

Contents

1	Executive Summary	2
2	Introduction	3
3	Understanding the drivers and challenges for the effective use of AI technologies	4
4	Rethinking traditional procurement practices and skillsets to enable effective adoption of AI	9
5	Trusted Providers	14
6	AI Regulation and Governance	16
7	Conclusion	22
	Appendix	23
	References	24

1 Executive Summary

Like no technology before it, artificial intelligence (AI) offers enormous potential to help solve global challenges, improve economic productivity and generally enhance the lives of ordinary people. Governments around the world are working at pace to understand the underlying technologies and landscape of opportunities they present for public office as well as for private citizens.

This paper addresses four key areas important for the public sector's assessment and procurement of AI technologies, specifically focusing on: (1) public sector AI use cases and the political, socio-economic and legal considerations; (2) flexible cloud and AI procurement practices; (3) the need to collaborate with trusted partners; and (4) an assessment of best practices in AI regulation and governance.

Formulating laws and policies that protect against potential risks while encouraging innovation is a difficult balance for governments to achieve and will necessarily depend on an appreciation for the unique drivers of AI use in their own jurisdictions. Governments that have clear, forward-thinking and flexible policies on the use and procurement of AI solutions will be in the best position to embrace the benefits and mitigate the challenges that AI technologies bring. Having an environment that supports safe and effective AI procurement is, however, only part of the story as engaging with these ground-breaking and advanced technologies will also require new skills and

approaches. Government stakeholders will need to consider and train for the competencies that will be required to make the best use of the exciting opportunities offered by current and future AI technologies.

Responsible AI procurement in the public sector will depend on government stakeholders and technology providers working together to best understand how to adopt and deploy AI in a compliant and ethical manner that minimizes the potential downsides. In addition, government endorsement of safe, ethical and responsible AI will result in not only gains for individual government departments (and the members of public they serve), but also, it will inspire confidence for wider market adoption across all sectors of the economy.

In this paper we explore the political, socio-economic and legal issues associated with the procurement of AI technologies by the public sector. Importantly, we share the work Microsoft is doing to promote the responsible use of AI as well as the work we are doing with our customer community to drive trust in AI.

We understand that much can be learned from the practical application of AI technologies, and in the spirit of collaboration, we share a number of use cases relevant to the public sector community. Finally, we share our thoughts on how modernizing procurement policies and approaches can help enable the public sector to embrace the advantages of cloud, AI and digital transformation technologies more easily.



2 Introduction

AI describes the use of computer technologies to perform tasks such as learning, reasoning and problem-solving that would ordinarily be thought of as requiring human intelligence. AI technology is already widespread in our society – from the algorithms that direct us to watch programs on our media streaming services through to complex models for assessing climate change and mapping protein structures. Recently, AI technology has attracted a surge of renewed interest as governments and citizens seek to understand the impact and potential uses of generative or ‘cognitive’ AI. Generative AI can be distinguished from other AI technologies (that primarily use pattern recognition to make decisions) by its ability to create new and unique content, such as text and art, using learnings from wider datasets. The accessibility of apps like ChatGPT has catapulted this powerful technology into the hands of ordinary citizens.

Recent developments in the field of AI are a leap forward in digital capability at least as significant as the arrival of internet browsers in the mid-90s or the widespread adoption of smartphones.

However, far more than the technologies that precede it, AI has the capacity to advance our understanding of the world and our ability to learn and express knowledge.¹ There is an unprecedented opportunity for governments to use AI to bring about positive change and innovation by driving advances in medicine and science, improving productivity and stimulating economic growth, and keeping societies healthier and safer. Microsoft has been committed to advancing AI in a responsible manner for many years, with such efforts being grounded in our company’s mission to empower every person and organization on the planet to achieve more.

Weighed against the huge potential for AI technologies to transform societies for the better are legitimate concerns around their possible use. We know, for example, there can be inherent risks in how AI technologies are designed and trained, as well as existential threats in how they might be deployed by reckless or malevolent players. Equally valid are citizen concerns about the impact on jobs and the potential infringement of individual rights.

As such, the public sector’s interest in advancing and harnessing the power of AI, in a responsible way, is paramount to society’s interests at large. Some of the biggest challenges for government customers using AI for public sector services will be in how they are able to justify and explain the techniques and algorithms used by these solutions to make decisions that affect the lives of real people. The potential for bias arising from training data as well as data protection and security concerns are issues government customers will need to address head-on in order to gain public trust in the deployment of AI solutions.² This paper addresses what governments can do to mitigate these risks, to enable the procurement, adoption and deployment of AI technology in a manner that allows them to continue to meet their regulatory and policy obligations, as well as the expectations of their citizens.

So how do customers in the public sector take advantage of the many benefits that AI technologies have to offer while effectively managing the risks and encouraging citizen confidence?

We propose the following four key steps for the safe and successful procurement, adoption and deployment of AI technologies by the public sector, which we explore in further detail below:

- Understand the drivers and challenges for effective use of AI technologies (see [Section 3](#)).
- Develop suitable procurement mechanisms and skills (see [Section 4](#)).
- Collaborate with trusted providers (see [Section 5](#)).
- Implement appropriate AI regulation and governance (see [Section 6](#)).





3 Understanding the drivers and challenges for the effective use of AI technologies

3.1 Use case considerations

AI has the potential to significantly improve almost all areas of public service and most governments acknowledge that they will, at least to some degree, have to embrace AI technology or risk being left behind. AI can make sense of data like no technology before it to help government and community leaders better understand demographic and behavioral trends that in turn can lead to smarter decision-making and more effective governing.

A further benefit of the use of AI in the public sector will undoubtedly be in its ability to eliminate inefficiencies to improve productivity and cost-effectiveness for public sector departments.

In healthcare, for example, applying AI technologies to improve coordination between sector specialists and to reduce healthcare waiting times and backlogs caused by bureaucracy will improve both the experience of service users and the working lives of overburdened public sector workers.

The World Economic Forum's AI Procurement Guidelines set out the expectation that leveraging the role of governments as embracers of AI will lead to an expansion and standardization of approach to the ethics and risk management associated with AI technologies in other sectors. In other words, where governments encourage and facilitate the responsible procurement and deployment of AI technologies, this inspires confidence and lays the groundwork for use in other sectors – to the benefit of the broader economy.

