



State and Local Government AI Roadmap

Executive Summary

Like no technology before it, artificial intelligence (AI) offers enormous potential to help governments increase efficiency, improve services and solve complex challenges. For years, governments have relied on AI for insights, such as providing early warnings of natural disasters and identifying potentially fraudulent benefits claims. Now the latest in AI, generative AI, can unlock even more opportunities to help governments improve efficiency and enhance services.

The process of evaluating and implementing new AI technologies can be informed by other changes in governments' digital transformation journey, such as the emergence of cloud computing, modern collaboration tools and remote work. Each brought new capabilities that could improve government performance, but each also had new challenges. Successful deployment required training technical, mission and elected leaders; careful strategies for adoption; responsible and innovative pilot deployments; and eventual full-scale adoption. Before adopting AI, governments need to understand how AI adoption impacts data privacy and security, equity and potential for bias.

Formulating policies that protect against potential risks while encouraging innovation is a balance governments need to achieve. Governments that have clear, forward-thinking and flexible policies on the use of AI will be in the best position to embrace the benefits and mitigate the challenges that AI technologies bring. However, having an environment that supports safe and effective AI innovation is only part of the story. Engaging with these ground-breaking and advanced technologies will also require employees to gain new skills and approaches. Government stakeholders will need to train for the competencies that will be required to make the best use of technology.

This paper defines a roadmap for government adoption of AI by identifying the government policies and operations most affected by the new technology and providing readiness guidance based on the work of leading governments that have already deployed AI-based solutions.

As AI technology and the government landscape continues to evolve rapidly, we encourage you to visit <http://microsoft.com/government> to learn about the latest technology and policy innovations in state and local government.



Roadmap for Government AI adoption

To download our practical guide on adopting AI in Government organizations, visit: <https://aka.ms/StateAIRoadmap>

Every government is on its own journey to adopting artificial intelligence solutions. This roadmap defines common steps governments should consider to start responsibly benefiting from AI innovations. They can be performed in any order.



1.

Set direction and prepare the workforce

- **Develop a plan for AI.** Issue executive orders, resolutions, state policies, or take other actions to direct government to assess and prepare for the responsible adoption of AI.
- **Implement an AI training program.** Provide officials and employees with a learning program that educates them on how to use the technology, the potential benefits, risks, and approaches for its responsible use.
- **Gain experience with generative AI.** Start using generative AI technology in low-risk situations, like creating draft documents using Microsoft Copilot and chatbots to facilitate government collaboration.



2.

Implement AI governance


- **Define a principles-based approach to AI governance.** Consider leveraging an AI governance board to guide the creation of AI principles. Refer to the Microsoft AI Principles for an example.
- **Create AI roles and responsibilities.** Determine the structure for governance and decision making centrally, within departments, and how they will relate to each other.
- **Develop and implement an AI use policy.** Define the rules by which government employees will be expected to request, develop, deploy, and use AI solutions – and leverage existing standards where they exist.



3.

Update procurement practices


- **Gain trust through certification and compliance practices.** Verify that technology providers obtain certifications and attest to compliance and governance standards.
- **Create a cloud and AI COE.** Bring together a virtual team with experience procuring cloud computing services to share best practices and train colleagues on how to procure cloud and AI services.
- **Create procurement guidance for cloud and AI services.** Update standards, templates, and guidance to reflect the technical nature of cloud and AI services.
- **Leverage AI for procurements.** Use AI to streamline the creation and review of proposals – and increase the diversity of bidders.



4.

Review data governance and security

- **Assess and automate data governance practices.** Ensure government data is properly labeled and secured, and that tools are used to automate the classification and protection of data as it is created.
- **Accelerate cloud migrations.** Move existing government data to the cloud to benefit from modern data protection and AI solutions.



5.

Foster AI innovation with trusted technology partners

- **Seek input on potential uses.** Encourage government departments and agencies to develop AI project ideas and use cases and seek guidance from technology and subject matter experts on how they can be responsibly implemented.
- **Hold an AI Discovery event.** With leading technology partners, convene technology and mission employees to gain hands-on experience with AI and develop a plan to incorporate it into government systems and services.
- **Create sandboxes for testing.** Stand up testbeds, or sandboxes, where employees can try using AI technologies in a controlled manner.
- **Encourage government pilots.** Encourage agencies to conduct pilot projects, starting with low-risk uses to build trust in the technology.
- **Scale with success.** Share use cases and pilot results across departments. Incorporate learnings, integrate and continue to innovate.



