February 2024

apolitical

Microsoft

Transforming Public Sector Services Using Generative AI

Global Case Studies

Foreword

Robyn Scott

CEO and Co-founder
Apolitical

Olivia Neal

Director
Public Sector Center of Expertise
Microsoft

If 2023 was an unprecedented year for AI, 2024 will be an unprecedented year for democracy. Nearly 4 billion people across 69 countries will vote. As the democratic process takes center stage, so will the issue of government delivery — from the quality of delivery influencing whether an incumbent wins, or the delivery sprints of new governments in their first 100 days.

AI presents a hugely important opportunity for governments hoping to deliver. Around the world public sector leaders are asking: How can the rapid gains in generative AI help civil servants raise the bar of public services? And how do we integrate AI into the machinery of government?

This report shares some early and encouraging case studies on the adoption of generative AI. From Canada to Tokyo to Portugal, governments have been leveraging AI quickly to deliver more citizen-centric services. These case studies are also evidence that governments can be leaders in technological innovation, challenging perceptions about slow adoption. Here are great examples of local and national governments seizing the opportunities AI offers without compromising on the management of its risks.

Delivering effectively in the public sector is both a responsibility and an opportunity like no other. It means working with taxpayers' money, delivering for all rather than a portion of a population, and safeguarding people's privacy and wellbeing. This means understanding and using the best tools available, in a responsible and secure manner.

For public sector organizations, the adoption of generative AI presents a unique opportunity to act and improve the quality of public services, increase operational efficiency, and meet the growing expectations of people and businesses. The case studies of civil servants and their teams shared here are leading examples to show how this can be done.

As with all tools, implementing generative AI needs leaders who are thoughtful, focused on the needs of their users and willing to embrace new opportunities. I'm grateful to those here for sharing their experiences, so that we can all learn from them. This time next year, there will undoubtedly be many more joining this group, to show how they too can go further and deliver better for all of us.

Reimagining public services with generative AI

The rise of generative AI

Over the last decade, AI has embedded itself into every aspect of our lives, whether we know it or not. But rapid advancements in generative AI and the release of already numerous groundbreaking tools have catapulted the technology to new levels of ubiquity. ChatGPT, for instance, attracted over 100 million users in its first month, making it the fastest-growing application in history at the time.

The sudden explosion of generative AI into the public sphere is down to many converging factors, including advances in deep learning technology, more data to train models, and the democratization of these tools. Its ease of use and ability to quickly generate new content has seen it grow hugely in popularity and we're already starting to see it reshape the way many sectors operate, from healthcare and education to business and government, the focus of this report.

How is generative AI impacting governments?

As technologies evolve, so do people's expectations of governments. Since Covid-19, it's become clear that public services urgently need to accelerate their digital and technological capabilities to meet public demand. Citizens increasingly expect faster, personalized and data-driven services, similar to those in the private sector. Many departments are also intrinsically motivated to take advantage of the efficiencies seen in other industries that have successfully harnessed the potential of technology.

This is where generative AI, when managed responsibly, presents such an opportunity. There's increasing evidence that new generative AI tools can help augment government productivity and improve the quality of public sector services.[1] If integrated carefully with appropriate guidance, these tools have the potential to reduce administrative workload, increase the efficiency of services, help public servants reach better, faster decisions and much more.

A brief history of AI

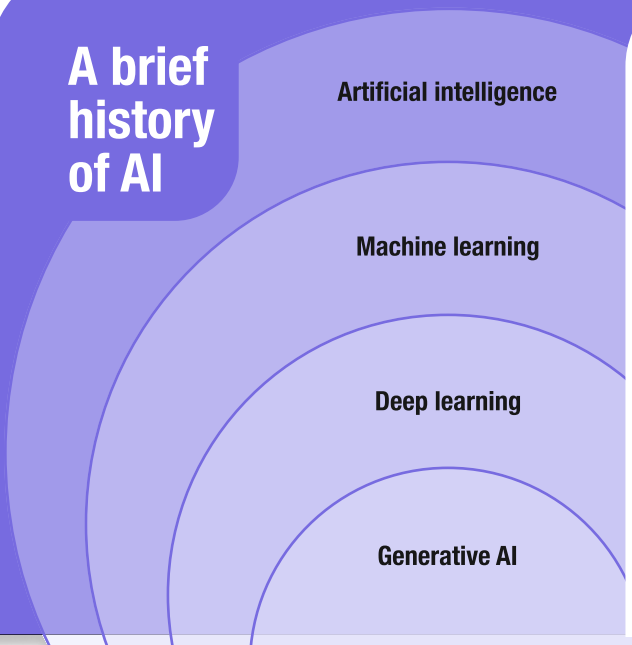

1950s
Artificial intelligence
A field in computer science focused on creating machines that can mimic human-like behavior. For example, speech recognition, decision-making and natural language understanding.

1959
Machine learning
A subset of AI that enables machines to sift through large amounts of data to find patterns and improve upon that data to make decisions or predictions.

