INSEAD

YANGO GROUP YANGO TECH

AI AS PUBLIC INFRASTRUCTURE:

Lessons from the UAE for Government Transformation

AUTHORS

Peter Zemsky, Pascale Balze, Victor Butenko

INSEAD, YANGO TECH (2026):

AI as Public Infrastructure: Lessons from the UAE for Government Transformation, Dubai, the United Arab Emirates.

DISCLAIMER:

No representation or warranty, either express or implied, is provided in relation to the information contained herein and with regard to its fitness, sufficiency, or applicability for any particular purpose.

The information contained in this report is provided for personal non-commercial use and information purposes only and may be reviewed and revised based on new information and data.

INSEAD and Yango Tech disclaim all liability relating to the content and use of the report and the information contained therein, and the report should not be used as a basis for any decision that may affect the business and financial interests of the reader or any other party.

The index's methodology and the rankings do not necessarily present the views of INSEAD and Yango Tech. The same applies to the substantive chapters in this report, which are the responsibility of the authors.

The information contained herein is proprietary in nature and no part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise without the prior per mission of INSEAD and Yango Tech.

EDITED

by Đurđica Gordić

PUBLISHED

by INSEAD, Fontainebleau, France; Yango Tech, Dubai, United Arab Emirates.

© 2026 by INSEAD and Yango Tech

TABLE OF CONTENTS

PREFACE 5

EXECUTIVE SUMMARY

From pilots to institutionalisation 7
Why the UAE case matters 7
A governance model of coordinated differentiation 7
From pilots to systems: execution, procurement, and evaluation 8
Implications for global policymakers 8
Conclusion 8

PART 1. INTRODUCTION & RESEARCH FRAMING

Research purpose and central statement 9
Methodology and scope 10
Analytical positioning and contribution 10
What this paper does not do 10

PART 2: GLOBAL BASELINE

The global overview 11
Strategic vision and leadership 12
Use cases: selection & prioritisation 13
Execution challenges and success factors 14
Procurement and ecosystem strategies 16
Evaluation, impact, and unintended consequences 17
What the literature establishes, and where it still struggles 18

PART 3: THE UAE EXPERIMENT

The uae case: context, method, and analytical lens 20
Strategic vision and leadership 21
Selection and prioritisation of ai use cases 25
Execution challenges and success factors 30
Procurement process & ecosystem strategy 36
Evaluation, impact, and learning loops 37
Lessons from the uae for global policymakers 38
Closing reflection 40

ACKNOWLEDGEMENTS 41

PROJECT TEAM 43

PREFACE: "FROM AI AMBITION TO DELIVERY"

A conversation that begins at an AI conference in Dubai's Museum of the Future is supposed to be futuristic.

Ours became practical.

When I met Peter Zemsky there a year ago, we kept circling one question: if AI is already so capable, where is it creating actual value?

In government, the answer is harder to see. Institutions move through budgets, rules, sovereignty choices, and public trust. A pilot can look alive on stage; a public system must actually perform later, with incomplete information, citizens, and someone accountable for the answer. Financial ROI has its place, but it doesn't capture the value of a faster service or a citizen who no longer has to fight the process.

Over the following months, with Peter, Pascale Balze, Daria Kishinskaia, and UAE AI leaders who spoke with rare openness, we carried that question into interviews across the country. They described the work after the demo, from getting to the data to changing processes enough for systems to scale.

The UAE made the question sharper. A country that appointed the world's first AI minister, with Abu Dhabi now moving toward AI-native government, can't afford for AI to remain a showcase. It has to become a service, a decision, a better Tuesday morning for someone dealing with the state.

As you read this analysis, I hope you can feel the human work behind it. AI in government is built by people: public servants carrying responsibility, technologists making systems usable, and researchers asking better questions. This paper began with curiosity and continued because people cared enough to share what worked, what didn't, and what they were still learning.

With gratitude to everyone who shared their experience and helped shape this research. And with warm memory of Sergey Loiter, whose energy, judgment, and humanity were part of this journey and will always stay with us.

VICTOR BUTENKO

INSEAD MBA '23J
Co-author, INSEAD x Yango Tech AI in Public Sector White Paper

“ AI SYSTEMS’ DEVELOPMENT CYCLES CAN TAKE 3 MONTHS TO CONTRACT AND 3 TO BUILD, **MAKING THE TRADITIONAL PROCUREMENT THE MAIN BOTTLENECK** ”

— Opinions of AI executives

EXECUTIVE SUMMARY

Artificial intelligence is no longer a peripheral experiment in government. It is becoming part of the machinery through which states deliver services, allocate resources, manage risk, and exercise administrative capacity.

Yet the global record remains uneven. Many governments have launched pilots, appointed AI leaders, and published strategies, but far fewer have turned AI into a durable institutional capability. The central challenge is therefore no longer access to advanced models. It is whether governments can organise themselves to use AI safely, effectively, and at scale.

This white paper argues that progress in government AI depends less on technology than on institutional design. The key shift is from treating AI as a portfolio of use cases to treating it as public infrastructure - embedded in leadership structures, data systems, procurement choices, operating models, and mechanisms of accountability.

FROM PILOTS TO INSTITUTIONALISATION

Across jurisdictions, AI is now widely deployed in service delivery, regulatory enforcement, internal operations, and policy analysis. However, outcomes remain highly uneven. The literature converges on a consistent diagnosis: failures rarely stem from model performance alone, but from deeper structural constraints.

These include:

- fragmented and poorly governed data environments;
- weak cross-entity coordination;
- shortages of policy–technology “translator” roles;
- rigid procurement systems misaligned with iterative AI development;
- limited capacity to evaluate impact, risks, and unintended consequences.

As a result, AI in government is best understood not as a technological challenge, but as a problem of state capacity and institutional readiness.

WHY THE UAE CASE MATTERS

The United Arab Emirates is analysed not as a model to replicate, but as a case that makes these institutional dynamics unusually visible.

Over more than a decade, the UAE has combined:

- sustained political commitment to AI as a national priority;
- significant investment in digital and data infrastructure;
- early moves from experimentation toward integration into core government functions.

This environment allows observation of how governments manage trade-offs that are often abstract in the literature, including:

- speed versus control;
- sovereignty versus access to global technology;
- central coordination versus local experimentation;
- automation versus human accountability.

The UAE therefore functions as a practical testbed for institutional approaches to AI at scale.

A GOVERNANCE MODEL OF COORDINATED DIFFERENTIATION

At the federal level, AI strategy operates primarily as a coordination and signalling framework, setting direction through soft-law instruments, standards, and shared priorities rather than detailed central planning.

Within this frame, two distinct but complementary execution models have emerged:

- Abu Dhabi pursues an infrastructure-led approach, centred on sovereign cloud, shared platforms, and domain-level redesign. AI is treated as foundational infrastructure, with an emphasis on reuse, lifecycle governance, and long-term control.
- Dubai emphasises execution speed and visible service impact, using structured pipelines, acceleration task forces, and controlled experimentation to move rapidly from pilot to deployment.

This pattern reflects not fragmentation, but a form of coordinated differentiation: a shared national direction combined with varied operating models.

FROM PILOTS TO SYSTEMS: EXECUTION, PROCUREMENT, AND EVALUATION

The UAE experience reinforces a broader shift from isolated AI use cases to system-level transformation.

Three features are particularly salient:

- Operating model redesign: AI is increasingly embedded into end-to-end workflows, with hybrid human -AI structures and clearer accountability at the process level.
- Procurement as strategy: partnerships, pilot-to-scale pathways, and ecosystem mechanisms are used to shape innovation, manage vendor dependence, and balance sovereignty with access to global capabilities.
- Emerging evaluation frameworks: monitoring is moving beyond technical performance toward multi-layered oversight, though challenges remain in measuring public value, trust, and long-term societal impact.

Despite progress, persistent tensions remain around human oversight at scale, decision accountability, and the governance of increasingly autonomous systems.

IMPLICATIONS FOR GLOBAL POLICYMAKERS

The central lesson is conceptual rather than prescriptive. Governments that treat AI as institutional infrastructure, rather than a collection of projects, are more likely to move beyond pilots to durable systems.

Transferable insights include:

- anchoring AI strategy in long-term data, infrastructure, and governance capacity;
- embedding human oversight within workflows, rather than relying on high-level principles;
- aligning procurement, partnerships, and ecosystem development with strategic objectives;
- designing clear ownership and accountability for AI systems across their lifecycle;
- treating evaluation as a feedback mechanism that can trigger redesign or decommissioning.

At the same time, the limits of transferability are significant. The UAE's high administrative capacity, centralised coordination, and regulatory flexibility differ from many other contexts. Its experience should therefore be understood as a source of insight, not a template.

CONCLUSION

AI is reshaping not only what governments do, but how they are organised to do it. The transition underway is from experimentation to institutionalisation.

The UAE case illustrates what becomes possible when leadership, infrastructure, governance, and execution are aligned. It also highlights the unresolved challenges that accompany this shift.

For policymakers, the implication is clear: the future of AI in government will be determined less by access to technology than by the ability to redesign institutions around it.

PART 1. INTRODUCTION & RESEARCH FRAMING

Artificial intelligence (AI) has moved from the periphery to the centre of government agendas worldwide. What began as a series of discrete experiments—chatbots, predictive analytics, automation pilots—has evolved into a broader re-examination of how public institutions design services, make decisions, and exercise state capacity. This shift has been accelerated by advances in generative AI, which, simultaneously, have lowered barriers to experimentation and has been raising new questions around accountability, sovereignty, trust, and long-term governance.

Yet despite the growing volume of AI initiatives across governments, outcomes remain uneven. Many public sectors report an accumulation of pilots without scale, fragmented ownership across agencies, limited internal capability, and persistent uncertainty over how to govern AI systems whose behaviour is probabilistic, adaptive, and deeply dependent on data and organisational context

> **AI IN GOVERNMENT HAS REACHED A STRUCTURAL INFLECTION POINT**
>
> the question is no longer whether it can be deployed, but how it is designed, governed, and institutionalised to deliver sustained value.

It is against this backdrop that the United Arab Emirates (UAE) has emerged as a particularly salient case. Over the past decade, the UAE has articulated a strong national ambition around AI, invested heavily in digital and data infrastructure, and moved beyond experimentation toward the integration of AI into core government functions. Unlike many jurisdictions where AI remains confined to isolated projects, UAE government entities are actively grappling with questions of leadership, procurement, sovereignty, evaluation, and organisational redesign—precisely the issues that the global literature identifies as decisive yet under-examined.

RESEARCH PURPOSE AND CENTRAL STATEMENT

This white paper seeks to address the following central statement:

The UAE's progress stems less from technology than from treating AI as public infrastructure, built around three institutional choices (density of leadership, domain-level redesign, procurement-as-strategy) that others may be able to adapt.

Rather than cataloguing technologies or use cases, the paper focuses on institutional choices: how AI strategies are translated into practice; how trade-offs between innovation, risk, and sovereignty are managed; and how public value is defined, measured, and sustained over time.

The study is structured around six interrelated themes that reflect both the global literature and the UAE experience:

- Strategic vision and leadership, including the role of senior political commitment, institutional mandate, and Chief AI Officers;
- Selection and prioritisation of AI use cases, and the mechanisms used to move from experimentation to scale;
- Execution challenges and success factors, spanning data, infrastructure, talent, organisational change, and responsible AI;
- Evaluation, impact, and unintended consequences, including how governments assess efficiency, citizen experience, trust, and broader social value;
- Procurement and ecosystem strategies, encompassing buy-versus-build decisions, public–private partnerships, open-source approaches, and vendor relationships;
- Lessons for global policymakers, with a focus on transferability, preconditions, and common pitfalls.

METHODOLOGY AND SCOPE

The analysis adopts a qualitative research methodology, grounded in in-depth interviews with Chief AI Officers, senior digital transformation leaders across key UAE government entities, and senior executives from the private sector partners. These interviews are complemented by a structured review of the global and regional literature on AI in government, drawing on policy reports, academic research, and practitioner studies.

The qualitative approach is intentional. While quantitative indicators can illuminate adoption rates or cost savings, they are less effective in capturing the organisational, governance, and political economy dynamics that shape AI outcomes in the public sector. Interviews allow for a closer examination of how decisions are made in practice, how constraints are navigated, and how success and failure are understood by those directly responsible for delivery.

ANALYTICAL POSITIONING AND CONTRIBUTION

This paper is positioned neither as a technology roadmap nor as a prescriptive "best practices" guide. Instead, it makes three distinct contributions:

1. It situates the UAE experience within the global AI-in-government literature, using existing research to frame the core design and governance challenges governments face.
2. It provides an empirically grounded deep dive into the UAE, offering rare insights into how these challenges are addressed in a high-capacity, fast-moving governmental context.
3. It distils lessons for global policymakers, emphasising design principles, institutional trade-offs, and conditions for transferability rather than replication.

Importantly, the paper does not assume that the UAE model can or should be copied wholesale. Differences in political systems, administrative capacity, data availability, and public trust mean that what works in one context may not translate directly to another. The objective is therefore not imitation, but learning: understanding which elements of the UAE experience are broadly applicable, which are context-dependent, and which raise new questions for future research and policy experimentation.

WHAT THIS PAPER DOES NOT DO

To avoid scope creep and over-generalisation, the paper does not attempt to:

- provide a comprehensive inventory of AI use cases across UAE government;
- evaluate individual AI systems from a technical or algorithmic perspective;
- offer normative judgments on the "optimal" level of AI adoption in government.

Instead, it focuses squarely on how governments organise themselves to use AI, and what this reveals about the future of public sector transformation.

PART 2. GLOBAL BASELINE

THE GLOBAL OVERVIEW

AI is now widely deployed across governments, spanning service delivery, regulatory enforcement, internal operations, and policy analysis. Across jurisdictions, public-sector organisations report tangible benefits from AI adoption, including improvements in processing speed, targeting accuracy, and operational efficiency. At the same time, outcomes remain highly uneven. While some governments have begun integrating AI into core administrative functions, many continue to struggle with fragmented initiatives, stalled pilots, and limited institutional uptake.