At a recent hearing in the US Congress to discuss how the Federal Government should update its procurement and acquisition policies to promote the responsible and effective adoption of artificial intelligence systems, US Senator Gary Peters, Chairman of the Homeland Security and Governmental Affairs Committee, spoke about the important role of AI technologies in providing more efficient public services and improving the capability of the federal workforce by automating routine tasks. He also emphasized the important influence of government procurement processes in shaping responsible and safe AI procurement in the private sector.³

There are already some interesting use cases that demonstrate how AI is bringing benefit to the public sector:

- The UK's Information Commissioner's Office (ICO), the UK's chief regulator of AI, is itself currently using AI technology powered by a Microsoft platform to help it process the huge volumes of user queries to its email inbox and to assist service users through the ICO registration process.⁴
- The state of Ceará in Brazil adopted the HMX Tax Intelligence System, which uses Microsoft Azure AI to monitor and identify any compliance issues regarding the taxation of retail vendors. This use has led to a 21% increase in tax revenues and an 84% increase in audit efficiency.⁵
- The City of Kelowna in 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. Citizen 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.⁶

We have set out, at [Appendix 1](#) below, some further examples of AI solutions and use cases currently being deployed in the public sector.

It is, however, important to recognize that just because an AI solution is readily available for use in the public sector does not mean that it should be used without an appropriate assessment of the relevant technology and its application. Before acquiring any AI solution, governments will need to carefully consider the potential use case and understand how the technology can be deployed in a compliant manner to achieve an expected outcome while causing the least harm. Governments will also need to assess how any AI solution will inform decision-making and what the role of humans will be in that process. We examine some of the common legal hurdles and concerns to AI adoption in [Section 3.4](#) below and discuss how Microsoft has worked to address these concerns for its customers.

3.2 Political considerations

Different governments (or even different states or provinces within nations) will have different priorities and problems they want to solve. They may be at different stages in their digital transformation journeys and have access to varying levels of investment. In addition, cultural and political differences in how AI is perceived mean governments may take differing approaches to deploying AI in their countries. This has led, and will continue to lead, to an uneven adoption of AI technology across different nations and socio-economic regions. Where governments are able to recognize the unique drivers and barriers to AI adoption that apply in their particular jurisdiction, they will be able to take an authentic outcomes-based approach to AI procurement and adopt solutions that meet their requirements while also providing practical and workable solutions.

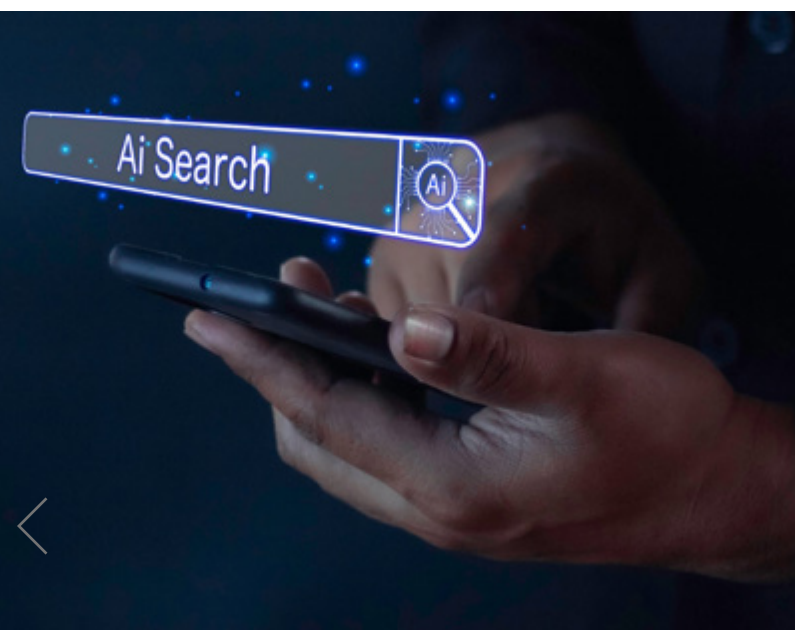
In [Section 6](#) below, we look at the public responses by various governments to AI technologies in the form of legislation, policies and guidance. Just as important, however, are the political pressures and drivers particular to each government that may sway the case for particular AI solutions. An interesting example of this is the current Federal Labor Government in Australia, which has expressed an aspiration to grow internal public-sector skills to place less reliance on the use of external consultants to manage public-sector projects. Employing AI technologies to realize efficiency gains is likely to hasten the re-capture of public services by civil servants and lessen reliance on external consultants. In addition, Australia's Digital Transformation Agency (DTA) and the Department of Industry, Science and Resources (DISR) have recently released 'interim', non-binding guidance for Australian public service staff on the use of generative AI platforms. Having these guidelines encourages staff to explore new and innovative opportunities that generative AI presents for government while carefully assessing potential risks and benefits.⁷ This all contributes to a receptive environment for the potential use of AI by Australian government agencies.

A particular political hot topic for governments looking to deploy AI solutions within the public sector is the fear that automation will lead to the loss of jobs. This concern is especially relevant to countries where there is a large unskilled workforce that might be more easily affected by automated technology tools. It would be difficult to argue that the advent of AI technologies will not have an impact on the labor market but the situation for workers is not necessarily only negative. According to a study by McKinsey, about half of all current work activities can be automated but only about 5% of occupations can be fully automated.⁸ While AI will be more efficient than humans at completing certain tasks, it is likely AI will take on the more monotonous tasks and free workers to concentrate on the more interesting aspects of their roles.

In addition, it is likely the loss or displacement of some roles will happen alongside the creation of additional jobs and opportunities (some of which may not yet have been conceived). McKinsey cites rising incomes and increased spend on health, infrastructure and technology as catalysts for growth of demand of between 21% and 33% of the global labor market up to 2030. This prediction is echoed by Gartner who predict that, 'by 2030, AI will become a net-positive job motivator, eliminating 1.8 million jobs while creating 2.3 million jobs'⁹. In other words, the jobs we worry about losing today may not necessarily be the jobs we will want in ten years, when there may be work and lifestyle options available that have not yet been considered. Contrary to fears about how AI will disproportionately pressure the labor market in low-income countries, McKinsey's view is that some of the biggest gains of the new job dynamism will be in developing nations such as India, where the working-age population is growing at pace.¹⁰

Governments will need to be forward-thinking in pre-empting labor and other political concerns about the use of AI technologies. Particularly relating to the public sector, they will need to put in place strategies for exploring the novel opportunities AI will bring in job creation and assisting with new skills training while providing realistic reassurance on the overall positive impact of AI on livelihoods. Leading by example in the public sector will encourage the business community to do the same.

