

The First 6 Weeks:

How Ministries Can Start Embedding AI Without Losing Control

A practical framework for moving from AI ambition to governed execution in ministry workflows.



1

AI starts to fail when it sits outside the workflow

Why many public-sector AI efforts create risk before they create value

Ministries and agencies are under pressure to deliver faster, manage growing complexity, and maintain trust at the same time. That is why AI has become so attractive. But public administration cannot use AI in the same way as a consumer app or a general enterprise assistant. Authorities must be able to explain what happened in a case, document how a result was produced, apply rules consistently, and keep sensitive data inside trusted environments.

This is where many initiatives begin to drift. Generic AI is often introduced outside the real case flow: outside the handovers, outside the documentation model, and outside the controls that make public administration governable. At first, that can look productive. Over time, it creates a second layer of work: one workflow for the formal case and another for the AI interaction.

Denmark's digital experience points to the underlying issue. Digitalization became more powerful only when compliance, case structure, communication, roles, and processes were embedded in one coherent administrative backbone. Without that coherence, more technology increased fragmentation instead of reducing it.



Generic AI can improve drafting, but that is not the same as governed case execution.



Public authorities need documentation, traceability, and consistent treatment.



AI outside the workflow often creates parallel processes.



Parallel processes weaken control over time.

2

Why early momentum fades

Why promising pilots so often stall before operational use

Many AI programs start too broadly and define tasks too vaguely. The result is early excitement without a workable model. Success depends on process discipline, trusted data, and clear legal context. Without those foundations, even strong models produce uneven results. Momentum also fades when AI sits outside the case flow instead of inside the governed workflow. Outputs may seem useful, but remain hard to explain, document, and control. In public administration, pilots that weaken transparency, consistency, or traceability rarely survive daily operations.

This pattern is familiar in government. Leaders are asked where AI can create value, and the answer is often framed too broadly: improve productivity, speed decisions, reduce backlogs. Those goals are reasonable, but too generic to guide implementation in ministries and agencies. Administrative work is structured around stages, rules, tasks, outputs, and exceptions. Without that structure, a pilot may impress in a workshop yet fail in real case handling.

The briefing material makes the same point in simpler terms: AI projects fail when expectations are too general and when AI is treated as the answer before the real operational problem has been mapped.

- Broad ambition is not an implementation model.
- A successful workshop is not the same as a working case process.
- Reliable results depend on workflow fit, data context, and review logic.
- In government, pilots fail when they cannot survive operational reality.

“AI projects fail when expectations are too general and when AI is treated as the answer before the real operational problem has been mapped.”

3

The shift: Treat AI as part of administrative production

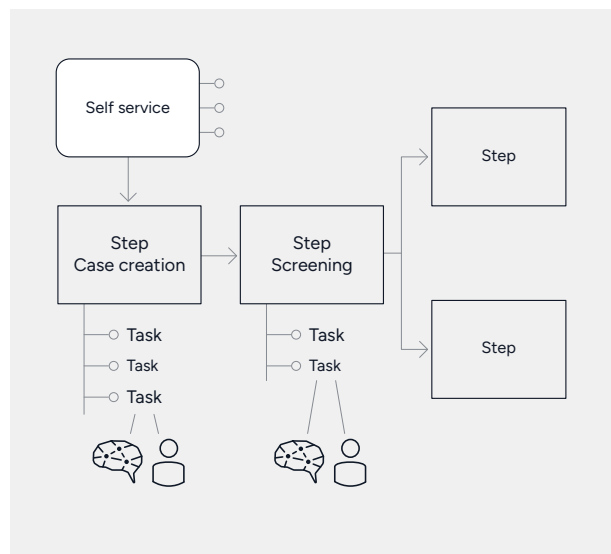
The right question is not “Where can we use AI?” but “Which tasks inside this process should AI perform?”

A more effective way to think about AI in government is to treat it as a specialized digital colleague that performs defined tasks inside a governed administrative process. That changes the implementation question completely. The issue is no longer where AI might be added in general. The issue becomes which concrete tasks in a case flow can be automated or supported, under which instructions, using which knowledge, with which controls.