How governments are using AI today

AI can make sense of data like no technology before it, improving productivity and cost-effectiveness. In customer service, chatbots powered by generative AI can understand and respond to complex questions using authoritative answers from existing government documentation. In case management scenarios, AI can quickly review years of case file interactions to create summaries that help new case workers onboard quickly. And in eligibility scenarios, AI can help individuals identify potential additional benefits to which they are entitled – and streamline the application process.

As previously noted in the Roadmap, we see governments primarily using AI in two ways:

- External use cases to simplify the constituent experience when engaging with government services. For example, a chatbot residents can engage 24/7 to obtain authoritative information about government services; and
- Internal use cases to empower government employees with insights at their fingertips. For example, summarizing a case file with years of notes to help a new case worker quickly get up to speed on a matter.

Governments are already realizing the power of new generative AI solutions for both use cases, including the following:

Simplifying the constituent experience

New York City released MyCity Chatbot, which uses information published by the NYC Department of Small Business Services to respond to constituents in a timely manner.

Families and individuals seeking food security resources rely on state SNAP programs. States dealing with backlogs of SNAP benefit applicants and are looking for ways to improve the efficiency of how they answer constituent questions about SNAP programs and reduce the volume of callers or visitors that are forced to wait for information. Benefits solutions provider YoungWilliams recently deployed Priya, which can answer customer questions about individual cases, take applications, assist with forms and questionnaires, and provide practical and policy advice. Users are able to get quality assistance for their SNAP, child support, TANF, disaster assistance, and childcare cases anywhere, anytime, and with no wait. For program employees, Priya acts as a copilot to help case managers quickly find information, get assistance, and do real work on routine and complex matters. For more information, visit <http://www.youngwilliams.com>

The Port Authority of New York and New Jersey (PANYNJ) created Beacon, a chatbot that helps minority, women-owned and small businesses navigate procurement processes and proposals. This tool offers easy access to information on Requests for Proposals, PA policies and procurement processes. It has a user-friendly chat and Q&A interface, with settings that let users customize their experience and ensure guidance is reliable and informative. This increased interaction, supported by the user-friendly and informative chatbot, is expected to result in a higher rate of MWBE participation in PANYNJ procurement opportunities. <https://beacon.panynj.gov/>

Empowering employees with insights at their fingertips

Microsoft 365 Copilot. One of the most anticipated AI-based tools for worker productivity, Microsoft 365 Copilot is now used by more than 10 million people worldwide, and will be available to all US state and local governments in summer 2024. Microsoft 365 Copilot combines the power of generative AI with your organization's data – all in the flow of work – to turn your words into one of the most powerful productivity tools on the planet. It works alongside popular Microsoft 365 apps such as Word, Excel, PowerPoint, Outlook, Teams and more. Microsoft 365 Copilot provides real-time intelligent assistance, enabling users to enhance their creativity, productivity and skills.

Government services provider RedMane incorporated AI functionality into their leading child welfare solution to help triage reporting and summarize cases. This functionality enables frontline workers to more quickly and accurately take action on a case, supervisors benefit from enhanced insights into workforce performance and practice needs, fostering a culture of continuous improvement, and improves outcomes for staff and families.

For all law enforcement agencies, Draft One is a valuable tool that will be an immediate force multiplier for their workforce. Reporting is a critical component of good police work; however, it has become a significant part of

the job that officers commonly refer to as “burdensome.” Axon found that every week officers in the U.S. can spend up to 40% of their time — or 15 hours per week — on what is essentially data entry. Additionally, prosecutors and defense attorneys have shared that they welcome the quality improvement of Draft One reports as they can help support and accelerate the judicial process.

“Every single officer in the U.S. writes police reports, often every day and normally multiple times a day,” said Axon CEO and Founder Rick Smith. “As we’ve done with Draft One, harnessing the power of AI will prove to be one of the most impactful technological advancements of our time to help scale police work and revolutionize the way public safety operates.” For more information, visit <https://www.axon.com/products/draft-one>

The City of Kelowna, Canada has started adopting the use of Microsoft Azure AI to improve its efficiency. The technology is able to search for and find specific laws and documentation at the click of a button. Resident concerns and queries are now dealt with using a combined human and AI service, something that has improved the time it takes to resolve issues.