> **DIFFERENCES IN OUTCOMES ARE NOT EXPLAINED BY ACCESS TO TECHNOLOGY**
>
> but by governance arrangements, institutional readiness, and organisational capacity

Governments with similar technical tools often experience markedly different outcomes depending on how AI initiatives are prioritised, governed, procured, and evaluated. As a result, AI in government is no longer understood as a question of technological adoption, but as one of institutional design and state capacity.

From early use cases to today

Early government applications of AI focused largely on automation—using rule-based systems and predictive models to streamline back-office processes and improve administrative efficiency. Over time, AI systems have evolved to support more complex forms of augmentation and decision support, assisting public officials in areas such as eligibility determination, risk assessment, and policy analysis.

Recent advances in generative and agentic AI mark an inflection point rather than a new category of use. These systems extend AI's role from supporting predefined tasks to interacting with users, synthesising information, and shaping workflows across organisational boundaries. The significance of this shift lies less in the underlying techniques than in its institutional implications: AI systems are becoming more deeply embedded in public-sector processes, raising new questions about governance, accountability, and long-term integration.

USE-CASE CATEGORY	ILLUSTRATIVE FOCUS
Citizen engagement and services	Chatbots, virtual assistants, service navigation, case handling
Compliance and enforcement	Risk scoring, fraud detection, inspection targeting
Internal operations	Process automation, workforce planning, document processing
Policy analysis and planning	Forecasting, scenario modelling, data-driven policy support

Table 1: The most common government AI use cases

Most common government AI use cases

The literature converges around a small number of dominant categories of AI use in government. These categories recur across jurisdictions and provide a useful typology for understanding where governments concentrate their efforts.

This typology serves as a reference point rather than evidence of impact. The prevalence of similar use cases across countries underscores the importance of governance and execution choices in determining outcomes.

The geography of AI in government

AI adoption in government is global. Advanced economies, emerging markets, and lower-income countries alike are experimenting with AI-enabled public services and administrative tools. While levels of maturity and scale vary widely, the literature points to increasing convergence around shared governance concerns, including data protection, transparency, accountability, and ethical risk management.

Differences across countries are therefore less about whether AI is being adopted, and more about how effectively it is institutionalised. Administrative capacity, regulatory frameworks, and organisational coherence shape readiness and outcomes far more than geography alone. As a result, geographic comparisons are most analytically useful when they illuminate variations in governance models rather than technological ambition.

Note on technology typologies

Detailed breakdowns of AI technologies (such as machine learning, natural language processing, computer vision, or automation) are intentionally excluded from this analysis. While these technologies underpin many government use cases, categorising them adds limited analytical value to a study focused on leadership, governance, procurement, and evaluation. The remainder of this paper therefore treats AI as an institutional capability rather than a set of discrete technical components.

STRATEGIC VISION AND LEADERSHIP

A dedicated, forward-looking strategic vision and strong leadership are widely recognised as foundational conditions for AI transformation in government. Senior leaders — supported by Chief AI Officers and cross-government governance mechanisms — shape whether AI remains a series of experiments or evolves into a system-wide capability for public value.

> **VISION ALONE IS INSUFFICIENT**
>
> without institutional mandate, sustained funding, and embedded governance structures, AI strategies risk remaining aspirational rather than operational (Deloitte ; World Bank)

How senior leadership and institutional commitment shape AI adoption

Effective senior leadership establishes a strategic direction for safe, innovative, and responsible AI, ensuring that AI priorities are embedded at the executive level rather than siloed as stand-alone digital projects. Institutional commitment translates this direction into durable structures: multi-year investment, capability development, procurement alignment, and mechanisms for public engagement. Where such commitment is fragmented or short-lived, AI initiatives tend to remain confined to pilots.

Executive leaders — including cabinet ministers and heads of digital or AI strategy — also play a signalling function. By articulating AI's purpose, risks, and societal implications, they help shape internal incentives and public trust, particularly in high-stakes domains (OECD). Yet the literature notes that communication without corresponding authority and resources can create expectations that outpace delivery.

The role of chief AI officers

The Chief AI Officer (CAIO) has emerged as a central coordinating role in many governments. Positioned at the executive level, CAIOs are typically responsible for aligning enterprise-wide AI strategy, technical standards, ethical frameworks, and procurement practices, while building institutional capability and ensuring responsible innovation (Government-Transformation.com).

In practice, CAIO responsibilities often include leading cross-entity working groups and advisory councils (AFERM). harmonising standards and procurement approaches across agencies (Government-Transformation), convening public–private and inter-ministerial partnerships (UK National Audit Office), and tracking risks related to explainability, fairness, and compliance. However, the effectiveness of the role depends heavily on mandate, budget authority, and integration into core decision-making processes. Where CAIOs lack formal authority or dedicated resources, the position risks becoming symbolic rather than transformative.

Cross-entity governance

Integrated AI governance — through innovation hubs, steering committees, digital authorities, or cross-ministry boards — aims to align priorities, minimise duplication, and enable knowledge-sharing across domains (AI Center for Government; World Bank). Strong oversight structures can support scaling by standardising risk frameworks, disseminating lessons learned, and maintaining whole-of-government alignment on ethics and technical standards.

At the same time, the literature cautions that governance mechanisms must balance coherence with flexibility. Overly centralised structures may slow experimentation, while fragmented arrangements increase duplication and weaken accountability. The institutional challenge is therefore not simply to create governance bodies, but to design them with clear authority, defined escalation paths, and sustained political backing.

USE CASES: SELECTION & PRIORITISATION

> **THE CORE CHALLENGE IS NOT GENERATING IDEAS**
>
> but reducing an abundance of plausible use cases to a small portfolio that can be executed and scaled

Across jurisdictions, a broadly similar four-stage logic has emerged:

- Ideation & discovery Agencies' and central units' surface potential applications through strategic planning cycles, problem statements, or national challenges, usually anchored in public-value objectives such as service quality, cost pressure, or resilience.
- Assessment & feasibility Use cases are evaluated for strategic alignment, data readiness, technical feasibility, legal and ethical risk, and potential impact — including both ROI and wider societal value.
- Strategic prioritisation Shortlisted initiatives are ranked so that limited capacity and funding are focused on a small number of projects that combine impact, feasibility, and policy relevance, rather than spreading effort thinly across too many pilots.
- Pilot implementation & scaling Selected use cases undergo pilot testing to validate effectiveness, user adoption, and manageability of risks before scaling. Only those that demonstrate measurable value, organisational buy-in, and replicability progress to system-wide or national deployment.

Within this shared logic, governments differ in where discretion sits, how tightly pilots are governed, and how quickly they are prepared to scale.

United Kingdom: "Scan → Pilot → Scale"

The UK formalises this logic in its AI Opportunities Action Plan, designed to systematically embed AI across public services. It structures selection and scaling as:

- Scan: Appoint AI leads for each mission, build horizon scanning and market intelligence, and partner with vendors and startups to anticipate future needs.
- Pilot: Decide whether to build, buy, or run innovation challenges; support rapid prototyping with in-house experts and user centred design; streamline access to data, infrastructure, and funding for low bureaucracy pilots.
- Scale: Provide central support and funding to help successful pilots reach national scale; launch national AI tenders; standardise infrastructure, open source code, and interoperability. Standardisation and reuse are treated as prerequisites for rollout, not afterthoughts.

Singapore: Cross agency, co created, governance heavy.

- Singapore applies a similarly systematic process, but with stronger cross agency coordination and public–private co creation. Use cases are prioritised through national level calls

for problem statements, assessed for strategic alignment, feasibility, scalability, and public value. Programmes like AI Singapore's "100 Experiments" pair agencies with researchers and developers to co create solutions, emphasising agile prototyping and iterative refinement. Deployment is grounded in explicit governance requirements, notably the Model A Governance Framework and AI Verify, which embed ethics, transparency, and explainability from the outset rather than as post hoc checks.

United States: Broad adoption and speed.

- The US federal AI Action Plan focuses on broad based adoption by lowering regulatory barriers, centralising leadership through the Chief Artificial Intelligence Officer Council, and providing agencies with toolkits for uniform AI procurement and integration. Key bodies such as the Department of Defence are tasked with identifying internal workflows for automation and forming priority partnerships with the private sector for computing infrastructure. Pilot programmes, especially for workforce development, are used, but the emphasis is on rapid deployment and extensive public private collaboration more than on a cautious, tightly sequenced pilot to scale funnel. (The White House).

European Union: Prescriptive and risk managed.

- The EU's AI strategies and legislation (including the AI Act and the Apply AI Strategy) promote AI in public services primarily where it improves service delivery and administrative efficiency, but under stringent regulatory conditions. Ethical evaluation, transparency, and citizen protection are required throughout the lifecycle; agencies must address risks and demonstrate safeguards before deployment. Pilot programmes and local innovation hubs are encouraged, but all activity must align with human centric principles and regulatory compliance, which can slow deployment but provides higher legal certainty and trust. (Live-in.eu).

China: Centrally coordinated, sovereignty focused model.

- China pursues a state driven, centrally coordinated model, as reflected in its Action Plan for Global AI Governance and earlier national plans. Strategic public sector use cases are identified at national level; sectoral guidance is issued and implementations mandated across agencies, backed by large state investment in infrastructure, open source standards, and data security. Pilots and scaling phases are tightly managed, often under the direction of the State Council and key ministries, with strong emphasis on software and hardware sovereignty, integration of AI into public management, and alignment with national security and social stability objectives.

  Across these models, a clear shift is visible: governments are moving away from opportunistic experimentation toward portfolio discipline. AI use case selection is no longer treated as a narrow technology exercise, but as a governance capability that requires prioritisation mechanisms, political backing, and institutional capacity to say no as well as yes.

The core trade-offs are consistent across jurisdictions:

- Speed vs control (for example, the US bias toward rapid deployment versus the EU's risk managed approach).
- Central mandate vs agency autonomy (for example, China's top down orchestration versus Singapore's co created, cross agency model).
- Standardisation vs contextual tailoring (for example, the UK's strong emphasis on shared platforms versus more federated systems).

EXECUTION CHALLENGES AND SUCCESS FACTORS

> **EXECUTION REMAINS THE DECISIVE BOTTLENECK**
>
> failures rarely stem from model performance, but from structural constraints in data, institutions, and governance

Core execution challenges

1. Data quality, interoperability, and governance. AI systems in government are only as effective as the data infrastructures on which they rely. Legacy IT systems, fragmented data ownership, inconsistent standards, and weak data hygiene routinely undermine deployment. These issues are compounded by stringent requirements around privacy, security, and lawful use of data. In practice,

many AI projects stall not because algorithms fail, but because data cannot be reliably shared, audited, or governed across agencies.

2. Talent and capability constraints. Governments face a persistent shortage of AI-literate talent, particularly at the intersection of technology, policy, and domain expertise. A 2024 Salesforce survey found that 60% of public-sector IT professionals cite AI skills shortages as their single greatest implementation challenge (Salesforce.com), exceeding private-sector levels. Recruitment lags, rigid pay structures, and formalised civil-service career paths make it difficult to attract scarce AI talent into permanent government roles. As a result, governments often depend heavily on external consultants and vendors. While necessary in the short term, this reliance can create power asymmetries, misaligned incentives, and long-term risks to institutional learning and in-house capability building.

3. Organisational change and operating models. AI adoption is not a technology upgrade but an organisational transformation. It requires redefining roles, decision rights, and accountability structures - yet change management is consistently underestimated. Public administrations are designed to privilege stability, procedural compliance, and risk avoidance. Deeply embedded practices, rigid hierarchies, and civil-service rules limit experimentation and slow iteration. Political cycles, leadership turnover, and competing departmental priorities further disrupt continuity, making sustained transformation difficult even when pilots succeed.

4. Cost, infrastructure, and legacy constraints. High upfront costs, long procurement cycles, and outdated infrastructure continue to delay or constrain AI deployment. Governments operate within fixed budget envelopes and multi-year investment cycles, making it harder to fund experimentation, absorb failure, or rapidly reallocate resources. Legacy systems often require extensive modernisation before they can support advanced AI workloads, adding time and cost that are rarely visible at the pilot stage.

5. Governance, regulation, and uncertainty. AI deployment in government unfolds within evolving regulatory and ethical frameworks. The need to ensure fairness, explainability, accountability, and legal defensibility adds layers of complexity absent from most private-sector deployments. Unclear or rapidly changing rules can slow execution, yet insufficient governance creates equally serious risks to public trust and legitimacy.

Public vs private

While private firms encounter many of the same obstacles, the public sector faces them under materially stricter constraints. Government data environments are more siloed, privacy obligations are higher, and failures carry social and political consequences rather than merely financial ones. Talent recruitment and budgeting are slower and more formalised, limiting flexibility. Organisational change must navigate entrenched hierarchies and public scrutiny, while private firms can restructure aggressively in pursuit of efficiency. Regulation in government prioritises ethical standards, transparency, and social accountability, whereas private firms typically focus on compliance, liability, and reputational risk.

Emerging success factors

Across jurisdictions, successful AI implementation in government shares several enabling conditions:

- Data foundations first: sustained investment in interoperable, well-governed data infrastructures before scaling AI use cases.
- Embedded capability: deliberate strategies to build internal AI literacy and ownership, reducing long-term dependence on vendors.
- Leadership continuity and mandate: clear political backing that survives electoral or administrative turnover.
- Adaptive governance: regulatory and ethical frameworks that enable experimentation while preserving accountability and trust.
- Change as design principle: treating organisational redesign and workforce transition as core components of AI programmes, not downstream afterthoughts.

Taken together, these factors suggest that the primary challenge of AI in government is not technological sophistication, but institutional readiness. Governments that treat AI as an operating-model transformation — rather than a portfolio of tools — are more likely to move from pilots to durable, trusted systems at scale.

PROCUREMENT AND ECOSYSTEM STRATEGIES



As AI systems become more data-dependent, adaptive, and embedded in core public functions, traditional public procurement models have struggled to keep pace.

Procurement beyond traditional IT logic

Unlike conventional IT systems, AI solutions are characterised by performance uncertainty, continuous iteration, and reliance on large volumes of data. Outputs cannot always be fully specified ex ante, and system behaviour may change over time as models are retrained, updated, or deployed in new contexts. The literature therefore highlights a growing mismatch between static procurement processes — designed around fixed deliverables — and the dynamic nature of AI systems.

