



Netcompany
Institute

AI is the key to
maintaining a modern
welfare state

Executive summary

Denmark can meet its national AI target with AI systems built for purpose.

The target is within reach. By 2035, the government, KL and Danske Regioner aim to free capacity in the public-sector corresponding to 30,000 full-time equivalents a year through AI systems. Assessing more than 16,000 public-sector tasks, we find a potential to automate work equivalent to about 90,000 FTEs, worth DKK 54 billion a year in wages. The potential reflects a scenario where every public employee uses AI systems for every relevant task at full efficiency.

But AI assistants alone do not get us far. We only count time-saving that is feasible under the standards that make public administration lawful and trusted. Meeting these standards means that more than two-thirds of the potential requires tailored AI systems embedded in public processes. Of the remaining third that could theoretically be realised with secure off-the-shelf assistants while meeting these standards, we estimate that 11,000-17,000 FTEs could be achieved cost-effectively.

Automation has different interpretation across occupations. Occupations in public administration hold the largest potential, 34,000 FTEs, and offers the greatest scope for strategically redirecting freed capacity to other uses. For occupations in health, social care, and education, automation will primarily reduce time spent on individual administrative tasks, freeing staff to spend more time with patients, citizens and students rather than creating capacity that can readily be redirected elsewhere. Policies should therefore be explicit about the intended use of the capacity released. Reaching the national target, however, requires addressing more than one of these groups.

Reaching the national target is an investment decision that must be taken soon. To reach the target, tailored AI systems must be scoped, built, tested, and in daily use across large parts of the public sector well before 2035. Such software embeds AI systems in structured workflows with rules, data and existing systems to ensure consistency, transparency and security, which underpin public-sector legitimacy.

Control over the infrastructure is an important part of the decision. Embedding AI in critical public systems at scale creates dependencies that raise questions of resilience, data sovereignty and strategic autonomy. Choices about provider diversification, substitutability, control over data and infrastructure, and investment in European capacity cannot be separated from the decision to build public sector capacity through AI.

AI offers a rare chance to help meet the compounding pressures on the welfare state by freeing public-sector capacity while preserving the quality and legitimacy on which public services depend. The technology to do so already exists, but realising the potential requires sound investments and careful implementation.

AI is the key to maintaining a modern welfare state

The potential time saving from AI in Denmark's public sector amounts to 90,000 full-time equivalents in 2035 – three times the national target of 30,000. Reaching the target by 2035, however, will require investment in tailored AI systems across the public sector to begin soon.

Key takeaways

- We estimate the potential of AI systems to free resources in the public sector at 90,000 full-time equivalents (FTEs) in 2035 – corresponding to a yearly wage bill of DKK 54 billion at current wage levels. The potential reflects a scenario where every public employee uses AI systems for every relevant task at full efficiency. What share of the potential will actually be realised depends on investment and implementation.
- Around 28,000 FTEs of the potential could theoretically be achieved using secure, compliant off-the-shelf AI assistants. Of the 28,000 FTEs, however, we estimate that around 11,000–17,000 FTEs could be realised cost-effectively, at an annual cost of DKK 900–1,900 million.
- Even in the most optimistic scenarios, reaching the Danish national target of 30,000 FTEs by 2035 would require tailored AI systems covering significant parts of the public sector to be scoped, developed, tested, implemented and used efficiently by employees before 2035.

By 2035, the national government, KL (Local Government Denmark) and Danske Regioner (Danish Regions) aim to free public-sector working time equivalent to at least 30,000 FTEs per year through AI systems.¹ It is one of the most concrete AI targets any government has set. This paper asks whether it can be met – and what meeting it would require.

Our approach goes a step further than most of the academic literature. We evaluate not only whether AI systems could technically save time on a task, but also whether that saving is feasible in a public-sector setting, and thus whether it can be achieved with secure off-the-shelf AI assistants or requires tailored AI systems to safeguard the good administrative practice on which public-sector legitimacy depends.² For off-the-shelf AI assistants, we go one step further by estimating how much of this theoretical potential can be realised cost-effectively.


In the coming years, a rising demand for services such as health and elderly care will continue to put pressure on the welfare state, while the bureaucracy required to keep public services transparent, secure and consistent grows year by year. With today's systems and processes, the Danish public sector is likely to employ a growing share of the labour force.