1 Set direction and prepare the workforce

Generative AI is arriving at a time of sustained change in government. The COVID-19 pandemic ushered in widespread changes in terms of new digital experiences. Overnight thousands of employees shifted to working remotely, shining light on manual and paper-based processes that required rework into online systems. At the same time, residents expected to engage with government services in the same way as commercial services. For example, think how easy it is to pull out your mobile phone and order groceries, check your online bank account or book a flight. In one survey, 85% of respondents expected digital government services to match or exceed those offered by the private sector.

Governments face these high expectations amid staffing shortages and resource pressures. According to the Rockefeller Institute of Government, there are nearly one million fewer state and local government employees now than in 2019. Finding budget to attract and retain staff is becoming more difficult. According to the National Association of State Budget Officers (NASBO), state government revenues are expected to decline in 2024 after four years of increases. At the same time, state and local governments have 928,000 fewer employees than before the COVID-19 crisis began. In a very real sense, government is expected to do more, with fewer people.

Many states have already implemented AI-based solutions, and in fact most government employees rely on AI technology in their daily work. Basic tools, such as spell check and email filters that identify junk mail, are powered by AI, as are the tools that help screen resident applications for benefits to suggest when answers may need to be corrected to speed up approvals. Generative AI technology will enable even more scenarios where government can both increase the efficiency of its employees and provide increasingly powerful self-service tools.

To realize the full benefit from AI tools, leaders should direct their government to plan for and begin adopting AI tools in an intentional manner.

Develop a plan for AI

Governors and mayors have taken proactive steps to create guidelines for AI technology use by government agencies. For instance, governors from more than a dozen states including California, Maryland, Oklahoma and Virginia have signed executive orders outlining guidance for the responsible use of AI while dedicating resources to continued innovation. New York City was one of the first cities to publish an "AI Action Plan" to build a robust governance framework, increase employee knowledge of AI and to support AI implementation. These orders often establish AI advisory panels responsible for oversight of AI deployment within state agencies, identifying potential use cases and promoting responsible AI technologies within the government.

Legislatures are also shaping government direction on AI. For example, legislatures in states like Texas and North Dakota have created advisory bodies to study and monitor AI systems used by state agencies. Such efforts ensure that AI adoption aligns with the state's strategic goals and public values.

Implement an AI training program

Although governments have likely used AI in some form for years, recent developments, including the advent of generative AI, are significantly changing how many employees understand the technology. Providing specific AI training will help employees understand both how to benefit from AI while minimizing potential risks unique to the technology.

Governors, mayors and other senior leaders can start the AI education journey by holding education sessions for their leadership teams. These sessions can explain how the technology works, its strengths and when it can be used in government. These sessions can include case studies and practical examples of usage – both by government employees to increase efficiency and when to consider incorporating AI into new and existing programs. Most importantly, these education sessions can help leaders understand the unique considerations and risks associated with AI technology and strategies to minimize them.

Governments should also provide AI training to all employees. Similar to annual cybersecurity or privacy training, a computer-based AI training program can help employees benefit from AI while educating users on how to avoid common mistakes. Training topics could include recognizing when AI tools are safe to use for work data, an introduction to generative AI including how to

Free AI training from LinkedIn Learning.

Microsoft and LinkedIn have developed two learning paths to help anyone understand the basics of generative AI and how to use the technology to increase productivity at work. Governments can leverage these learning paths to train their employees. Free Learning Paths for Top Jobs (opportunity.linkedin.com)

For additional training content, check LinkedIn Learning's catalog of AI training courses: Artificial Intelligence (AI) Online Training Courses | LinkedIn Learning

write "prompts" that aid in the completion of daily tasks and the limitations of AI technology, such as its ability to make mistakes.

By educating their teams, leaders can foster a culture of informed decision-making and responsible experimentation with AI solutions.


Gain experience with generative AI

Leaders should also gain first-hand experience using AI tools themselves in low-risk scenarios, demonstrating a commitment to understanding and leveraging the technology. By immersing themselves in the functionalities and potential of generative AI tools, leaders can gain firsthand insights into the technology's practical applications and limitations. Moreover, when government leaders actively use generative AI technologies, they set a powerful precedent for their administrations, signaling the importance of adopting and adapting to technological advancements.


Among the types of generative AI tools that government leaders could experiment with are assistants like Microsoft Copilot (<http://copilot.microsoft.com>) and ChatGPT (<http://chat.openai.com>), which can draft documents, answer queries and provide summaries of complex information. Additionally, tools like Copilot and DALL-E offer capabilities to generate visual content from text prompts, which can be used for public communication and educational materials. As with any technology, be sure to understand what data privacy and security commitments are made by an AI tool before entering sensitive data into the tool.