This mismatch has governance implications. Procurement decisions increasingly determine not only cost and functionality, but also how risks related to bias, explainability, data protection, and accountability are managed. In high-stakes public-sector contexts, contracts must embed requirements for transparency, auditability, and contestability — features that extend well beyond traditional service-level agreements.

Vendor dependence and sovereignty considerations

A recurring concern in the literature is the risk of vendor dependence. AI systems often rely on proprietary models, specialised talent, and compute-intensive infrastructure, creating asymmetries of information and bargaining power between governments and suppliers. Where internal capability is limited, procurement can inadvertently lock public institutions into long-term dependencies that constrain future policy choices.

As a result, procurement is increasingly linked to questions of digital sovereignty. Decisions about where models are trained, how data is stored and accessed, and who controls system updates have implications for national autonomy, resilience, and trust. The literature does not prescribe a single solution, but it consistently emphasises the need for governments to retain sufficient control over foundational layers to govern AI systems over their full lifecycle.

Ecosystem-oriented procurement

Beyond individual contracts, procurement is also framed as a tool for shaping broader AI ecosystems. Governments use procurement to encourage competition, support domestic innovation, and foster collaboration between public agencies, academia, and private-sector actors. Approaches such as innovation partnerships, challenge-based procurement, and framework agreements aim to balance experimentation with accountability.

However, the literature cautions that ecosystem-oriented procurement introduces additional coordination challenges. Without clear pathways from experimentation to operational deployment, such approaches risk reproducing the very fragmentation they seek to avoid. Effective ecosystem strategies therefore depend on strong governance, clear ownership, and mechanisms to transition successful initiatives into core public-sector systems.

Emerging global frameworks

Recognising these challenges, international and national bodies have begun to adapt guidance and regulation to the realities of AI procurement. The OECD 'AI Procurement in a Box' guidelines promote risk based procurement with strong transparency and explainability expectations. The European Union AI Act imposes binding requirements on high risk AI systems, which public procurers must ensure are met when acquiring such systems. The US federal government has issued AI related procurement guidance and relies on tools such as the NIST AI Risk Management Framework and OMB AI Guidance to emphasise risk management and vendor accountability. Singapore's Model AI Governance Framework, and its AI Verify testing toolkit are increasingly embedded into public sector procurement standards, seeking to balance innovation with robust safeguards and collaborative development with vendors.

While these frameworks vary in form and legal force, they reflect a shared recognition that procurement is inseparable from AI governance.

EVALUATION, IMPACT, AND UNINTENDED CONSEQUENCES

Evaluation

> **EVALUATION HAS BECOME A GOVERNANCE FUNCTION**
>
> legitimacy depends not only on performance, but on demonstrable impact and the management of unintended consequences

Across advanced jurisdictions, evaluation frameworks are becoming more sophisticated, but they diverge in how prescriptive, transparent, and enforceable they are.

The UK Government has long emphasised rigorous monitoring and evaluation through the Magenta Book and related guidance. AI initiatives are expected to incorporate structured piloting, ongoing evaluation, and stakeholder engagement throughout the lifecycle, with methods ranging from qualitative assessment to randomised control trials where appropriate. This reflects a strong tradition of evidence-based policymaking, but also imposes relatively high procedural overhead.

In contrast, the United States Government has recently shifted toward a more flexible, decentralised approach. While agencies remain responsible for impact assessments, testing, and public consultation, the 2025 federal guidance reduces mandatory procedural requirements and places greater weight on innovation, agency discretion, and competitiveness. Transparency and public input persist, but through more consultative and less binding mechanisms than under earlier federal policy.

The European Union adopts the most formalised and legally binding model. Under the AI Act, public-sector uses classified as “high-risk” must undergo extensive pre-deployment and ongoing impact assessments, accompanied by documentation, auditability, and citizen redress mechanisms. This approach prioritises rights protection and trust, but can slow deployment and raise compliance costs.

Government of Singapore occupies a distinctive position. Its Model AI Governance Framework is voluntary yet comprehensive, combining technical testing (via AI Verify) with process checks, advisory oversight, and public consultation. Importantly, this framework has evolved to encompass generative AI, allowing both self-assessment and third-party validation. Singapore’s model is widely regarded as pragmatic: it enforces discipline without freezing innovation, and treats evaluation as a learning mechanism rather than a purely regulatory one.

Impact

Where evaluation frameworks are mature, evidence of impact is increasingly visible - though uneven across domains.

Operational efficiency gains are among the most consistently documented outcomes. For example, AI-enabled systems deployed by Singapore’s Ministry of Manpower reportedly reduced call-centre workloads by around 50%, largely through automated transcription and knowledge retrieval (NCS). In other sectors, predictive AI models have significantly improved detection accuracy in areas such as seismic monitoring, with reported gains of up to 70% in Japan and China (Indiaai.gov) (Deccan Herald).

Beyond efficiency, global reviews by institutions such as the World Bank suggest that AI can improve service quality, transparency, and responsiveness, while generating measurable cost savings through reduced errors and faster processing. (The World Bank).

Unintended consequences

As AI systems scale within public administration, evaluation efforts increasingly surface second-order effects that are difficult to quantify yet central to public legitimacy. Beyond efficiency metrics, assessments consistently point to risks such as job displacement in administrative functions, widening digital divides, disruptions to public-sector revenue models, and heightened exposure to algorithmic bias, privacy breaches, and workforce dissatisfaction — patterns documented across global analyses by the World Bank, the UN, and leading policy think tanks.

Crucially, these risks are not theoretical. Several high-profile failures over the past decade demonstrate that harm in government AI deployments rarely stemsfrom model performance alone, but from breakdowns in data governance, oversight, and organisational accountability. The UK Home Office Windrush scandal revealed how inaccurate data, opaque automated decision systems, and weak appeal mechanisms led to wrongful detentions

and deportations, causing lasting damage to public trust. Similarly, the Dutch Tax and Customs Administration childcare benefits scandal (Toeslagenaffaire) showed how poorly governed algorithmic risk scoring, combined with inadequate oversight and institutional inertia, resulted in systemic discrimination and ultimately the resignation of the national government.

These cases illustrate a central lesson for policymakers: in government, execution failure is not neutral. Unlike the private sector, where errors may be absorbed as commercial losses, public-sector AI failures generate legal harm, political fallout, and long-term erosion of institutional legitimacy. As a result, many governments are now reframing unintended consequences as a core evaluation dimension, treating them as early warning signals of governance and organisational weakness rather than as peripheral externalities to be managed after deployment.

The productivity paradox in public administration

A further insight emerging from recent research is a public-sector productivity paradox. Studies published between 2024 and 2025 show that while automation reduces repetitive tasks, it can simultaneously increase cognitive load and stress for public servants, who must monitor, interpret, and override AI outputs. This shift from execution to supervision can reduce job satisfaction and, in some cases, offset expected productivity gains. (Public Management Review).

AI changes the nature of public-sector work before it improves aggregate productivity. Without deliberate role redesign, training, and organisational adaptation, AI risks creating new bottlenecks even as it removes old ones.

Implications for government ai strategy

Taken together, global experience suggests three implications:

- Evaluation must be continuous, not limited to pre-deployment approval.
- Impact must be defined broadly, encompassing trust, workforce effects, and social outcomes alongside efficiency.
- Unintended consequences are governance signals, revealing institutional weaknesses rather than marginal technical flaws.

Governments that treat evaluation as a learning system — rather than a compliance exercise — are better positioned to sustain AI adoption at scale while maintaining public legitimacy.

WHAT THE LITERATURE ESTABLISHES, AND WHERE IT STILL STRUGGLES

Across academic research, multilateral policy work, and practitioner guidance, there is broad agreement that AI adoption in the public sector is no longer primarily a technological challenge. It is increasingly framed as an institutional, organisational, and governance problem, shaped by leadership structures, incentive systems, procurement choices, and mechanisms of accountability.

The literature is particularly strong in its diagnosis of a first generation of challenges governments face. It shows that political leadership and strategic vision are important but tend to be insufficient without formal mandates, sustained funding, and institutional continuity. It documents how unstructured experimentation and project-based funding can lead to fragmented pilots that struggle to scale. It repeatedly finds that execution failures are predominantly socio-technical rather than algorithmic, rooted in data fragmentation, organisational silos, and misaligned incentives. And it highlights the limits of narrow efficiency-based evaluation, emphasising the importance of trust, legitimacy, and long-term public value in democratic settings.

In this sense, the field has moved beyond early optimism and simplistic narratives of technological transformation. These core obstacles are well documented, widely acknowledged, and observed across a range of jurisdictions, even if depth of evidence varies by context.

Where the literature continues to struggle is less in recognising these challenges than in resolving a second set of structural and institutional questions that have become more acute as AI systems grow more autonomous, interconnected, and embedded in public administration. First, while the need to scale beyond pilots is widely recognised, there are comparatively few detailed, empirical accounts of the mechanisms of scale. Existing work offers only scattered examples of how governments redesign funding models, operating structures, and governance arrangements to support AI as a sustained, cross-cutting capability rather than a sequence of isolated projects.

Second, although organisational dynamics are frequently discussed, questions of institutional ownership over time remain only partially specified. As AI systems evolve, are repurposed, or operate across administrative boundaries, issues of authority, liability, and stewardship become increasingly complex, and the literature offers limited guidance on how responsibility is allocated and maintained as systems shift from decision support toward more agentic forms of action.

Third, procurement is increasingly recognised as central to AI governance, yet it is still only unevenly theorised as a strategic design choice. Concerns about vendor dependence and loss of sovereignty are well noted, but practical guidance on balancing build, buy, and hybrid approaches — particularly in the context of frontier models and rising compute costs — remains fragmented.

Fourth, despite growing attention to sovereignty and data control, AI in government is still more often analysed as a collection of applications than as a form of foundational digital public infrastructure, although a small but expanding body of work has begun to adopt this lens. This leaves partly unresolved questions about long-term stewardship, lifecycle governance, and the institutional arrangements required to manage shared AI capabilities that underpin multiple policy domains.

Finally, although evaluation frameworks have expanded to include fairness, transparency, and impact assessment, the evaluation–governance feedback loop remains under-examined. There is limited comparative evidence on how evaluation findings translate into policy adjustment, system redesign, or decommissioning, as opposed to remaining largely formal compliance exercises.

These gaps do not reflect a lack of awareness or conceptual effort so much as the limits of literature-based analysis in the absence of detailed empirical insight into how governments actually attempt to resolve these tensions in practice. It is at this boundary - between relatively well-established diagnosis and under-specified institutional resolution - that the empirical analysis of the UAE begins.

PART 3. THE UAE EXPERIMENT

THE UAE CASE: CONTEXT, METHOD, AND ANALYTICAL LENS

This section turns from literature-based diagnosis to empirical analysis. Building on the institutional questions identified in Part I, it examines how AI is being designed, implemented, and governed in the public sector of the United Arab Emirates. The UAE is not presented as a benchmark to be replicated, nor as an outlier insulated from constraints faced elsewhere. Rather, it is analysed as a case of concentrated policy intent and administrative capacity, offering insight into how governments attempt, in practice, to address the second-wave institutional challenges highlighted by the literature.

Why the UAE is analytically relevant

The global literature increasingly converges on the view that the primary barriers to effective AI deployment in government are institutional rather than technical. Yet empirical accounts of how governments operationalise leadership intent, manage procurement trade-offs, assign ownership over time, and integrate evaluation into governance remain limited. The UAE provides a particularly useful setting to explore these dynamics, for three reasons.

First, the UAE has articulated AI as a strategic state priority over a sustained period, embedding it within broader agendas of public-sector reform, digital government, and economic transformation. This continuity allows examination of how long-term ambition interacts with organisational realities beyond the pilot phase.

Second, the UAE public sector operates within a governance context characterised by comparatively high state capacity, strong central coordination, and the ability to mobilise resources rapidly. This does not eliminate institutional challenges, but it alters their form and visibility, making trade-offs around scale, control, and accountability more explicit.

Third, the UAE's approach combines domestic capability-building with extensive engagement with global technology firms and research communities. This hybrid orientation creates a relevant setting in which to examine questions of procurement, vendor dependence, and sovereignty that are increasingly salient across jurisdictions.

These characteristics do not make the UAE representative of all governments. They do, however, make it analytically useful for understanding how the institutional problems identified in the literature are confronted under conditions that are relatively favourable to execution.

> **THE UAE IS NOT A MODEL TO REPLICATE**
>
> but a case through which institutional choices, trade-offs, and conditions for transferability can be understood.

Methodological approach

The UAE analysis is based on a qualitative research design centred on in-depth interviews with senior leaders responsible for digital transformation, data, and AI strategy across key government entities, complemented by interviews with senior private-sector partners operating at the government interface. These interviews are supplemented by a review of publicly available policy documents, strategies, and regulatory materials to situate interview evidence within its formal institutional context.

The qualitative approach is intentional. While quantitative indicators can illuminate adoption levels or headline efficiency gains, they are less effective in capturing how decisions are made, how authority is exercised, and how trade-offs are managed over time. Interviews allow closer examination of how AI initiatives are prioritised, governed, and adapted in response to organisational constraints and evolving risk perceptions.

The analysis does not seek to evaluate individual AI systems from a technical perspective, nor to provide a comprehensive inventory of AI use cases across the UAE government. Instead, it focuses on institutional arrangements and decision processes: how leadership intent is translated into operational structures, how procurement and partnerships are managed, how responsibility is allocated, and how impact is assessed.

Analytical structure

The UAE case is organised around the same core themes developed in Part I, allowing direct comparison between the literature and observed practice. Subsequent sections examine:

- how strategic vision and leadership are institutionalised;
- how AI use cases are selected, prioritised, and scaled;
- how execution challenges related to data, talent, and organisational change are addressed;
- how procurement and ecosystem strategies shape capability and control;
- how evaluation and risk management are incorporated into governance.

This thematic structure is designed to foreground design choices and trade-offs, rather than outcomes alone.

