type validation:** verified that $R_t$and $C_t$are non-negative integers and that $D_t$is a positive numeric value.
4. **Consistency checks:** flagged any month where $C_t > R_t$for manual review (possible if closed includes cases received in prior months). No records were removed; all computations are reproducible from the published series.

7.5. KPI Construction:

Let $R_t$denote received complaints in month t, $C_t$closed complaints in month t, and $D_t$average closure duration (days) in month t.

(1) Throughput ratio (closure rate)

$$TR_t = \frac{C_t}{R_t} \times 100$$

(2) Net flow (pressure proxy)

$$NF_t = R_t - C_t$$

Given monthly aggregation, $NF_t$is interpreted as a pressure proxy rather than a true backlog measure.

(3) Period-level throughput

$$TR_{\text{overall}} = \frac{\sum_{t=1}^{T} C_t}{\sum_{t=1}^{T} R_t} \times 100$$

(4) Month-over-month change in demand

$$\Delta R_t = R_t - R_{t-1}, \Delta R_t^{\%} = \frac{R_t - R_{t-1}}{R_{t-1}} \times 100$$

Short-series diagnostics ($N = 9$):

Given the short monthly series, we do not emphasize anomaly detection. We focus on descriptive plots and (optionally) non-parametric association checks (e.g., Spearman correlation) between timeliness $D_t$and pressure signals (R_t, NF_t).

7.6. Results and Interpretation

Across January–September (9 months), the dataset reports 10,071 received complaints and 9,822 closed complaints, yielding an overall throughput ratio $TR_{\text{overall}}$of 97.53%.

Monthly demand increased over the period (range: 578 in January to 1,497 in August). Pressure signals are visible in the net flow proxy NF_t: early months show near alignment between intake and closures ($TR_t \approx 100\%$), while later months exhibit widening divergence. The month with the peak in net flow is September, with a total of 146 cases, while it is also the month with the minimum throughput ratio at 90.18%. This condition represents a late-period pressure regime, which means that throughput is less than intake. The timeliness varies on a month-by-month basis, with a mean duration of closure between 7 days in May and 15 days in July. Notably, July combines high intake with the longest closure duration, while September combines sustained high intake with the largest net flow proxy suggesting that triage quality, routing efficiency, and workload balancing become more critical under elevated demand.

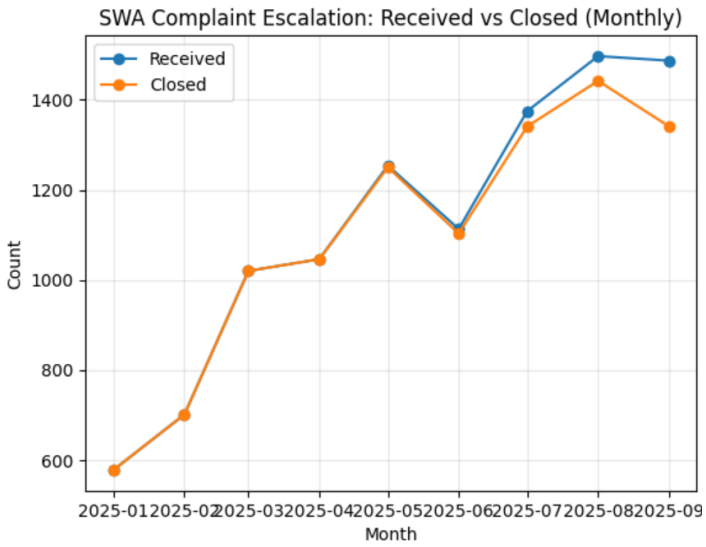

Figure 5. Received vs. Closed (Monthly). Source: Author's own work.

The figure compares monthly received and closed complaints. Close alignment indicates operational throughput stability, whereas divergence toward the end of the period signals emerging service-pressure conditions.

