


AI in government works best when it stays inside the process

Michael Wätjen, Vice President, cBrain, Klavs A. Holm, Diplomatic Advisor and Ambassador, ret., cBrain, and Adam Mollerup, Head of GovTech



AI needs administrative order

Public authorities are under pressure to do two things at once: accelerate delivery and strengthen control. That is why AI has become so attractive. Yet the public sector cannot use AI in the same way as a consumer app or a general enterprise assistant. Ministries and agencies must be able to document what happened, explain why a case moved forward, apply rules consistently, and keep sensitive data inside trusted environments.

This creates a paradox. The more serious the administrative task, the less useful generic AI becomes on its own. General-purpose copilots may be impressive in conversation, but public administration is not a conversation. It is a sequence of defined tasks carried out within case processes, legal frameworks, access rights, and documentation requirements. If AI is introduced outside that structure, it often adds a new layer of risk rather than a new layer of capacity.

The opportunity is therefore not to replace bureaucracy with AI. It is to make administrative production more effective by embedding AI in the existing workflow, at the level of concrete tasks. When AI is trained, instructed, and governed within a case process, authorities can increase productivity without giving up compliance, transparency, or human accountability.

Why the challenge persists

Many public-sector AI initiatives stall because they begin with technology enthusiasm instead of administrative reality. Leaders understandably ask where AI can create value, but the answer is often framed too broadly: improve productivity, speed up decisions, help knowledge workers, or reduce backlogs. Those ambitions are reasonable, yet they are too generic to guide implementation in ministries and agencies.

Administrative work is highly structured. Cases move through defined stages. At each stage, staff perform specific tasks based on legislation, policy, instructions, and professional knowledge. Some tasks can be standardized heavily; others require judgement, escalation, or exception handling. In that context, a general AI tool quickly runs into practical limits. It may produce plausible language, but the authority still needs traceability, documentation, role-based access, and consistency across thousands of cases.

A second reason the challenge persists is that many organizations try to apply AI before they have mapped the process in sufficient detail. Without clarity on steps, tasks, data sources, outputs, and handovers, it is impossible to decide where AI should be used, what knowledge it needs, and what a good result actually looks like. The result is familiar: pilots that impress in demonstration but fail in daily operations.

Why early momentum fades

Many AI programs start too broadly and define tasks too vaguely, creating excitement without a workable model. Success depends on process discipline, trusted data, and clear legal context. Momentum also fades when AI sits outside the case flow, making outputs harder to explain, document, and control. In public administration, pilots that weaken transparency, consistency, or traceability rarely survive operations well.

Why generic AI falls short

The dominant approaches to public-sector AI rely on assumptions that are only partly true. The first assumption is that a powerful general model can simply be placed in front of complex work and start producing value. That works for low-risk drafting or experimentation, but it is insufficient for

administrative production. A ministry does not need a clever answer in isolation; it needs the right action, in the right case, based on the right sources, recorded in the right place.

A second assumption is that the central challenge is model quality. In practice, the bigger issue is operational context. Authorities need to know which data the AI can use, which task it is performing, what instruction applies, when a human must review the outcome, and how the result becomes part of the case record. Without that context, even a strong model becomes unreliable in production.

Why assumptions fail

Many current AI initiatives in government are guided by conclusions that sound plausible, but do not hold up in operational case handling:

- Fluent drafting does not mean administrative work can be performed safely at scale.
- A successful workshop pilot does not guarantee success in a live case process.
- A reasonable-sounding answer is not the same as reliable decision support.
- Giving staff access to chat does not mean the process itself has been modernized.
- Data that exists somewhere in the organization is not automatically usable in context.

These assumptions ignore how public administration actually creates value: through repeatable tasks, documented handovers, and controlled exceptions. In real case handling, one task may involve screening ownership records, checking servitudes, identifying protected species, or preparing a draft assessment. The challenge is not merely to generate text; it is to execute the task in a way that fits the workflow and produces an auditable result.

A further weakness in current approaches is that they blur the boundary between automation and judgement. Public authorities should absolutely use AI to reduce manual work. But they should not let AI bypass the moments where accountability, legal interpretation, or exception handling requires human responsibility. Productive AI in government therefore depends on a disciplined allocation of work between automated tasks and human decisions.

What ministries need instead

A solution that works at task level, inside the case process, with explicit instructions and governed access to data.

A setup where AI can support both screening and case handling, while keeping documentation attached to the case.

A model for continuous improvement, so prompts, skills, and workflows can be tuned over time as experience grows.

AI must work inside the process

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 perspective changes the implementation question completely. The issue is no longer, "Where can we use AI?" but rather, "Which concrete tasks in this process can be automated or supported, under which instructions, using which knowledge, with which controls?"

This is the logic behind F2 AI Task Specialist. The product is designed to automate manual tasks in case handling without changing the fundamental principles of administration. Instead of acting as a general agent detached from the workflow, it operates within cBrain F2® and alongside F2 Service Builder, where authorities define and maintain the case process, including steps, tasks, and the conditions under which work moves forward. Because the solution is offered on-premise, data does not have to leave the authority's own environment.