Scope and limits

It is important to clarify what this case analysis does not claim. The UAE experience does not offer a single model of AI governance, nor does it resolve all of the institutional questions identified in the literature. Differences in political systems, administrative traditions, and public accountability mechanisms limit direct transferability. The purpose of the analysis is therefore not prescription, but learning: to identify which elements of the UAE approach appear contingent on context, which may be adaptable elsewhere, and which raise new questions for policymakers and researchers.

The sections that follow examine these issues in detail, beginning with the role of strategic vision and leadership in shaping AI adoption within the UAE public sector.

STRATEGIC VISION AND LEADERSHIP

Global literature highlights strategic vision and senior leadership as essential for AI adoption in government, yet notes that vision alone often fails to drive lasting change without institutional support. The UAE demonstrates how federal direction integrates with emirate-specific approaches, creating a layered governance system that balances coherence and flexibility.

Federal level: strategic direction, economic imperatives, and cross-government alignment

At the federal level, AI leadership is anchored in the UAE National Artificial Intelligence Strategy 2031, launched in 2018 as a long-term framework to embed artificial intelligence across government, the economy, and society. Interview evidence suggests that the strategy functions less as a technology roadmap than as a strategic coordination device, aligning AI deployment with national priorities related to resilience, economic diversification, and public- sector transformation.

While the rationale for AI adoption mirrors global narratives around efficiency, transparency, and service quality, interviewees also framed AI as carrying additional structural significance in the UAE context. In particular, AI is viewed as a lever for reducing long-term dependence on a large, predominantly foreign, manual workforce, and for accelerating the transition toward more knowledge-intensive, locally embedded state capabilities. This economic and demographic dimension lends urgency to AI adoption that extends beyond service optimisation alone.

The Strategy articulates a comprehensive set of objectives spanning ecosystem development, priority-sector deployment, talent and research capacity, secure infrastructure, and responsible governance. Rather than prescribing detailed implementation pathways, it establishes a shared reference framework within which emirates and federal entities can align initiatives while retaining discretion over execution. This approach reflects a deliberate balance between coherence and flexibility, a trade-off highlighted in the global literature as central to scaling AI beyond pilots.

> **SUCCESS IN GOVERNMENT AI IS NOT ABOUT HOW MANY PILOTS ARE LAUNCHED**
>
> but how deeply AI improves the daily citizen experience

Implementation and coordination are led by the UAE Office of Artificial Intelligence, Digital Economy and Remote Work Applications, which plays a convening role across ethics, talent development, research partnerships, and international engagement. Interviewees described this federal office as providing continuity, signalling political commitment, and arbitrating cross- cutting issues, rather than acting as a central delivery unit. Its mandate reinforces AI as a cross-government concern, helping to mitigate fragmentation while avoiding excessive centralisation.

Taken together, the federal layer positions AI as both an economic transformation instrument and a governance challenge, setting direction and expectations while enabling differentiated implementation across emirates. This role is critical in shaping the conditions under which Abu Dhabi and Dubai pursue distinct but complementary leadership approaches, examined in the sections that follow.

Abu Dhabi: infrastructure-led leadership and systemic integration

Leadership in Abu Dhabi is characterised by an infrastructure- and ecosystem-led approach, with a strong emphasis on sovereign AI capacity, foundational research, and long-term institutional resilience. Interview evidence suggests that AI is treated not merely as a tool for improving individual services, but as a strategic layer underpinning government operations, economic development, and security-related functions.

This orientation reflects Abu Dhabi's broader role within the federation, where long-term investment horizons and stewardship of critical state assets shape governance priorities. Leadership decisions therefore privilege control, reliability, and lifecycle governance over rapid experimentation, particularly where AI systems interact with sensitive data or core regulatory functions.

These priorities are formalised in the Abu Dhabi Government Digital Strategy 2025–2027, which articulates the ambition of positioning the emirate as an "AI-native government." Rather than framing AI as an add-on to existing digital services, the strategy treats it as a foundational capability embedded across public service delivery, policy design, and internal operations.

Central to this approach is the vision of a "government without walls," which seeks to integrate services, data, and institutional boundaries into coherent citizen journeys. Strategic priorities include hyper-personalised digital services, more proactive policy intervention, and the integration of emerging technologies such as generative AI, sovereign cloud infrastructure, and digital identity systems. While these ambitions are forward-looking, interview evidence indicates that they are pursued through staged implementation, with particular attention to data readiness, governance safeguards, and system interoperability.

Abu Dhabi's infrastructure-led strategy is supported by significant institutional and financial commitments. A landmark partnership between Microsoft and G42 underpins the development of secure cloud and AI infrastructure designed to support scale while preserving data sovereignty. An allocation of AED 13 billion signals the depth of political and fiscal commitment behind the transformation agenda. Importantly, interviewees framed these investments not as ends in themselves, but as enabling conditions for sustained capability-building and cross-government reuse.

Governance of this transformation is structured around three pillars: personalised citizen journeys, AI-enabled policymaking, and agile, future-ready government operations. Together, these pillars reflect an attempt to align technological innovation with institutional safeguards related to trust, privacy, and resilience. Rather than prioritising speed alone, leadership in Abu Dhabi appears to emphasise system integrity and long-term stewardship, consistent with the literature's emerging focus on AI as a component of digital public infrastructure.

In analytical terms, Abu Dhabi's approach illustrates one resolution of the scale–control trade-off identified in the global literature. By investing heavily in shared infrastructure, research capacity, and sovereign control, the emirate seeks to reduce fragmentation and enable AI deployment at scale across policy domains.

Dubai: executive sponsorship and execution-oriented leadership

Dubai's leadership model places greater emphasis on **execution speed, executive sponsorship, and service-level impact**. Interviewees described a governance environment in which senior leaders actively champion AI initiatives tied to specific operational or citizen-facing outcomes, often with shorter delivery cycles and clearer performance signals.

Institutions such as the Dubai Centre for Artificial Intelligence and Dubai Digital Authority play a convening role, supporting experimentation, rapid prototyping, and collaboration with private-sector partners. AI initiatives are frequently advanced through pilots that are tightly linked to service delivery metrics, with leadership prepared to iterate or discontinue projects based on observed results.

This execution-oriented approach reflects Dubai's long-standing administrative culture of innovation and responsiveness. Rather than prioritising full system integration from the outset, leadership accepts a degree of fragmentation in exchange for learning, speed, relying on governance mechanisms to manage risk and prevent uncontrolled diffusion.

UAE AI ECOSYSTEM: COORDINATED DIFFERENTIATION ACROSS FEDERAL, ABU DHABI AND DUBAI LAYERS

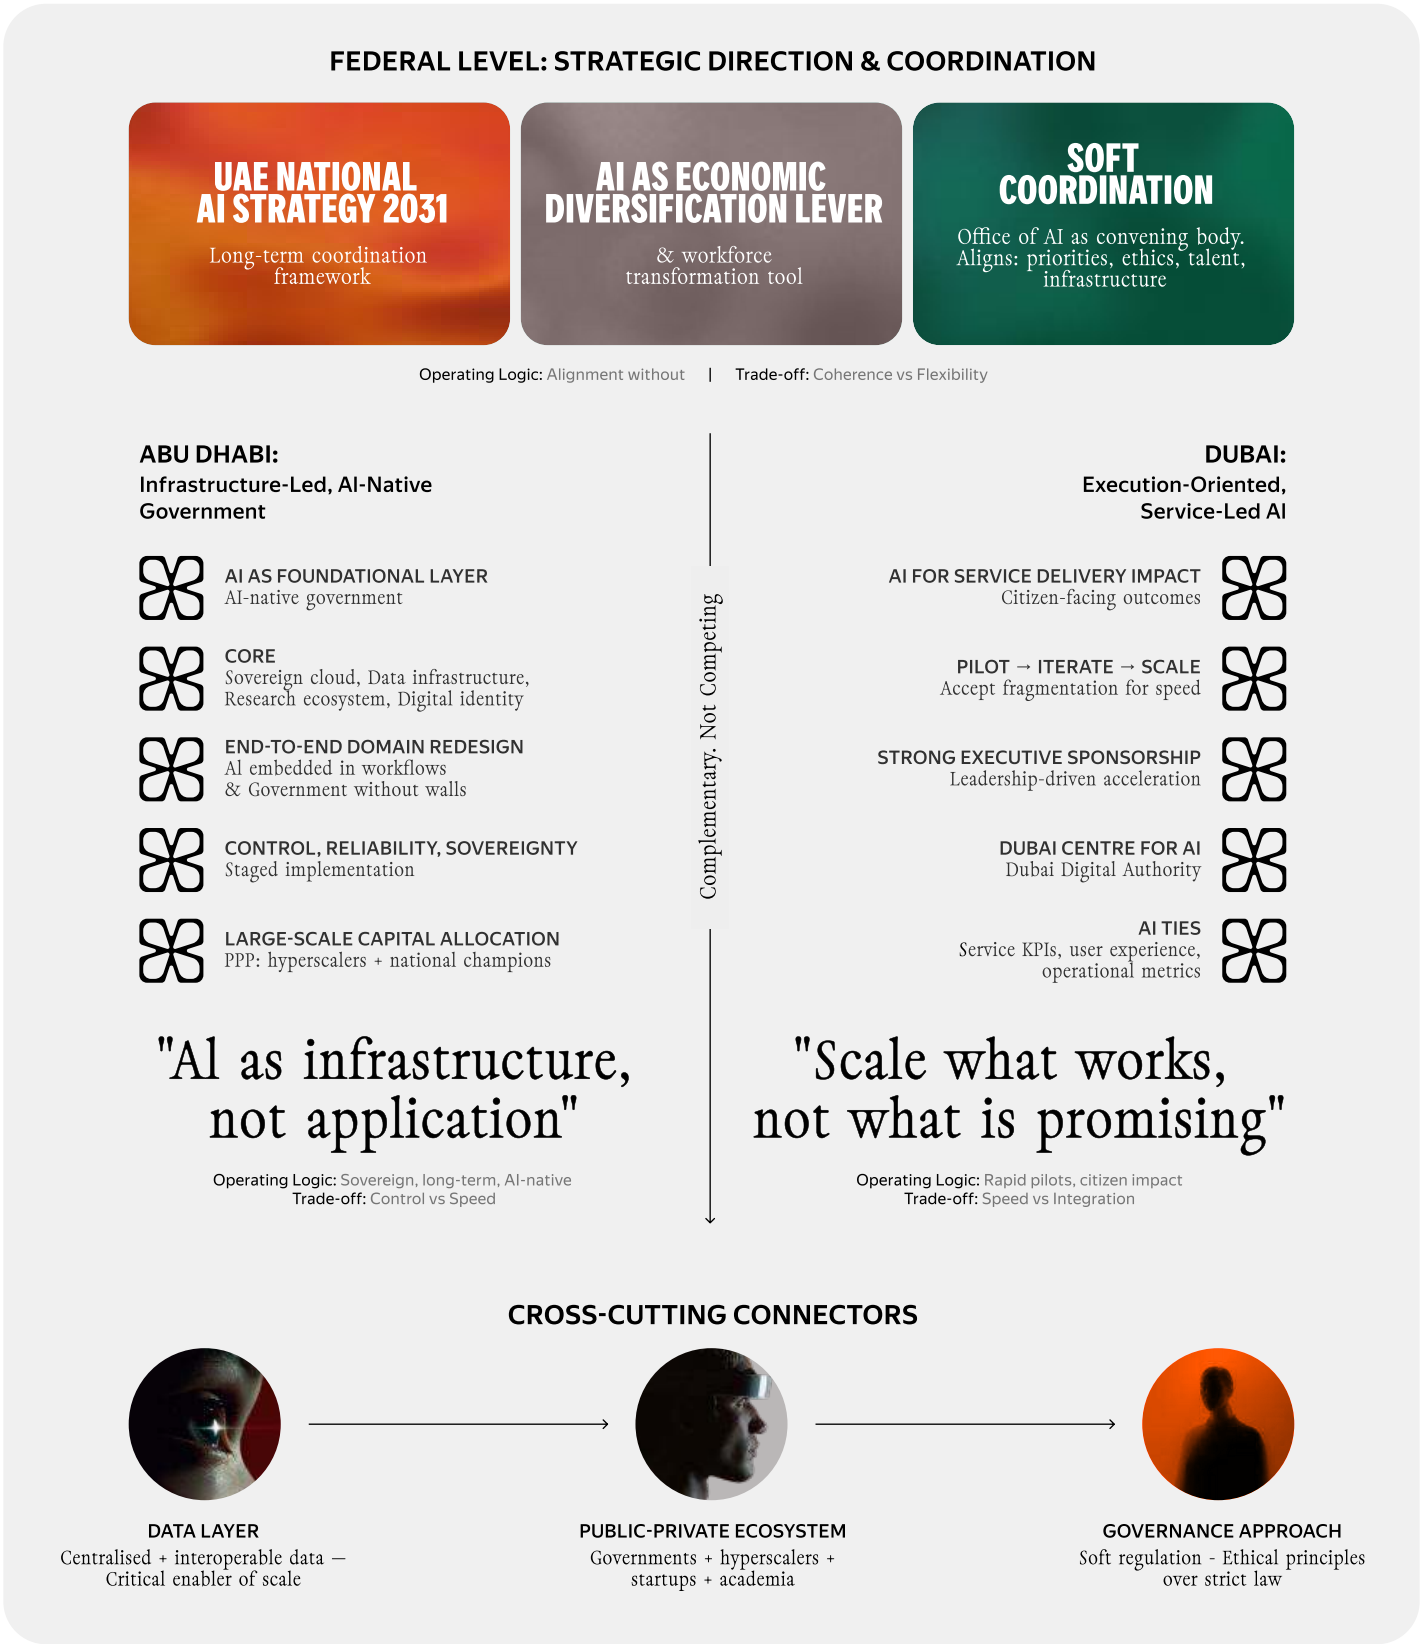

Image 1: The UAE AI Ecosystem

Leadership density and coordinated differentiation

Across federal and emirate levels, the United Arab Emirates exhibits a notable density of senior leadership roles related to artificial intelligence, data, and digital transformation. The early appointment of the world's first Minister of Artificial Intelligence, H.E. Omar Sultan Al Olama, signalled an explicit intent to elevate AI from a technical or innovation agenda to a matter of executive governance. Since then, Chief AI Officers, digital government leaders, and sectoral executives have operated with overlapping but distinct mandates across institutions.