Implementing AI systems in the public sector presents a unique opportunity to reverse this trend, but unlike in the private sector, where competition makes AI adoption self-reinforcing, adoption in the public sector requires deliberate policy action. Early adoption would not only free considerable public-sector capacity to help meet these compounding pressures; it could also shift considerable demand towards AI systems aligned with European values such as democracy and individual freedom.

Our results show that the national target of 30,000 FTEs in 2035 is achievable – but only if tailored AI systems across large parts of public administration are in use by then. Assessing the more than 16,000 tasks carried out by Danish public employees, we find a feasible automation potential equivalent to about 90,000 FTEs in 2035, three times the target. Over two-thirds of that potential, however, is in tasks where automation must meet the standards of good administrative practice; about half of the remainder lies in occupations such as childcare and education, where the gains could be difficult to consolidate and redirect.

Potential time savings are often much higher than realistic gains on the ground. We set up four scenarios to give a ball-park figure of the costs and gains of rolling out off-the-shelf AI systems to public employees. Of the 28,000 FTE potential, we find that the cost-effective level of investment yields between 11,000 and 17,000 FTEs at an annual cost of 900–1,900 million DKK.

Freed time can create value in several ways



Note: Time saved through automation can translate into different outcomes across areas of public-sector work. Depending on the context, it could support higher output, better service quality, lower workloads, the ability to meet future demand with the existing workforce or, in some cases, lower staffing requirements. How the freed capacity is ultimately used is a political and managerial choice.

Meeting the target of 30,000 FTEs therefore depends on tailored AI systems being scoped, built, tested and in daily use across large parts of the public sector by 2035 – no small feat. Such systems embed existing AI models within structured workflows that combine them with rules, data and other systems, with individual components or agents responsible for clearly defined steps. They are required not because today's off-the-shelf AI assistants lack the technical capacity, but because a high level of consistency, transparency and security is required to maintain public-sector legitimacy.

Reaching the target is thus, above all, an investment decision – and one that must be taken soon. The paper proceeds as follows. Section 1 reviews what AI systems are capable of today, and why capability alone is not enough in public administration. Sections 2 and 3 present our approach and our results, and section 4 compares them with previous studies. Section 5 turns to the questions of costs, infrastructure and sovereignty implications of implementing AI across the public sector at scale, and section 6 concludes.

AI in the public sector requires careful implementation

The capability of AI models is increasing fast. Current AI systems hold considerable promise for raising productivity. Across a growing range of tasks usually associated with public-sector employment, AI systems can both make work faster and improve the quality of output, as several randomised controlled experiments have shown. Professionals given an AI assistant for writing tasks finished 40 per cent faster, while quality rose by 18 per cent.³ Customer-support agents resolved 15 per cent more enquiries per hour, with the least experienced staff gaining about 30 per cent.⁴ Another study showed that performance in legal drafting and analysis rose by between 50 and 130 per cent, with the largest gains on more complex tasks.⁵

There is reason to believe that the task-level gains could translate into productivity gains on a larger scale, even if they have yet to show up in macro data. Studies on economy-wide AI exposure generally show that public-sector work contains a high share of automatable tasks. The IMF estimates that around 60 per cent of employment in advanced economies is to some degree exposed to AI.⁶ Both the IMF and the ILO find that administrative and clerical work is particularly exposed,⁷ while the OECD finds that AI models are currently most useful in standardised, codifiable and information-processing tasks.⁸ At the same time, the OECD notes that the gap between current AI capabilities and those required for more complex and specialised tasks is narrowing.

Public-sector automation demands more than technical capability

Technical capability, however, is not enough to establish what AI systems can feasibly do in the public sector. As AI systems advance, public institutions will be expected to capture productivity gains. But they must do so while meeting the standards that make public administration lawful and trusted.

For administrative processes, this translates into concrete safeguards. Decisions must rest on documented reasons that citizens can understand and courts can review. Comparable cases must be hand-

led consistently. The path from input to decision must be traceable, and personal citizen data must be handled lawfully and securely. These requirements follow from administrative and data-protection law and are reflected in the Ombudsman's guidance for public IT systems.⁹ They are fundamental to legitimate public administration, and this paper assumes that any new AI systems used in public services must adhere to the same standards.

