

Agentic AI Use Cases for Government



Choosing useful tasks and defining safe authority for Saudi public services

The Sohob Arabia perspective

Government leaders should evaluate agentic AI by the service work it can complete within a defined mandate. The most useful starting point is often a bounded administrative task: preparing a case file, collecting missing information, or coordinating an approved maintenance request. These tasks have an observable outcome and a clear point at which authority remains with an official.

Agentic systems can select actions and use tools over multiple steps. That capability makes task design, system access, and escalation rules central to the investment decision. An assistant that drafts advice and an agent that changes a government record to create different exposures, even when they share the same underlying model.

Sohob Arabia recommends selecting government agents by the service outcome and the authority the entity can safely delegate. This applies to the AI Architecture & Platforms and Responsible AI, Governance & Compliance dimensions of our enterprise framework. The four use cases below are design proposals, not reporting agency deployments; each keeps a named official accountable for consequences.

Choose a task that needs an agent

Do not add autonomous planning to a predictable process

An ordinary workflow can execute a known sequence of steps reliably. An agent becomes useful when the next step depends on information discovered during execution, and the organization can constrain the choices available. Anthropic's engineering guidance distinguishes predefined workflows from systems in which a model directs its process and tool to use. [1]

For a public entity, the distinction should shape the use case. Validating whether a required field is present usually needs a deterministic rule. Investigating an incomplete service request may require several searches and clarification steps. The latter may justify a bounded agent, provided each tool is authorized for the task, and the result can be checked.

Compare candidates by public service value

Application preparation is useful when requirements depend on the applicant's situation. An agent can consult approved guidance, ask for missing evidence, and assemble a cited case summary. It should not infer eligibility from missing documents. Measure complete submissions and repeat requests, while retaining the official's eligibility decision.


Policy research has a different boundary. An agent can follow references across approved circulars, compare effective dates, and identify conflicting provisions. Its deliverable is a source-linked evidence brief, with unresolved conflicts visible. An authorized policy specialist interprets the material and approves advice. OECD research highlights governance, data, and skills as enablers for government AI. [2]

For internal procurement, an agent could assemble specifications and flag missing justification across approved records. Keep supplier selection and commitments with authorized staff. Use a conventional workflow whenever the next step is already known; agentic search is justified only where discovery changes the task.

Candidate task	Permitted contribution	Human boundary
Application preparation	Find missing documents and prepare a case summary.	Decide eligibility or approve the application.
Municipal requests	Collect details and propose a work order.	Approve priority and dispatch exceptions.
Policy research	Assemble cited material from approved sources.	Interpret policy and approve advice.
Internal procurement support	Collect specifications and flag missing evidence.	Select suppliers and authorize commitments.

Design the service journey

A municipal maintenance request with controlled execution



Imagine a resident reporting a damaged pavement through an approved service channel. The agent checks that the request has a usable location and description, searches the permitted maintenance records, and asks for clarification when information is missing. It then prepares a draft work order with supporting references.

The system should confirm that the location belongs to the entity's service area using an authoritative geographic service. It should not invent coordinates from an ambiguous description. A duplicate check can suggest that two reports concern the same issue, but uncertain cases should remain separate for an employee to review.

Separate preparation from commitment

The agent may collect and organize information. A rules-based service can verify mandatory fields and routing rules. During an initial release, an authorized dispatcher approves the order before it reaches an operational queue. The agent should not change urgency classifications, create spending commitments, or close to a resident case outside its delegated scope.

Once an action is approved, use a transaction identifier and an idempotency control, so a retry does not create a second order. Confirm the result in the destination system before telling the resident that the request was submitted. A successful tool call is not enough if the record was rejected downstream.

Keep the resident informed

Show a plain language status, the next expected step, and a route to a person. The service must distinguish between a draft, a submitted request, and completed maintenance. Avoid an overly confident conversational response that implies work has been done when only a ticket exists.

The municipal case succeeds when a complete, authorized work order reaches the correct queue exactly once. Track rejected submissions, duplicate orders, and inaccurate status messages alongside speed. This is distinct from application preparation, research, and procurement: the agent participates in an operational transaction, so destination verification is essential.