2017
Deep learning
A subset of machine learning based on neural networks that allows a machine to train itself to perform a task.

2021
Generative AI
A subset of artificial intelligence that uses techniques (such as deep learning) to generate new content like text, images and videos.

Generative AI is set to enhance how governments operate

AI is already helping governments and other organizations tackle some of society's biggest challenges. From innovations like predictive healthcare to personalized job-seeking advice[2] and wider access to education[3], the technology is powering new solutions to improve people's lives.

$1.75 trillion per year estimated in productivity gains from Generative AI for public services by 2033.[4]

Generative AI could similarly usher in a new wave of transformation for government operations, bringing significant economic and productivity benefits. But adoption has been understandably tentative as officials debate the appropriate guardrails needed to ensure security and reliability when using these tools. Nevertheless, early pilot studies in a few governments, and "bottom-up" innovations driven by individual employees and small teams are showing some promising results.

Trends in government

- Improve speed and efficiency of public services and engagement
- Optimize employee time by reducing manual processes
- Boost research and writing skills
- Draw insights to empower decision-makers
- Improve internal knowledge management
- Support code writing to accelerate development

Whilst use cases for generative AI in government are still fairly new, there are trends emerging that tell us a lot about how public services could derive value from these tools.

Improve speed and efficiency of public services and engagement

Chatbots powered by generative AI are being trained to handle routine public and citizen enquiries and route more complex enquiries to the right department, relieving pressure on overstretched staff and providing personalized, 24/7 services.

Optimize employee time by reducing manual processes

Generative AI can free up time by assisting with repetitive tasks and augmenting workflows like document processing or case management. For example, the US Department of Defense is prototyping a tool to streamline their procurement process.

Boost research and writing skills

Tools like Copilot are acting as creative aides, helping public servants ideate and produce first drafts of briefings, speeches, memos and citizen guides. They can also find and summarize large amounts of information for faster research.

Draw insights to empower decision-makers

AI helps governments rapidly draw insights from reports and data by identifying patterns and anomalies. In turn, these evidence-based insights can inform policy formation and even provide evaluations of policy impacts.

Improve internal knowledge management

Chatbots don't just have to be citizen-facing. Some governments are using internal chatbots to answer employee questions about department policies and find information created by other teams.[5]

Supporting code writing to accelerate development

Tools like Github Copilot are being used to help produce, check and debug programming code. This is particularly useful as many governments are focusing on improving their digital services and retiring legacy software.[6]

Governments are rapidly addressing privacy and security concerns

Despite generative AI's potential advantages, its use in government raises valid concerns, particularly regarding privacy, security and transparency. This includes carefully considering how sensitive government data is processed and stored on third-party websites.

To address this, some governments have started releasing guidance documents that lay out best practices for working with the technology. These principles empower public servants to take advantage of generative AI whilst understanding how to handle data in a responsible way. However, these guidelines will need to evolve continually – incorporating the emerging best practices on what's working around the globe – as the uses and capabilities of generative AI develop.

Work between governments is continuing to build momentum, and in December 2023 the G7 leaders endorsed the 'Hiroshima AI Process Comprehensive Policy Framework'.[6] This is the first international framework to include guiding principles and a code of conduct aimed at promoting safe, secure and trustworthy advanced AI systems.

Governments are also increasingly developing secure, customizable applications using foundational models from external providers, like OpenAI, whilst working with providers to create clarity and trust in how data is used.. This reduces the chances of information leaks and means governments can ensure their data and prompts won't be used to train public generative AI tools.

Perspectives from government

We interviewed experts from diverse locations and levels of government to provide a comprehensive perspective on the integration of generative AI into government work. By featuring public servants spanning state and city governments from Japan, Portugal and Canada, we've gathered a holistic understanding of the challenges, successes and unique insights associated with generative AI implementation.