How Generative AI is benefitting governments

Simplify the constituent experience
when engaging with government services



Empower government employees
with insights at their fingertips



2 Implement AI governance

Governance is the process of making and implementing decisions about the design, development, deployment and use of IT systems. Governance ensures that IT investments and usage align with the organization's goals, values and policies, and that they comply with relevant laws, regulations and ethical standards. Governance also involves monitoring and evaluating the outcomes and impacts of IT systems and adjusting them as needed to address any issues or risks.

A first step to adopting AI is providing guidance to departments and leaders on how AI should be used to benefit the public and guardrails for its proper usage. When implemented correctly, governance should both accelerate adoption by encouraging responsible innovation while decreasing risks.

Government and industry organizations are establishing standards for the use of AI in various industries and settings. This includes the International Organization for Standardization (ISO) and the National Institute of Standards and Technology (NIST). When developing AI governance, consider referencing existing standards that can guide the configuration, deployment and testing of AI technology.

Define a principles-based approach to AI governance

As we have seen with previous innovations, technology evolves more rapidly than the policy developed to govern it – and AI is no different. Therefore, it's important to develop governance that can consistently guide government decision-making.

One way to do this is by developing a set of outcome-based principles that are durable over time and will continue to apply, even as granular details for configuring and implementing a technology may change. Government may also choose to leverage an existing framework, like the NIST AI Risk Management Framework (RMF) or the ISO/IEC Risk Management Standard: ISO/IEC 23894:2023 - Information technology — Artificial intelligence — Guidance on risk management.

At Microsoft, we created our Office of Responsible AI in 2019 to foster a comprehensive approach to developing and deploying AI responsibly. This office has dedicated significant resources to understanding and setting out the principles that should form the basis for a responsible approach to AI, and the standards that should inform development and implementation of AI systems in accordance with the principles. From these principles, we created aligned rules and a governance process, which include conducting impact assessments of an AI tool or service prior to development and release. This approach is detailed in our Responsible AI Standard, which along with a sample impact assessment and transparency notes, is available for download at <http://www.microsoft.com/rai>

Create AI roles and responsibilities

Given the potential for AI to affect all government departments and stakeholders, establishing clear roles and responsibilities for who will guide AI governance and implementation efforts is critical.

In August 2023, Maryland became the first state to appoint a Senior Advisor for AI, who is responsible for guiding the state's AI strategy, including ethical guidelines, business pilots and coordination with the Federal Government. New Jersey recently expanded the Chief Innovation Officer role to also serve as the State's Chief AI Strategist. More commonly, we see existing chief information officers/IT directors leading on the development of AI policies and the review of new AI solutions.

As governments consider what role should take on responsibility for AI, governments should consider who is responsible for the following:

- Convening an AI governance body comprised of AI and data experts; department leaders representing law, privacy, administration, IT, HR and enterprise funds; representative leaders from general and enterprise fund departments; and executive and legislative leader representatives. These bodies should advise the AI leader on the direction of AI technology, potential benefits and risk mitigation strategies, feedback on how implementations are progressing across government and public sentiment on solutions.
- Leading the development of AI principles, use policies and processes to ensure AI solutions are developed in accordance with the direction of elected officials and the AI governance body.
- Increasing the knowledge of AI across government and creating a culture that encourages the ideation and responsible use of AI systems.
- Depending on the organizational structure, creation of an AI leader network across departments in a hub-and-spoke model, creating competency and awareness of AI governance close to where line-of-business/mission decisions are made.
- Maintaining visibility into AI implementations, leveraging that context to measure progress and evolve policies and processes as needed.

Develop and implement an AI use policy

While most governments have robust policy frameworks in place to govern the selection, implementation and operation of new technologies, many are finding the approach to implementing AI solutions and the relevant risks to be sufficiently different as to require specific AI use policies. These policies can define the expectations for how AI is to be used by government staff and the process to follow for requesting new AI tools and solutions.

Cities, including Boston and San Jose, and the State of Kansas have implemented AI use policies that establish processes by which AI solutions can be considered and adopted. We also see governments working together to create model policies that can work across governments,

AI Principles



Fairness

AI systems should be designed to treat all individuals and groups fairly and equitably, without bias or discrimination.



Reliability and safety

AI systems should be designed to operate reliably and safely, minimizing unintended harm to people and property.