That logic is consistent with how public administration already works. Cases move through defined steps. At each step, staff perform tasks based on rules, instructions, and professional knowledge. AI becomes useful when it helps with that production model rather than bypassing it. The product materials make the analogy explicit: just as industrial robots automated defined tasks in a production line, public-sector AI should automate defined tasks in the administrative process.

The practical implication is important. Ministries do not need AI that acts independently of the operating model. They need AI that works within the process, writes back into the case context, respects role and access structures, and supports human judgment where judgment is still required.

- The right unit of implementation is the task, not the chatbot.
- Administrative production depends on steps, handovers, and controlled exceptions.
- AI should reduce manual burden without weakening legal integrity.
- Human responsibility must remain explicit where judgment is required.



*Figure 1:
Example of a simple workflow strip
with highlighted AI-supported tasks*

4

Weeks 1–2: Choose the right case flow

Start where the workflow is visible, the burden is real, and the benefit can be measured

The first mistake leaders make is to start with enterprise-wide ambition. A better starting point is one bounded case flow with visible volume, repetitive manual work, and clear ownership. The goal is not to prove AI in theory. The goal is to identify one operational context where process steps, tasks, handovers, data sources, and decision points can be mapped clearly enough to support governed execution.

The cBrain® phased method begins here: process mapping, process re-engineering where needed, and identification of tasks suitable for initial AI automation. This is also consistent with the broader lesson from Denmark that sustainable digitalization starts with practical exploration of workflows that mirror how government actually works.

A good first use case is not necessarily the most strategic process in the organization. It is the process where the logic is clear enough to support fast learning without compromising control.



Checklist for a good first use case

- Pick one bounded case flow, not a ministry-wide promise.
- Map the current flow step by step.
- Identify tasks, handovers, data inputs, outputs, and exceptions.
- Prefer workflows with repetitive manual burden and visible friction.
- Choose a case flow where progress can be measured quickly.

5

Weeks 3–4: Define guardrails before automation

Separate what AI can do from what must remain explicitly human-led

Once the process is mapped, leaders need to decide where AI fits and where it does not. That means breaking the flow into tasks and distinguishing clearly between automation, support, and judgment. Some tasks are repetitive, document-heavy, and rule-informed. Others involve interpretation, exception handling, or formal responsibility. Public authorities should absolutely use AI to reduce manual work, but they should not let AI bypass the moments where accountability remains human.

This is where governance becomes operational. For each AI-supported task, teams should define the instruction, the knowledge required, the data required, what completion looks like, how the output is documented, and when review is mandatory. The briefing paper recommends exactly this kind of decomposition: break the

process into steps and tasks, then specify the knowledge, data, and expected result for each AI task before choosing tools.

The product material reinforces the same principle through three prompt types: role, instruction, and skills. That structure matters because it turns vague AI usage into a governed task model that can be tuned, versioned, and maintained over time.

- Define which tasks are suitable for AI.
- Keep judgment, escalation, and accountability explicit.
- Specify data, knowledge, instruction, and output per task.
- Decide how outputs will be reviewed and recorded.
- Build governance into the workflow from day one.

