


REGION / OCTOBER 2026

AI for UAE government at scale

Complete the public-service journey across federal and emirate institutions



THE DECISION IN BRIEF

Government AI should be funded around a completed public-service journey, with accountable handovers and accessible recourse across every institutional boundary.

- 01

Make the citizen’s completed journey the unit of delivery. An intelligent answer creates limited value if the resident still carries evidence between agencies or cannot resolve an exception.
- 02

UAE federal ambition, Abu Dhabi investment programmes and government-reported deployments are different evidence. Track programme delivery separately from measured completion, correctness, effort and access.
- 03

Give one journey owner responsibility for coordination while statutory authority stays with each institution. Reuse identity and records where lawful, and fund the human route as part of the service.

INSIDE THIS PERSPECTIVE


The argument at a glance	3
01 / Fund the completed journey	4
Exhibit 01 / A permit journey needs an owner at each handover	5
02 / Own every handover	6
Exhibit 02 / UAE evidence answers three different questions	7
03 / Make public value observable	8-9
04 / Scale what the service proves	10
The leadership agenda	11
Method and references	12

THE MAYDEN VIEW

The public-service outcome is fewer obligations for the resident to manage.

ONE DECISION, FOUR CONNECTED QUESTIONS

Connect institutions around a service



CIVIC SERVICE NETWORK / CONCEPTUAL ARTWORK

01 Fund the completed journey

Move the burden from the resident to the service

02 Own every handover

Make every handover an owned obligation

03 Make public value observable

Reuse infrastructure without obscuring statutory authority

04 Scale what the service proves

Contract for the part each party can deliver

The public-service outcome is fewer obligations for the resident to manage.

CHAPTER 01

01

Fund the completed journey

Move the burden from the resident to the service

A resident needs one permit. Several agencies hold the evidence and authority. An AI assistant can explain each requirement while leaving the resident to reconcile conflicting instructions and chase the approval. The redesign opportunity lies in those handovers: reduce the obligations the resident must manage and make completion verifiable. That is the service mayden proposes funding, testing and scaling.

There is a substantial international learning base, but it should not be mistaken for proof of a universal delivery model. The OECD's 2025 review examines 200 use cases across 11 government functions and identifies common enabling conditions. ^[1] These cases differ in maturity, purpose and context. They support a structured search for patterns; they do not establish a pooled causal return on government AI investment. A procurement business case still needs evidence from the specific service.

The UAE offers accessible primary records of federal ambition and emirate-level programmes. These records provide the setting for this paper; Bahrain, Kuwait, Oman, Qatar and Saudi Arabia are not assessed or ranked. The proposed journey design must be tested against the institutions, laws and users of each service.

Keep ambition, investment and outcomes in separate accounts

On 23 April 2026, the UAE Cabinet announced a framework aiming to introduce agentic AI across 50% of federal government sectors, services and operations within two years. ^[2] The announcement also describes a phased approach and employee training. It is evidence of stated policy direction and delivery ambition. It is not an independently audited measure of how many services have become autonomous or how much public money has been saved.

Abu Dhabi's January 2025 announcement describes an AED 13 billion digital strategy for 2025-2027, including sovereign cloud, a unified ERP platform and more than 200 planned AI solutions. ^[3] This is an emirate-level programme, distinct from the federal framework. Its significance for service design is the combination of infrastructure and organizational change. Treating the budget as an AI-only expenditure figure, or the planned solution count as a verified result, would go beyond the source.

TAMM provides a named service example beyond a policy announcement. In its January 2026 account of the previous year, Abu Dhabi's Department of Government Enablement describes the application and the launch of AutoGov for recurring needs including licence renewals and medical appointments. ^[9] This is the programme owner's account of deployment, not an independent evaluation. Its published headline outcomes do not provide the case-level definitions, observation windows and comparison design needed here to estimate an AI effect, so this report does not reproduce them as performance evidence. The operational question is how each journey confirms completion, handles exceptions and preserves the authority of the responsible institution.

Maintain two executive accounts. The first tracks delivery against the announced programme. The second tracks observed service outcomes. Adoption targets can mobilise institutions, but only the outcome account can show whether the resident experiences less effort, correct completion and usable recourse. Investigate any gap between the two before multiplying the portfolio.