Privacy and security

AI systems should be designed to protect the privacy and security of individuals and their data.



Inclusiveness

AI systems should be designed to empower everyone and engage people in the design and development process.



Transparency

AI systems should be designed to provide transparency into their operation and decision-making processes.



Accountability

AI systems should be designed to be accountable for their operation and decision-making processes, and to enable responsible oversight.

such as the National Association of Counties (NACo) forming an AI Exploratory Committee to produce a model AI policy and toolkit.

While policies will vary to meet specific government needs and structures, they generally should address:

- Limited use of consumer AI tools in the workplace. The policy should set limits on the use of consumer and non-approved AI tools in the workplace, such as freely available ChatGPT, Bing Copilot and Google Bard services. Governments do not have visibility into how employees use these tools, and therefore cannot ensure the protection of resident and critical operational data provided to them.
- How AI solutions should be developed and deployed. This can include requiring development environments (sandboxes) that have been pre-configured to meet industry-leading privacy and security standards, plus additional guidance on how to request the procurement and deployment of additional tools and services that meet industry-leading privacy and security requirements, and do not require highly customized local requirements.
- Ensuring that the AI systems are transparent, accountable and fair. For solutions built in-house by government teams, this can include testing strategies relevant to AI, such as "red teaming" which may not be familiar to staff today, impact assessments, when to engage resources like an AI governance board, and defining who makes final decisions on AI system readiness before deployment. For purchases of vendor solutions, policies should consider requiring the completion of impact

assessments. The policy should also require the disclosure when users are interacting directly with AI technology and that a human is always accountable for the AI solution and its output.

- Privacy, security and compliance. This should include reviews of the terms and conditions for AI services to understand how the service provider limits use of data for reasons other than providing the contracted service, prohibits sharing of data for development of foundation models unless explicitly approved, and the government's data classification and related requirements align to the AI services' terms.
- Accessibility. Policies should require that new AI-based solutions consider how the technology will be used by individuals with accessibility needs and opportunities to use AI to make government services more accessible to the public.
- Monitoring and evaluating the performance of AI systems. This should include the requirement for new systems to have plans for how periodic testing should be performed to ensure AI systems are reliable, accurate and effective.

For additional insight into what should be included in an AI use policy, or to understand the role of potentially new processes necessitated by AI like "red teaming" testing, consider leveraging existing standards, like the NIST AI Risk Management Framework or ISO 42001, the AI Management System (AIMS). These are strong foundations for a consistent, impactful and process-based approach to mapping, measuring and managing risk across AI systems.

3

Develop suitable procurement mechanisms and skills

Procurement is a critical government function, facilitating the competition and providing the transparency necessary to ensure governments maximize value and desired outcomes. Many procurement rules for technology were written in an era when distinct copies of software were sent to a customer in a shrink-wrapped box and installed in government data centers. The advent of cloud computing a decade ago began to change how governments consumed technology. Instead of "boxed" software, governments mostly rely on services – software installed in a technology company's data center, or cloud, which is billed on a unit of utilization, such as per gigabyte of storage or per user per month.

AI relies on the vast compute and storage capacity that can only be provided by cloud computing services. As governments prepare for AI adoption, they should consider what procurement rules and processes may need to be updated to facilitate timely and cost-effective adoption of cloud services, and to consider industry standard terms based on how AI and cloud computing services are delivered, thus minimizing risk and facilitating timely service delivery to internal stakeholders.

Identify opportunities to gain trust through certification and compliance practices

Cloud computing operates at hyperscale, providing millions of customers with a common service to realize the cost benefits of scale. Technology providers and governments have worked together to develop scalable risk mitigations that reflect how cloud technologies work. For example, government

standards like FedRAMP Authority to Operate, define security and compliance practices that cloud providers must implement and continually operate. Microsoft maintains compliance with more than 90 standards and certifications to ensure governments can trust its operations are secure, available and operating in accordance with our contractual commitments. Governments can obtain evidence of our compliance at the Microsoft Trust Center (<http://aka.ms/trustcenter>).

Governments should review their existing procurement guidance and templates to consider that they:

- Require contractors to include terms that require industry trust certifications, such as FedRAMP Authority to Operate or StateRAMP certification and responsible AI usage as noted by conformance to the NIST AI Risk Management Framework.
- Recognize when the intended use case is for standardized hyper cloud services and be empowered to accept those standardized product terms, identifying any specific minimum requirements or standards for the government's commercial terms that do not contradict the use case and contractor's operational, product terms.