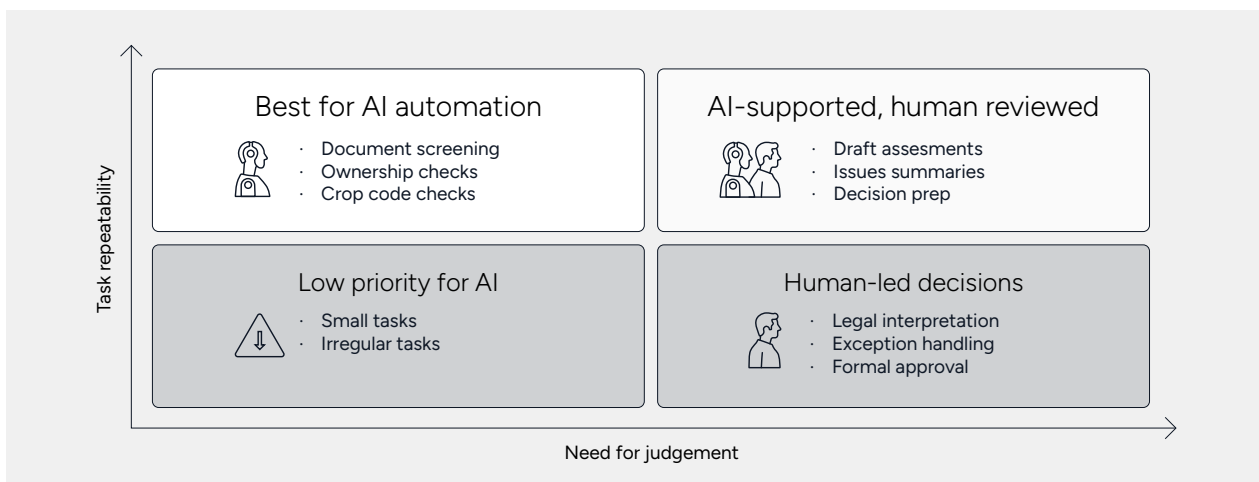


Figure 2: Which work should AI automate, support, or leave human-led?

6

Weeks 5–6: Test in context, then decide

Move from pilot theatre to operational evidence

The final step in the first six weeks is not broad rollout. It is controlled testing in context. That means validating whether the AI-supported tasks fit the actual workflow, whether prompts and instructions are precise enough, whether needed integrations are clear, and whether users can work with the result without creating new exceptions or workarounds.

cBrain’s phased approach is useful here because it is explicitly built around go/no-go decisions. The first phase typically combines workshops, process mapping, initial task automation, and an alpha version. The second phase completes the workflow, tunes prompts, clarifies integrations, and tests the solution further before go-live. Each

phase is intentionally bounded, so leaders can learn before they scale.

This matters because ministries do not need more symbolic adoption. They need a repeatable method for turning specialist knowledge into controlled productivity gains. A six-week start is credible when the scope is narrow, the process is clear, and the success criteria are operational.

- Test the workflow, not just the model.
- Tune prompts and task logic in context.
- Clarify integration dependencies early.
- Use go/no-go decisions before expanding scope.
- Measure readiness by control and fit, not novelty.

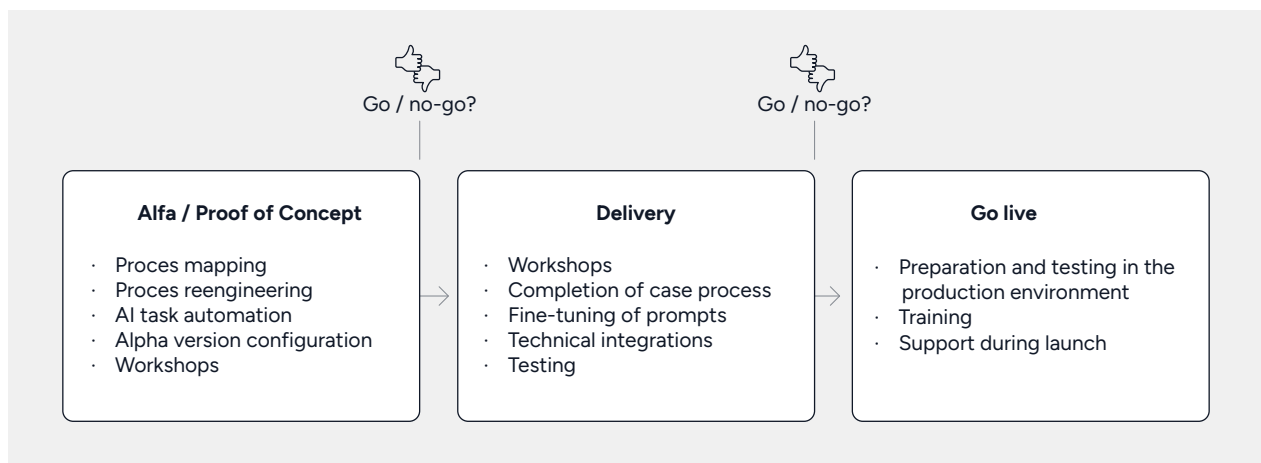


Figure 3: Example of our project model

7

Five mistakes that lock ministries into fragmentation

Common errors that make AI harder to govern and harder to scale

Most public-sector AI mistakes are not technical. They are implementation mistakes. They happen when the operating model is ignored and AI is treated as something separate from the administrative process.