Off-the-shelf AI assistants are highly useful for a range of tasks, but on their own offer limited control over the consistency and predictability of their output. Where public administration requires consistency, transparency and oversight throughout the process, additional safeguards may therefore be necessary. The use of personal citizen data brings further requirements for data protection and security.

Denmark has already seen the consequences of automated systems falling short of these standards. In 2025, the Ombudsman found that an automated process at the property-valuation agency had wrongly closed around 700 cases. Later that year, the Ombudsman opened a separate inquiry into whether the agency's use of AI allows citizens sufficient insight into the calculations behind their property valuations.¹⁰ In 2026, one municipality acknowledged a caseworker's use of a general-purpose chatbot in a disability case as a clear error,¹¹ and the Danish Data Protection Agency opened an investigation into another municipality's use of off-the-shelf AI assistants without adequate risk assessments.¹² These cases show that using AI systems for tasks clearly within their technical capability can go wrong if not implemented carefully.

The technical potential of AI systems is a poor measure of what the public sector can achieve

Technical potential captures the task time that current AI capabilities could in principle automate. In this paper we focus on the narrower concept of "feasible potential": what can be automated while maintaining good administrative practice, which is fundamental to public legitimacy.

For some tasks, the two potentials are close. Summarising documents, drafting text or extracting information can often be supported by a secure off-the-shelf assistant with limited additional infrastructure. Other public processes place much greater demands on consistency, traceability, data protection or inte-

gration with existing systems. Here, an off-the-shelf assistant would typically not be enough.

That does not put these tasks beyond automation with AI systems. A task that cannot safely be handed to an off-the-shelf assistant might still be handled by an AI embedded in a system that grounds its output in authoritative data, records the steps taken, constrains what the model can do and keeps personal citizen data within the required environment. This is an example of what we call tailor-made AI systems in this report.

Our results ultimately distinguish between two routes to automation: tasks that can be reached with secure off-the-shelf AI assistants and little additional investment, and tasks that require tailored AI systems to meet public-sector standards. How we measure these two routes is the subject of the next section.

Our approach – assessing feasibility task by task

To estimate the feasible automation potential of the Danish public sector, we first measure the technical potential using the predominant method in the academic literature: we map all Danish public-sector occupations onto more than 16,000 separate tasks using a US Department of Labor database.¹³ We then assess automation potential task by task and use Danish public-sector employment data to estimate FTEs.¹⁴ To translate the estimated hours saved into the equivalent wage bill, we apply publicly available occupation-level earnings data.¹⁵

To do this, each task receives two assessments. First, we evaluate the share of the task that AI systems could currently perform given the right environment, tools and data – and while meeting the specific public-sector requirements discussed in the previous section. Second, we categorise the automatable tasks into two broad groups based on what level of investment would be required for reliable use in a public-sector setting:

Off-the-shelf AI assistants. This category covers tasks that a trained employee could carry out with a compliant and secure off-the-shelf AI assistant, using public information or material that the employee supplies to it. This could include a caseworker turning meeting notes into a draft case record, a

civil servant preparing the first draft of a committee paper or a teacher preparing material for a lesson.

Tailored AI systems. This category covers tasks that require properties that off-the-shelf AI assistants do not inherently possess. Here, software is developed for the specific process in question with the aim of maintaining security and good administrative practice. Examples include handling record-based citizen casework, coordinating appointments across systems and processing permits for citizens and enterprises.

Tailored AI systems combine existing AI models with rules and data in a structured workflow, with different components or agents performing clearly defined steps. This allows AI systems to be used while meeting requirements for security, consistency and transparency across the process. The software will generally need to be adapted to the authority's data, systems, procedures or professional safeguards. This may involve configuring and integrating an existing product or developing a purpose-built solution.


Our main estimate applies conservative assessment rules. It credits only capabilities that current AI systems demonstrably possess, while at least maintaining the current quality of public service. We find a conservative estimate much more useful from a policy perspective than speculative estimates of future AI capabilities, given that there is no theoretical limit to what AI systems might eventually be able to do.

The technical appendix provides further detail on the assessment rules, calibration, task weighting, occupational mappings, statistical tables, wage calculations and limitations.

Two types of AI systems


Initiated by employees' prompts

Secure off-the-shelf AI assistant



Same predefined process every time

Tailor-made AI system



90,000 FTEs: the feasible potential for freeing public-sector capacity

We estimate the feasible automation potential in 2035 at approximately 90,000 FTEs, or 11.2 per cent of the civilian general-government workforce. This is well above the Danish national target: reaching 30,000 FTEs by 2035 requires realising around a third of the potential (see figure 1). Over two-thirds of the potential, however, requires tailored AI systems adapted to public-sector data, workflows and safeguards.