Set authority at the tool boundary

A prompt is not an access control mechanism

Define an allowlist of actions for each agent's role. Permissions should specify which records it can read, which fields it can propose, and which changes require approval. Enforce those rules in identity, application, and integration controls. A statement inside a prompt cannot replace a system of permission.

Treat citizen messages, attachments, and retrieved documents as untrusted content. An instruction hidden in a document must not gain authority to change the task, reveal protected records, or invoke another tool. Limit the data returned by tools and validate their inputs against the service contract.

Protect data across the complete path

SDAIA's PDPL guidance distinguishes controllers and processors and explains that outsourcing does not remove controllers' responsibilities. [3] For an agentic service, map data through the intake channel, model service, retrieval store, logs, integration layer, and support access. Confirm the applicable processing basis, retention arrangements, and transfer requirements with the responsible functions.

The NIST Generative AI Profile describes risks including false outputs and component integration. [4] In this design, separate those risks: verify the factual content of a case summary, and independently verify whether an action was permitted and completed correctly.


Default to escalation when the agent cannot establish authority, resolve contradictory sources, or recover from a failed transaction. Limit the number of steps and the time available for a task, so an unproductive loop cannot continue consuming resources.

Action class	Suggested initial control
Read approved service guidance	Permission is aware of retrieval and source of recording.
Read a resident record	Verified identity and case specific authorization.
Draft a proposed update	Validation before any system is written.
Commit an operational action	Named approval or a narrowly delegated rule.
Change entitlement or impose a penalty	Retain authorized official decision making.

Prove the use case before expansion

Test complete tasks including the difficult exceptions

A municipal request separates three responsibilities



Measure accepted task completion, time to a correct handoff, effort and citizen follow up. Also measure unauthorized action attempts, duplicate transactions, and cases in which the agent reports success incorrectly. A fluent response should never substitute for a verified service result.

Build realistic evaluation cases

Use separate tests for each shortlisted service. For applications, introduce a missing document that must not imply ineligibility. For policy research, supply two circulars with conflicting dates. For procurement, request an unauthorized commitment. For municipal requests, test incomplete addresses, Arabic variants, duplicate reports, and a withdrawn request.

Review the execution trace and simulate an unavailable destination system. The operator should identify the information used, approvals obtained, and final state without exposing unrelated personal information. Retain evidence proportionately. A failed transaction must remain visibly pending or failed until reconciled, rather than being reported as completed.

Launch inside a narrow service boundary

Start with one request category, one operating team, and a clear working schedule. Keep a manual route available and establish who can pause the agent. Expand only normal cases and exceptions to meet the agreed service standard, and the team can manage the escalation of load.

If the agent spends most of its time following a stable sequence, replace that portion with a conventional workflow. If it regularly needs broad human interpretation, it prepares evidence. The purpose is dependable on public service, not maximizing the number of autonomous steps.

At the review, choose between keeping the scope, expanding a specific permission, simplifying the design, or stopping the use case. Record what evidence would justify the next choice.

Build a portfolio around service outcomes

Select the next use case from operational evidence



For the first selection workshop, bring service managers, frontline staff, data owners, architecture, and the relevant control teams. Ask each group to describe one recurring delay and one failure that would be unacceptable. This produces a more useful shortlist than starting with a catalogue of agent products.

The selection output should be a brief for each candidate: public outcome, reason for using an agent, permitted tools, human decision boundaries, and acceptance tests. Prioritize the candidate with a meaningful service benefit and a boundary the operating team can enforce. Its first milestone is verified upon completion within that boundary.

Discuss your government use case with Sohob Arabia

Bring one service journey and the systems it crosses to Sohob Arabia. The useful design output is a permitted-action map, a tested handoff to staff, and an evaluation set tied to service outcomes. These connect architecture and governance to a task the entity can operate.

[Contact Sohob Arabia](#)

Evidence and further reading

[1] Anthropic | [Building Effective Agents](#)

[2] OECD | [Governing with Artificial Intelligence Executive Summary](#)

[3] SDAIA | [Guide to the Saudi Personal Data Protection Law for Controllers and Processors](#)

[4] NIST | [Generative Artificial Intelligence Profile NIST AI 600-1](#)

Original Sohob Arabia editorial analysis. Examples and proposed methods are illustrative, not reported on client outcomes. Sources inform the discussion; international references do not replace applicable Saudi requirements.