1. Treating AI as a standalone tool

Why it happens: It is faster to deploy chat access than redesign work.

Why it is risky: It creates a second workflow outside the case context.

What to do instead: Embed AI in defined tasks inside the case flow.

2. Launching pilots before mapping the process

Why it happens: Teams want momentum quickly.

Why it is risky: Nobody can define what good output looks like.

What to do instead: Map steps, tasks, handovers, and exceptions first.

3. Separating AI from legal and procedural requirements

Why it happens: Governance is treated as a later-stage concern.

Why it is risky: The pilot cannot survive operational reality.

What to do instead: Build documentation, accountability, and review logic in from day one.

4. Optimizing for experimentation instead of execution

Why it happens: Demonstrations are easier than delivery.

Why it is risky: Workshops look good, but case handling does not improve.

What to do instead: Test the process in context and use go/no-go gates.

5. Creating parallel processes that weaken accountability

Why it happens: AI outputs live in one place; the formal case lives elsewhere.

Why it is risky: Traceability, consistency, and control erode.

What to do instead: Ensure outputs remain attached to the case and reviewable in context.

- Most failures come from workflow design, not from model weakness.
- Fragmentation returns when AI sits outside the process.
- Governance should shape implementation from the beginning.
- The best defense against failure is operational clarity.



What good looks like


The target state for operational AI in ministries and agencies

Effective implementation does not look like a generic assistant rolled out across the organization. It looks like a governed operational model. AI is embedded in the case workflow. Tasks are clearly defined. Outputs are written back into the case context. Review points are explicit. Human accountability is visible. Legal and documentation requirements remain intact. Leaders can see where AI is being used, what work it supports, and how value is being created.

The broader architecture behind this outcome: automatic compliance, one authoritative source of truth, role-based control, and processes embedded in the administrative model itself. The AI materials extend that same logic to task-level automation. Together, they point to a target state where AI strengthens coherence rather than weakening it.

In practical terms, good looks like this: organizations start with one workflow, build confidence through governed use cases, then expand into adjacent processes that share the same logic. Over time, instructions, prompts, and skills become reusable assets instead of one-off pilot artifacts.

- AI is part of the workflow, not outside it.
- Legal integrity and documentation remain intact.
- Roles, review logic, and accountability are explicit.
- Architecture supports scale without creating new silos.
- Value is measured through controlled productivity gains.



9

A better way to move from AI ambition to execution

Start with one workflow. Build evidence fast. Keep control intact.

The real question for ministries is not whether AI belongs in government. It is how to implement AI without weakening the case workflow that legal certainty, consistency, and accountability depend on. The answer is not a generic bolt-on tool. It is a workflow-centered implementation approach: start with one bounded process, define tasks clearly, embed guardrails from the start, test in context, and expand only when the operating model holds.

That is the practical lesson behind both Denmark's administrative digitalization journey and cBrain's approach to AI-supported case handling. Sustainable gains come from embedding capability into the production model of government itself.

If you are assessing where AI can create value in your ministry or agency, start with one workflow, one defined burden, and one realistic six-week exploration. That is often enough to see whether AI can support the operating reality you actually need.

- Do not start with a broad AI promise.
- Start with one governed workflow.
- Test operational fit before scaling.
- Use structure to create speed without losing control.

Discuss one workflow, one use case, one 6-week start

You are welcome to contact us for discussion and dialogue via www.cbrain.com.

Kalkbrænderiløbskaj 2
DK-2100 Copenhagen
Denmark

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

CBRAIN®