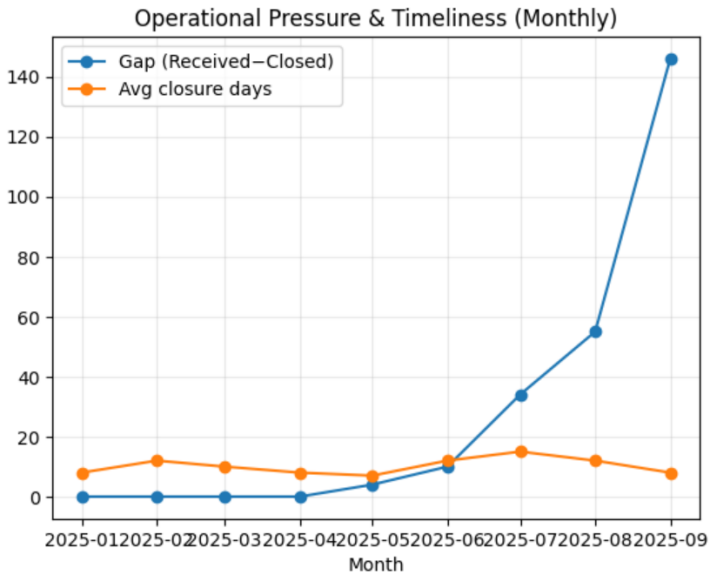

Figure 6. Operational Pressure and Timeliness (Monthly). Source: Author's own work.

The figure overlays the net flow proxy $NF_t = R_t - C_t$ and average closure duration D_t. Simultaneous increases suggest operational stress regimes where triage quality and routing efficiency become more critical.

7.7. How KPIs Map to Bounded-Agent Value and Evaluation

In a regulated deployment, a bounded agent would:

- standardize intake summarization and classification,
- execute allowlisted routing and ticketing actions under RBAC,
- trigger HITL gating for high-risk categories,
- maintain auditability and PDPL-aligned minimization.

Operational impact would be evaluated using the same KPI family: reductions in NF_t, improvement in TR_t, and reductions in D_t, complemented by governance compliance indicators (audit completeness, HITL rates for high-risk cases, and access-control enforcement logs).

7.8. Bounded-Agent Impact: Evaluation Design and Effect Size Targets

To assess whether bounded-agent support improves operational performance without weakening governance controls, we propose an outcomes-and-safeguards evaluation framework aligned with the

KPIs defined in Section 7.5. The primary operational outcomes are pressure reduction and timeliness improvement, measured by decreases in NF_t, increases in TR_t, and reductions in D_t. Because the present analysis is based on aggregated monthly open data and does not observe agent deployment, impact assessment should be conducted once the bounded agent is implemented using a pre-specified design that supports credible inference.

If such a rollout is possible, a stepped-wedge design or difference-in-differences can be used to compare treated units with bounded agent triage/routing enabled to control units with business-as-usual operation. In the event of having only one deployment site, an interrupted time series can be used to evaluate changes in level and trend of TR_t, NF_t, and D_t pre- and post-deployment. Evaluation should jointly report governance and safety metrics—audit completeness, HITL compliance for high-risk categories, RBAC enforcement, and PDPL-aligned minimization—to ensure efficiency gains do not come at the cost of accountability or privacy.

8. Interpreting the Saudi Pathway as Regulated Agentic Transformation

The findings support a Saudi pathway where scaling is shaped by the co-evolution of capability-building and enforceable governance. Advanced service-experience maturity provides an execution environment for high-volume language workflows (e.g., summarization, routing, standardized drafting, multilingual guidance). Meanwhile, institutional transformation compliance indicates the presence of standardized digital practices required for auditable tool-use, controlled access, and accountable human gating. This suggests that regulated agentic transformation is not only a technology trajectory but a governance trajectory: autonomy becomes feasible when compliance-by-design and institutional orchestration constrain behavior deterministically through access rules, auditability, and approval workflows.

9. Practical Application: A Governance-Ready Bounded Agent Prototype for Citizen Case Triage and Routing

9.1. Use Case Definition

One of the most common high-impact workflows that can be found in various public sector organizations includes citizen request intake, triage, summarization, routing, prioritization, and the composition of responses that align with the policies of the government. Delays are usually introduced if the process is done manually; thus, a bounded agent can aid the workforce in standardizing the process while speeding it up within the government's rules and regulations.