Interview evidence suggests that this leadership density is not managed through rigid hierarchical control or uniform reporting lines. Instead, coordination relies on frequent cross-entity interaction, shared strategic reference frameworks, and informal escalation channels that allow issues to be surfaced and resolved quickly. This approach appears to mitigate fragmentation risks typically associated with distributed ownership, while preserving sufficient autonomy for entities to adapt AI initiatives to their operational contexts.

SELECTION AND PRIORITISATION OF AI USE CASES

While pipeline formation determines how AI initiatives enter the system, prioritisation determines which initiatives receive sustained attention, resources, and a pathway to scale. Across interviews, prioritisation in the UAE is not described as a competition between projects, nor as an open-ended innovation exercise. Instead, it is treated as an institutional filtering process, shaped by federal steering instruments, emirate-level operating models, and differentiated risk tolerance.

Interview evidence points to a consistent set of criteria used — explicitly or implicitly — to judge whether an AI use case is "strategic." These include: expected public value and service impact; data readiness and integration cost; legal, ethical, and reputational risk; scalability and cross-entity reuse potential; and alignment with long-term infrastructure and sovereignty investments. How these criteria are weighted, however, varies across governance levels and emirates, producing distinct prioritisation logics rather than a single national template.

Federal: strategic framing and soft law steering

At the federal level, prioritisation operates primarily through strategic framing and soft-law steering, rather than through direct project selection. Strategic priorities articulated in the National AI Strategy 2031 define where AI should matter — government transformation, priority sectors, ecosystem growth, and ethical governance — without prescribing a fixed portfolio of initiatives.

The UAE Office of Artificial Intelligence, Digital Economy and Remote Work Applications operationalises this framing through instruments such as AI ethics guidelines, AI maturity and readiness self-assessment tools, and an AI Adoption Guideline in Government Services. Together, these tools function as early-stage prioritisation filters, encouraging entities to assess alignment with national priorities, governance expectations, data readiness, and reuse of shared platforms before initiatives progress.

Abu Dhabi: AI-native, platform-centric pipelines

In Abu Dhabi, prioritisation is shaped by a governance logic that places system integration, reuse, and long-term control ahead of rapid expansion in the number of initiatives. Pipeline design is tightly coupled to the emirate's ambition to build an AI-native government and to avoid the proliferation of uncoordinated AI stacks across entities. Officials describe a two-phase journey: initially, each entity identified efficiency and experience-oriented use cases under leadership KPIs that emphasised the number of initiatives; subsequently, the Department of Government Enablement overlaid an AI-native blueprint that treats entire domains as "white spaces" and designs agentic systems end-to-end, with human roles added only where technology falls short or risks are high. Proposed initiatives are now assessed both on immediate value and on their ability to be embedded in shared platforms and infrastructure — particularly TAMM, which has evolved into an AI-enabled "crown jewel" hosting more than a thousand services and modular containers reused across government-to-government and government-to-business channels. Abu Dhabi is in the process of embedding governance through a new AI policy, AI-nativity scale and maturity benchmarks, and an 'AI-in-a-box' stack with built-in technical and ethical guardrails, alongside validation agents and targeted human-in-the-loop oversight in high-stakes domains such as health and large

infrastructure investments. Taken together with the stated plan to prepare a sovereign deployment path in parallel for any promising solution trialled on external infrastructure, these instruments are intended to ensure that, as the model matures, only use cases that can meet Abu Dhabi's requirements on sovereignty, governed data sharing, and risk control will be scaled across government.

Dubai: impact-driven, experimentation-led pipelines

By contrast, prioritisation in Dubai places greater weight on speed, visibility, and measurable service impact, with pipelines explicitly designed to surface and test high-potential use cases quickly. Rather than anchoring prioritisation in platform consolidation, Dubai's approach emphasises orchestrated experimentation, supported by structured governance mechanisms to manage risk and scale success. Through the Dubai Centre for Artificial Intelligence (DCAI), launched in 2023, 33 government entities collectively identified operational challenges and strategic opportunities suitable for AI, generating a repository of 183 potential generative-AI applications that reflected both service needs and the emirate's wider innovation agenda. These candidates' use cases then passed through successive filters - assessment of technical feasibility and data readiness, alignment with Dubai's strategic priorities, and expected, measurable impact on citizen services, efficiency, and resilience - narrowing the pipeline to 75 pilots and ultimately to 15 high-impact exemplars selected for scaling.

According to interview testimony, departments are asked a simple question — "If this went live tomorrow, how would it impact people's lives?"— and use cases are scored on citizen impact, financial pain points, and cross entity scalability, with examples ranging from dynamic ambulance repositioning to customs image classification and early detection of depression to support clinicians. Governance design shapes how these ideas move through the pipeline: a "strategy in a box" playbook gives Chief AI Officers a standard path from idea to pilot (leadership buy in, stakeholder mapping, risk and KPI design), while an AI acceleration task force brings together finance, data-security, infrastructure, and strategy decision-makers to resolve blockers such as budget or data-sharing constraints in real time. Instruments such as the Dubai Future Accelerators sandbox and the Dubai AI Seal further ensure that promising pilots can access external innovation and move toward scale under vetted providers and agreed ethical criteria, making public-private partnerships and procurement design integral components of the pipeline rather than ex-post constraints.

Across both emirates, interview evidence suggests a deliberate shift away from unstructured experimentation toward curated portfolios of AI use cases that can be governed, funded, and sustained over time. Increasingly, initiatives are expected to demonstrate not only immediate efficiency or experience gains, but also clear alignment with data readiness constraints, cross-entity reuse, sovereignty requirements, and long-term operating models. In this sense, use-case prioritisation functions as a mechanism for translating ambition into institutional capacity, directly addressing the "pilot trap" and adoption-to-value gap highlighted in the global literature.

DUBAI DCAI – 15 AI USE CASES IN GOVERNMENT

Nº	USE CASE TITLE	DOMAIN & SUMMARY
1	AI in Customer Services	Chatbots, virtual assistants, and avatars streamline citizen interactions
2	Personalized Healthcare	Genomic and lifestyle data used for tailored diagnosis and treatment
3	Empowering Inclusion	Digital avatars and speech-to-sign translation support people of determination
4	Transforming Information Access	Automated cataloguing, metadata, and personalised recommendations in libraries
5	Smart Urbanism	AI for urban planning, sustainable architecture, and resource optimisation
6	Transforming Lost & Found	Image recognition and ML speed up retrieval of lost items in airports/police
7	AI Text Recognition for Ancient Manuscripts	OCR digitises and preserves complex Arabic-script manuscripts
8	AI in Real Estate	Predictive analytics for valuations, transactions, and market insights
9	AI in Logistics and Customs Security	Intelligent surveillance and risk detection in border and trade operations
10	AI-Powered Workforce Management	Algorithms for recruitment, training, and adaptive workforce allocation
11	AI-Powered Investment Tools	Predictive analytics to guide FDI and economic competitiveness
12	AI for Spend Management	Automated compliance, anomaly detection, and budget optimisation
13	AI in Advertising	AI-driven content creation and targeted public communication campaigns
14	AI for Digital Historical Reconstructions	VR/AR and 3D modelling to recreate heritage sites for culture and education
15	AI Chatbots for Government Transactions	Unified service chatbots for forms, payments, and information queries

Table 2: Dubai DCAI:15 most common government AI use cases

AI IN ACTION: GOVERNMENT SERVICES & SMART CITIES IN THE UAE

AI deployment in the UAE spans a wide spectrum — from citizen-facing digital assistants to city scale infrastructure systems in mobility, safety, and utilities