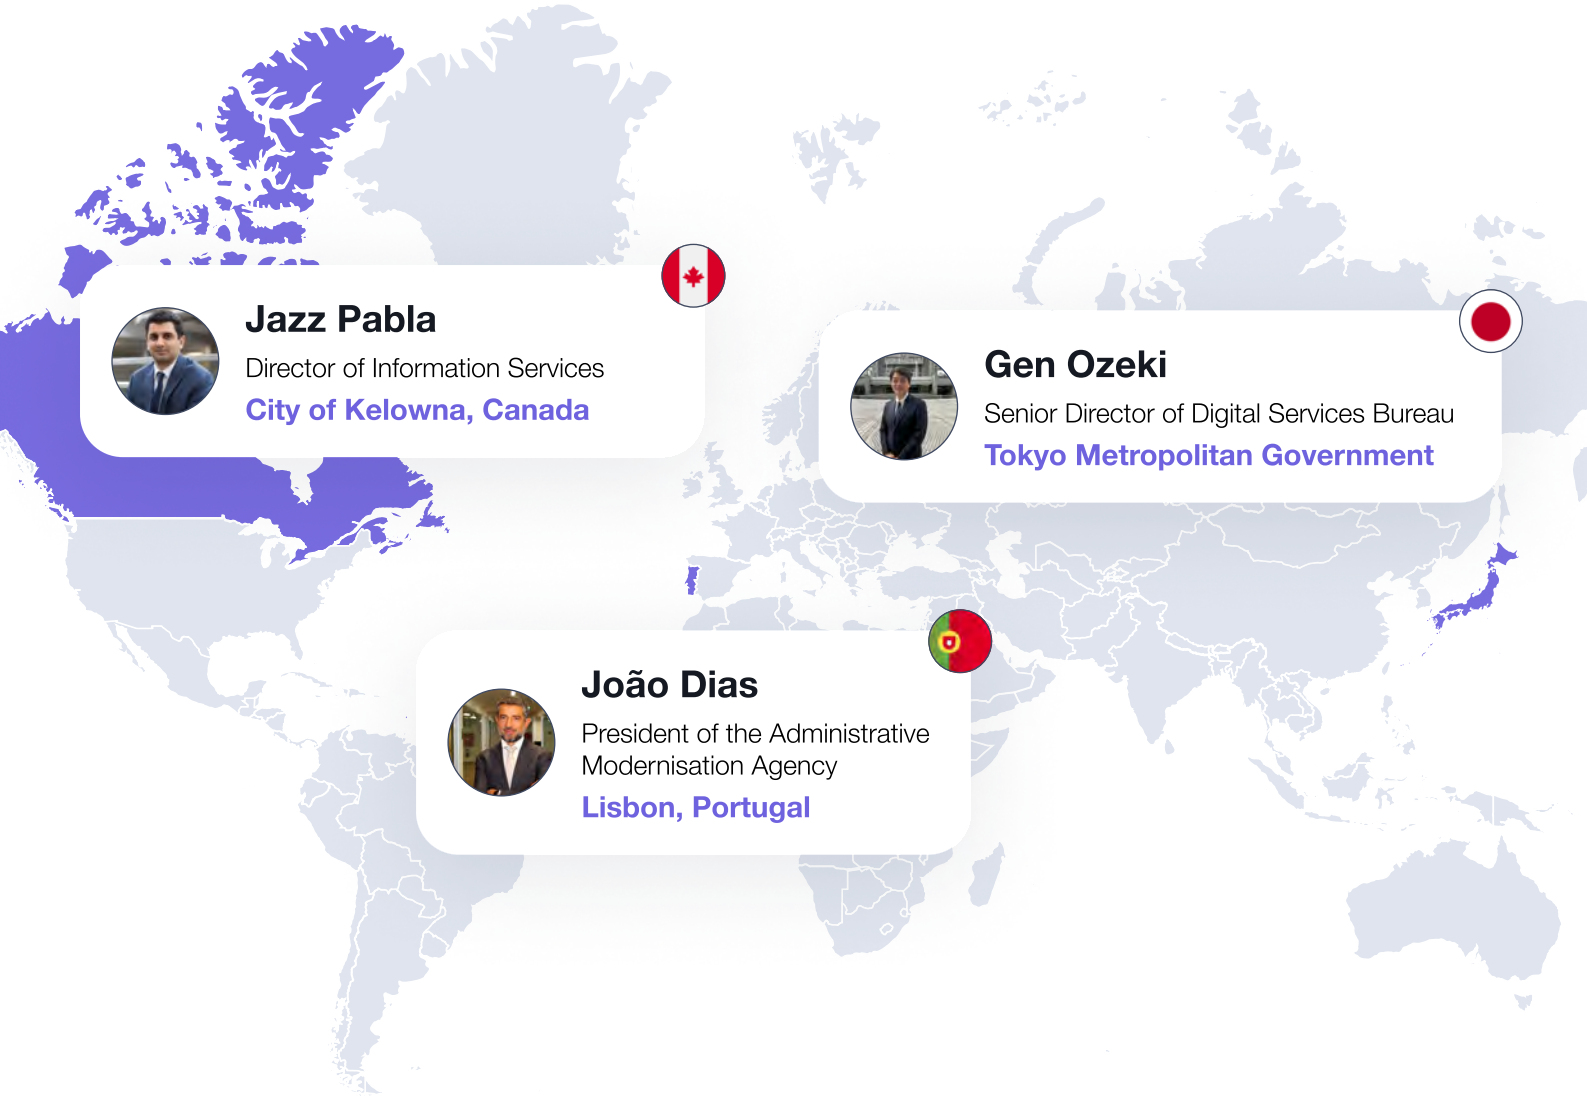

Jazz Pabla
Director of Information Services
City of Kelowna, Canada

Gen Ozeki
Senior Director of Digital Services Bureau
Tokyo Metropolitan Government

João Dias
President of the Administrative Modernisation Agency
Lisbon, Portugal

Transforming city services in Kelowna

How innovative partnerships and a commitment to responsible AI have rapidly accelerated Kelowna's permitting process

The city of Kelowna is on a mission to become 'the best digitally managed city in all of Canada'. Throughout 2022, the Intelligent City team has focused on automating online services for citizens. Now, after receiving a $350,000 grant from the government of British Columbia, the city is developing a generative AI chatbot that accelerates the planning permissions process for new homes and apartments.