The model is powerful precisely because it is narrow where it needs to be narrow. F2 AI Task Specialist is configured per task and guided by three prompt types: role, instruction, and skills. Role defines how the specialist should behave. Instruction defines the specific task it must carry out. Skills provide the domain knowledge required to perform that task. This mirrors how experienced organizations train employees: not with general ambition, but with a role, a work instruction, and the knowledge needed to perform it well.

The practical implication is that AI can be introduced where ministries feel the pressure most clearly: in repetitive, document-heavy, rule-informed work. A concrete example is the forestation case flow described in the product material. There, AI is used to support screening of

AI creates value in government when it works inside the process, not outside it.

ownership conditions, protected species, crop codes, and servitudes, while also assisting the caseworker during subsequent handling of flagged issues. The process remains intact; AI reduces the manual burden within it.

This reframing matters because it preserves control. The case flow still determines what happens next. Exceptions still become visible for manual handling. Human caseworkers can still question the AI, ask for elaboration, and make the final assessment where necessary. In other words, AI becomes a production capability inside bureaucracy, not an uncontrolled alternative to it.

A simple operating model for public-sector AI

Map the process first. Break the case flow into steps, tasks, data inputs, outputs, and decision points.

Then identify which tasks are repetitive, knowledge-intensive, and suitable for task-bound AI support or automation.

Why phasing matters

cBrain's phased method is relevant here because it recognizes that operational AI must be tested in context. The recommended path starts with an alpha phase built around workshops, process mapping, re-engineering where needed, and initial AI task automation. It then moves into preparation, where workflows are completed, prompts are tuned, and integrations are clarified, before go-live in the operational environment. This go/no-go structure lowers risk and gives ministries a practical way to learn before scaling.

The broader lesson is clear: successful AI adoption in government depends less on bold vision statements and more on disciplined process design. Once authorities accept that, AI becomes easier to govern, easier to explain, and far more likely to produce measurable value.

AI in government should automate the work around decisions, not obscure the basis for decisions.

How to put AI to work

For public-sector leaders, this perspective leads to a different set of priorities. The goal is not to launch the largest possible AI initiative. The goal is to introduce AI where the administrative logic is already sufficiently understood, the tasks are explicit, and the control model is clear.

1. Start with one bounded case flow, not an enterprise-wide promise.
2. Break that flow into tasks and distinguish clearly between automation, support, and judgement.
3. Define the data, instructions, and knowledge each AI-supported task requires.
4. Ensure outputs are written back into the case context and remain reviewable.
5. Use phased testing with go/no-go decisions before scaling into live operations.

Leaders should also reconsider how they define success. In the public sector, value is not only measured in speed. It is measured in speed with consistency, productivity with traceability, and better service without loss of control. That means implementation teams must include process owners, operational specialists, and compliance-minded stakeholders from the beginning, not only technologists.

There is also an organizational implication. Ministries often talk about lacking enough experienced specialists to keep up with growing case volumes and rising complexity. Task-bound AI offers a practical answer: it allows organizations to encode selected parts of specialist knowledge into reusable skills and instructions that can support staff across similar tasks. The aim is not to replace expertise, but to extend its reach.

What leaders decide early

Which case flow has the clearest backlog, the strongest process maturity, and the most repetitive manual burden.

Which tasks are appropriate for AI automation, and which must remain explicitly human-led.

Which integrations, data-quality issues, and governance constraints must be resolved before go-live.

How prompts, skills, and workflows will be maintained as the solution learns from operational experience.

If leaders make these choices early, AI moves from abstract ambition to practical capability. The organization gains a repeatable method for turning specialist knowledge into controlled productivity improvements - without losing sight of legal certainty, documentation, or accountability.

A practical starting point

Treat AI as part of administrative production. Begin where the process is visible, the task is clear, and the benefit can be measured.

Build confidence through governed use cases, then expand across adjacent processes that share the same logic.

What leaders should do next

The next phase of public-sector AI will separate symbolic adoption from operational adoption. Organizations that continue to deploy generic AI outside the workflow will struggle with trust, auditability, and uneven results. They may still generate excitement, but they will find it hard to convert experiments into stable production capacity.

By contrast, authorities that embed AI within defined case processes can create a compounding advantage. They will shorten handling times on selected tasks, free specialists for the most complex work, and build reusable libraries of instructions and skills that improve over time. Just as importantly, they will do so in a way that remains compatible with public administration's core obligations.

For ministries and agencies under pressure to deliver more with limited capacity, that is the real promise of AI: not a break with bureaucracy, but a more productive form of it.

If your organization is interested in learning how AI can be embedded within case processes to strengthen both productivity and control, please reach out to the authors:

Michael Wätjen, Vice President, cBrain,
mwa@cbrain.com

Klavs A. Holm, Diplomatic Advisor and
Ambassador, ret., cBrain
kho@cbrain.com

Adam Mollerup, Head of GovTech, cBrain
ada@cbrain.com

Kalkbrænderiløbskaj 2
DK-2100 Copenhagen
Denmark

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

CBRAIN®