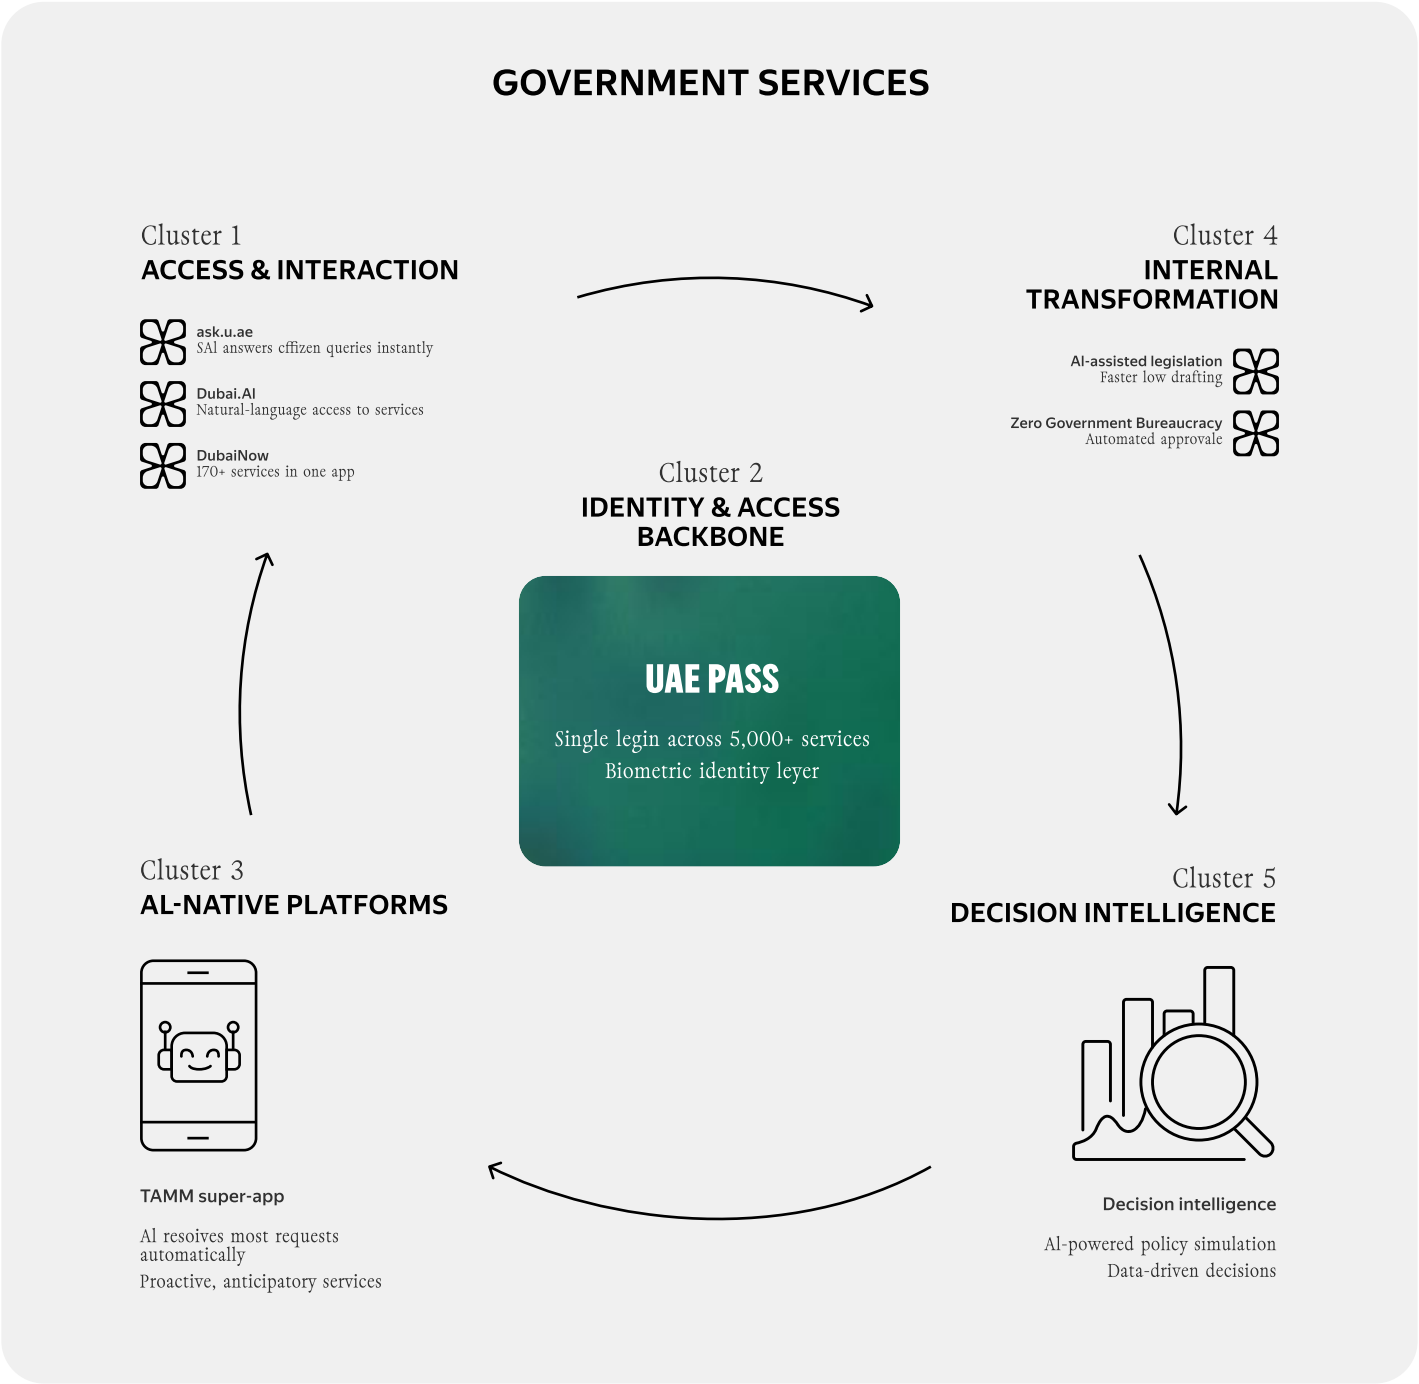

Image 2.1: AI in Action: Government Services & Smart Cities in the UAE

AI IN ACTION: GOVERNMENT SERVICES & SMART CITIES IN THE UAE

From Interaction ⟶ Infrastructure

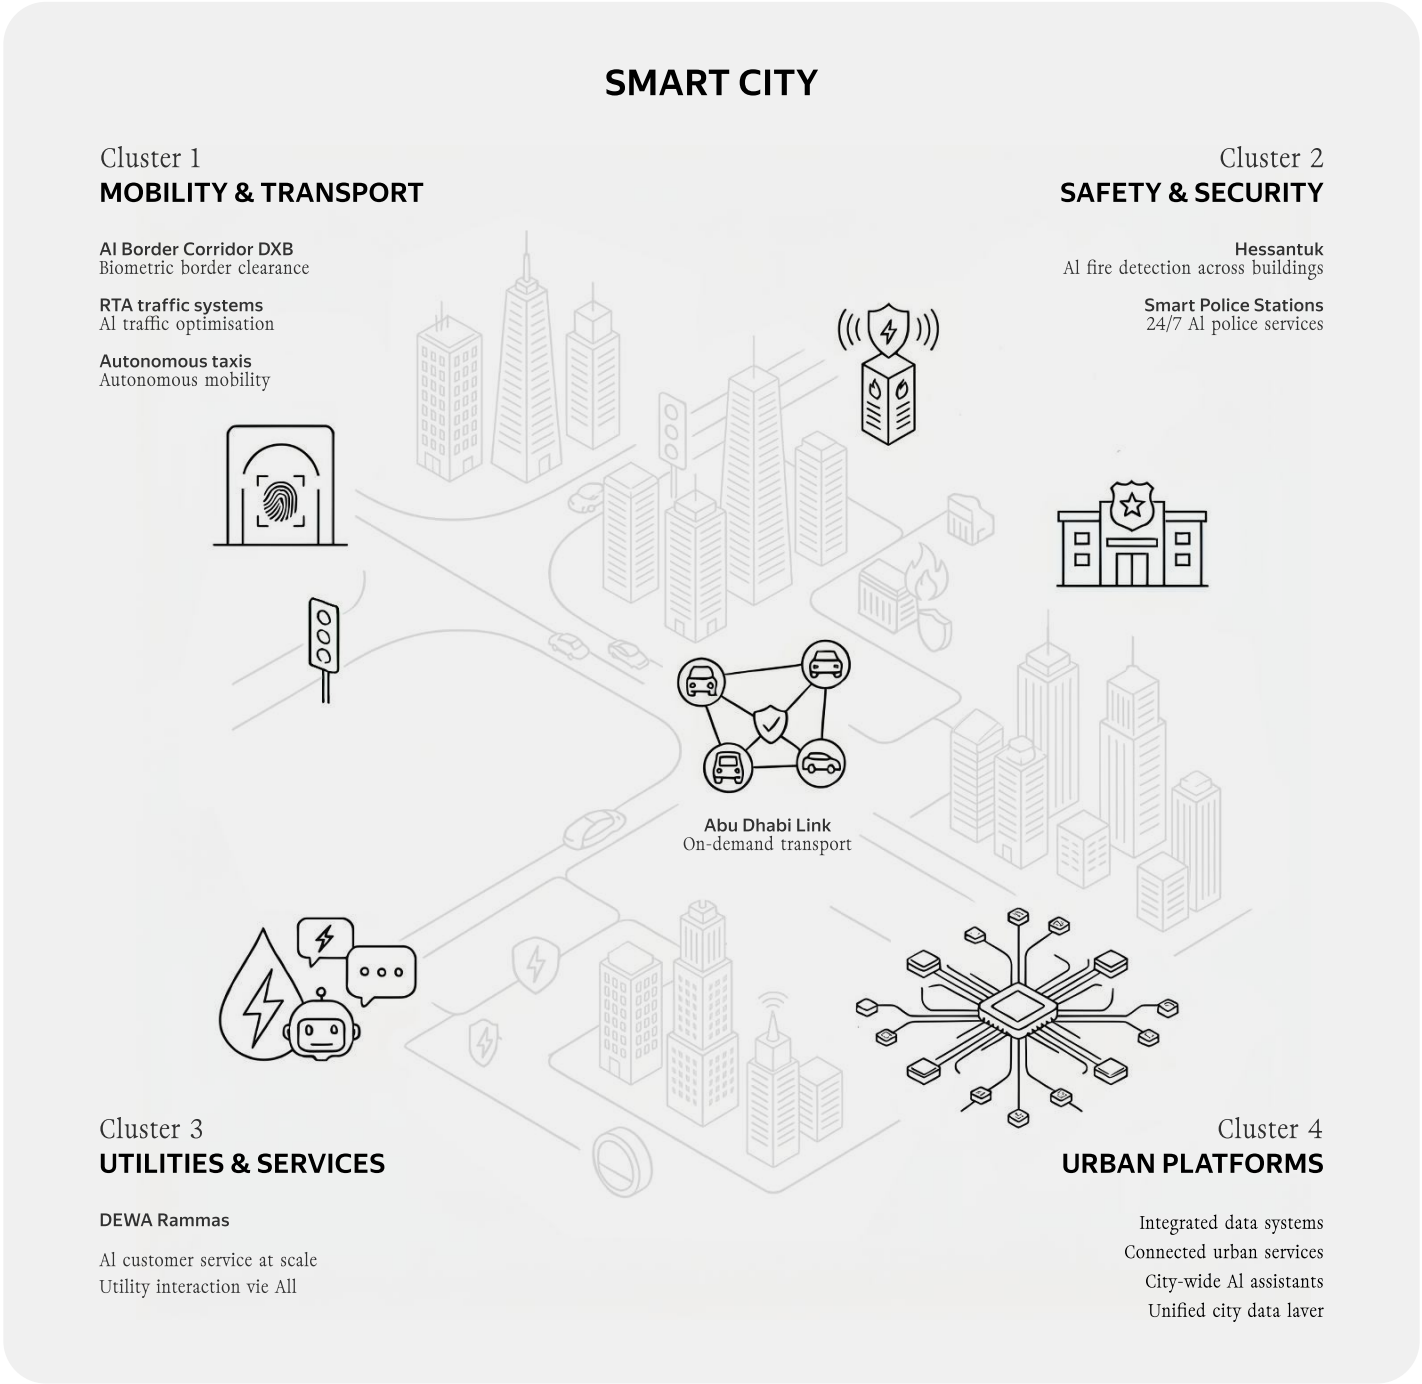

Image 2.2: AI in Action: Government Services & Smart Cities in the UAE

EXECUTION CHALLENGES AND SUCCESS FACTORS

As in other jurisdictions, the expansion of AI initiatives across the United Arab Emirates has made execution - not technological ambition — the decisive bottleneck. Interview evidence confirms that many of the challenges identified in the global literature are present in the UAE: fragmented data environments, capability gaps at the intersection of policy and technology, organisational inertia, infrastructure and cost constraints, and uncertainty around governance and accountability. What distinguishes the UAE case is not the absence of these constraints, but the ways in which they are surfaced, prioritised, and addressed within a multi-level governance system.

This section examines execution through the same analytical lens as the global baseline, focusing on structural constraints rather than individual projects or technologies. It treats execution as an institutional challenge embedded in data architectures, operating models, talent systems, and governance arrangements. Rather than assessing outcomes or performance, the analysis asks how execution conditions are shaped, and how different governance choices enable or constrain progress over time.

Across subsections, a consistent internal logic is applied. Each execution challenge is examined in four steps: first, the core constraint as identified in the global literature; second, the factors that make this constraint particularly salient in the UAE context; third, the contrasting ways in which Abu Dhabi and Dubai seek to address it, reflecting different operating models and risk tolerances; and finally, the residual tensions and trade-offs that remain unresolved. This structure allows the UAE case to be read not as an exception to global patterns, but as a practical exploration of how governments attempt to work through the execution challenges that the literature consistently identifies.

Taken together, the sections that follow show that execution success in the UAE is less a function of technological sophistication than of institutional readiness. Progress depends on how deliberately governments invest in data foundations, distribute capability, redesign organisations, manage cost and infrastructure constraints, and adapt governance mechanisms as systems move from experimentation toward sustained operation.

Data quality, interoperability, and governance

Across interviews, data governance — rather than model capability — emerges as the primary execution bottleneck for AI deployment in the UAE. AI initiatives rarely fail because algorithms cannot perform; they stall because data is fragmented across entities, governed under incompatible rules, or too sensitive to be accessed in live environments. Practitioners repeatedly emphasised that while technical solutions can be built in weeks, deploying them securely in government contexts can take months due to unresolved data-access, classification, and security constraints. In this sense, data problems surface not as technical glitches, but as institutional barriers embedded in public-sector operating models.

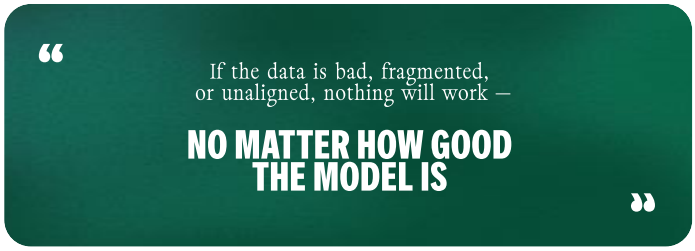

Interviewees also identified unstructured content as the most difficult and consequential data challenge. Estimates suggest that grand majority — 80, 90, or more percent of organisational data falls into this category (IBM, Deloitte, IDC). Making such “dark data” usable for AI requires slow, labour-intensive work on business glossaries, shared definitions, metadata standards, and document curation. While generative AI and large language models are valued for their ability to search, index, and reason over unstructured material, several interviewees cautioned that without prior governance and tagging, these tools simply accelerate confusion rather than create value. In practice, unstructured data governance was described as foundational to any serious AI programme, even though it attracts far less attention than model development.

Abu Dhabi: sovereignty-led integration as a precondition for scale

In Abu Dhabi, data execution challenges are addressed as part of the emirate’s broader AI-native and sovereign-AI strategy. Interviewees highlighted the “inherent compartmentalisation of data ownership,” noting that resistance to sharing is often rooted in institutional control.

Digital sovereignty is operationalised through a layered framework — distinguishing data sovereignty, operational sovereignty, and technology sovereignty — and through a national classification system that determines whether data and AI workloads can be deployed on standard public cloud, sovereign public cloud, or fully private or air-gapped environments. Shared stacks such as 'AI in a box' and observability tools like '42 insight' are intended to make it easier for entities to apply national sovereignty and security rules in practice, by embedding common ethical and technical guardrails and giving CIOs visibility into where data resides, how it is encrypted, and whether it aligns with the National Cloud Security Policy and classification scheme. Interview evidence points to a deliberate sequencing strategy in which experimentation happens first on non sovereign infrastructure using synthetic or non sensitive data, while plans for a sovereign setup are prepared in parallel so that, once a concept is proven, sovereignty becomes a mandate for scale; this operationalises the view that data sovereignty is 'not a choice' while avoiding the cost and rigidity of treating all workloads as 'secret' by default. Centralising AI on shared, sovereignty aware stacks can make it harder for individual entities to move fast in the short term, but it improves control, auditability, and confidence when scaling AI in sensitive domains, because data location, security, and guardrails are managed consistently rather than on a project by project basis.

Dubai: pragmatic experimentation with downstream consolidation

In Dubai, data constraints tend to surface through pilots and service-level initiatives rather than upfront architectural design. Interviewees described projects that stalled despite available budgets and technical capability because finance, security, and data owners had not been aligned early enough. In response, organisational mechanisms have been introduced to resolve data-access and governance issues more dynamically. An AI Acceleration Task Force brings together security, infrastructure, finance, and strategy stakeholders to address blockers in real time, while a strategy-in-a-box playbook requires AI leaders to anticipate data and governance risks before committing to delivery.

Talent and capability constraints

Across the project, interviews converge on a clear finding: talent and capability constitute a binding execution constraint — but not in the form most commonly assumed. The primary challenge is not the availability of engineers or access to advanced technical skills. Rather, it lies in the scarcity — and misallocation of hybrid profiles able to operate at the intersection of technology, public policy, sovereignty requirements, and service design.

More specifically, the most acute bottleneck is not generic "AI talent" but what several interviewees described as "translators": individuals who combine deep domain understanding with sufficient grasp of AI capabilities and limitations to shape, prioritise, and sequence transformation initiatives. This gap is compounded by a second, distinct deficiency: the lack of "orchestrators" — programme leaders capable of coordinating complex, cross-functional AI efforts and driving them through to execution.

Taken together, these gaps explain a recurring pattern: AI execution falters less because of models or compute, and more because of people, authority, and organisational design.

Where execution breaks down

These constraints become particularly visible at key points in the execution chain. At the leadership level, decision-making is often slowed by limited familiarity with core frameworks such as data classification or national cloud policies. Faced with complex trade-offs, senior leaders may default to overly restrictive positions — treating all data as highly sensitive, for instance — not because of formal regulatory barriers, but due to uncertainty and risk aversion. Execution is therefore delayed at the point where strategic decisions need to be made.

At the organisational level, the issue is less one of individual capability than of role design and accountability. Dual-hatted Chief AI Officers without dedicated budgets or mandates frequently struggle to move beyond coordination, while ownership of adoption and scaling remains diffuse. By contrast, the presence of fully empowered AI leaders — with autonomy, protected time, and budgetary control — consistently emerges as a critical enabler. Capability, in this sense, is inseparable from governance.

A further point of failure lies in the persistent divide between those who build AI systems and those expected to use them. Interviewees described repeated cases where technically sound tools failed to gain traction because they were not embedded in workflows or co-designed with users. Bridging this divide requires embedding AI expertise directly within service delivery and policy teams, rather than positioning it as a separate technical function.

Capability as an architecture problem

These observations point to a deeper conclusion: the constraint is not primarily a hiring problem, but a capability architecture problem. It is generally easier to upskill strong business or strategy profiles in AI than to teach technical specialists the complexities of sovereign or policy-heavy environments. The critical task, therefore, is to design and position roles — translators, orchestrators, and empowered domain leaders — in ways that align expertise with authority and accountability.

Most organisations are expected to follow a staged evolution, in which these capabilities are initially concentrated in central units to build momentum, before progressively embedding into functions as AI becomes part of the operating model.