The remaining third, or 28,000 FTEs, consists of work that can be automated using secure, compliant off-the-shelf AI assistants. About half of this lies in education, childcare and research, where gains are hard to measure and redirect, as we discuss below.

Feasibility does not mean that all of this potential will be realised by 2035. In fact, realising the full poten-


tial is highly unlikely. Figure 1 presents the full potential along with two illustrative realisation scenarios, where 60 and 80 per cent of the potential is reached. Meeting the national target of 30,000 FTEs by 2035 – and a significant share already by 2030 – will require large-scale implementation of tailored AI systems in addition to making secure and compliant AI assistants available for a large share of public employees.

Time savings vary across occupational groups – and lead to different outcomes

The automation potential spans four broad groups of public employment, of which “Administration”¹⁶ holds the greatest potential for AI automation at 34,000 FTEs, followed by sizeable potential in “Education & Childcare” and “Health & Social Care”, and a residual “Other” group (see figure 2). While the three main groups contribute similar magnitudes of potential, the interpretation of their contributions differs markedly.

Figure 1

The national target of 30,000 FTEs requires investment in tailored AI system




Note: Achieving the national target of 30,000 FTEs requires investment in tailored AI systems, even at highly ambitious degrees of realisation.

Source: Own calculations based on O*NET, the US Bureau of Labor Statistics and Statistics Denmark. Totals may not sum due to rounding.

Figure 2


Administrative occupations have the largest potential, and most of it requires tailored AI systems



Source: Own calculations based on O*NET, the US Bureau of Labor Statistics and Statistics Denmark. Totals may not sum due to rounding.

Figure 3

Administrative occupations stand out by being highly automatable by AI



Note: Marker area is proportional to automatable FTEs.

Source: Own calculations based on O*NET, the US Bureau of Labor Statistics and Statistics Denmark.

“Administration” consists of occupations such as office clerks, secretaries and legal professionals, which represent a relatively small share of public employment but a very high automation potential (see figure 3). The tasks carried out in “Administration”, in particular casework, are often procedural in nature and generally require processes to be auditable and transparent. Three-quarters of the potential therefore requires tailored AI systems to maintain good administrative practice.

“Education & Childcare” spans childcare and schooling through to teaching and research at universities. It accounts for more than 250,000 public-sector FTEs, although a smaller share of its work can be automated. Most of the potential lies in preparation and other supporting tasks in education, as well as in research. Many of these tasks do not carry the same documentation and transparency requirements as administrative casework. As a result, the group accounts for around half of the total potential that can be realised with secure, compliant off-the-shelf AI assistants. The resulting time savings are likely to be spread across many employees and supporting tasks, while scheduled commitments to teaching, supervision and care remain.

“Health & Social Care” represents another large share of public employment with relatively low automation potential. But unlike in “Education & Childcare”, automatable work in this area is primarily documentation and administration, governed by strict rules and procedures. As in “Administration”, three-quarters of the potential therefore requires tailored AI systems to maintain the quality of service.

Time saved through automation can translate into different outcomes across areas of public-sector work. Depending on the context, it could support higher output, better service quality, lower workloads, the ability to meet future demand with the existing workforce or, in some cases, lower staffing requirements. How the freed capacity is ultimately used is a political and managerial choice.

In “Administration”, automatable tasks account for a relatively large share of work, so the resulting time savings are more likely to add up to meaningful capacity at team or unit level. That capacity could be used to handle more cases, redirected towards other tasks or used to reduce overall FTEs.

In “Health & Social Care”, much of the automatable work consists of documentation and other support-

ing tasks carried out alongside patient and citizen contact. Time saved on these tasks is unlikely to be large per employee but could help staff spend more time with patients and citizens or fit existing responsibilities more easily into the working day. The gains would probably appear primarily as more face-to-face time with citizens and higher job satisfaction rather than redirected workload or possible staffing cuts.

The same goes for “Education & Childcare”, where time saved on preparation and other supporting tasks is unlikely to allow more students or children to be served or to reduce staffing needs in the short term. The gains are therefore more likely to appear as improved quality than as measurable extra capacity.