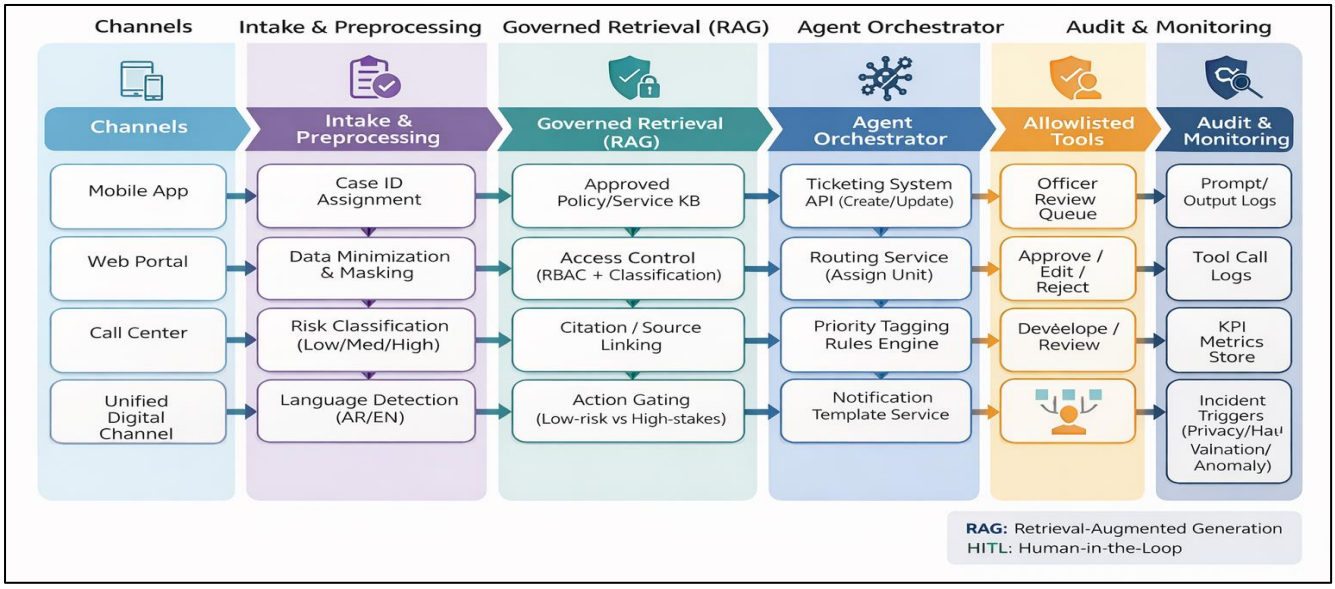

Figure7. Governance-ready bounded agent architecture for citizen case triage and routing. Source: Author's own work.

This description outlines the end-to-end system components that facilitate a controlled integration of large language models and agents, which includes governed retrieval augmented generation, tool usage, human-in-the-loop gating, and comprehensive audit and monitoring instrumentation.

9.2. End-to-End Agentic Workflow (Bounded Execution)

The prototype operationalizes bounded autonomy through a controlled sequence of steps:

- Intake & minimization: Ingest the request; mask unnecessary identifiers; assign a case ID; classify sensitivity.
- Governed retrieval (RAG): Retrieve policy content from an approved knowledge base with access control by case type/classification.
- Plan generation: Produce a structured plan (task decomposition); store the trace for auditability.
- Tool use: Execute allowlisted actions (create ticket, attach summary, route unit, assign priority tags).
- HITL gating: Require approval for high-risk cases and any actions affecting rights/enforcement outcomes.
- Feedback loop: Capture correction reasons, resolution time, and satisfaction; update templates/rules under governed change management.