By 2030, AI will become a net-positive job motivator, eliminating 1.8 million jobs while creating **2.3 million jobs**⁹



3.3 Socio-economic considerations

There is a concern that an easier path to AI deployment in high-income or middle-income nations could strengthen their economic advantage over low-income nations. As an example, although AI advancement may lead to a net gain in jobs globally, it is likely that, at least initially, there may be a disproportionately higher number of jobs affected in developing nations than in developed nations.

The developed world will simultaneously be likely to see benefits in increased automation of those administrative and computer coding roles that have traditionally been outsourced to workers in developing countries.

On the other hand, a recent report by the World Bank argues that AI offers the potential to 'lower costs and barriers to entry for businesses and deliver innovative business models that can leapfrog traditional solutions and reach the underserved'. The report goes on to suggest that the ability to harness AI technology might be the key to ending poverty and boosting economic development in many nations.¹¹ The M-Pesa mobile phone app, currently used by more than half of all Kenyans for money transfer services, uses algorithms and artificial intelligence to detect and prevent fraud. It is a good example of AI technology being deployed successfully (in a highly regulated sector) to not only improve the lives of citizens living in isolated rural communities but also to stimulate wider economic activity. M-Tiba, another Kenyan mobile phone app, uses similar technology to offer health services, which again is beneficial both for citizens and the national health service.

3.4 Legal considerations

The perceived risks associated with AI range from discussions about privacy and accuracy of data to ethical and human rights considerations. We reflect below on some of the commonly cited legal issues associated with AI adoption and provide details as to how Microsoft is addressing these concerns; for example, in how our AI solutions are architected, as well as in how we contract with our customers.

Microsoft is committed to ensuring the AI applications users deploy on our platforms meet the legal and regulatory requirements for responsible AI. For example, and as described in more detail below, we are implementing the AI Risk Management Framework recently published by the US National Institute of Standards and Technology.¹² This commitment applies to any future relevant international standards.¹³

3.4.1 Intellectual property (IP)

A key legal consideration for governments using AI technology is whether or not the particular solution utilizes data that is protected by copyright and therefore whether government use infringes third-party intellectual property rights. Generative AI solutions require significant volumes of data in order to learn. The data used to train the AI models can be sourced directly from the internet, often in a manner that creates copies (although there are AI models that are trained only with pre-assembled datasets).

Because IP laws differ across jurisdictions, whether or not there has been any infringement of copyright or other IP rights may depend on where the copying takes place or from where the data originates. Some jurisdictions provide exceptions that may circumvent issues of copyright for data used in AI systems. The Singapore Copyright Act 2021, for example, includes a broad exception for text and data mining (which is the automated analysis of large amounts of data to discover patterns, trends and other useful information), allowing both commercial and non-commercial organizations to identify, extract and analyze information from the copyrighted work using a computer program.¹⁴ This exception potentially supports the processing of copyrighted material for AI model development and use.

Microsoft is committed to ensuring our customers can use our AI technologies in compliance with applicable laws and mitigates potential infringement of intellectual property rights as follows:

The large language models (LLM) that Microsoft makes available have been trained using vast datasets, including copyrighted works. However, training has been done in compliance with copyright law to make the knowledge in those works more accessible and useful.

Microsoft provides contractual commitments to defend customers against third party claims via its Customer Copyright Commitment. Specifically, if a third party sues a commercial customer for copyright infringement arising from the output content from generative AI services, including Copilot for Microsoft 365 or the Azure OpenAI Service, Microsoft will defend the customer and will pay the amount of any adverse judgments or settlements that result from the lawsuit, as long as the customer used the required guardrails and safety measures relevant to the service.

Microsoft has taken steps in the design process for its AI solutions to mitigate the risk that the output from these technologies will violate copyright laws. Customers can further take steps to mitigate risk through the ways in which they use Microsoft tools and assessing the commercial useability of the output that is generated.



3.4.2 Confidentiality and data security

When assessing cloud and AI solutions, we know that confidentiality, security and integrity of data are key areas of consideration for all our customers. We appreciate it is especially important for governments to know exactly how the data they input is treated, including where it is hosted and how it is used. AI tools, by design, adapt to become more efficient by retaining and further processing the data entered

into them to produce better and more tailored results with each new use. Governments will want to be comfortable that information inputted into the AI solutions they use remains confidential and secure, particularly where such information is sensitive information (for example, relating to citizen health) or security classified.

Microsoft ensures the confidentiality and security of customer data as follows:

Microsoft's AI solutions leverage the entire complement of cybersecurity and data security tools available in Azure and/or Office 365. This best-of-breed protection is applied to the entire application dataflow.

On top of technical controls, Microsoft commits that 'your data is your data' and that it will handle customer data in accordance with the customer's agreements with Microsoft and as directed by the customer. For example, customers wishing to leverage their own data to tailor more specific or domain appropriate responses from the LLM may create a prompt engineering flow or fine tune a model, which remain private and confidential from other customers and from Microsoft, and under the customer's control.

Azure OpenAI Service foundation models, which Copilot for Microsoft 365 calls on, do not store customer data. For Azure Open AI Service customers, prompts and responses are stored for up to 30 days for abuse monitoring, but certain customers can opt out of Microsoft's abuse monitoring process, and thereby prevent even this short-term storage from occurring.

Although data can be secured in any region with proper configuration, controls, and monitoring, many public sector customers nonetheless prefer to also have data residency. This means that their data stays within the cloud region in their own country, where this is available.

Copilot for Microsoft 365 is a Core Online Service and is protected by Microsoft's data security and privacy safeguards.

Copilot for Microsoft 365 use is bound by an organization's data access policies and permissions, which determine what users and Copilot for Microsoft 365 can access.

Copilot for Microsoft 365 presents only data that a user can already access using the same underlying controls for data access used in all Microsoft 365 services.

Azure OpenAI Service is considered one of the Azure Cognitive Services and is protected by the same contractual and policy safeguards as any other Azure Cognitive Service. This means data will be stored in the same region in which the service is instantiated. Microsoft is making a massive global investment with the goal of bringing Azure OpenAI Service to all of our Azure regions.

However, in the interim, government customers can rest assured that the 'stateless' nature of the application models and the opt-out of abuse monitoring can still provide data residency at rest in the customer's preferred region even if the transitory processing is done outside of the country or countries from which they are provisioned.

Note that if the customer is in a country where Azure OpenAI Service is not available, the customer would also have to refrain from using fine-tuning capabilities of the service because that would require storing the training data in the location where the fine-tuning is done.