We spoke to Jazz Pabla, Director of Information Services for the City of Kelowna, about the team's approach to innovation, the project and the lessons learned along the way.

Leveraging generative AI to tackle housing shortages

Even before ChatGPT became mainstream, Kelowna had started to use AI to create better citizen-facing solutions. "Generative AI was just the next solution and we walked right into it," says Pabla.

The city's latest venture to accelerate the housing permit approval process is part of an effort to address Kelowna's growing housing shortage. "It's not the role of government in Canada to build new houses," Pabla explains, "But if we accelerate the process, then developers aren't waiting six months for a permit. That plays a large role in increasing the housing supply."

To accelerate the process, the city has developed two solutions. The first solution – the 'explore journey' – is a chatbot from *zammo.ai* which answers people's permitting questions and helps them understand what they can and can't do. The second solution – the 'apply journey' – uses a Copilot chatbot to guide people through the application process.

> "Once you find the pain points, you can start from there."

Kicking off the project

Kelowna's IT team started by collaborating with the Planning and Development department to understand the needs and common questions citizens had about permitting.

The teams ran two-week Agile sprints to start developing solutions, sharing feedback and working side-by-side. "This project wouldn't have been

possible without the planning team," reflects Pabla. "They were constantly giving us feedback on how the existing process worked and what the new process needed to look like." Together with IT and Planning, Pabla included non-technical members in his team with a deeper knowledge of processes and user behavior. "This helped us know what was important from that side of the experience too."

A final component in these early stages was stakeholder buy-in. Both the internal planning team and external property developers were given the opportunity to test generative AI solutions as they were being developed. "We wanted businesses who owned these processes to feel they were really getting a say in what the solutions looked like." It also allowed non-IT stakeholders to see the benefits these tools could offer. "There can be a sort of public fear around generative AI and what it's going to do," says Pabla. So the team focused on helping others make a connection with the technology by showing how it could help.

How did the team develop the right infrastructure, data and skills to work effectively with AI?

According to Pabla, there was initially a steep learning curve for the team as they migrated relevant business services to the cloud. But it became easier once development cycles were up and people could see how the technology worked. "Now the team can't imagine a world without us being in the cloud and without using generative AI."

It was also important for the team to approach data in a structured way as they moved it to the cloud infrastructure. "You shouldn't strive for perfection in your data to start this journey," he stresses. "But have a plan about where you want to be and then start to slowly move data to each of the generative AI solutions you want to use."

The team took a trial-and-error approach to developing skills. "When we started using generative AI, there wasn't any training available because it was so new." So part of the permitting solution was just the team learning. Pabla explains how they collaborated with Microsoft's AI team. "If Microsoft had learned something and multiple people agreed that this was the way to do it, they would pull us in and explain it to us," he says. "There wasn't any sort of formal training – we were literally building the plane as we were flying it!"

But constantly adapting as generative AI's capabilities grow isn't easy. Pabla admits that it was an inherent challenge, especially when dealing with other levels of government — "It's tough to get everyone on the same page." This challenge won't be solved overnight, but Pabla says focusing on the purpose of the project rather than just the technology helps. "When you have everyone focusing on the same problems and generative AI is simply a solution, things start to move a little bit quicker."

How was the sensitive issue of data privacy and security handled in the project?

The team took multiple approaches to ensure the new tools handled data in a responsible way. To start with, Kelowna's internal procurement, legal, risk and IT teams came together with the Montreal Institute of AI to create their first iteration of a Responsible AI Framework. "This

includes the privacy aspect of working with personal information," says Pabla. The framework will be shared with internal staff and in 2024, they will release an AI registry to show people exactly where the Kelowna government is using generative AI, how they're using it and where the data is coming from.

There was also a big focus on cybersecurity. "The thing with generative AI is, you know it might be used for nefarious purposes," says Pabla. "So having a good cybersecurity framework that's agile and has continuous improvement is vital." For example, the team runs several security tests with all their chatbots before they go live.

But ultimately, when it comes to building trust in AI, Pabla believes that having a strong, collaborative relationship with the AI provider is paramount.

How did the team launch the tool and how are they measuring success?

The first version of the permitting 'explore' tool was soft-launched earlier in 2023. This allowed the team to see the types of questions people were asking and assess the responses the chatbots were giving. Instead of a grand reveal, they shared updates early and often throughout the development process. This allowed plenty of feedback and kept the project on track.

While the permitting project is still in its infancy, automated hours are a good early indicator of success which the team plans to measure when they have more data. Measuring automated hours starts by talking to frontline staff about how long a particular process takes, then using a new generative AI solution to automate the task. The team can then calculate how much time has been saved in automated hours.