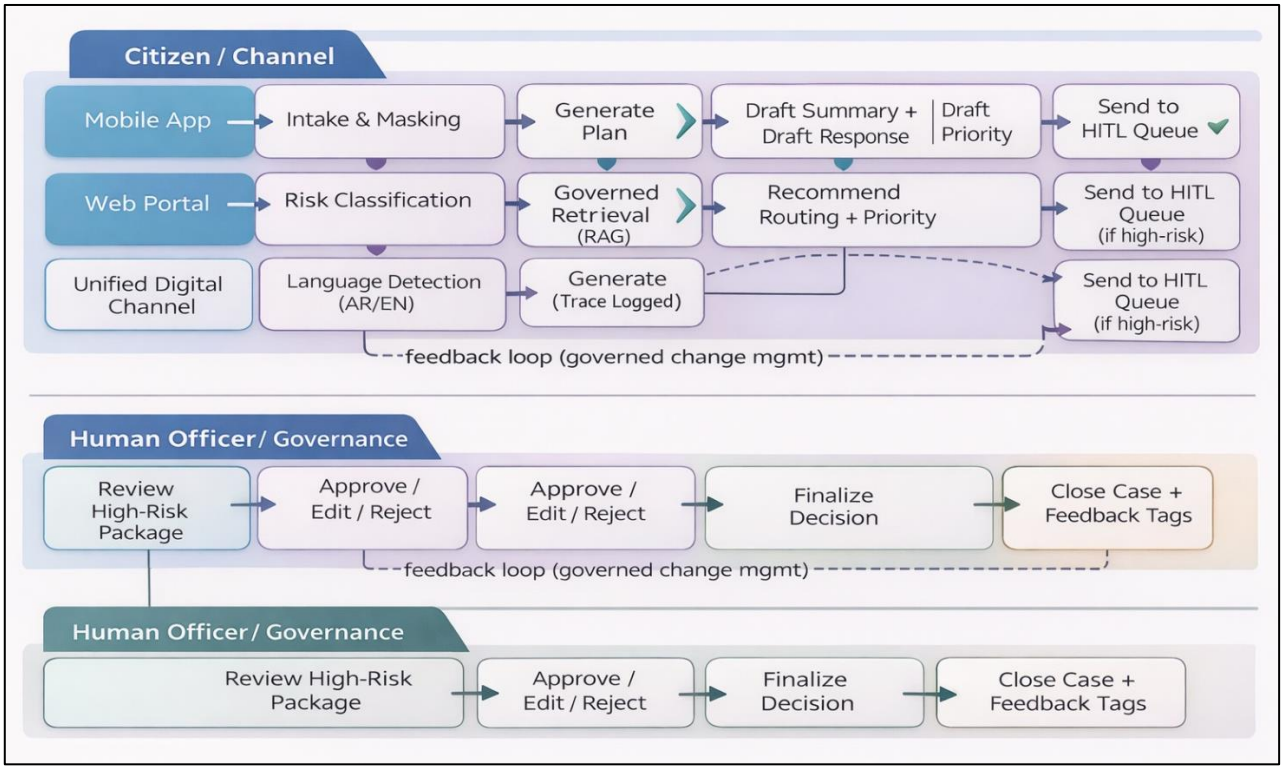

Figure 8. indicates that autonomy is limited by risk classification, deterministic policy checks, and approval workflows. Source: Author's own work.

9.3. Risk-Based Decisioning and Autonomy Constraints

To prevent unsafe automation, the bounded agent applies risk-based decisioning: low-risk cases can proceed with allowlisted tool actions, while high-risk cases are automatically escalated to human review with a recommendation package.

Figure 9. Risk-based decision flow for bounded autonomy (low-risk auto-actions vs. high-risk HITL escalation). Source: Author's own work.

Figure 9 shows how risk classification determines whether the agent executes allowlisted actions or routes the case to the HITL queue. As autonomy increases from generation to tool-executing agency, risks expand across three domains. The risks that governance faces include opaque decision paths, unregulated tool actions, lack of accountability, and traceability issues. The risks that privacy faces include overcollection, over retention, leakage between cases through memory or state, and unauthorized access during retrieval and tool usage. The risks that ethics and safety face include unfair routing or prioritization, unfair treatment in high-stakes situations, rationales that hallucinate and drive action, and overreach by automation itself. In the Saudi context, these risks are best mitigated through enforceable controls aligned with PDPL and national data-governance instruments: PDPL-aligned data minimization and masking, strict purpose limitation, role-based access control (RBAC), allowlisted tool-use with deterministic policy checks, comprehensive audit logging (prompts, retrieval, tool calls, and outputs), and mandatory human-in-the-loop (HITL) gating for high-risk cases. Together, these controls operationalize regulated agentic transformation by enabling bounded execution, traceability, and accountable oversight while supporting scalable modernization to position adoption scope under enforceable governance, we define four Agentic Readiness Levels (ARL) that map increasing autonomy to required controls. The levels are not a capability ladder only; they are a governance ladder advancing only when auditability, access control, and risk-gated oversight are demonstrably in place.

- **ARL-1: Assistive LLM (generation-only).**

Function: drafting/summarization/classification with no tool execution.

Controls: PDPL-aligned minimization, content safeguards, logging of prompts/outputs.

Success signal: reduced drafting time; high acceptance rate of summaries.

