

From Pilot to Scale: A Practical Path to AI-Native Government in the UAE

How to turn agentic AI ambition into operational AI Task Specialists in weeks, not years

Christoffer Tejs-Knudsen, Vice President, cBrain



The UAE's AI ambition is set – the method is the question

Across the world, governments are moving quickly to understand how artificial intelligence can support public administration. The United Arab Emirates is moving faster than most.

The ambition is to become an AI-native government: a government where AI is part of how services are delivered, operations are run and public institutions work. The question is no longer whether AI can support government. The question is how ministries and public entities move from ambition to operational impact – safely, visibly and fast.

Many AI initiatives begin with the same idea: take a powerful AI model, connect it to existing systems and let it operate across the organisation. It promises speed and autonomy. But in government, this approach often meets a harder reality. Government work is rule-bound, documented and accountable – and an AI agent cannot simply be placed on top of this landscape and be expected to operate responsibly.

The faster route begins not with the AI model, but with the work of government itself. This is the practical route to scale: building AI task specialists. The ambition is AI-native government. The method is to build it through real work: task by task, process by process, on a governed platform.

The UAE moment: from digital government to AI-native government

The UAE is already recognised as one of the world's most advanced digital governments. Citizens and residents experience many services as fast, digital and seamless. The next opportunity lies behind the front end, in the operational work that makes those services run in the back office.

This is where the shift to agentic AI matters. Ministries and federal entities are not being asked to discuss AI in the abstract. They are being asked to make AI part of real services and operations.

A digital service is not the same as an automated operation. A citizen may experience a service as digital because the front end is online, while the work behind it still depends on manual checking, routing and approvals. These tasks are often where the cost, delay and complexity sit.

AI-native government is therefore not only about better citizen interfaces. It is about the operating machinery of government: how cases move, how decisions are prepared and how repetitive work can be handled without losing accountability.

For the UAE, this creates a significant opportunity. The country does not need to begin with basic digitization. It can move directly to the next question: how to make AI operational inside a modern public administration.

Why AI pilots often stall before production

The first wave of government AI pilots has produced many useful lessons. One of the clearest is that AI demonstrations are easier than AI operations.

A model can summarise a document, draft a reply or classify information. In a demonstration, these capabilities can look impressive. But public administration requires the AI to understand the context in which the output will be used.

The AI must know which documents belong to the case, who may access them and what should be recorded. It must also understand when a human must review or approve the result. Without that context, the output may be useful as text, but difficult to use as official government work.

This is where generic AI pilots often struggle. They begin outside the operational structure of government, either as a separate tool beside existing systems or as a layer on top of them. Meanwhile, the real work continues to happen across case systems, document stores and established procedures.

The AI may be powerful, but it lacks a governed place to work.

This creates a paradox. The more ambitious the AI use case becomes, the more dependent it becomes on the underlying administrative architecture. To perform real work, the AI must interact with processes, data, roles and decisions. If these are fragmented, the AI project must first solve the fragmentation.

The result is delay. Before AI creates operational impact, the ministry must handle integrations, security, access control and auditability. What began as an AI pilot becomes a systems integration programme.

The problem is not the AI model. The problem is that the AI has not been embedded in the work of government.

The problem is not the AI model. The problem is that the AI has not been embedded in the work of government.

Government work has a structure

Government work is often complex, but it is not shapeless.

A permit, an inspection or a grant case will usually follow an underlying structure. Information is received, the case is reviewed, a decision is

prepared and the outcome is recorded. The details vary, but the administrative logic is recognisable.

This is why a task-based approach matters. Government workflows are built from process steps, tasks and clear task instructions. AI automation should respect that structure.

A task is the smallest meaningful unit of government work. It may involve:

- checking documents
- comparing information against a rule
- drafting a response
- routing a case to the right unit.

These are the places where AI can help immediately.

The task already has an instruction. That instruction can be translated into a prompt, defining the AI's role, the expected result and the data it may use. The AI does not need to understand the entire ministry; it needs to perform a defined task inside a defined process.

This is the first step towards agentic AI in government. The AI is not a free-floating

intelligence. It is an autonomous colleague assigned to a specific task.

The fastest route to scale: automate tasks inside real processes

The fastest route to AI-native government is to embed AI into the actual flow of government work.

That work is made up of processes, steps and tasks. Each task has a purpose, an owner, an instruction and a required output. This is where agentic AI can create measurable value quickly: by drafting, checking, classifying or validating work inside a real process.

The caseworker remains accountable for the case, while the AI performs defined tasks inside the workflow.

Each automated task becomes a reusable building block. As more tasks are automated, the process becomes more efficient; as more processes are added, the ministry builds a scalable foundation for AI-native government.

This is operational AI. It does not wait for a full transformation programme before impact appears. A ministry or agency can select one process, configure it, automate selected tasks, test the result and then decide how to expand.

The first step is practical, but it is also strategic. It creates the pattern that can be repeated across the organisation.




Figure 1: One task at a time, F2 AI Task Specialists execute defined tasks at each step of a governed process, with a human still in control.

Why the platform is what makes it scale

Task-by-task automation only works at scale if the process itself is structured.

This is where a digital platform becomes important. The platform provides the governed architecture in which cases, records, workflows and audit trails already belong together. AI then works inside this environment, not around it.