3.4.3 Personal data privacy

The application of personal data law is another key consideration in the procurement, adoption and deployment of AI products. Where personal data is used in an AI system (whether to develop, train or test an AI system or to provide prompts or input), developers and users will need to ensure its use is fair and lawful under the relevant data protection laws. In turn, many governments looking at AI legislation are focusing on ensuring AI developers and users adhere to a robust data governance regime that implements core principles such as accountability, data protection by design and default and data minimization.

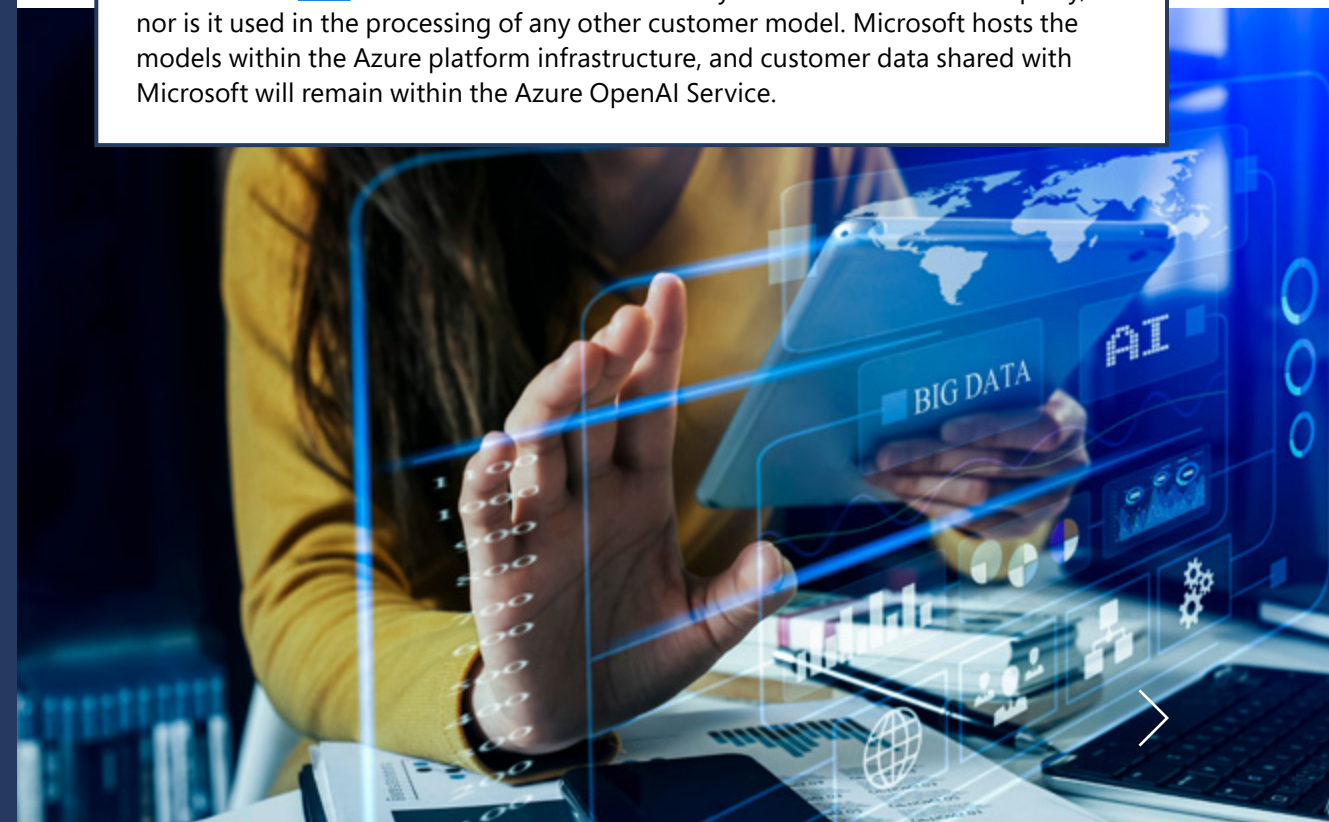
Governments have a balance to strike in permitting the controlled use of personal data in AI applications to improve service delivery (including by further developing and training the relevant AI solutions) while concurrently creating a legislative environment that protects the individual rights and freedoms of their citizens.

Microsoft safeguards personal data in the following manner:

Microsoft adheres to the best practice core data governance principles described throughout this paper.

Copilot for Microsoft 365 complies with Microsoft's security, General Data Protection Regulation (GDPR), and European Union (EU) Data Boundary commitments. Customer data is not used to train Azure OpenAI Service foundation models.

Azure OpenAI can use data that the customer uploads to create a customized model for that particular customer only, other forms of data processing by Azure OpenAI can be found [here](#). This data is not shared with any other customer or third party, nor is it used in the processing of any other customer model. Microsoft hosts the models within the Azure platform infrastructure, and customer data shared with Microsoft will remain within the Azure OpenAI Service.



3.4.4 Content regulation

A significant challenge posed by AI is the potential for misuse by bad actors. For instance, generative AI can be used to fabricate images and videos, which can be used to spread hate speech and misinformation or manipulate public opinion. Generative AI could potentially also be used to create highly sophisticated phishing attacks that aim to install malware onto a victim's device.

The potential for misuse highlights the importance of implementing appropriate guardrails and safety measures against abuse and unintended harm in the use of AI. Several governments have introduced laws relating to platform regulation, content regulation and online safety to try to address these concerns. Some examples of these include:

European Union – Digital Services Act

The European Parliament has set rules requiring internet platforms to implement procedures designed to take down illegal material such as hate speech, incitement to terrorism, and child sexual abuse.¹⁵

United Kingdom – Online Safety Act

Passed in September 2023 with the primary aim of protecting children and adults online, the Act seeks to make social media companies more responsible for the safety of users on their platforms and tackles a wide range of illegal and harmful content.¹⁶

Singapore – Protection from Online Falsehoods and Manipulation Act in 2019, Online Safety Bill in 2022 and Online Criminal Harms Act in 2023

These were enacted to guard against misinformation and to tackle online content that facilitates crimes such as scams, online incitement of mass public disorder, and malicious cyber activities.¹⁷

Microsoft is committed to ensuring content is managed to the extent necessary to ensure its AI systems are used safely and for the correct purpose. Microsoft supports the regulation of online harms as follows:

- By default, the Azure OpenAI Service includes a content management system that works alongside the LLMs to minimize harmful content. It has two components: (i) content filtering; and (ii) abuse monitoring.
- Azure OpenAI Service content-filtering component is processed in real time, running both the input prompts and generated completions through a detection system to review if content should be filtered. No data is stored as part of this process.
- Azure OpenAI Service abuse monitoring retains prompts and outputs are logged and retained for up to 30 days to allow for the application of the abuse detection tool. This data is stored within the Azure OpenAI Service boundary in the same geography as the customer's Azure OpenAI resource; it is encrypted at rest, subject to access controls and deleted after 30 days.
- For Copilot for Microsoft 365, content filtering occurs in real time, with no data stored as part of this process. Copilot for Microsoft 365 has not enabled Azure OpenAI Service abuse monitoring storage and human review.