ILLUSTRATIVE DECISION TOOL

A permit journey needs an owner at each handover

Fictional cross-agency permit journey

Illustrative handover	Accountable decision	Record that proves progress
Applicant to intake	Journey owner confirms service and usable request	Confirmed request; unresolved language/access need retained
Intake to eligibility	Authorized agency resolves disputed status	Versioned evidence and agency disposition
Eligibility to execution	Executing owner accepts a permitted transaction	Transaction identifier and observed state, including partial completion
Decision to applicant or review	Responsible agency explains outcome and accessible recourse	Delivered decision, open exception or review case; no silent closure

THE MANAGEMENT IMPLICATION

Count completion under an agreed definition; a handover or generated answer is not that definition.

Original mayden design, not a reconstruction of TAMM or a measured government process. The journey owner coordinates; legal authority remains with the responsible institution. Sources: [1], [4], [5]

CHAPTER 02

02

Own every handover

Make every handover an owned obligation

Mayden's proposed service-completion model has five stages: understand the need, establish eligibility, assemble permitted evidence, execute the authorized actions and resolve the outcome. Each stage needs an owner and a handoff that can be verified. A conversational system can support all five, but its fluent response is not itself proof that an obligation has been fulfilled. For example, generating a payment instruction is different from confirming payment and updating the authoritative record.

A practical design workshop should begin with a real journey and its difficult cases. Identify every place where the resident must carry information between agencies. Ask whether the information can be reused lawfully, whether a single authoritative record exists and whether the receiving agency trusts the record's meaning. Some obstacles are policy choices, not technical deficiencies. If two agencies legitimately apply different eligibility rules, a language model should explain the difference rather than silently harmonize it.

The World Bank's 2025 GovTech update covers shared systems, service delivery, engagement and enabling institutions. It expressly says its index is not a performance ranking and highlights gaps in monitoring actual uptake. ^[4] Our inference is that capability inventories need a service-level companion. Count successful completions, unresolved exceptions and repeated requests for evidence. Those measures reveal where shared infrastructure is helping and where it exists on paper without changing the citizen's experience.

Design the route for the difficult case

Most public-service designs are tested first on straightforward cases. The difficult ones expose whether the service is fair and operable: an expired document, a name represented differently in two languages, a disability, a disputed status or a resident who cannot complete a digital identity step. These examples are proposed test cases, not claims about prevalence in Gulf populations. Their importance comes from the consequences of exclusion, not from an assumed failure rate.

An accessible service should make the next step clear when automation cannot resolve a case. The resident should know what is missing, which authority can decide and how to reach that authority without repeating the entire story. Human assistance must be funded as part of the service. Otherwise a programme can reduce digital processing cost while transferring work and frustration to call centres, front desks or the resident's family. That transfer can disappear from a narrow technology business case.

UNESCO's AI ethics recommendation emphasizes human rights, inclusion and human oversight. ^[5] The management implication developed here is to make access and remedy design requirements from the beginning. Evaluate the service in the languages and channels it will actually use, including Arabic terminology and local administrative meaning. Translation accuracy is only one dimension: an answer can be grammatically correct while naming the wrong document, authority or procedure. Local practitioners and affected users need a role in testing that cannot be supplied by an imported benchmark.

SCOPED PRIMARY-SOURCE COMPARISON

UAE evidence answers three different questions

UAE federal and Abu Dhabi evidence; not GCC-wide

Announced federal direction	Reviewed example April 2026 agentic-AI framework	Limit on the claim Not completed adoption or audited savings
Emirate investment programme	Reviewed example Abu Dhabi AED 13 bn digital strategy 2025-27	Limit on the claim Not AI-only spending or verified return
Government-reported deployment	Reviewed example DGE January 2026 account of TAMM and AutoGov	Limit on the claim Named service activity; not an independently identified AI effect
Local outcome evidence to obtain	Reviewed example Completion, corrections, access and recourse for a defined journey	Limit on the claim Needs denominators, observation period and credible comparison
THE MANAGEMENT IMPLICATION Keep programme delivery and measured public value in separate records.		