Finally, some of the capacity freed by AI will be offset by new tasks and roles created by its adoption. Operating AI systems at scale will increase the need for cybersecurity, compliance, data monitoring and other forms of oversight that also require staff time.

It is useful to consider what policymakers want to achieve with the national target of 30,000 FTEs by 2035. A policy meant to reduce the public-sector wage bill is likely to find the largest potential among administrative workers, while policies meant to free up time for citizen contact could focus more broadly. Whatever the aim, reaching the stated target in even the most optimistic scenario requires addressing more than one of the groups – and a considerable degree of software development beyond simply rolling out existing AI assistants.

The full potential of 90,000 FTEs corresponds to wages of DKK 54 billion a year

Applying current occupation-level wages to the automatable potential translates the 90,000 FTEs into DKK 54 billion in annual wage expenditure, equivalent to 11.8 per cent of the civilian general-government wage bill, only slightly higher than the 11.2 per cent of hours automated. Around DKK 37 billion of this potential requires tailored AI systems, with the remainder accessible through secure, compliant off-the-shelf AI assistants.¹⁷ The DKK 54 billion can be interpreted as the cost of achieving the same potential through hiring at current wage levels.

Measured by the wage bill of automatable hours, the broad employment groups in figure 4 closely mirror the FTE picture in figure 2. The share of “Adminis-

tration” becomes somewhat larger, as occupations in the group are better paid than average, while the opposite is true for “Education & Childcare”.


Our estimate is a conservative upper bound – realised gains will depend on implementation

The estimated potential of 90,000 FTEs represents our best estimate of the working hours that could feasibly be automated using current technology within current processes. It is an upper bound, as it assumes that AI systems are used 100 per cent effectively for every relevant task across the public sector – an overly optimistic scenario. The scenari-

os in figure 1 show how much capacity could be released at different levels of adoption and effective use, recognising that implementation will be gradual – see section 5 for more explicit scenarios for rolling out our off-the-shelf AI assistants to public employees. In another sense, the results are merely a snapshot of a fast-moving target. By the time AI systems have been broadly adopted in the public sector, the technology is likely to be far more advanced than today, and our results might therefore be understated. We also assume that public processes and employment structures remain unchanged over the period. Given how disruptive a technology AI is, a reorganisation of public services could well reap significant additional benefits on top of the per-task savings we identify.

Figure 4

The wage costs associated with time freed skew heavily towards administrative occupations



Note: The wage costs associated with automatable hours closely mirror the FTE potential.

*Source: Own calculations based on O*NET, the US Bureau of Labor Statistics and Statistics Denmark. Totals may not sum due to rounding.*

Our feasibility criterion places our estimate below most previous estimates

We employ a method that has been widely used in the academic literature and by think tanks. Because of our feasibility criterion, our results differ from previous findings in interpretation, but they can nonetheless be compared.¹⁸

Our method mirrors that of the task-level automation literature, where two academic studies, ILO-NASK (2025) and Eloundou et al. (2024), are relatively closely aligned with ours. Applying their occupational automation rates to the employment composition of Denmark's public sector results in much larger implied automation rates. An important caveat here is that both studies measure technical potential without addressing whether automation would be lawful and desirable in the Danish public sector. Applying the ILO-NASK data produces an estimate of 29.4 per cent, or over 200,000 FTEs.¹⁹ Eloundou et al.'s 2024 occupational automation rates yield 32.5 per cent, or more than 240,000 FTEs.²⁰

These figures measure occupational exposure to generative AI in the public-sector workforce. They do not explicitly estimate how much working time AI systems would save or how many posts could be released. An IMF analysis similarly reports exposure rather than automation: 63.6 per cent of Danish employment overall and 65.0 per cent of public-sector employment are classified as highly exposed to AI.²¹

Other Danish analyses also find substantial potential. Statistics Denmark (2024) applies Felten et al.'s occupational exposure scores to Danish register data but does not estimate time savings.²² Danmarks Nationalbank (2026) uses the same scores to classify 39–61 per cent of Danish work as AI-exposed and converts this into a projected increase of 0.1–1.0 percentage points in annual productivity growth – the high end of which is close to our results.²³ Working at task level, Kraka-Deloitte (2024) constructs its own exposure measure and estimates that 38 per cent of working time in Danish public administration could technically be automated.²⁴ None of these studies estimates feasible public-sector hours or FTEs released.