- **ARL-2: Tool-augmented under strict permissions.**

Function: LLM produces outputs + invokes read-only or low-risk tools via allowlists (e.g., retrieve policy passages, prefill templates).

Controls: RBAC + allowlisted tools + retrieval governance (approved KB/RAG) + audit logs.

Success signal: improved consistency and faster retrieval-to-draft cycle without policy violations.

- **ARL-3: Governance-constrained (bounded) workflow agent with HITL approvals.**

Function: executes bounded actions (e.g., create ticket, attach summary, route to unit) with risk-based escalation and mandatory HITL for high-stakes steps.

Controls: risk classification, HITL gating, full trace of tool calls/decisions, change management for prompts/rules.

Success signal: improved closure rate, reduced backlog gap, reduced average closure days—while maintaining audit completeness and zero unauthorized actions.

- **ARL-4: Semi-autonomous with continuous oversight and audits.**

Function: broader autonomous sequencing across steps with continuous monitoring and periodic audits.

Controls: continuous evaluation, anomaly monitoring, independent audit review, strict incident response, and governance sign-off for scope expansion.

Success signal: sustained KPI improvement at scale with stable compliance performance and low incident rates.

Gate criteria for progression. The transition from ARL-1 to ARL-4 requires evidence of: (i) boundedness of tool actions, ensuring enforceability; (ii) access control using RBAC with sensitivity-awareness; (iii) end-to-end traceability of all actions using audit logs; and (iv) HITL, ensuring high-impact decisions are human-authorized, in conformance with PDPL and national data governance regulations.

9.4. Practical Governance Matrix (Controls-by-Component)

This section specifies enforceable controls aligned with each workflow component to ensure compliance-by-design.

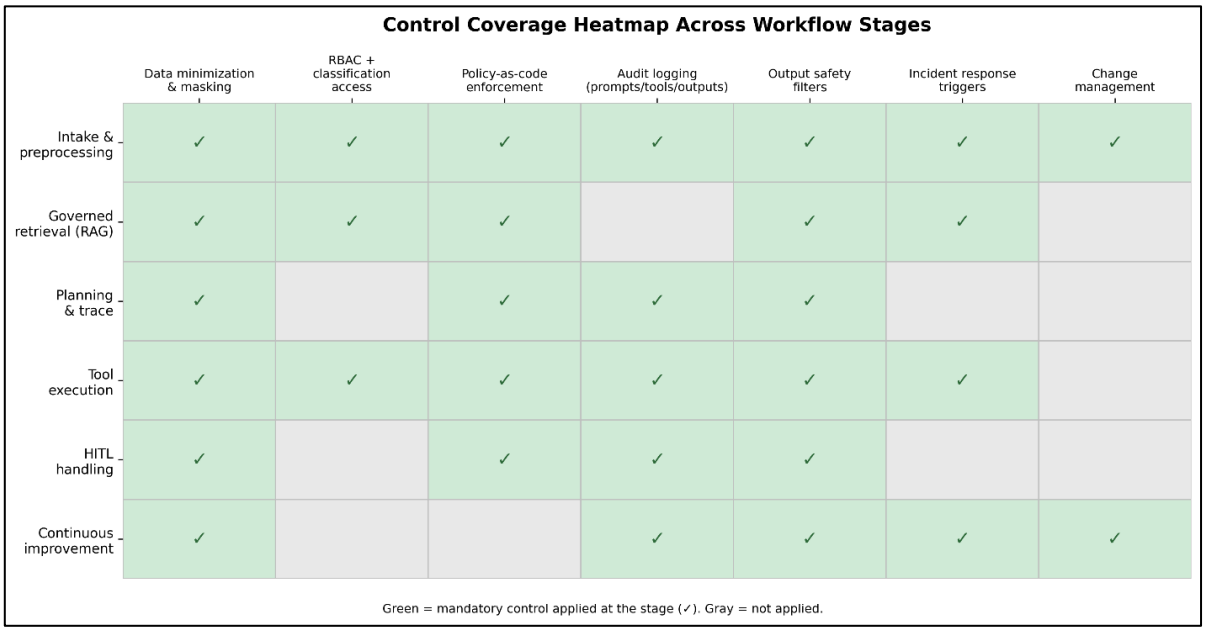

Figure 10 Practical governance matrix for the bounded agent workflow (capability–risk–control mapping). Source: Author's own work.