To take full advantage of the benefits of AI, governments will need to work with AI providers to find solutions and routes to adoption that meet their particular needs. Microsoft continues to collaborate with individual governments in a manner that is thoughtful, substantive and broad-minded, to use cooperation with the public sector as an opportunity to encourage responsible innovation, and strengthen the way we work in individual jurisdictions and across borders.



4 Rethinking traditional procurement practices and skillsets to enable effective adoption of AI

As described above, there is widespread public recognition by many governments that AI technologies will play a pivotal role in digital transformation in the public sector. Although governments have legitimate concerns about potential hostile applications of AI technologies, there is no avoiding the enormous potential these technologies have in aiding governments to meet their own diverse challenges. While governments work to develop their legislative responses (which we touch on briefly in [Section 6.3](#) below), many are simultaneously encouraging a curious and innovative approach to the use of AI technologies, which in turn requires a review of traditional public sector procurement practices.

Our ongoing engagement with governments across the globe has highlighted several procurement blockers that currently affect public sector customers looking to embrace the benefits of cloud and AI technologies. Factors such as outdated procurement practices, the absence of appropriate costing models and lack of appropriate skillsets all play a role in impeding the ability of public sector customers to successfully deploy new AI technologies.

At this point, it is worth recognizing there would be no AI without hyperscale cloud and, to an extent, some of the groundwork for AI procurement has been laid already by previous procurement policy work related to the use of cloud computing. Regardless of where a nation is at on its digital transformation journey, there are common learnings from the procurement of cloud services that governments can take advantage of when acquiring and deploying AI technology in the public sector.

4.1 Procurement principles

Looking to the future, it is critical that procurement processes and principles remain up to date and fit for purpose – for instance, embedding innovation and sustainability as requisite criteria in the procurement process. Procuring the most innovative or sustainable digital solution may not be the cheapest option but we see the move towards a more holistic procurement approach (most advantageous versus most economically advantageous) as both positive and inevitable.

Advanced technologies like AI have a huge role to play in reducing redundancy and wastage and thereby supporting sustainability initiatives. However, considering innovation and sustainability also means seeking to ensure the deliverables provided as a result of the procurement are themselves innovative and sustainable. Consider, for example, a national electricity grid that is looking to benefit from a more targeted distribution of power based on demographics, weather predictions and particular renewable energy inputs. AI has the capability to rapidly assess multiple datasets to make accurate and targeted predictions and real-time decisions to optimize energy flow. In this situation, the AI technology is itself more sustainable than traditional IT solutions but will also deliver a more sustainable outcome overall. Similarly, entrenching principles like innovation within the procurement process will necessarily result in procurers seeking out more capably advanced solutions like AI to solve complex challenges.

4.2 Budgeting and cost/benefit analysis

Although one of the most touted benefits of AI technologies is cost efficiency, through our engagement with governments, we understand that the upfront expense of assessing and deploying AI technologies and upskilling staff can, ironically, be one of the biggest hurdles to its procurement. The initial outlay in acquiring AI technologies coupled with traditional budgeting and revenue recognition practices in government can lead to a cost/benefit analysis that swings unfavorably away from AI procurement. The inflexible financing practices of some governments is something we have explored previously in the context of cloud procurement.¹⁸ Arguably, given the novelty of AI technologies, it is even more relevant here and needs to be addressed.

To make the most of AI technology, public sector organizations will need to intentionally and thoughtfully address rigid financing rules that can inhibit successful digital transformation and innovation. The traditional financing model, which is based on a yearly allocation of a fixed sum of money for IT spend, encourages a preference for shorter-term individual projects with pre-set requirements that are not easy to adjust down the line. In addition, the typical cost/benefit models used by governments are generally not fit for purpose in assessing the long-term gains in cost efficiency and the service or productivity improvements AI technologies are likely to deliver.

Like cloud solutions, the acquisition of AI solutions requires a more holistic, flexible view that reflects the reality of how they will likely be consumed (i.e., standardized, scalable subscription-based services) and where and when their benefits are likely to be realized. This likely will require cooperation between governments and providers of AI technologies to develop new models for assessing the impact and whole-of-life benefits of these emerging technologies.

As we set out in our [Building Blocks](#) paper, Microsoft believes it is advantageous for governments to have a clear financial plan that sets out technology spending in a targeted and controlled manner. There are benefits in having a central budget authority to take on the responsibility of coordinating the requirements of individual departments and agencies to set a cohesive strategy for digitalization that does not leave individual agencies behind due to their separate budget constraints. Digitalization strategies and budgets should, in our view, be reconsidered in the context of digital expansion to:

1. Facilitate multi-year commitments to consumption.
2. Allow for rapid deployment of extra funds to scale up existing services or provide access to new services as required.
3. Cover not just the acquisition of AI technology but also the upskilling and support services required to optimize the benefit of such technology.
4. Take into account not just the acquisition of an AI platform or solution, but potentially also the costs of smaller projects that could benefit from such AI technologies.¹⁹





4.3 Flexible procurement practices

Whole-of-government framework agreements are widespread across digitally mature governments. These agreements are usually structured as a main framework agreement between the entity representing the government and the supplier, with the possibility of call-off terms for individual government departments.

As we noted in the [Building Blocks](#) paper, Microsoft considers this type of centralized approach to contracting is an important contribution to the effective implementation of a digital transformation strategy. However, our observation on the ground is that an inflexible application of framework agreements can also engender a degree of inflexibility and a one-size-fits-all approach to contracting that is not compatible with the acquisition of innovative technologies like AI.

Avoidable challenges may arise in the government migrating to AI solutions where the government fails to consider the inherent attributes of such solutions and is fixed on adhering to or replicating the mandatory requirements used for bespoke outsourcing solutions.