Closest in scope and results is a 2024 BCG study prepared for the Danish Employers' Confederation

(DA), which adjusts for expected adoption and finds a generative-AI productivity potential equivalent to 84,000–96,000 FTEs in 70 per cent of public employment in 2040.²⁵ That result is numerically very close to our 90,000 FTEs, though the two measure different things: a future-oriented, adoption-adjusted projection against a feasibility-filtered snapshot of today.

Finally, the Tony Blair Institute's 2024 study of the UK comes closest to our own design, as it divides the potential according to the technology needed to realise it. It estimates that around 40 per cent of UK public-sector tasks could be partly automated, saving 19.8 per cent of working time – a much higher share than our estimated 11.2 per cent. Free or low-cost tools account for 10.8 percentage points (about 55 per cent) of this potential. Bespoke AI systems, sensory hardware and high-cost equipment account for the remaining 9.0 points (about 45 per cent). In both studies, then, a large share of the potential depends on more than off-the-shelf tools – around 45 per cent for the Tony Blair Institute, over two-thirds in our case.²⁶

Using AI in public processes raises questions of cost and control

This section turns to two questions that any government implementing AI at this scale should ask: What does it cost, and what does it mean for control and sovereignty?

We estimate the former using simple scenario-calculations for rolling out off-the-shelf AI assistants. Their relatively standardised nature makes it possible to estimate implementation costs within a reasonable range. Tailored AI systems, by contrast, depend on the specific processes, data and systems in which they are deployed. Costing them would require a comprehensive mapping of all public-sector processes and is therefore beyond the scope of this analysis.

The latter question; realising the productivity gains while retaining sufficient control over critical functions, data and infrastructure is a complex matter. A comprehensive analysis of the strategic implications of these dependencies is beyond the scope of this paper, but policymakers should recognise that laying the foundation of future public processes in the age of AI is deeply intertwined with geopolitics and questions of control and sovereignty.

Optimal investment in off-the-shelf AI-systems: Four scenarios

To get from a theoretical potential to the optimal investment in off-the-shelf AI assistants, two data-points are needed: The number of public employees for whom a license would be cost-effective and the annual cost per license. Our data already shows the potential working time freed per public employee as well as their wages. We therefore only need the price of a licence for a secure and compliant off-the-shelf AI assistant and some assumptions about how effectively it is used. A licence is cost-effective when the wage-equivalent value of the working time it frees exceeds its price.


To provide a ball-park estimate of actual implementation of off-the-shelf AI systems in 2035, we set up four scenarios as a two-by-two along the axes

of “license price” and “efficiency of use”. Based on publicly available data on price-setting by major suppliers of AI-systems and market research from Netcompany, we arrive at an annual licence cost of between 2,000 DKK and 4,000 DKK. For efficiency of use, we consider it unrealistic to assume that the average public employee would use AI at maximum efficiency for all relevant tasks. Instead, we assume 80 per cent efficiency in the generous case and 60 per cent efficiency in the conservative case.

The two assumed prices, 2,000 and 4,000 DKK, and the two assumptions about efficiency of use, 60 and 80 per cent, gives us what we need for four scenarios. For all scenarios we assume that users need two days a year to learn how to use the systems, corresponding to about 20 minutes a week, covering both on-the-job training and formal courses.²⁷

Figure 5.

Off-the-shelf AI assistants could get Denmark about halfway towards the national target at the price of 900-1,900 million DKK per year.



Note: The four points in the figure correspond to the four scenarios described in the main text – that an off-the-shelf AI-system license costs 2000/4000 DKK and that the systems are used 60/80 per cent effectively by public employees.

Source: Own calculations based on O*NET, the US Bureau of Labor Statistics and Statistics Denmark.

With these assumptions, we calculate the number of public employees for whom the value of the working time freed exceeds the cost of a licence. The results are shown for each scenario in Figure 5 below, with the range between them highlighted to show the uncertainty. Across the four scenarios, a licence is cost-effective for approximately 435,000–516,000 public employees (full-time equivalents), at an annual licence expenditure of DKK 0.9–1.9 billion.

Of the theoretical potential of 28,000 FTEs from secure and compliant off-the-shelf AI assistants, we estimate that 11,000–17,000 FTEs of working time could be realised cost-effectively. The working time freed has an estimated wage-equivalent value of DKK 7–10.5 billion annually, net of training time.