9.5. Evaluation Plan (IPO-Aligned KPIs)

To support evidence-based scaling, the prototype includes an evaluation plan aligned with the IPO outcome layer:

- Productivity: processing time per case; cases handled per staff/day; % auto-summaries accepted without edits.
- Quality: routing accuracy; staff-rated summary fidelity; hallucination rate.
- Governance: audit completeness; privacy leakage incidents; unauthorized access attempts blocked.
- User experience: citizen satisfaction score; resolution time reduction; repeat-contact rate.

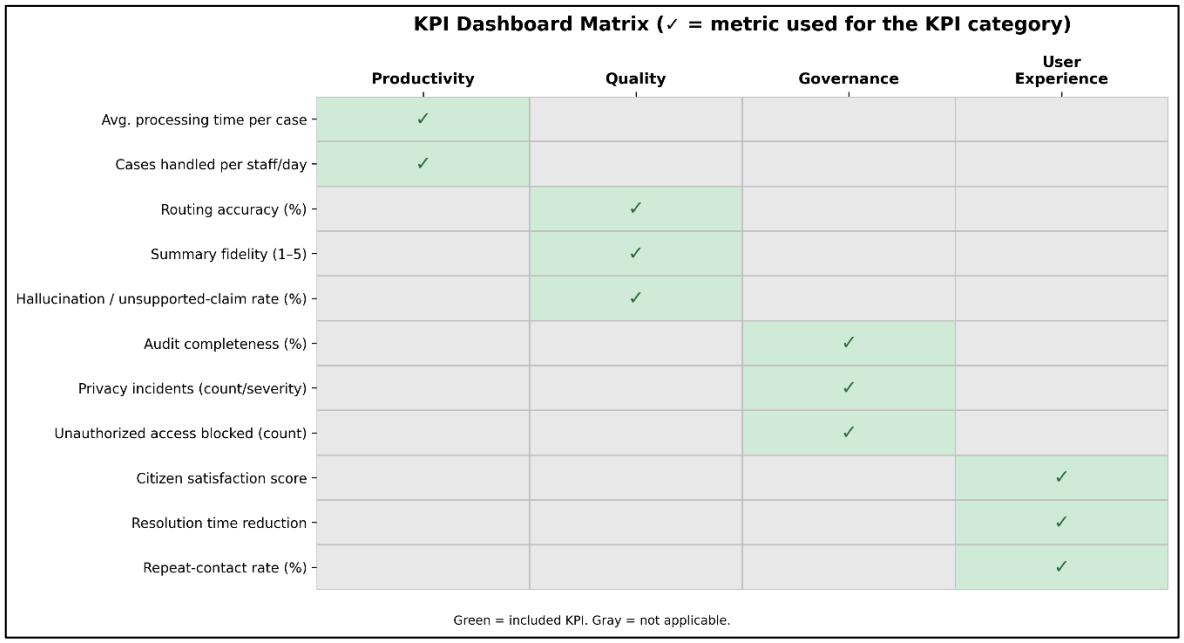

Figure 11. IPO-aligned KPI dashboard for evaluating the bounded agent prototype. Source: Author's own work.

Figure 11 shows the measures of productivity, quality, governance, and user experience to facilitate evidence-based scaling.

10. Discussion:

The analysis supports a Saudi pathway of regulated agentic transformation in which scaling LLM-enabled and agentic capabilities is shaped primarily by enforceable governance and institutional orchestration rather than model capability alone. Within the IPO framing, national inputs—Vision 2030/NSDAI strategic direction, SDAIA-led coordination, and binding governance instruments (NDMO and PDPL)—define the governance-ready envelope that conditions what systems may access,

what they may execute, and how accountability is maintained at scale. At the process layer, the most defensible route in high-stakes public services is bounded autonomy: LLMs function as a scalable language layer for drafting, summarization, and intake triage, while agentic value becomes feasible when planning and tool use are constrained by deterministic controls (allowlisted actions, controlled state/memory, verification loops, RBAC, comprehensive audit logging, and HITL gating for high-risk decisions). Readiness is anchored in triangulated signals rather than asserted abstractly: official maturity reporting provides quantitative anchors for the output layer (e.g., DXMI at 86.71% and Qiyas at 87.14%), complemented by macro-level benchmarking (GTMI) and a flagship vignette (Tawakkalna) illustrating a platform-scale execution environment with standardized workflows and auditable integration surfaces (see Section 6.3 for full indicator mapping).