Create a Cloud and AI Center of Excellence (CoE)

Procurement officials have a strong accountability role to ensure that procurements are transparent and competitive. Procurement officials must have the capacity to handle many different procurement needs, which may require the same procurement personnel to procure more than specialized IT procurement.

Governments should consider bringing together a team of people dedicated to identifying and leading the adoption of best practices in a CoE. This can be a virtual team of procurement professionals from across departments and agencies who have helped facilitate the implementation of these transformative technologies. Some of the functions of a CoE should include promoting cross-government collaboration, identifying training needs,

Strong policies set a roadmap for the era of AI



Clear directives to develop government use cases can help ensure that read benefits and impacts reach local communities



Sandbox environments can enable governments to experiment with AI and discover new applications



Specific timelines lay out clear direction for government, from ethics establishment to research and development to responsible deployment

providing customized training and influencing cultural change. The CoE should offer insights beyond the technical products and services that the government intends to procure. Professional associations and technology providers are a good source for technology training content. Microsoft training can be found at <http://aka.ms/MSLearn> and <http://aka.ms/AzureAI>

Create procurement guidance for cloud and AI services

The shift from “boxed” software to cloud computing services raises important questions for government procurement offices. Is the government buying software or a service? Who is responsible for the privacy and security of the government’s data? What can the cloud provider do with the data in their custody? Ten years into the cloud computing revolution, many governments now run many, if not a majority, of their technology services in the cloud and have developed governance processes and procedures aligned with the hyper-scale nature of cloud computing and its standardized contract terms.

We suggest governments consider how cloud and AI procurements affect the three general procurement phases: research, select and contract.

Research

This is the premarket engagement to learn. When software was installed and run in government data centers, government employees could install and test trial copies of software with limited assistance. In the

cloud-first world, governments partner with technology providers to test and build proofs of concept on cloud computing platforms. This level of partnership is especially critical as government technology teams gain experience with new tools like generative AI. Procurement rules and processes should be reviewed to consider:

- How to allow government stakeholders to gain an understanding of technology offerings through demonstrations, proofs of concept and similar engagements in a manner that complies with equal treatment and transparency rules.
- Whether the government has what it needs to deploy the solutions. For example, cloud adoption and data requirements.
- How technology providers’ approaches to data security and other legal and ethical standards will be evaluated. For example, by using a list of screening questions.
- Consideration for delivery and contractual models, including pricing.

Select

At this stage the business decisionmakers and procurement officials have collaboratively developed their specifications based on the business’s needs. It will be important for the procurement officials to ensure that the specification is output based for that intended use. Understanding the key result is essential for procurement to align with the business stakeholder’s objective

for the procurement. Some useful points to consider when creating selection criteria that enables innovative procurement include:

- Collaborate between stakeholders and learn from the business decisionmaker about the intended use, its specifications and KPIs.
- Seek training from or create a Center of Excellence to learn more about standardization in cloud computing in addition to stakeholder collaboration.
- Consider whether there are existing industry standards to require as selection criteria, for example requiring cloud providers to have a FedRAMP Authority to Operate (ATO) or StateRAMP certification, and conformance to the NIST AI Risk Management Framework.
- Acknowledging that one of the benefits of cloud computing is hyperscale and standardization of underlying technical infrastructure ensure selection criteria focus on business needs and compatibility with existing systems and recognize that some legacy selection criteria may require updates to align with the intended use of cloud computing.
- Consider the available compliant procurement options, which may be an existing framework amendment.

Recognize the differences between operational, product terms that are often not negotiable and commercial terms that may have more flexibility.

Contract

At this final stage in the public procurement, the business decisionmaker, procurement officials and lawyers should be collaborating to determine what additional commercial terms are required to address local requirements that align with contracting for the standardized, hyperscale cloud computing that has been researched and selected.

Leverage AI for procurements

Along with training in cloud and AI for procurement, we have seen many examples of how procurement teams can use them in their own work. It may make sense for procurement teams to improve the procurement request and change their procurement template for better alignment, as discussed here before using generative AI to produce reusable material (Requests for Proposals, specifications, formal correspondence, summarizations and more). Generative AI could be used to base their business decisionmaker’s information to determine the use case with more informed questions. Procurement teams may also want to use generative AI for external purposes for transparency and competition, such as the Port Authority of New York and New Jersey’s Beacon chatbot referenced on <https://beacon.panynj.gov>



4 Review data governance and security

Data has become the lifeblood of every organization but, with the shift to hybrid work and unprecedented levels of information dissemination, much of governments' data now lives outside of the traditional borders of data infrastructure, including cloud computing services. Because AI relies on data, the availability and quality of data made available to AI models directly affects the quality of its output.