Another possible indicator is measuring the volume of questions answered by the chatbot service. While there isn't sufficient data yet, insights can be drawn from a comparable case at City Hall. For instance, in 2023, their phone line digital assistant efficiently addressed 79,000 citizen queries. As a result, City Hall answered around 96,000 citizen queries in total — almost 9 times more than the previous year.

How do you see this AI initiative evolving in the future?

"I think any enterprise system that's used by internal people will have the ability to have generative AI attached to it," says Pabla. "In the future, rather than logging into a system to input data, you'll be asking Copilot to input it for you — I want to start that process."

Beyond that, Pabla says he's excited about "Continuing to provide digital systems and generative AI solutions that will benefit our citizens and, equally importantly, our internal staff."

What advice would you give to other public sector organisations considering similar Generative AI solutions?

"Don't think that just because you're in the public sector, you can't innovate and you can't do these things. Just find the most painful or mundane process and show people there's possibility here."

Boosting administrative efficiency in Tokyo Metropolitan Government

How ChatGPT has improved government operations without compromising security

In 2016, the Japanese government introduced its Super Smart Society initiative (Society 5.0) as an aspirational vision for the future of Japan. This society would leverage AI, big data and other high-tech innovations to improve citizens' standard of living and address various social issues. With this long-term goal in mind, as was stated in the 'Integrated Innovation Strategy 2023', Japan has been keen to explore how generative AI could be used in government.[7]

In June 2023, Governor Yuriko Koike announced that the Tokyo Metropolitan Government (TMG) planned to use ChatGPT to assist in tasks like text creation, idea generation and other clerical work. The TMG introduced its own private version of ChatGPT to do this.They also create a project team to test the efficiency of the tool and create guidelines for its use.

We spoke to Gen Ozeki, Senior Director for Planning and Coordination at the TMG's Bureau of Digital Services, to explore how ChatGPT has been received, their approach to data privacy and recommended best practices.

Integrating ChatGPT into government operations

When asked why the TMG made the decision to introduce ChatGPT, Ozeki had a simple answer: "We ultimately want to enhance the level of services provided to citizens."

As they work towards enhancing services, TMG is determined to innovate without compromising on security. This commitment can be reflected in a wider, pressing question: How can governments take advantage of generative AI's capabilities without compromising on security or risking improper use?

To answer this question, the Digital Services team focused intently on security, particularly in the early stages of integration. This included running pilot studies, interviewing research groups and consulting with experts to create detailed usage guidelines before releasing their ChatGPT tool to staff.

Building on these security measures, the TMG implemented an additional safeguard: a mandatory declaration form for staff intending to use ChatGPT. It serves as a reminder of responsible usage and outlines four foundational guidelines:

1. Don't input highly confidential information
2. Don't generate text that infringes copyright
3. Always verify and validate responses given by AI independently
4. If you directly use responses generated by AI, indicate these responses were created using AI

This approach ensures that staff are fully aware of how to interact with generative AI at work, aligning with the TMG's commitment to responsible and secure technology use.

How did the TMG choose a generative AI provider and train their staff to use ChatGPT?

TMG's commitment to data protection was a major factor in choosing a provider. They decided to partner with Microsoft for enhanced security, using Azure OpenAI Service and their GPT-3.5-Turbo model (a large language model optimized for conversational interfaces). Key to this decision was the assurance that employee prompts would neither be used to train Microsoft's AI models nor stored on Microsoft's servers, adhering to TMG's strict information security standards.

As part of their commitment to security, the Digital Services team is ensuring all employees receive comprehensive training on ChatGPT. They've developed an online portal offering training on digital tools, including lectures on ChatGPT, and have prompt engineers who work as instructors, hold study sessions and run workshops for an in-depth understanding of ChatGPT's functionality.

The team also prioritized continuous feedback on tool usage. Two months after the introduction of ChatGPT, they conducted an extensive survey among 50,000 employees. Now, they've established a portal for employees to regularly submit opinions and requests about the usability of digital tools provided by the Digital Services Bureau.

What frameworks and guidelines have been used to oversee the responsible use of ChatGPT in the TMG?

Before releasing the ChatGPT tool to staff, the team wanted to create detailed, beginner-friendly guidelines to optimize the use of ChatGPT and reduce the chance of security breaches.

“We sought advice and consensus from an expert group of scholars, legal experts, developers and service providers,” Ozeki explains. This allowed the team to assess the strengths and weaknesses of the tool, identify potential safety breaches and gather diverse perspectives on the topic. The Digital Services team also ran an ‘ideathon’ to share thoughts about the kind of work ChatGPT could be used for. This session identified about 200 use cases, with idea generation and task efficiency emerging as the most prevalent among them. The importance of precise prompt wording was also emphasized. Additionally, a pilot study involving 300 employees was conducted to collect feedback and various usage examples.

From these insights, the team have started creating a detailed set of guidelines for staff. It includes advice about how to handle confidential information safely and examples of tasks that ChatGPT could be used for. Templates for effective prompts are also being developed.