Practically, these results imply that safe scaling should prioritize bounded agent augmentation in high-volume, language-intensive workflows while treating compliance-by-design and measurement as prerequisites: PDPL-aligned minimization and purpose limitation, RBAC, allowlisted tool use, end-to-end auditability, and mandatory HITL gates for rights-affecting cases. The paper contributes (i) a governance-first conceptualization of regulated agentic transformation, (ii) an IPO model linking drivers, guardrails, mechanisms, and measurable outcomes, and (iii) evaluation-ready measurement through indicator triangulation, an open-data KPI baseline, and a governance-ready bounded-agent prototype, while acknowledging that the desk-based design supports readiness evidence rather than causal estimates of deployed agent impact.

10.1. Final Contribution Statements

This study provides a policy-grounded and empirically informed synthesis of how generative and agentic AI particularly large language models may reshape digital governance and public-service delivery in Saudi Arabia, aligned with Vision 2030 and SDAIA-led transformation. It proposes and applies an Input–Process–Output (IPO) framework linking national drivers to bounded agentic mechanisms and measurable outcomes. The paper further strengthens feasibility claims through official maturity indicators and flagship implementation evidence, showing advanced service-experience maturity and institutional transformation readiness that support scaling LLM-enabled service augmentation and governance-constrained agentic execution under regulated constraints.

11. Implications:

The paper has four three implications which are:

- **Implications for practice:** Digital-government organizations should prioritize bounded-agent deployments in high-volume, language-intensive workflows (intake triage, routing, summarization,

and reporting) where measurable gains are plausible without delegating rights-affecting decisions to automation.

- **Implications for policy/governance:** Scaling should be treated as compliance-by-design: PDPL- and NDMO-aligned constraints must be translated into enforceable technical and organizational controls, including data minimization, RBAC, allowlisted tool use, end-to-end audit logging, and mandatory HITL gates for high-risk actions.
- **Implications for capability and accountability:** Agencies should define accountable roles, training, and escalation procedures for HITL decisions and audit review, ensuring that responsibility remains traceable and contestability is preserved.
- **Implications for research and evaluation:** Indicator triangulation and open-data KPI baselines should be complemented by deployment-stage empirical designs that jointly report operational outcomes (timeliness, throughput) and governance outcomes (audit completeness, HITL compliance, unsupported-claim/hallucination rate).

12. Limitations and Future Work

Limitations. This study relies on documentary evidence and publicly available outcome indicators rather than operational metrics; therefore, it does not estimate causal effects of deployed agents on outcomes. Publicly available information may not capture all implementation details (e.g., control flow configurations or incident tracking). The open-data KPI analysis is descriptive and is intended to provide an operational baseline for measurement and design targeting rather than to claim realized effects from an implemented intervention.

Future work. Future research should validate the proposed framework through bounded-agent deployments in specific high-volume public-service workflows, using rigorous evaluation designs (e.g., interrupted time series, stepped-wedge rollouts, or quasi-experiments). Evaluation should jointly report operational outcomes (throughput, timeliness, pressure signals) and governance outcomes (evidence-linking rate, unsupported-claim/hallucination rate, audit completeness, RBAC enforcement, and HITL compliance), with particular attention to bilingual (Arabic/English) service settings and PDPL/NDMO-aligned compliance-by-design controls.

13. Conclusion

This paper argues that LLMs and agentic AI can reshape Saudi digital government through language automation, workflow orchestration, and governance-aware decision support, with scalability conditioned by enforceable governance and institutional orchestration. Official maturity indicators

(DXMI/Qiyas) and macro benchmarking (GTMI) provide convergent evidence that Saudi Arabia is progressing from digitization to digital maturity and is increasingly governance-ready for scaling LLM-enabled services and bounded agentic execution under regulated constraints aligned with Vision 2030. The proposed IPO model and its prototype demonstrate the viability of pursuing productivity, decision-quality, and service-modernization benefits without compromising privacy, auditability, and accountable human oversight in alignment with PDPL. Overall, the Saudi case represents an instance of regulated agentic transformation, where sustainable adoption is achieved through the simultaneous maturation of AI capability, digital-government readiness, and enforceable governance.