TAMM statements are attributed to DGE. No headline performance rates are reproduced because the measurement detail is insufficient for this comparison. Sources: [2], [3], [9]

CHAPTER 03

03

Make public value observable

Reuse infrastructure without obscuring statutory authority

Shared infrastructure offers economies of reuse, but it also concentrates failure. A common identity service, document exchange or model gateway can support many agencies. If its permissions or availability fail, the consequences can span the portfolio. mayden recommends centralizing reusable technical capabilities while keeping each statutory decision and service obligation attached to its accountable institution. A central AI team should make safe delivery easier; it should not accidentally become the unrecorded decision maker for every agency.

The foundation should distinguish identity, data access and authority to act. Knowing who a resident is does not automatically authorize every agency to use every record about that person. Nor does permission to retrieve a record

imply permission to alter it. These distinctions become essential when an agent can invoke several services in sequence. A transaction should carry a scoped authorization and a clear record of completion, with safe handling of partial failure. The resident should not be told that a journey is complete when one agency's step remains pending.

NCSC guidance treats secure AI operation as a lifecycle responsibility. ^[6] Our service-design extension is to rehearse failure across institutional boundaries. Test a revoked permission, a stale record and an unavailable downstream agency. Preserve a safe manual route for consequential services and establish who communicates delays. Resilience is not an invisible engineering detail: when access to a public entitlement depends on a digital chain, the recovery experience is part of the service promise.

CHAPTER 03

03

Make public value observable

Measure the completed journey and displaced effort

A credible evaluation starts before release. Record the current journey's completion time, error and rework patterns, user effort, operating cost and complaint route. Specify which part AI is expected to change and what else is changing at the same time. A simplified form or revised eligibility rule may account for much of the improvement attributed to an assistant. Where feasible, compare phased rollouts or matched groups; where that is impractical, report the weaker inference honestly.

The UN E-Government Survey assesses all 193 Member States and identifies persistent differences between national and local digital-government development. ^[7] This is a reminder to avoid letting a strong national portal stand in for every service experience. In the proposed outcome ledger, show distributions and exceptions rather than only an average. A lower median completion time can coexist with longer delays for people needing human review. A rising digital-completion rate can coexist with exclusion if difficult applicants abandon the process.

A UK government trial illustrates the measurement problem. It provided Microsoft 365 Copilot licences to 20,000 employees from September to December 2024. The published

estimate of 26 minutes saved per day came from self-reported ranges; the report says it could not establish how saved time was used. ^[8] This is useful operational evidence, but it does not establish equivalent cash savings or completed public-service outcomes.

For a Gulf service programme, specify the conversion mechanism before multiplying a time estimate by headcount. Capacity might be used to reduce a queue, handle exceptions, improve access or avoid future hiring. Each route needs a separate measure. mayden proposes pairing staff task time with the citizen's completed journey, rework and appeal outcomes. A faster internal step has public value only when the service system can make use of the capacity released.

mayden proposes a balanced scorecard with four outcomes: correct completion, avoidable effort, equitable access and recoverability. Establish local baselines and tolerances, then publish the methodology appropriate to the service. Do not produce an invented regional benchmark or aggregate unrelated services into a spurious score. Citizens should be able to understand the measures. Internally, agencies should be able to reconcile them to operational records and explain which improvement came from process redesign, AI assistance or a policy change.

CHAPTER 04

04

Scale what the service proves

Contract for the part each party can deliver

The commercial contract can reinforce a narrow interface measure even when leadership wants a better public service. In an illustrative permit programme, a supplier might be paid for handling conversations while the agency bears the cost of correcting applications. That arrangement rewards activity before the outcome is known. mayden proposes linking acceptance to a defined journey and a jointly reviewed exception sample. The supplier should demonstrate accurate information, lawful handoffs and reliable completion signals; the agency remains responsible for the policy and administrative decision.

Do not transfer risks the supplier cannot control and expect the resulting price to be efficient. A contractor cannot guarantee completion when another agency has an unresolved approval backlog. Separate supplier-controlled performance from dependencies owned by government, and give both a deadline and an escalation route. When a journey misses its target, the operating record should identify which obligation failed. This makes a commercial dispute less likely to be mistaken for a model-quality problem and keeps the public outcome visible across institutional boundaries.

Retain service knowledge in a form the next delivery team can use. The procurement pack should include the authoritative content map, interface specifications, test cases, exception classifications and the method for measuring completion. Test whether those materials can be exported and maintained without the original implementation team. Government may choose a long relationship with a supplier, but it should make that choice with a credible alternative. This is particularly important for shared components: a difficult exit from one agency's assistant can

become a constraint on many services when the same component is reused.