According to the Ozeki, creating usage guidelines for staff turned out to be the most challenging aspect of introducing ChatGPT because of the need to satisfy various stakeholders. The team tackled this by seeking advice from external experts and gathering knowledge from academic societies like the Japan Deep Learning Association. “We now believe the quality of our usage guidelines are actually what set us apart from other cities,” he reflects. However, Ozeki acknowledges that there are still concerns about the potential inaccuracy of some ChatGPT answers and the influence of unconscious bias. Mitigating these risks will be a continuing focus.

> "We share our implementation journey on the Japanese Blog Site ‘note’ as a reference for other cities.”

The TMG's guide has been published online for wider access. They're also actively engaging in speaking events to share their experiences.

How is the TMG measuring the effectiveness of ChatGPT and how do they see its usage evolving in the future?

The TMG is focused on the value of qualitative feedback to measure success in its early stages. The Digital Services team are conducting surveys to gather feedback about how much time staff members feel was saved by using ChatGPT for processing tasks. But they're also interested in the quality of the generative AI output.

> "We want to make sure staff feel the quality of their work has improved as well as saving time."

If there becomes a need to invest in expanding the capabilities of generative AI, the team say that the Financial Bureau, which is responsible for the budget, would request an explanation of the efficiency gains and measure the impact of streamlining operations.

When asked about future plans, Ozeki says the TMG is currently considering two directions to evolve its use of ChatGPT. First, they want to enable the use of ChatGPT in a wider range of tasks by promoting services that incorporate ChatGPT into various business software. Secondly, they want to enhance ChatGPT by incorporating "Specific administrative regulations, white papers, web pages and other information to gain more accurate answers."

What advice would TMG give to other public sector organizations considering similar generative AI solutions?

Whilst the TMG hopes to maximize AI's positive impact on society, privacy and security are critical concerns for them. What enabled them to take this step was building trust and confidence by working with a provider who could effectively eliminate the risk of information leaks.

They also recommend implementing system-level measures and creating clear guidance to reduce the chances of privacy breaches. Data security is a crucial foundation for any government, and setting up these safety measures means officials can have greater peace of mind when using ChatGPT.

Supporting citizens in their digital interactions with Portuguese public services

How a virtual assistant has helped citizens navigate their digital authentication portal

Portugal is no stranger to technological transformation. In fact, OECD indices have recognized it as one of the global leaders in digital government.[8] Guided by the Administrative Modernization Agency (AMA), the government's journey with generative AI began by exploring its potential to improve digital services for citizens. This led to the launch of a pilot program in late 2023. The program introduced a generative AI-powered virtual assistant to help citizens navigate their digital authentication services.

We spoke to João Dias, President of AMA's Board of Directors, about what he learned from the pilot, the importance of collaboration and the benefits of starting small.

Using generative AI to reduce friction and bring clarity to citizen services

At its core, AMA is driven by a simple guiding principle: a belief in the power of technology. The agency works cross-governmentally to promote a culture that's curious about technology and enable the digital transformation of Portuguese public services.

It was this principle that motivated AMA to start exploring uses of generative AI. "We saw the potential of generative AI in two big domains," explains Dias. "Improving internal efficiency and improving the relationship between the citizen and the state."

A pain point AMA wanted to tackle was customers queuing for in-person government services. The process was frustrating and often, citizens would have to visit a service desk multiple times after bringing the wrong information. "They'd have to go through the same process just because they'd missed one document," says Dias, "It was so painful." AMA wanted to create a virtual assistant to guide citizens so they knew what to do and what documents they'd need.

> "We have to make use of the best tools around us."

How did AMA kick off the project and ensure stakeholder collaboration?

"It was important for us to start small and focus on one service before scaling up," explains Dias. "That way, we could see what worked and what didn't." The team decided to start with their digital

authenticator Mobile Key Solution, an e-portal that gives citizens access to their digital public services. They planned to train the virtual assistant using common questions and then pilot the avatar on the platform.

Collaboration was a key factor when designing the pilot. Dias emphasizes that it's a crucial part of AMA's culture. "It's in our DNA to work with other entities — we are a cross-governmental organization." He explains that if you want to put the citizen at the center of your services, you have to start by breaking down silos because their problem might span multiple agencies. Dias also recognizes the value of working with the private sector to make the most of their cutting-edge technology. AMA worked with Microsoft's Azure OpenAI services and collaborated with AI startups *Defined.AI* and *Daredata* to leverage their expertise in voice processing and converting text to speech. This helped create a more sophisticated avatar that feels like you're talking to a real person.

But achieving this collaboration wasn't always easy. Dias admits that one of the biggest challenges was ensuring cohesion between internal and external stakeholders with different priorities, legacy systems and cultures. There wasn't a simple solution, but AMA found that developing relationships was crucial, as was ensuring each department felt that working together was a rewarding experience.