Equally important is how access to data is controlled, including access permissions. While some government data is public, governments hold sensitive data on its residents where access should be limited to only those government employees who need information derived from the data to perform their role. Government data should be classified and "labeled", or tagged in such a way that government can track the security, privacy and regulatory requirements relevant to the data. New technologies can label both existing data and new data as it is created, increasing the likelihood that only those with a need to access the data can do so, that the data is used in a compliant manner and that AI-based tools also comply with these rules.

One of the most anticipated AI-based tools for worker productivity, Microsoft 365 Copilot, will become available to US state and local governments in summer 2024. Microsoft 365 Copilot combines the power of generative AI with your organization's data – all in the flow of work – to turn your words into one of the most powerful productivity tools on the planet. It works alongside popular Microsoft 365 apps, such as Word, Excel, PowerPoint, Outlook, Teams and more. Microsoft 365 Copilot provides real-time intelligent assistance, enabling users to enhance their creativity, productivity, and skills. Copilot and many new AI-based solutions envisioned by governments will rely on

existing government data. Therefore, it is essential to ensure that the data used by AI solutions is of high quality, well structured and secure. This will not only facilitate the proper use of AI, but also protect the privacy and rights of residents and stakeholders.

The actions in this section are about technical changes to a government's environment and may require additional technical assistance. Government may already have access to tools that can help perform the steps below and have access to support resources necessary to configure and deploy them. For additional assistance, governments can engage their Microsoft account team or support partner for additional help.

Assess and automate data governance practices

Data governance is the process of defining and implementing policies, standards, roles and responsibilities for the collection, management and use of data within an organization. Data governance aims to ensure that data is consistent, reliable, secure and aligned with the organization's goals and values. For governments, data governance is especially important as they deal with sensitive and personal information of residents, businesses and other entities. Moreover, data governance can help governments comply with legal and ethical obligations, improve public trust and accountability, and enable data-driven decision-making.

Most governments have an existing data governance program, or aspects of one such as a process for classifying data based on compliance requirements and risk. Yet many governments rely on trust and other manual processes to ensure data is consistently and accurately handled in accordance with program requirements.

Prior to implementing new AI solutions, governments should consider implementing tools to automate data classification and labeling. In other words, ensuring data can only be used by those individuals who should have access to it and for intended purposes. Governments should also audit current data access to ensure that confidential data is only accessible to intended users.



Accelerate cloud migrations

AI solutions require massive amounts of data and compute resources, which are often not available or affordable in on-premises government data centers. Cloud computing services offer scalable, flexible and cost-effective access to large volumes of data and powerful AI tools and frameworks.

On-premises infrastructure has served governments well for decades, but it is now holding them back. When we talk to governments that still primarily rely on data centers, we often hear about their desire to benefit from technology innovation only available in the cloud, like AI. Unlike cloud-based systems, on-premises infrastructure does not enable flexible computing capacity, which can lead to governments paying for more than you need during slow times while also running the risk of system overload during high demand events.

In 2020, many states had an on-prem data center for their labor departments and had been slowly modernizing apps one at a time. In March

2020, when the pandemic struck, Americans filed for unemployment in record numbers, overloading many states' legacy systems. This led to claimants experiencing significant delays in receiving their benefits, sometimes waiting over six months to collect needed payments.

In an investigation of the failure, one audit found that the incremental approach their labor department had been following did not prepare its unemployment system for the demand placed on it when it was needed most. Despite having begun the modernization process, the on-prem system lacked efficient claims processing, automated communications, and data reporting and analytics commonly found in cloud-based systems. This resulted in delayed service for constituents.

Having your data and applications on cloud infrastructure enables you to scale computing as needed, while benefiting from a more secure environment. Once migrated to the cloud, governments free up the resources devoted to an on-prem data center, creating more capacity for modernization and innovation work.

5 Foster AI innovation

AI is a game-changing technology. However, to harness its full potential, governments need to foster a culture of innovation that encourages experimentation, exploration and learning. This means both providing the technical infrastructure and the programmatic support for employees to try out new ideas and solutions using generative AI.