In response, governments are beginning to reshape how capability is built and deployed. Operating models are shifting towards hybrid human – AI teams, with cross-functional collaboration replacing the traditional separation between IT and business. At the same time, system-wide efforts to raise AI literacy — through academies, leadership programmes, and structured playbooks — aim to create a shared baseline of understanding and reduce reliance on ad hoc judgement.

Partnership models are also evolving. Rather than relying on transactional outsourcing, organisations increasingly favour embedded “one-team” approaches, with a stronger emphasis on knowledge transfer and internal capability development to avoid long-term dependency.

Despite these efforts, a structural trade-off remains. Reliance on external partners enables rapid experimentation and early wins but risks weakening institutional ownership, particularly in sovereign or high-stakes contexts. Conversely, investing in internal capability — through embedded teams, empowered roles, and stronger governance — slows initial delivery but is widely seen as necessary for long-term control, accountability, and value realisation.

Taken together, the evidence reframes talent and capability not as a supply constraint in the labour market, but as a design problem for the state. The UAE has access to engineers, vendors, and advanced tools; the more difficult challenge lies in embedding hybrid skills, authority, and institutional memory within government itself. Without this, AI remains a series of pilots. With it, it becomes an operating system.

Organisational change, operating models, and accountability

The global literature consistently identifies organisational change as one of the most persistent barriers to effective AI adoption in government. While AI systems may perform as intended, execution often falters when they encounter established workflows, decision rights, and accountability structures. Public administrations are designed around clear hierarchies, procedural compliance, and risk containment. AI systems — particularly those that are agentic, cross-functional, or continuously operating — challenge these arrangements by redistributing tasks, compressing decision cycles, and blurring traditional boundaries of responsibility.

> “AI IS NOT AN ADD-ON TO GOVERNMENT PROCESSES
> it becomes the operating system through which the state thinks and decides”

Interview evidence suggests that organisational friction is especially pronounced in the UAE because AI ambitions extend beyond incremental automation toward reconfiguring how government operates. As AI systems increasingly move from analysis and recommendation to orchestration and action initiation, questions of ownership, liability, and auditability become unavoidable. Interviewees emphasised that hesitation around AI deployment rarely reflects opposition to technology itself; instead, it stems from uncertainty about who is accountable

when decisions are distributed across humans and machines, and how such decisions can be justified within existing administrative and legal frameworks. Insights from the Machines Can Think discussions reinforce this point: trust and clarity of ownership, rather than model accuracy, emerge as the limiting factors once AI enters live workflows.

Abu Dhabi: domain-level redesign and platform-led accountability

In Abu Dhabi, organisational change is pursued through deliberate, system-wide redesign led by the Department of Government Enablement (DGE). Interview evidence suggests a clear shift from a fragmented, use-case-driven approach toward domain-level transformation, in which AI is used to reconfigure entire operational areas end to end.

Early efforts focused on building a large portfolio of AI use cases, but this approach proved insufficient to capture meaningful value. As one interviewee noted, "patching AI" onto existing processes rarely delivers structural impact. In response, Abu Dhabi adopted a different logic: prioritising a limited number of domains — such as human capital, compliance, or investment for Sovereign Wealth Fund Mubadala — and redesigning them as coherent, AI-enabled systems. These domains are treated as integrated journeys, with AI embedded across the full lifecycle rather than applied to isolated tasks.

This shift is reflected in the design approach. Transformation begins with structured design sessions in which domain teams define an unconstrained "ideal" end state - reimagining entire workflows (for example, the full hire-to-retire cycle in HR) before mapping pain points, user needs, and AI interventions across each stage. Execution then follows a staged roadmap: a two- to three-year target architecture is defined upfront, but delivery is sequenced through a small number of high-impact, high-readiness building blocks, prioritised using feasibility–impact criteria and constrained by data readiness, budget, and delivery capacity.

Within this model, AI is positioned not as a tool but as an organising layer. Target operating models assume that AI systems perform orchestration, monitoring, and decision preparation by default, with human roles inserted selectively where judgment, risk mitigation, or accountability are required. This logic underpins constructs such as hybrid AI–human pods, where civil servants and AI agents operate as a single unit with clearly defined responsibilities.

Accountability is embedded at the workflow level rather than the task level. Validation mechanisms surface confidence and risk indicators, escalation paths are predefined, and high-stakes decisions remain explicitly human-in-the-loop. The objective is to make responsibility auditable and defensible even as systems become more autonomous.

Dubai: workflow embedding and incremental organisational adaptation

In Dubai, interviewees described a more incremental approach to organisational change, in which AI is embedded into existing workflows and services, and organisational adaptation follows demonstrated value. Rather than upfront domain redesign, AI initiatives are often deployed within current structures, with operating models evolving as friction points emerge. This approach prioritises speed, service impact, and learning-by-doing, allowing departments to test how AI interacts with real operational constraints before committing to deeper reorganisation.

Governance mechanisms in Dubai focus on coordination and unblocking rather than structural overhaul. Cross-functional acceleration task forces bring together policy, operations, finance, security, and digital teams to resolve accountability and decision bottlenecks in real time. Playbooks and shared guidelines standardise how initiatives move from concept to deployment, while leaving organisational forms relatively flexible. Interviewees noted that this model preserves continuity and buy-in, but can also result in parallel processes if successful AI systems are not subsequently integrated into core operating models.

Across both emirates, interviews highlight unresolved tensions between coherence and adaptability. Abu Dhabi's platform-led, domain-level redesign strengthens accountability and auditability but requires sustained leadership commitment and tolerance for disruption. Dubai's incremental embedding lowers organisational resistance and accelerates uptake, but places greater demands on later consolidation to avoid fragmentation. More broadly, the transition from human-in-the-loop to more supervisory, on-the-loop models raises difficult questions about liability, escalation, and public legitimacy that neither approach fully resolves.

Taken together, the UAE experience illustrates that organisational change is not a secondary implementation issue, but a central design choice shaping how AI systems are trusted, governed, and sustained. Differences across emirates reflect alternative pathways for aligning workflows, authority, and accountability with increasingly capable AI systems, rather than divergent levels of ambition or maturity.

Cost, infrastructure, and legacy constraints

Across jurisdictions, cost, infrastructure, and legacy systems operate as structural constraints on government AI execution, not secondary implementation details. The practical barrier is rarely "lack of compute" in the abstract; rather, it is the interaction between infrastructure choices, security and compliance requirements, and institutional readiness that determines whether AI systems can move from experimentation to sustained deployment. As AI capabilities become more agentic and integrated into live services, governments face a shift from one-off project costs toward ongoing operating commitments — including secure environments, monitoring and observability, integration layers, and lifecycle maintenance — that are often underestimated at the pilot stage.

In the UAE, the salience of this constraint is heightened by ambitious targets, high expectations of reliability and service continuity, and an explicit emphasis on sovereignty. Interviews and the literature suggest that the UAE's agenda is enabled by substantial investment in digital infrastructure, yet execution frictions persist where legacy IT, interoperability gaps, and security constraints complicate system-wide deployment. Several interviewees underline that compute is "almost never" the primary bottleneck; the harder problem is secure access to data, especially in sensitive government-adjacent environments where role-based access control, classification rules, and "need-to-know" boundaries increase both the complexity and the cost of AI deployment. In practice, building a model-enabled solution can be relatively fast; integrating it into production-grade, securely governed environments — compatible with classification requirements and legacy systems — can take months, and becomes the dominant driver of time and cost.

Abu Dhabi: infrastructure-first and sovereign capacity as an 'intelligence grid'

In Abu Dhabi, interviews describe an infrastructure-first approach that treats sovereign AI capacity as a foundational layer — an "always-on" capability that must be invisibly available to host sensitive workloads under clear data-classification rules. This logic reframes "sovereign AI" as an infrastructure problem: building shared sovereign cloud and compute capacity, standardised workload placement mechanisms, and monitoring/observability tools so that entities can deploy AI within controlled environments without bespoke, project-by-project arrangements. The strategic intent is to internalise lifecycle costs - security hardening, redundancy, monitoring, compliance - within shared platforms rather than pushing them entirely onto individual entities or vendors.

> "AI INFRASTRUCTURE SHOULD BECOME AS INVISIBLE AND RELIABLE AS ELECTRICITY"
>
> always on, always trusted, and not something users think about

This approach can reduce duplication and strengthen predictability in the long run, but it also introduces a coordination requirement: entities must align to shared stacks, standards, and classification-driven deployment models. The trade-off is therefore not simply cost versus speed, but short-term flexibility versus long-term sustainability and control.

Dubai: service-led deployment with downstream consolidation pressure. In Dubai, interviews and the literature point to a service-led model where AI is layered into existing digital platforms, portals, and operational command environments, with infrastructure scaling as value and demand become clearer. This approach supports faster experimentation and reduces upfront infrastructure commitments for uncertain use cases. At the same time, it pushes some costs downstream: as pilots proliferate, integration burden, monitoring requirements, and security controls accumulate across a growing portfolio of services. Over time, this creates pressure for consolidation — standardised integration patterns, shared monitoring and compliance tooling, and clearer platform boundaries — so that operating costs do not escalate through fragmentation and duplicated technical stacks.

Across both emirates, the interviews point to two persistent tensions. First, governments must manage a shift from capital-style digital project investment to recurring operating expenditure for cloud services, data platforms, and AI workloads — requiring new budgeting disciplines and stronger visibility into total cost of ownership. Several interviewees explicitly invoke a FinOps/TCO mindset: infrastructure should be sized for the organisation and shared services, not repeatedly rebuilt around each use case, to avoid duplicated demand for high-cost resources and inconsistent security practices. Second, sovereignty and classification choices can unintentionally inflate costs: treating all workloads as maximally sensitive and defaulting to private environments may protect risk at the margin, but can sharply increase operating cost and limit access to advanced capabilities that could otherwise be deployed under differentiated workload classifications.

Taken together, the UAE evidence reinforces a broader point from the literature: cost and infrastructure are not merely enabling conditions; they are steering mechanisms that shape what is feasible, what is scalable, and what is sustainable. In this context, the UAE illustrates how differing pathways — shared sovereign foundations versus service-led layering — each attempt to reconcile sovereignty, reliability, and cost transparency with the practical realities of legacy systems and governed deployment.

Governance, regulation, and trust

Across jurisdictions, governance and trust consistently emerge as binding constraints on government AI execution. While regulatory frameworks, ethical principles, and risk-management toolkits have proliferated, the literature shows that formal rules alone rarely enable deployment. Instead, uncertainty around accountability, explainability, and legal defensibility often slows or stalls implementation — particularly in high-stakes public-sector domains.

A recurring finding in the global literature is that governments face a dual risk: under-governance, which erodes public trust and legitimacy, and over-governance, which creates paralysis by treating all AI use cases as equally risky. The execution challenge lies not in defining principles, but in translating them into operational decision rights, escalation paths, and enforceable accountability embedded within workflows.

In the UAE, governance and trust carry heightened importance for three reasons. First, AI systems are increasingly deployed in core public services — from licensing and welfare administration to health, policing, and infrastructure oversight — where errors or opacity can have immediate social and reputational consequences. Second, the country's ambition to position itself as a global AI testbed amplifies external scrutiny around ethical use, sovereignty, and data protection. Third, the use of advanced and increasingly agentic systems intensifies questions of responsibility when decisions are automated or partially autonomous.

> **TECHNOLOGY IS NEVER THE PROBLEM**
>
> everything comes down to people and how you orchestrate influence across them

Interviewees repeatedly stress that trust is not assumed; it must be actively designed into systems. Governance is therefore treated less as a compliance overlay and more as a prerequisite for scale - without credible safeguards, promising pilots cannot be authorised to operate at population level.

How governance is operationalised: contrasting but aligned approaches

Rather than relying on a single, centralised regulatory model, the UAE exhibits multi-layered governance, combining federal principles with emirate-level operationalisation

At the federal level, governance is framed through national AI ethics principles, risk-based guidance, and alignment with international standards. These instruments establish a shared baseline on fairness, transparency, accountability, and human oversight, while deliberately avoiding prescriptive rules that would constrain experimentation. Federal guidance functions as soft law: it sets expectations and reference points, but leaves implementation choices to operating entities.

In Abu Dhabi, governance is tightly integrated into system architecture and operating models. Interviewees describe an approach in which policy, technology, and oversight are co-designed. AI systems are expected to comply with data-classification rules, national cloud security policies, and auditability

requirements by design, rather than through ex post controls. Governance mechanisms such as validation agents, observability layers, and structured human-in-the-loop checkpoints are embedded directly into AI-native platforms, particularly for health, infrastructure, and other high-risk domains. Accountability is therefore treated as a system property, not merely a managerial responsibility.

In Dubai, governance is more explicitly tied to use-case pipelines and execution processes. Rather than enforcing uniform controls upfront, governance is applied progressively as initiatives move from experimentation to scale. Mechanisms such as ethics reviews, AI seals, sandbox participation, and cross-functional acceleration task forces allow risks to be surfaced and addressed early, while preserving speed. Interviewees emphasise that trust is built through visible service impact combined with clear guardrails, ensuring that governance does not become detached from citizen-facing outcomes.

Across both emirates, a shared principle is evident: governance follows risk, not technology category. Low-risk monitoring and back-office applications may operate with lighter oversight, while systems affecting rights, entitlements, or safety trigger tighter controls and explicit human accountability.

Despite these advances, interviewees highlight several unresolved tensions. One concerns the boundary between human-in-the-loop and human-on-the-loop oversight. While current practice favours active human review for high-stakes decisions, the growing use of agentic systems raises questions about scalability, fatigue, and real-world effectiveness of continuous human supervision.

A second tension lies between legal certainty and adaptive governance. As AI systems evolve through updates, retraining, or repurposing across domains, responsibility becomes harder to assign, particularly when multiple entities share platforms or data. Interviewees note that accountability frameworks often lag behind technical realities, creating grey zones rather than clear liability.

Finally, there is an inherent trade-off between trust through caution and trust through performance. Overly restrictive governance can delay benefits and undermine confidence in government capability, while rapid deployment without visible safeguards risks public backlash. The UAE's experience suggests that trust is not maximised by either extreme, but by credible, transparent mechanisms that allow systems to be challenged, corrected, and, if necessary, withdrawn.