The team created this environment by laying a foundation of clear and open communication channels. They arranged regular collaborative meetings and workshops to promote idea-sharing and joint problem-solving, fostering a sense of shared purpose. Tailored stakeholder engagement plans were crucial in recognizing and celebrating the unique strengths of each entity involved. This included highlighting specific contributions, milestones and areas of focus to instil a sense of pride and accomplishment in each department and private organization, reinforcing the notion that individual effort was key to overall success. Finally, through regular updates, transparent reporting on challenges and inclusive decision-making, they created an environment where everyone felt "Informed, heard and valued".

How did AMA build trust in generative AI?

Building trust was another undeniable challenge. Many people saw the generative AI technology as a black box and were concerned about the risks of hallucinations and privacy breaches.

Dias says that AMA tried to address these concerns topic by topic to build stakeholder confidence. They started by creating their own curated database to train the algorithm. This was built using manuals and technical data from the Electronic Identification team to create an FAQ-structured database. The FAQs were assessed by senior members of the Digital Content and Usability team and Citizen and Business Contact Centre to ensure a citizen perspective. The database was then put through rigorous testing and some answers were tweaked for clarity and to accommodate how citizens referred to concepts. The results were instantly visible. "We quickly moved from 40% accuracy in the answers to 90% accuracy in two weeks! Everyone was very impressed and could see that it actually worked".

> "Every answer was one less citizen queuing or waiting for a phone call."

AMA also made it clear that this was a pilot, that citizens were talking to a machine and that it might make some mistakes. This level of visibility helped manage citizen expectations. At the same time, AMA assembled a team to monitor the quality and accuracy of the avatar's answers which added another level of precaution. These interventions

and AMA's decision to address concerns one by one significantly changed the perception and trust of the technology during the pilot.

What frameworks were used to oversee the responsible deployment of generative AI during the project?

Before starting the pilot, the Portuguese government launched their ethical guide for applying AI in the public sector [9] along with a wider collection of AI guidance.[10] This was AMA's north star when designing the virtual assistant project.

AMA created an internal knowledge base to reduce the risk of privacy breaches and ensured that Microsoft Azure was certified by the Portuguese Cybersecurity National Center before working with them. In the future, Dias says they may make the virtual assistant transactional, not just informational. This would mean using ID authentication to generate more personalized responses. While this could be a major milestone in citizen engagement, Dias explains that it would bring new privacy challenges they'll need to solve before they reach that point.

Does the Portuguese government have plans to expand the use of generative AI?

After the success of the pilot, AMA will be expanding the virtual assistant to other government services. They're particularly interested in using it to help Portugal's immigrant population. Due to language barriers and complicated systems, many new residents find interacting with Portuguese public services difficult. AMA hopes to use their virtual assistant to provide information about relevant services like tax, social security and health. The avatar will speak between 15 and 20 languages, making communication far clearer. "Imagine what this will mean for inclusiveness and integration," reflects Dias.

AMA also plans to use generative AI to better understand user needs. "We have a huge amount of information from citizens", says Dias. "Imagine using generative AI to read all this feedback, cluster it, provide insights and then write first draft responses."

> "If we use the technology in the right way and increase efficiency, we can focus people on more complex questions, on creating and innovating."

What were AMA's biggest lessons from the pilot and what advice would they give other public sector organizations considering similar solutions?

The first piece of advice is to choose your use case carefully. Find a trade-off between something that's interesting for the user but also has a balanced risk level. From there, start small and be explicit that this is a pilot. "It's very important to be transparent that we're learning, we're testing the technology and we might need to change things." Gather feedback throughout the pilot phase and test, test, test in a controlled environment.

"We also learned you have to take a very careful approach when announcing the tool publicly", Dias adds. AMA decided to share a press release the night before launching their tool. This generated a lot of buzz but meant that when they ran a demo, the system was overloaded and crashed. "It was a bit embarrassing," Dias chuckles, "So now we'd recommend launching a demo within an hour of any press releases."

Finally, Dias' says attitude is one of the most critical factors. "You need to start with curiosity and not be dogmatic. Having blind faith in technology can be dangerous. But being too suspicious isn't helpful either. Of course, there are concerns about risks. So address them and test solutions but don't throw away the baby with the bathwater." Once people see the results, attitudes change and there's much more support to keep testing new use cases.

Takeaways for future initiatives

At this early stage, the results from Kelowna, Tokyo and Portugal suggest that generative AI has the potential to be a leading driver of positive change for the future of government. Whether that's streamlining a laborious permitting process, or helping navigate a nationwide portal, these small changes are big wins for governments and citizens alike. A large slice of the success of these initiatives can be credited to the thoughtful and agile approaches of their project teams.

These teams all saw generative AI as an opportunity to improve productivity and develop better services. Beyond a shared appetite for innovation in government, these case studies reveal a set of principles that can be useful for future generative AI initiatives.

1

Start with a focused approach

The best results come from starting small and focusing on a single, manageable initiative before expanding. This allows project teams to experiment, test and iron out errors in a controlled environment. If they're successful, data from these initiatives can be used to build internal confidence and gather support for expanding use cases.

Identifying pain points before choosing an initiative and working backwards can be a good place to start. By drawing attention to the problem that's being solved, it can reduce friction and frame the technology as a useful tool.