Microsoft has been engaging with government customers about the impracticality of requiring bespoke terms to be negotiated for standard 'off-the-shelf' AI solutions running on cloud platforms that are, by definition, non-customized, shared services that rely upon economies of scale achieved through standardization. The design of scale solutions, like cloud and AI, means all customers receive the same high-standard product and related service levels. Although 'as-a-service' solutions can to an extent be configured to meet a customer's particular needs, it is often not possible to add bespoke features and/or services not already part of the specification 'menu'. Contractual terms designed for either non-digital transformation projects or traditional outsourcing deals are therefore usually fundamentally incompatible with 'as-a-service' solutions. Government-mandated commercial terms that are not compatible with the product being procured can create insurmountable barriers to participating in a procurement process by vendors who would otherwise have participated and may lead to the relevant government customer acquiring an inferior product compared to what may otherwise have been possible through greater flexibility in the procurement process.

Our observation is that a strict insistence on government-mandated terms and/or non-acceptance of vendor commercial terms is often culture- or policy-driven and often has little to do with actual risk. The inclusion of vendor terms in framework agreements need not increase the risk profile for public sector

customers; in fact, it may decrease risk as the vendor terms will most accurately reflect how the actual services are operationalized and delivered by the vendor. This is particularly true for advanced AI technologies where the particular vendor terms may help to deliver the most appropriate or innovative outcome, reflecting what is being provided. We note that the procurement agencies in some countries have been open to considering the adoption of vendor terms (e.g., on data protection, SLAs, and so on) as part of their framework contracting, though we believe there is an opportunity for government customers to accommodate vendor terms more readily.

We encourage the continued use of framework agreements, but with a change of approach to encourage greater flexibility through collaborative dialogue with the relevant technology providers, a focus on the desired outcome of the procurement, and departing from a tendency to mandate requirements (some of which may not be commercially beneficial and/or technologically feasible). Having a flexible approach to contracting (such as accepting vendor terms that reflect the evolution of AI and cloud technologies and maintain quality/continuity of service) will enable government customers to gain the maximum benefit from the most innovative AI products on the market. It is critical that framework agreements remain agile and adapt along with the advancements in technology.



4.4 Centralized procurement function or central purchasing entity

Individual government departments or agencies will have different views on AI, different budget constraints and their own range of existing individual vendors, procurement practices and technology needs. These challenges are often cited as reasons for maintaining a decentralized procurement approach that places the procurement function closer to the needs of the final user. However, this decentralized, siloed approach has downsides in that it limits the efficiency possible with a central decision-making purchasing agency, increases bureaucracy and may result in more expensive, less transparent and less accessible tenders. These factors in turn can limit the government's access to the benefits of AI technologies. A centralized, 'whole-of-government' procurement model has several benefits. For example, coordinated budget approval and volume purchases make it possible to obtain significant cost savings and/or receive better services at lower cost. This not only contributes to cost reduction but also allows a greater use of shared resources on the government side to manage those terms. The prospect of contracting with a single AI provider or reseller, with the potential

to provide services across government, is more appealing to suppliers than having to navigate the requirements of complex and disparate procurement processes.

As discussed in our [Building Blocks](#) and [Public Sector Procurement](#) papers, we consider the establishment of a centralized procurement function is fundamental in developing and implementing a digital strategy, which in turn will almost certainly include the procurement of AI technologies. While no centralized procurement function will be perfect (for example, processes such as renewals and enrolments can be time-consuming and the procurement function may still be impeded by legacy approaches to procurement), our view is a centralized approach has the potential to improve transparency and efficiency and be more attractive to a wider range of potential AI technology suppliers for government, therefore increasing the quality of AI products available to governments, and ultimately to the citizens they serve. A further benefit of wider access for vendors to public sector purchasers is the cycle of economic growth and innovation it supports in the market more generally.

Two examples of this centralized procurement model are Australia and the UK:

Australia

Australia's Digital Transformation Agency (DTA) has streamlined its digital sourcing process by consolidating its Digital Marketplace with the BuyICT platform, so all types of ICT procurement are now in one place. The DTA is responsible for whole-of-government IT contracts, with the possibility of centralized enrolments for the benefit of all agencies as well as individual enrolments to be signed by individual agencies for specific needs they may have for which they can leverage the same framework. This approach has the effect of reducing the risk of inter-departmental friction and inconsistency in the approach to procurement.

UK

The UK Government's Crown Commercial Service (CCS) is the largest public procurement organization in the UK and is responsible for procuring and managing the central government framework agreements used across the public sector. Public sector employees have access to a Digital Marketplace (managed by the Government Digital Services), which, depending on the relevant framework agreement (between government and supplier), is used by the public sector organization to purchase various digital services. The government has set up three new bodies to support the use of AI, build the right infrastructure and facilitate public and private sector adoption of these technologies.

These three new bodies are the:

- **AI Council**, which will be an expert committee providing high-level leadership on implementing the AI Sector Deal.
- **Office for AI** (part of the Department for Science, Innovation and Technology), which works with industry, academia and the third sector to coordinate and oversee the implementation of the UK's AI strategy.
- **Centre for Data Ethics and Innovation**, which identifies the measures needed to make sure the development of AI is safe, ethical and innovative.

It is important that any centralization simplifies the procurement of AI services through policy and practice rather than increases the complexity of procedural layers that officials need to work through to deliver innovation. One further example of the benefit of a centralized function is the Canadian Federal Government's list of pre-qualified ethical AI providers who can provide the Government of Canada with responsible AI services (published by Public Services and Procurement Canada and the Treasury Board of Canada Secretariat)²⁰. The pre-approved list of AI vendors, spread across three tiers of anticipated spend, aims to streamline the procurement of AI solutions.



4.5 Rethinking the skillset

Public procurement on standard framework agreements can sometimes be seen as a scale exercise and may be carried out by staff who may (understandably) not be fully aware of the intricacies of what they are procuring, the practical application of the technology or the legal terms they are seeking to impose. In our experience, officials can sometimes be more focused on the procedural aspects of procurement rather than on the outcomes the technology should drive for future public policy and services. At times, this makes it challenging for providers to engage in meaningful discussion concerning what they perceive as unreasonable or unnecessary requirements. Given the complexity of AI technologies and the increased focus on their transparency and explainability (as discussed above), developing the right skills in this area is an important procurement consideration.

A lack of outcome-focused practice can lead to officials focusing on mandating requirements rather than embracing the innovative solutions AI can offer.