Taken together, governance, regulation, and trust in the UAE function not as static constraints but as dynamic execution levers. The central challenge is no longer defining ethical principles, but embedding accountability, explainability, and human judgment into the daily operation of AI systems. How effectively governments resolve this challenge will largely determine whether AI remains confined to controlled pilots - or becomes a durable, trusted component of public-sector decision-making at scale.

> **SCALE WHAT IS PROVEN**
>
> not what is promising

PROCUREMENT PROCESS & ECOSYSTEM STRATEGY

In most public-sector contexts, procurement is treated as a downstream administrative function — focused on compliance, vendor selection, and cost control. In the UAE, however, it has progressively become a strategic layer of AI capability, shaping not only how solutions are acquired but what can realistically be built, scaled, and governed.

This shift reflects a structural feature of AI systems: their development cycles are iterative, data-dependent, and tightly coupled with organisational workflows. As a result, traditional procurement models - designed for well-specified, static solutions - struggle to accommodate the uncertainty and speed inherent to AI deployment. This misalignment is visible even in the UAE, where interviewees consistently pointed to procurement as one of the main bottlenecks in otherwise fast-moving environments. One executive summarised the imbalance starkly: projects can take "three months to contract and three months to build", highlighting how administrative timelines can rival or exceed technical delivery cycles.

> **ORGANISATIONAL READINESS**
>
> not compute or models — is the primary bottleneck to AI deployment

Earlier phases of Dubai's AI push were similarly constrained by procurement rules that limited engagement to pre-approved or locally registered vendors, effectively setting a "ceiling on innovation." This has led to the emergence of parallel procurement pathways designed to reconcile experimentation with compliance. Initiatives such as the Dubai Future Accelerators illustrate this logic: rather than specifying solutions upfront, government entities define problems and engage a broad pool of global firms to co-develop prototypes. Successful pilots can then transition into formal contracts, effectively reversing the traditional sequence of procurement. In this model, procurement becomes less about selecting a predefined solution and more about discovering, validating, and scaling viable approaches.

At the same time, the role of external partners is being redefined in Abu Dhabi. Instead of acting as vendors delivering fixed outputs, firms are increasingly integrated into co-creation arrangements, working alongside government teams to design and implement solutions iteratively. This reflects a broader transition from transactional procurement to ecosystem orchestration, where the state curates a network of partners, capabilities, and platforms rather than simply contracting for services. Interviewees repeatedly emphasised that successful projects depend on shared ownership, continuous interaction, and the presence of internal "receivers" capable of taking over and scaling solutions — conditions that cannot be achieved through arm's-length vendor relationships alone.

However, this shift introduces new challenges. The rapid expansion of the AI vendors — particularly in hubs such as Dubai — has created a highly crowded and uneven market, where differentiating credible providers from opportunistic entrants becomes increasingly difficult. In response, instruments such as the Dubai AI Seal have emerged as filtering mechanisms, helping government entities navigate this complexity by signalling a baseline level of technical and organisational credibility. Procurement thus takes on an additional function: not only selecting suppliers, but structuring and stabilising the ecosystem in which innovation occurs.

As procurement becomes more central to AI capability, it also becomes the site of key strategic trade-offs. Efforts to accelerate delivery through pilot-to-scale pathways must be balanced against the need for robust governance, auditability, and long-term accountability. Similarly, access to cutting-edge global technologies must be reconciled with national priorities around data sovereignty and operational control, often resulting in hybrid architectures that combine sovereign and external infrastructure. Finally, the flexibility required for iterative development can come into tension with the need for standardisation across entities, creating risks of fragmentation and duplication that central bodies such as the Department of Government Enablement or Dubai Centre for AI seek to mitigate.

Taken together, these dynamics suggest that the UAE's progress in AI is not only a function of leadership or infrastructure, but of its ability to reposition procurement as an instrument of execution. In this model, procurement defines the boundaries of innovation, determines the speed of deployment, and shapes the structure of the ecosystem through which AI systems are built and scaled. At the same time, this approach remains contingent on strong central coordination, sustained leadership alignment, and the careful management of trade-offs between speed, control, and coherence. Its effectiveness in the UAE does not automatically guarantee transferability to more decentralised or regulation-heavy systems, where similar adaptations may prove more difficult to implement.

EVALUATION, IMPACT, AND LEARNING LOOPS

Across the global literature, evaluation is widely identified as a weak link in government AI programmes. While performance frameworks have expanded beyond narrow efficiency metrics to include fairness, transparency, and risk, evaluation only partially feeds back into decision making and often remains a compliance or reporting exercise rather than a steering mechanism. Many governments measure outputs (cost savings, processing time, volumes handled) but struggle to assess broader public value, distributional effects,

or institutional learning, and few have codified criteria for when to continue, redesign, or decommission AI systems. As a result, AI initiatives frequently accumulate in an ad hoc way, with governments at risk of scaling systems whose impacts are poorly understood or, conversely, shelving promising pilots because benefits are not visible through conventional metrics.

In the UAE, evaluation carries particular weight because AI initiatives are deployed at speed, often across multiple entities, and in domains with high visibility and public consequence. Interviewees emphasised that traditional public sector KPIs — focused on throughput, compliance, or service delivery volumes— are insufficient for AI systems that reshape workflows, redistribute decision authority, or introduce automated judgment. Several practitioners noted that early AI deployments were evaluated primarily on technical performance and user satisfaction, which proved useful for pilots but inadequate for governing systems at scale. As AI systems become more embedded and agentic, the question shifts from “Does this work?” to “Is this still the right system, operating in the right way, for this context?”, which requires evaluation frameworks capable of capturing risk, trust, and long term institutional effects rather than just immediate efficiency gains.

How evaluation is evolving in practice

Interview evidence suggests that the UAE is gradually moving toward more multi layered evaluation approaches, even if these remain unevenly institutionalised and are not yet described as a single, formal framework. At the operational level, AI systems are increasingly supported by ongoing monitoring and technical oversight—through performance dashboards, error and exception tracking, and logging of human overrides — to surface drift, bias, or operational stress before failures become externally visible. In higher stakes settings, explainability tools, audit trails, and federal instruments such as ethics self assessments and maturity tools contribute to a traceable decision record and more structured scrutiny of how AI recommendations are used.

At the organisational level, interviewees describe emerging efforts to link evaluation to portfolio and platform decisions rather than judging initiatives in isolation. In Dubai, for example, use case pipelines and “scale what’s proven” disciplines push entities to assess AI systems based on replicability across agencies, contribution to shared platforms, and alignment with long term infrastructure investments, echoing FinOps and total cost of ownership logics. This shifts evaluation away from short term ROI toward strategic fit and sustainability, including whether AI is reinforcing already strong processes or compensating for deeper organisational weaknesses.

At the governance level, evaluation is increasingly linked to risk sensitivity, with higher risk domains subject to more intensive oversight. Federal and emirate level instruments — such as ethical guidelines aligned with OECD principles, Dubai’s AI related data protection rules, and the UAE AI Seal — define expectations for transparency, human oversight, and risk assessment that go beyond simple performance checks. In practice, this means low risk systems can be assessed primarily on reliability and efficiency, while systems affecting rights, safety, or entitlements trigger deeper review, including ethical scrutiny and formal approvals for continued operation, even if a full, cross government risk classification ladder is still emerging rather than fully codified.

Despite these advances, both the literature and interviews highlight several unresolved challenges. First, linking evaluation to authority remains difficult: even when metrics or audits reveal underperformance or rising risk, it is not always clear which actor has the mandate to pause, redesign, or retire a system — especially when platforms or data pipelines are shared across entities. Second, there is a persistent tension between learning and accountability; adaptive evaluation frameworks encourage experimentation and iteration, but public sector environments demand clear responsibility and defensibility, and fear of reputational or political fallout can discourage honest reporting of failures, weakening learning loops. Third, measuring public value remains conceptually and practically difficult: equity, trust, and legitimacy are recognised as critical outcomes in global and UAE debates, yet they resist simple quantification and unfold over long time horizons, raising the risk that evaluation frameworks privilege what is easily measurable over what is institutionally important.

LESSONS FROM THE UAE FOR GLOBAL POLICYMAKERS

The UAE case does not offer a single model to be replicated. Instead, it provides a set of practical resolutions to institutional challenges that the global

literature consistently identifies but rarely observes at scale. The value of the case lies less in specific technologies than in how strategic intent, governance choices, and operating models are aligned under real-world constraints.

> **AI FAILS AT THE BOUNDARY OF A DECISION**
>
> without trust, ownership, and explainability, it is used for drafting — not for action

Several transferable lessons emerge, alongside limitations that reflect the UAE's particular political and regulatory context.

Treat AI as institutional infrastructure, not a portfolio of tools

A central lesson from the UAE is the shift from viewing AI as a collection of applications toward treating it as public-sector infrastructure. The result of this reframing changes both investment logic and governance design. Rather than funding isolated projects, the UAE increasingly invests in shared platforms, data foundations, and operating capabilities that can support multiple policy domains over time.

For global policymakers, the implication is that AI strategy should be anchored in long-term institutional capacity — including data governance, security, monitoring, and human oversight - rather than in headline use cases. This approach reduces duplication, clarifies accountability, and makes lifecycle costs more visible, even if it slows initial deployment.

Resolve coherence — flexibility trade-offs through differentiated governance

The UAE demonstrates that national coherence does not require uniformity. Federal alignment around ethics, sovereignty, and strategic direction coexists with emirate-level differentiation in operating models and execution tempo. Abu Dhabi's platform-led, infrastructure-first approach and Dubai's service-driven, iterative deployment illustrate alternative pathways within a shared national frame.

For other governments - particularly federated or decentralised systems — the lesson is that coordinated differentiation can be more effective than enforced standardisation. Clear reference frameworks and interoperability expectations matter more than identical organisational forms, though transferability will depend on each country's constitutional setup, political culture, and degree of central control.

Make procurement, infrastructure, and cost explicit strategy choices

Across interviews, procurement and infrastructure emerge not as downstream implementation issues but as strategic levers shaping sovereignty, vendor dependence, and sustainability. The UAE case shows that choices about cloud models, data classification, and shared compute directly determine what kinds of AI systems can be governed and scaled.

For global policymakers, this underscores the importance of adopting a FinOps/TCO mindset: aligning funding models, procurement rules, and infrastructure architecture with long-term operating realities. Governments that fail to surface these trade-offs early risk either escalating costs or over-constraining innovation through overly defensive choices, especially where fear of sovereignty loss leads to "everything is secret" classifications and unsustainable private-cloud dependence.

Design human oversight into operating models, not policy statements

The UAE experience reinforces that trust in government AI is built operationally, not rhetorically. Rather than relying solely on ethical principles, oversight is increasingly embedded into workflows through defined human-in-the-loop checkpoints, validation mechanisms, and audit trails - especially in high-stakes domains.

The broader lesson is that accountability must be executable. Global frameworks that specify values without clarifying who intervenes, when, and with what authority are unlikely to sustain trust as AI systems become more agentic. At the same time, the UAE still relies heavily on soft-law instruments and voluntary self-assessment tools, which - absent stronger civil-society oversight and independent audits - carry a real risk of ethics-washing.

Treat evaluation as a learning loop, not a reporting exercise

Finally, the UAE case highlights the limits of traditional performance metrics for governing AI. As systems reshape workflows and decision authority, evaluation must capture risk, adaptability, and public value—not just efficiency gains. More importantly, evaluation only matters if institutions are empowered to act on it: to redesign systems, adjust governance intensity, or decommission tools when contexts change.

For policymakers elsewhere, this points to evaluation as a test of institutional maturity. AI programmes succeed not when metrics are perfect, but when feedback loops genuinely inform prioritisation and governance over time - recognising that, in the UAE itself, many tools remain voluntary, systematic impact studies are scarce, and independent scrutiny is still emerging.

> “THE CHALLENGE OF AI IN GOVERNMENT IS NOT TECHNOLOGICAL SOPHISTICATION
>
> but institutional transformation”

CLOSING REFLECTION

Taken together, the UAE experience suggests that the central challenge of AI in government is not technological capability, but institutional readiness. Governments that align strategy, infrastructure, governance, and operating models are better positioned to move beyond pilots toward durable, trusted systems. Those that treat AI as a sequence of projects— however sophisticated - are likely to remain trapped in cycles of experimentation without sustained impact.

The UAE does not resolve all the tensions identified in the global literature. Its reliance on soft-law tools, limited civil-society participation, and uneven independent evaluation mean that questions of accountability and public trust remain open. But it offers a rare, empirically grounded view of how governments attempt to navigate these tensions in practice — making it a valuable reference point, and at times a cautionary one, for policymakers confronting similar trade-offs worldwide.

ACKNOWLEDGEMENTS

Participants in the research on AI in the UAE public sector

GOVERNMENT LEADERS & INTERVIEWEES



SAEED AL FALASI
Executive Director, Dubai Centre for Artificial Intelligence
Dubai Future Foundation



MOHAMMED ADNANE RETMI
Vice President, Sovereign Cloud
Core42



NAIMA AL FALASI
Senior Vice President, AI Strategy & Transformation
Mubadala



TAREQ ASHMAWY
Advisor
Department of Government Enablement (DGE)



AIDAN MILLAR
Head of AI Enablement
Mubadala



ROBIN MILLS
Head of AI Research
Inception

ACKNOWLEDGEMENTS

Participants in the research on AI in the UAE public sector

INSEAD PARTICIPANTS



PETER ZEMSKY
Eli Lilly Chaired Professor
of Strategy and Innovation
INSEAD



PASCALE BALZE
Research Lead
INSEAD

PROJECT TEAM

Contributors from Yango

YANGO PARTICIPANTS



DANIIL SHULEYKO
CEO
Yango Group



SERGEI LOITER
CEO of AI, AdTech, and Search
Yango Group



VICTOR BUTENKO
Director of Global Partnerships
Yango Tech



POLINA PESTOVA
Head of PR
Yango Group



DARIA KISHINSKAIA
Head of PR
Yango Tech



ALEXANDER MERKUSHEV
Head of AI
Yango Tech



YANGO GROUP YANGO TECH

© 2026 by INSEAD and Yango Tech