2

Emphasize collaboration and transparency

By fostering early and frequent collaboration, project teams can build stakeholder trust and deepen their knowledge of user needs. This includes involving multiple teams in conversations to ensure all perspectives are considered and letting stakeholders experience generative AI tools firsthand.

Aligning teams with competing priorities and pace is difficult, but strong project management and communication channels can make a big difference. As can focusing on the purpose of the project, or the problem you're solving together, rather than the technology itself.

3

Adopt an agile, iterative approach

Any AI initiative will have a steep learning curve. But a lot can be gained from adopting a trial-and-error approach, iterating and refining strategies based on regular citizen, staff and stakeholder feedback. This could include developing pilots in agile sprints or running a soft launch to assess results. It's also helpful for project teams to avoid striving for perfection and instead embrace transparency as they experiment and iterate.

4

Address risk and establish trust

Privacy and security risks are always important considerations in the public sector. Being transparent and acknowledging these risks early can make a big difference to public and stakeholder trust. As can taking a thoughtful, step-by-step approach to mitigate them.

Developing clear user guidelines can build an understanding of how to use the tools safely. It's useful to bring in a diverse mix of internal and external stakeholders like academics, industry experts and internal legal and risk teams during the guideline development phase. Creating regular review cycles will help ensure the guidelines stay updated as the technology evolves. Alongside this, a strong, agile cybersecurity network and collaboration with trusted AI providers will reduce the risk of data breaches and ensure that sensitive information isn't used to train AI models.

5

Invest time in staff training

Generative AI tools become more effective when users are well-trained. Training ensures employees can fully and safely use these tools, preventing misuse that could risk sensitive information. It also empowers people who might otherwise not use these tools.

Investing in detailed training increases confidence and responsible use of generative AI tools. Training can include usage guidelines, study sessions, workshops, online courses, or peer-led training. Regular feedback on tool usage helps monitor its adoption and identifies areas for more training.

6

Build support through results

It can be challenging to build support when initiatives are still in their infancy and results are limited. But sharing early-stage feedback and efficiency gains can make a big difference to stakeholder confidence and broader interest in generative AI. Early measurements of success could include the volume of citizen questions answered, the accuracy of these answers, the amount of time saved for staff or wider feedback from surveys. Only by testing the technology and achieving tangible results can we better understand both its capabilities and limitations in a government context.

References

Sources

1. **Strauss, Steven (2023).** "Some Emerging Hypotheses about Using Generative AI in Public Sector Operations." Accessed November 25, 2023.https://ssrn.com/abstract=4544943 or http://dx.doi.org/10.2139/ssrn.4544943

2. **Keio University (2020).** "Artificial Intelligence for Social Good." Accessed November 25, 2023. https://hdl.handle.net/20.500.12870/5671.

3. **Deloitte Insights.** "Use Cases of AI Adoption in Governments across the Globe." Accessed December 18, 2023. https://www2.deloitte.com/us/en/insights/industry/public-sector/global-government-ai-case-studies.html?id=us:2sm:3li:4diUS176002:5awa:6di:MMDDYY:author&pkid=1011069.

4. **Deloitte Insights.** "DHS Australia chatbot answers internal as well as external queries." Accessed December 18, 2023. https://www2.deloitte.com/us/en/insights/industry/public-sector/global-government-ai-case-studies.html?id=us:2sm:3li:4diUS176002:5awa:6di:MMDDYY:author&pkid=1011069#dhs-australia-chatbot-answers-internal.

5. **Actuaries in government.** "Leveraging Generative AI in Our Work." Accessed December 18, 2023. https://actuaries.blog.gov.uk/2023/10/27/leveraging-generative-ai-in-our-work/.

6. **Hiroshima AI Process.** "About the Hiroshima AI Process." Accessed January 22, 2024. https://www.soumu.go.jp/hiroshimaaiprocess/en/index.html.

7. **Cabinet Office Home Page.** "Science,Technology and Innovation. Council for Science and Technology Policy." Accessed January 22, 2024. https://www8.cao.go.jp/cstp/english/index.html.

8. **CAF.** "Portugal's Journey Towards Digital Progress." Accessed January 22, 2024. https://www.caf.com/en/knowledge/views/2021/07/portugals-journey-towards-digital-progress/.

9. **AMA.** "Guia Para a Inteligência Artificial, Ética, Transparente e Responsável." 2023. Accessed January 15, 2024. https://bo.mosaico.gov.pt/api/assets/mosaico-site/7127c349-4491-424f-927f-2ac1202f4766.

10. **Inteligência Artificial.** "Inteligência Artificial." Accessed January 22, 2024. https://mosaico.gov.pt/areas-tecnicas/inteligencia-artificial.

Resources

For more conversations, insights, ideas and resources on the use of AI in government from fellow public servants, you can explore Apolitical's AI in Government Community.

For the latest thinking about digital transformation and AI in the public sector, we'd recommend Microsoft's Public Sector Centre of Expertise and Public Sector Centre of Digital Skills.