Scale a completed journey, then its reusable components

Cross-agency redesign can take longer than launching an information assistant. Permit a bounded assistant where it reduces effort, uses authoritative content and hands off effectively. Keep the unresolved journey in the programme's scope. A useful first release and a completed transformation are different milestones.

Choose one service with material user effort and a manageable institutional boundary. Appoint a journey owner, agree completion and exception measures with participating agencies, and test ordinary and difficult cases with frontline staff and representative users under lawful research arrangements. Record the handover that fails, its responsible institution and the action needed to resume the case.

Expand when the journey shows correct completion, usable recourse and an affordable support model. Reuse shared components where the evidence supports reuse; retain local tests for eligibility, language and consequence. Report the work citizens no longer have to do alongside the obligations government has taken on. That is a more meaningful account of progress than the number of assistants launched.

FROM INSIGHT TO ACTION

Decisions to take forward

JOURNEY OWNER AND PARTICIPATING AGENCIES

Resolve who owns each incomplete obligation before adding another interface.

Evidence to require: Handover map with eligibility, execution, exceptions and recourse

SERVICE MEASUREMENT OWNER

Separate adoption commitments and provider statements from measured outcomes.

Evidence to require: Defined denominator, observation window and completion/correction/access results

PROCUREMENT AND DELIVERY OWNERS

Scale the journey whose evidence supports completion; retain smaller scope where it does not.

Evidence to require: Exportable service knowledge, agency dependencies and support costs

mayden connects senior judgement, research and engineering to turn strategic choices into working capabilities.

EVIDENCE YOU CAN INSPECT

Research notes

Method and boundaries

Comparative desk research of nine primary institutional sources, followed by service-design synthesis; cutoff 10 October 2026.

UAE federal and Abu Dhabi announcements and a named DGE service account define the geographic cases; international sources provide methods, comparative context and safeguards.

No fieldwork, citizen interviews, service data or independent audit of announced programmes. UAE federal and Abu Dhabi evidence cannot be generalized to the entire GCC. Proposed measures and model require local validation. TAMM is an attributed programme-owner deployment account; no causal impact estimate or inadequately specified headline rate is reproduced.

Primary sources and author research

[1] Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions

OECD / 2025-09-18. Executive summary, Key findings: 200 use cases across 11 core functions.

[Read the source](#)

[2] New UAE government framework for agentic AI

UAE Cabinet / 2026-04-23. Announcement paragraphs 24,33-35.

[Read the source](#)

[3] Abu Dhabi Government launches Digital Strategy 2025-2027

Abu Dhabi Media Office / 2025-01-21. Strategy funding and infrastructure paragraphs 83-86.

[Read the source](#)

[4] 2025 GovTech Maturity Index Update

World Bank / 2025 edition; accessed 2026-10-10. Methodology and reporting on uptake; paragraphs 28-33,50.

[Read the source](#)

[5] Recommendation on the Ethics of Artificial Intelligence

UNESCO / 2021 recommendation; current page accessed 2026-10-10. Human rights, inclusion and human oversight principles.

[Read the source](#)

[6] Guidelines for secure AI system development

UK National Cyber Security Centre and partners / 2023-11-27; version 1.0. Secure operation and maintenance.

[Read the source](#)

[7] UN E-Government Survey 2024

United Nations DESA / 2024-09. Publication overview and local-government findings.

[Read the source](#)

[8] Microsoft 365 Copilot Experiment: Cross-Government Findings Report

UK Government Digital Service / 2025-06-02; trial 2024-09-30 to 2024-12-31. Sections 2.2, 3, 4.2 and 6. Not a Gulf field study; no causal cash-saving estimate; licence population is not the response denominator.

[Read the source](#)

[9] Department of Government Enablement Records Landmark Year Under AI-Native Vision

Abu Dhabi Department of Government Enablement / 2026-01-08. TAMM Sets Global Benchmarks; TAMM and AutoGov description. Self-reported programme account, not an independent audit or causal evaluation. Headline outcomes lack the definitions and comparison detail needed here and are not reproduced.

[Read the source](#)