14. References:

Abou Ali, M., & Dornaika, F. (2025). Agentic AI: A comprehensive survey of architectures, applications and future directions. *Artificial Intelligence Review*. https://doi.org/10.1007/s10462-025-11422-4

Bandi, A. (2025). The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges. *Future Internet, 17*(9), 404. https://doi.org/10.3390/fi17090404

Brynjolfsson, E., Li, D., & Raymond, L. (2024). *Generative AI at work*. arXiv. https://arxiv.org/abs/2304.11771

Digital Government Authority. (2024). *Digital transformation measurement 2024 (Qiyas 2024)*. https://dga.gov.sa/sites/default/files/2025-05/Digital%20Transformation%20Measurement%202024.pdf

Digital Government Authority. (2025). *Digital experience maturity index for government services (2025)* (Version 1.0). https://dga.gov.sa/sites/default/files/2025-09/Digital%20Experience%20Maturity%20Index%20for%20Government%20Services%20%282025%29-V1.0.pdf

Filippucci, F., Gal, P., Jona-Lasinio, C., Leandro, A., & Nicoletti, G. (2024). *The impact of artificial intelligence on productivity, distribution and growth: Key mechanisms, initial evidence and policy challenges*. OECD. https://www.oecd.org/content/dam/oecd/en/publications/reports/2024/04/the-impact-of-artificial-intelligence-on-productivity-distribution-and-growth_d54e2842/8d900037-en.pdf

IBM. (2024). *SDAIA's Arabic large language model "ALLaM" is now fully operational on IBM watsonx*. https://mea.newsroom.ibm.com/sdaia-launches-allam-on-watsonx

National Data Bank. (n.d.). *National Data Bank: Home*. https://data.gov.sa/en

National Data Lake. (n.d.). *National Data Lake: Home*. https://lake.data.gov.sa/en

National Data Management Office. (2020). *National data governance interim regulations* (Version 1). https://openresearch-repository.anu.edu.au/bitstreams/c3146c87-743a-48b6-aed9-c99411867c14/download

National Data Management Office. (2023). *The implementing regulation of the Personal Data Protection Law*. National Competitiveness Center. https://istitlaa.ncc.gov.sa/en/Transportation/NDMO/IMPLEMENTINGPDPL/Pages/default.aspx

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf

Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). *AI agents vs. agentic AI: A conceptual taxonomy, applications and challenge*. arXiv. https://arxiv.org/abs/2505.10468

Saudi Data & AI Authority. (2020). *Saudi Arabia: National strategy for data and AI (NSDAI): Realizing our best tomorrow*. https://wp.oecd.ai/app/uploads/2021/12/Saudi_Arabia_National_Strategy_for_Data_and_AI_2020.pdf

Saudi Data & AI Authority. (2023a). *Personal Data Protection Law (PDPL)* (English version 2). https://sdaia.gov.sa/en/SDAIA/about/Documents/Personal%20Data%20English%20V2-23April2023-%20Reviewed-.pdf

Saudi Data & AI Authority. (2023b). *Guide to the Saudi Personal Data Protection Law* (Version 1.0). https://dgp.sdaia.gov.sa/wps/wcm/connect/85a664fd-6e4e-4edb-bab4-1cbdb02ba6d5/Guide%2Bto%2Bthe%2Bsaudi%2BPDP%2Blaw%2BHANDBOOK__%D9%90EN.pdf

Saudi Data & AI Authority. (2024). *Regulation on personal data transfer outside the Kingdom* (Version 2.0). https://dgp.sdaia.gov.sa/wps/wcm/connect/e5bbede0-1119-4f70-b4ef-f043ce58d780/Regulation%2Bon%2bPersonal%2bData%2bTransfer%2bOutside%2bthe%2bKingdom..pdf

Saudi Data & AI Authority, & National Data Management Office. (2020). *National data governance interim regulations*. https://sdaia.gov.sa/ndmo/Files/PoliciesEn.pdf

SaudiPedia. (n.d.). *Tawakkalna application*. https://saudipedia.com/en/tawakkalna-application%E2%80%8E

Saudi Water Authority. (2025). *Complaint escalation platform* [Dataset]. Open Data Library. https://www.swa.gov.sa/ar/open-data-library

World Bank. (2025). *GovTech Maturity Index (GTMI)*. https://www.worldbank.org/en/programs/govtech/gtmi

Doi: https://doi.org/10.52133/ijrsp.v7.80.2