The results show four points of optimal investment in off-the-shelf AI-systems given the assumptions in the four scenarios. The main take-away is that even the most optimistic scenario only gets about half-way to the national target of 30,000 FTEs. Reaching it thus entails use of tailor-made systems in large shares of public processes already in use by 2035.

AI in critical public processes creates sovereignty dilemmas

Operational resilience requires that critical public functions do not depend on technologies, models or compute that external actors can unilaterally restrict or withdraw; data sovereignty, that governments retain control over sensitive public data; and autonomy, that they can switch technologies or providers when circumstances change.

Dependencies take several forms. Systems hosted outside Europe may be exposed to decisions by foreign governments, and even infrastructure in Europe may be operated by companies under foreign jurisdiction. Limited competition and vendor lock-in can make switching costly, and access to the models themselves can be restricted, as the temporary US export controls on Anthropic's leading models in June 2026 illustrate.

Control over advanced chips, computing power and AI models is increasingly intertwined with economic and foreign policy, which makes these dependencies geopolitical: a commercial dependency can become a strategic vulnerability when regulation or international relations change. The safeguards discussed in this paper must therefore be accompanied by

scrutiny of the infrastructure underlying AI systems.

European providers hold only around 15 per cent of the European cloud market, and Europe depends on foreign providers of leading AI models and advanced chips. Expanding Europe's options requires substantial investment in computing, datacentre and energy capacity; growing public-sector demand could help create a market for European alternatives.

Such capacity will take time to build and will not remove the need for global technology. Governments must therefore continue to assess where control is essential and where dependencies can be accepted and managed – diversifying providers and keeping models and infrastructure substitutable where possible. Infrastructure, access and control will thus be central to how Denmark captures the productivity gains while maintaining a legitimate, trusted and resilient public sector.

30,000 FTEs is an achievable target – if Denmark builds for it

Denmark faces decades where demand for welfare services grows faster than the workforce that supplies them. Our results suggest that AI systems can close much of the gap: a feasible potential of 90,000 FTEs – a capacity corresponding to DKK 54 billion a year in wage terms – against a national target of 30,000. The target does not require technological optimism; it requires realising around a third of what is already feasible with today's technology under today's administrative standards.

It does, however, require the right kind of effort. Over two-thirds of the potential depends on tailored AI systems built to the standards that make public administration lawful and trusted. While the remainder is technically reachable with off-the-shelf assistants, going from theoretical potential to realised gains reduces this further. We estimate that cost-efficient implementation of off-the-shelf AI-assistants would realistically increase capacity by about 15,000 FTEs for an estimated investment of 900 to 1,900 million DKK.

Most of the gains from off-the-shelf AI assistants are in education, research and health where time saved is inherently harder for management and policymak-

ers to consolidate and redirect. This emphasises the point that policymakers should be explicit about what they want freed capacity to achieve. Redirecting capacity to new tasks or reducing staffing needs may point to automating administration, whereas more face-to-face time with citizens could also be gained in education, care or health. Either way, large parts of public-sector employment must be covered in any policy aimed at reaching the national target, and the necessary systems must be scoped, built and in daily use well before 2035.

Finally, because these systems will sit inside critical public processes, control over the infrastructure beneath them becomes part of the policy decision. Overreliance on vendors outside Europe puts Denmark at a strategic disadvantage and creates new geopolitical pressure points. While this reliance cannot be avoided completely, a decision to invest in the underlying infrastructure can increase control over the technology stack on which public processes will run in the future.

Welfare states are under pressure everywhere, and Denmark's is no exception. AI technology offers a rare chance to shrink the time absorbed by administration while preserving the quality and legitimacy of public services. Seizing it requires no technological breakthrough – only broad and sensible implementation, in which AI systems are made to meet the standards of good administrative practice rather than allowed to redefine them.

Endnotes

[1] Danish Ministry of Digital Affairs, “Kunstig intelligens skal frigøre mere tid til det vigtige”, 10 June 2025. The announcement sets out a joint ambition to free up at least 50 million hours, equivalent to at least 30,000 full-time equivalents (FTEs), by 2035, with a significant share realised by 2030.