In our observation, there is also at times a tendency to think flexibility in approach to procurement equates to an increase in risk when, as outlined above, a flexible approach may more accurately memorialize what is being sought. Empowering officials to use procurement to innovate is as much an issue of culture, organization and skills as it is of policy and procedure.

Employee communication, engagement and transition strategies are key. Governments must implement communication plans that help employees understand the changes and approaches that may need to be made to implement AI technology. An understanding of AI, cloud and other digital technologies is not only important for the use of such technologies but is also necessary to ensure stakeholders across government can collaborate on the policies and strategies associated with a government's digital transformation objectives. In other words, it is critical that procurement officials understand the technology solutions they require and how to contract appropriately for the best outcomes.

To truly embrace the benefits AI has to offer the public sector, government agencies will need talented and forward-thinking officials who feel supported by leadership when they take decisions that break new ground and find new

opportunities. These officials will require different skills from those of their predecessors and there should be an increased focus on building digital skills in procurement professionals as well as growing an awareness of emerging AI technologies and their impact. Governments should consider performing an analysis of skill gaps, including a review of the current state of the government IT workforce and projection of the future skills that will be required to make the most of what AI technologies can offer. Training should not be constrained to technical but extended to contracting and procurement. A workplace AI training program (akin to the scheme launched by Microsoft with the UK Cabinet Office in 2017) should enable the government to attract, train and support workers with the skills needed for the next stage of digital growth.

As a practical step, governments should consider implementing a Digital Transformation Centre of Excellence (CoE) comprising a team of people dedicated to the creation, spread and institutionalization of best practices, structures and governance for the evolution of AI, cloud and digital transformation technology. Some of the functions of a CoE should include promoting cross-government collaboration, identifying training needs, providing customized training and influencing cultural change.





4.6 Practical considerations for AI procurement in the public sector

Some of the key practical considerations for government departments procuring and using AI technology are as follows:

4.6.1 Safe use

- (a) Understand AI and equip staff with the relevant skills to understand and manage AI projects.
- (b) Factor the use of AI into broader risk management frameworks. Plan and prepare for artificial intelligence implementation.
- (c) Consider if it is the right solution.
- (d) Conduct proper testing.
- (e) Provide ongoing human supervision and the use of circuit-breakers.
- (f) Identify edge and high-risk cases and act accordingly.

4.6.2 Data issues

- (a) Ensure the AI system is trained on sufficient, high-quality data.
- (b) Ensure data is compatible with confidentiality obligations.
- (c) Ensure the use of personal data complies with GDPR and/or other applicable data protection regimes.
- (d) Consider the use of regulatory or development sandboxes.
- (e) Conduct a data protection impact assessment where necessary.

4.6.3 Ownership/IP issues

- (a) If the AI project involves collaboration with another party, consider who will 'own' the outcome of that project.
- (b) Identify what elements of the project they are most interested in – the algorithm, the data, the output of the algorithm and so on.
- (c) Clarify commercial aims for each element – that is, positive rights to use and negative rights to prevent others from using.

4.6.4 Regulation and liability

- (a) If provided under a contract, address the standards the system must meet and include appropriate limitation and exclusion provisions.
- (b) If not provided under a contract, consider the use of appropriate disclaimers with end users.
- (c) Align with wider government and regulatory policies on AI.

4.6.5 Ethical challenges

- (a) Be open and transparent about the use of AI.
- (b) Evaluate potential impact.
- (c) Where systems help make decisions about individuals, evaluate whether the process is fair, lawful and avoids discriminatory outcomes.
- (d) Work with AI providers to understand the data used to train generative AI tools and consider how results derived from such data may disproportionately affect individuals or communities.
- (e) Ensure there is human involvement and oversight in the decisions being made by AI technologies.



5 Trusted Providers

5.1 Relationships of trust

Regulators and the public alike will expect the gold standard in compliance and risk management from government departments that use AI technology. Unlike the cloud services market, which is represented by well-established, trusted vendors, the AI landscape is nascent, rapidly evolving and disparate by comparison. Government procurers are understandably cautious about entrusting large volumes of personal, confidential or classified information to new technologies and vendors.

As discussed, a key principle in safe AI procurement is accountability, and government customers that use AI will be required to explain the legitimacy of the decisions and impacts of the AI technologies they employ. To offer this transparency to service users, government customers in turn will need to rely on their technology providers to adequately articulate how the AI solutions they use make decisions (or support decision-makers) and generate outputs. In addition, delivering public sector services often requires the processing of sensitive or classified data and government customers require additional assurances and risk mitigations by AI technology providers for the safeguarding of that data.

Regular, meaningful discussion between industry providers and government is key to successful long-term partnerships that will ultimately drive wider participation by vendors and result in better end products. Microsoft is already a trusted provider of AI technologies to the public sector, but further inspiration can be drawn from Microsoft's track record of successful partnerships for AI solution deployment within the financial services sector. We have seen in the private sector the uptake of cloud and AI technology as a move away from a 'vendor-customer' relationship to a partnership of trust. Financial services, including institutions such as banks, are heavily regulated in most jurisdictions and regions and yet fintech partnerships are currently revolutionizing the banking world.

Ultimately, responsible AI begins with creators and developers who hold themselves to account. Microsoft has invested for many years in a program that ensures the AI systems we develop are responsible from the design phase upwards.

Microsoft understands that, with all AI systems, it is critical to understand their limitations as well as their capabilities. Using this knowledge and understanding has allowed Microsoft to identify common risks, such as the generation of content that perpetuated stereotypes, as well as the ability of AI systems to generate responses that, despite being convincing, were factually incorrect.²¹

The benefit of obtaining these key insights early in the development of AI systems is it has allowed corrections and additional steps to be added as the technology has developed. At both Microsoft and OpenAI, this has led to rapid progress and reinforced the depth and breadth of expertise needed to advance the state of the art on responsible AI, in addition to highlighting the growing need for new norms, standards and laws.²²

To further help customers on their responsible AI journey, Microsoft recently launched the [Microsoft AI Customer Commitments](#).²³

These commitments focus on three areas:

1. Sharing and providing AI resources to help customers deploy AI responsibly;
2. Creation of an AI assurance program; and
3. Developing responsible AI partner programs. For example, most recently, Microsoft launched a partnership with PwC and EY to leverage responsible AI expertise to help mutual customers deploy their own responsible AI systems.