In combination with the other efforts described in the Roadmap, from developing use policies to training employees, government leaders can use their power to convene and set the direction of government to increase excitement, trust and momentum for AI adoption. Below are several actions leaders can take to help employees consider beneficial uses of AI.

Seek input on potential uses

Advances in AI create vast new potential use cases across government departments, with many ideas coming from government employees who operate specific programs and processes on a daily basis. Leading governments are encouraging departments and individual employees to suggest use cases and define their potential benefit. For example, The State of North Dakota formed AI learning cohorts through which employees learned about potential uses for AI, and each cohort pitched a potential AI use case to increase government efficiency and improve citizen service. In the State of California, the Governor's AI Executive Order called for departments to develop potential AI use cases, with high-value use cases selected for pilot development. Encourage government departments and agencies to develop AI project ideas and use cases, and seek guidance from technology and subject matter experts on how these solutions can be responsibly implemented in government.

Hold an AI Discovery event

While more people are discovering the power of AI through publicly available tools like ChatGPT and Microsoft Copilot, even more tools are available to technology and mission professionals responsible for government programs and operations. Holding an "AI Discovery Day" with employees and representatives from AI technology providers like Microsoft can help employees gain hands-on experience with new AI tools, understand their benefits and also the controls to facilitate safe integration into government systems and processes.

Create sandboxes for testing

One way to help employees feel supported when ideating how to use AI is to create a safe space for testing – what is commonly called a "sandbox". An AI sandbox is configured with the proper security and controls to help users comply with use policies, including security and privacy requirements, but still allows flexibility to access testing data and try building new AI solutions. AI sandboxes can enable rapid prototyping and iteration of AI solutions, as well as facilitate collaboration and knowledge sharing across teams and departments. Promising solutions developed in a sandbox can be further developed into solutions ready to be used within a department or across government and help identify best practices that can broadly benefit government employees.


Encourage government pilots

Another way to foster AI innovation is to stimulate the ideation of AI use cases that can solve real-world problems and create value for residents. Governments can do this by holding AI workshops, hackathons or competitions that invite employees to brainstorm and pitch AI ideas, as well as provide feedback and mentoring from experts and peers.

Governors from states including California, Oklahoma and Washington have issued executive orders requiring departments to develop potential use cases for AI to

The adoption journey for generative AI use cases

See how these horizons play out with a DMV assistant



Example queries

Horizon 1:
What is the process for a vehicle title transfer?

Horizon 2:
I just bought a motorcycle; what are the license requirements in this state?

Horizon 3:
Can you give me an update on my commercial license application?

help encourage further exploration of AI's benefits. Local governments, including New York City and Montgomery County, Maryland, are taking a similar approach.

Scale with success

One positive aspect of generative AI compared with previous waves of innovation is its relatively lower barrier to entry. Organizations don't need to make a huge upfront investment and can start their journey with rapidly deployable use cases that have already been implemented successfully in other contexts. For example, the NYC conversational assistant we looked at was implemented in 45 days, so some use cases can be implemented very quickly with the right circumstances.

We've presented this journey through the lens of McKinsey's Three Horizons Model, which groups projects into three "horizons" of investment.

- Here, Horizon 1 might represent the first 60 days where you are developing a policy foundation and experimenting with low-risk entry level use cases.
- At Horizon 2, 2-9 months in generative AI is becoming more normalized and governments might be starting to foster emerging use cases.
- Horizon 3, at 1-2 year, governments are creating genuinely new, innovative, bespoke services powered by generative AI.

Initial AI services can also evolve over time to incorporate more advanced and complete functionality as the government and the public gain trust in the solutions performance. For example, a chatbot that was launched to answer questions on how to obtain a benefit service could be refined over time to answer more contextual questions about complex or uncommon scenarios and integrate with other systems to allow users to access their specific account information and make changes from the chatbot interface.

Conclusion

The increasing use of AI technologies to drive economic efficiencies, improve lives and help solve challenges presents great promise to governments that constantly need to expand their capacity from within existing resources.

Technology is a proven pillar of competitiveness and growth for governments across the country. Governments that are able to harness the power of technology, especially new

and emerging technologies like AI, stand to increase productivity, foster innovation and realize cost savings, all while better engaging with and supporting their residents.

The journey to digital transformation is a constantly evolving one, punctuated by pivots and altered by individual national responses to an ever-changing global policy climate. At Microsoft, we look forward to continuing our engagement with our public sector customers on this exciting journey to see what we can accomplish together when we harness the capability of ground-breaking technologies for the common goal of bringing benefit to all.