[2] This analysis distinguishes between off-the-shelf AI assistants and tailor-made AI systems. The former help public employees with general tasks using information provided by the employee, requiring minimal system integration. The latter address tasks requiring greater consistency, transparency and oversight throughout the process to comply with public-administration requirements. The distinction is elaborated on p. 4.

[3] Shakked Noy and Whitney Zhang (2023), “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence”, *Science* 381(6654), pp. 187–192.

[4] Erik Brynjolfsson, Danielle Li and Lindsey Raymond (2025), “Generative AI at Work”, *Quarterly Journal of Economics* 140(2), pp. 889–942.

[5] Daniel Schwarcz, Sam Manning, Patrick Barry, David R. Cleveland, J. J. Prescott and Beverly Rich (2026), “AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice”, *Journal of Law and Empirical Analysis* 3(1).

[6] IMF (2024), *Gen-AI: Artificial Intelligence and the Future of Work*, Staff Discussion Note SDN/2024/001.

[7] Paweł Gmyrek, Janine Berg and David Bescond (2023), *Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality*, ILO Working Paper 96.

[8] OECD (2026), *The OECD AI Exposure Measure: Mapping the OECD AI Capability Indicators to Occupations*.

[9] Folketingets Ombudsmand, *Myndighedsguiden, overview no. 13, “Forvaltningsretlige krav til offentlige it-systemer”*.

[10] Folketingets Ombudsmand (2025), *FOB 2025-13 (Vurderingsstyrelsen); opening letter, case no. 25-05712, October 2025; and Ejendomsvurderinger – Vurderingsstyrelsens mulighed for at forklare boligejere om bagvedliggende beregninger mv. for en vurdering*.

[11] DR, “Jonas var ikke i tvivl, og han fik ret: Kommunen brugte AI og brød reglerne”, 5 July 2026.

[12] Ingeniøren, “Kommune brugte AI i smug med lukkede øjne: Vagthund gransker nu sagen”, 23 June 2026.

[13] US Department of Labor, O*NET Resource Center (2023), O*NET Database, release 27.2.

[14] Since time studies of Danish public-sector employment are almost non-existent, we assume that the distribution of time spent on tasks within occupations follows the pattern in the US dataset. Where possible, we use Danish time studies to calibrate our results.

[15] Authors’ calculations using Statistics Denmark’s LONS20 data for 2024.

[16] “Administration” covers administrative and clerical occupations across the public sector—office clerks, secretaries, economists, lawyers, HR staff and similar occupations—regardless of the subsector in which they work.

[17] Authors’ calculations using Statistics Denmark’s LONS20 data for 2024.

[18] McKinsey & Company (2023), *Det økonomiske potentiale af GenAI i Danmark*, figure 10, pp. 26–27.

[19] Authors’ calculations applying the open occupational scores from Paweł Gmyrek et al. (2025), *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*, ILO Working Paper 140, International Labour Organization and NASK.

[20] Authors’ calculations applying the open occupational task-exposure classifications from Tyna Eloundou, Sam Manning, Pamela Mishkin and Daniel Rock (2024), “GPTs are GPTs: Labor Market Impact Potential of LLMs”, *Science* 384(6702), pp. 1306–1308.

[21] Théodore Renault (2025), *The Impact of Artificial Intelligence on Denmark’s Labor Market*, IMF Selected Issues Paper No. 2025/119.

[22] Ole Teutloff, Johanna Einsiedler and Fenja Søndergaard Møller (2024), *Large Language Models and the Danish Labour Market*, Statistics Denmark Analysis No. 2024:02.

[23] Kim Abildgren and Rasmus Mose Jensen (2026), *Artificial Intelligence Can Boost Productivity in the Danish Economy*, Danmarks Nationalbank Analysis No. 6.

[24] Ninja Ritter Klejnstrup and Anders Gotfredsen (2024), *Stort potentiale for automatisering af danske jobs*, Kraka-Deloitte.

[25] Boston Consulting Group (2024), *GenAI – et væsentligt potentiale for den danske offentlige sektor i 2040*, prepared for Dansk Arbejdsgiverforening.

[26] Isabel Atkinson and James Browne (2024), *The Potential Impact of AI on the Public-Sector Workforce*, Tony Blair Institute for Global Change.

[27] The Tony Blair Institute made the same assumption in 2024. AI courses offered to public employees in Denmark range from a few hours to full days. Two workdays per year would cover these courses as well as self-directed learning on the job.