This architecture removes much of the work that normally delays AI projects. If the AI is embedded in the platform, it can use the platform's data, roles and logging instead of recreating them in a separate layer.

The process is configured, the tasks are defined and the AI operates within that structure. This gives ministries a way to move quickly without turning every process into a custom development project.

In a custom-build model, each process risks becoming its own solution, with its own integrations and maintenance burden. Over time, that creates a new form of fragmentation.

A platform approach works differently. The first process is configured on a common foundation, and the next process can reuse that foundation. Workflows, prompts and skills can be improved over time while governance remains consistent.

This is how individual AI use cases become an operating model.

From prototype to real solution: a practical 10-week model

A common weakness of pilots is that they are not built to become real operations.

They demonstrate a possibility, but they are often not production-ready. They may use sample data, simplified rules or temporary workarounds. When the pilot ends, the ministry still needs to build the real solution.

This creates a gap between innovation and implementation.

A platform-based approach closes that gap. The first version is not a throwaway prototype; it is the first version of the real process.

A phased 10-week project can support this:

1. An Alpha/PoC phase
2. A Delivery/Beta phase
3. A Final/Go Live phase, with go/no-go decisions after the early phases.

The early phase can be limited in scope, but still real in architecture.

This matters for government buyers because it gives them a way to explore AI without taking unnecessary risk. The ministry can see the process working, users can respond to something tangible, and leadership can assess whether the approach creates value.

If the result is not strong enough, the ministry can stop. If it is strong enough, the same configuration can be matured and taken forward.




Figure 2: The practical 10-week model. Three phases - Alpha (weeks 1-4), Beta (weeks 5-7) and Final (weeks 8-10) - turn a first version into a production-ready process.

Make the AI a caseworker – not a chatbot.

The pilot is therefore not a side experiment. It is a production ready version one.

The right starting process should be bounded, visible and repetitive. It should matter enough to create value but not be so broad that the first project becomes a large reform programme. A licensing case, a grant process or an internal approval flow can all be good starting points if the ownership is clear.

The point is not to prove that AI can do something in theory. The point is to prove that AI can perform useful work inside a real government process.

This is where the phrase weeks, not years, becomes credible. The promise is that a real AI-enabled process can be created, tested and matured quickly, then used as the basis for scale.

Compliance by design

In government, compliance cannot be added afterwards.

If a human performs a task, the organisation must know what happened. If AI performs a task, the requirement is even stronger. The organisation must be able to see the data used, the output produced and the human decision that followed.

This is why AI governance must be operational, not only policy-based.

One key lesson from Denmark's digitalisation journey was that compliance had to become automatic. Archiving, traceability and record-keeping could not depend on individual discipline; they had to happen as a natural by-product of daily work.

The same principle applies to AI. If AI works outside the process, compliance becomes a separate

problem. If AI works inside the process, compliance can follow the work.

The AI uses the same case context, operates under defined access rights and leaves its output in the record. The human caseworker remains accountable, but the administrative burden around the task can be reduced.

This is the practical meaning of governed AI. Control is not only described in a policy; it is built into the way the AI performs work.

Make the AI a caseworker - not a chatbot.

That does not mean the AI makes final decisions independently without human control. It means the AI works within the same administrative structure as a caseworker: assigned to a role, given an instruction and visible in the audit trail.

This is what makes AI usable in daily government operations.

Scaling from one process to many: using process libraries

A successful pilot is not the end point. It is the beginning of the scaling model.

The risk with AI is that every ministry or agency launches separate pilots with separate tools and governance assumptions. That may create innovation activity, but it does not create an AI-native government. It creates a new generation of fragmentation.

The better route is to treat each AI-enabled process as part of a growing process library.

A process library allows government to define, reuse and adapt operational processes. A permit process can be adjusted for another permit type, and a document review task can be reused in another workflow. The organisation does not start from zero every time.

This is where task-by-task automation becomes strategically powerful. AI may begin by supporting a few tasks in one process, then expand to similar tasks in related processes. Over time, entire categories of administrative work can become AI-enabled while remaining under human accountability and institutional control.

The organisation learns, the platform improves and the process library grows. Scale is not achieved by

making one giant AI system. It is achieved by repeating a controlled pattern.

The choice ahead: scattered pilots, or an AI-native administration

The UAE has the ambition, urgency and institutional focus to define the next generation of government AI. The opportunity is to avoid the slow path: scattered pilots, generic AI layers and delayed operational value.

The faster path is more practical. Start with one process, map the tasks, configure the process on a governed platform and automate selected work. Once value is proven, add more tasks, more processes and a reusable library that can scale across the organisation.

This is how AI becomes part of the operating model of government.

It is not achieved by asking AI to run everything at once. It is achieved by giving AI the right task, in the right process, on the right platform.

For governments, that is the difference between an AI pilot and AI-native administration.

For the UAE, it is a practical path from ambition to operational impact - in weeks, not years.

Exploring where to start

For ministries and federal entities, the question is no longer whether AI can support government – it is which first process to begin with. For governments interested in how to move from ambition to operational impact, or in discussing how a first AI pilot could work in their own context, please reach out:

Christoffer Tejs Knudsen, Vice President, cBrain,
ckn@cbrain.com

Kalkbrænderiløbskaj 2
DK-2100 Copenhagen
Denmark

+ 45 7216 1811
info@cbrain.com
www.cbrain.com