5.2 Microsoft's AI Principles

As discussed, Microsoft has been on a responsible AI journey since 2017 when we established our research-led AI, Ethics, and Effects in Engineering and Research (AETHER) committee in order to explore AI technologies and the standards that should apply to their development. Since then, Microsoft has amassed a wealth of relevant expertise and resource, including policy experts dedicated to implementing a robust governance process that guides the design, development and deployment of AI in safe, secure and transparent ways.²⁴

In 2019, Microsoft established the Office of Responsible AI in order to foster a comprehensive approach to Responsible AI. 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.

Microsoft has identified six principles that should guide AI development and use:



1. Fairness

Mitigating unfairness starts with understanding the implications and limitations of AI predictions and recommendations. It is important for developers to recognize how different types of bias can potentially be introduced into data, machine learning models, or systems that leverage multiple models, and for those developers to use tools, governance methodologies and other resources to help detect and mitigate unfairness. Developers should have robust governance processes that promote fairness and continually monitor for drift or deterioration in relation to the same. In addition, decisions made by AI models will need to be supplemented by the sound judgment of humans who are accountable for the conclusions of the AI models they have assessed.

2. Reliability and safety

AI systems can make a negative impact if they do not operate reliably, safely and consistently. To understand the consequences of this, one needs only to imagine the role of AI in making decisions such as who benefits from government aid – here, reliability and safety will be critical in managing genuine citizen needs in relation to health and poverty issues. Rigorous testing is essential during system development and deployment to prevent unexpected performance failures and to ensure systems do not evolve in ways that are inconsistent with original expectations. After testing and deployment, it is equally important that organizations properly operate, maintain and protect their AI systems over the lifespan of their use so the systems do not become unreliable or inaccurate.

3. Privacy and security

Access to data is essential for AI systems to be able to make accurate and informed decisions. As AI technologies become more prevalent, protecting privacy and personal data is becoming an increasingly complex and critical area of focus. In addition, new and improved authentication factors such as facial or voice recognition and behavioral biometrics are leading to new ways to collect and process personal data. It is essential that the benefits of AI, such as efficiency and accuracy gains, do not outweigh citizen rights to privacy and security.

4. Inclusiveness

At Microsoft, we firmly believe everyone should benefit from intelligent technology and, in the public sector in particular, there is huge potential for AI to be used to address a broad range of human needs and to level the access for people with different backgrounds, skill levels and perspectives. Inclusive design practices help system developers understand and address potential barriers in a product environment that could unintentionally exclude people.

5. Transparency

Since AI relies on extrapolated logic rather than hard-coded rules, it can sometimes function as a 'black box' where users do not necessarily understand how the system's outputs were derived from its inputs. However, regulation for AI increasingly demands that organizations provide transparency about how their AI technologies work. There are three components of transparency:

- Transparency relies on a foundation of traceability, with developers clearly documenting their goals, definitions, design choices, and any assumptions made in development of the particular AI model.
- Transparency requires effective communication from providers and customers as to how, when and why a particular AI model has been developed and deployed. An important component of this communication is honesty as to the limitations and downsides of a particular model.
- Transparency depends on intelligibility – that is, people should be able to fully understand and monitor the technical behavior of the relevant AI technology. This understanding helps data scientists evaluate and debug models and make informed decisions about how to improve the models over time. It also helps governments and customers to assess how much to trust to put in a model's predictions or recommendations.

6. Accountability

Anyone who deploys AI technologies will need to be accountable for how those technologies operate. For government users of AI technologies, it will be especially important to establish an internal governance framework that provides oversight and guidance to ensure that humans, and not AI systems, are the final authority on any decision affecting citizens.

Based on the above principles, Microsoft has launched its 'Responsible AI Standard', which is a multi-year effort to define product development requirements for responsible AI.²⁵ This standard consists of goals and requirements for each of the six principles and is intended to function as a checklist or scorecard for companies developing and implementing AI systems.

6 AI Regulation and Governance

6.1 Microsoft's position on AI policy

Microsoft's position on AI policy can be found in [Governing AI: A Blueprint for the Future](#) (25 May 2023), at the heart of which is a mission to advance a globally coherent framework for AI safety and security regulation that fosters trust. Microsoft was one of the first companies to have committed to implement the National Institute of Standards and Technology (NIST) AI Risk Management Framework²⁶ across our AI development and deployment practices, and to attest to this with our customers. (Further detail on the NIST framework is set out in [Section 6.3.5](#) below.)

Microsoft also welcomed the White House Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence (30 October 2020).²⁷

Microsoft's blueprint for the public governance of AI procurement considers the policy, law and regulation surrounding AI and recommends five key ways governments should manage AI implementation within their own systems:

6.1.1 Implement and build upon new government-led AI safety frameworks

One of the most effective ways to move quickly is to build on recent advances in governmental work that improve AI safety. Microsoft's AI Risk Management Framework provides a strong foundation that companies and governments alike can immediately put into action to ensure the safer use of artificial intelligence. While no single effort can answer every question, the immediate adoption of this framework will accelerate AI safety momentum around the world.²⁸

6.1.2 Require effective safety brakes for AI systems that control critical infrastructure

With AI becoming increasingly powerful and sophisticated, it is important to ensure sufficient checks and balances are in place for high-risk use cases; for example, where AI systems control critical infrastructure. Governments should require system developers to ensure safety brakes are built by design into the use of AI systems for the control of critical infrastructure and to ensure they are tested rigorously to certify their efficacy.²⁹

6.1.3 Develop a broader legal and regulatory framework based on the technology architecture for AI

There is a strong need to develop laws and regulations that are suited to the technology architecture for AI to ensure those involved in the development and operation of AI have clear legal obligations imposed upon them. In the development of AI, software companies build a 'stack', creating layers of functionality that work together to create AI technology. (See the table below setting out the full Microsoft AI technology stack.) Microsoft believes the most effective way to regulate AI is by focusing on the following three layers of the AI technology stack, with different obligations at each level: (i) Applications; (ii) pre-trained AI models; and (iii) AI data infrastructure. The corresponding obligations for each level of focus are set out in the table below.³⁰

AI Architecture Stack

Stack level	Description	Obligations
Applications	Software programs where the output of an AI model is put to work.	Ensure the use of AI in the application complies with all existing and evolving laws and regulations.
API services	APIs (Application Program Interfaces), or endpoints, through which applications access pre-trained models.	
Powerful pre-trained ai models	Pre-trained models like GPT-4 that can be used to solve similar problems without starting from scratch.	Regulate through the pre-release safety and security requirements, then license deployment for permitted uses in a licensed AI data center with post-deployment safety and security monitoring and protection.
Machine learning acceleration software	Software that speeds up the process of developing and deploying large AI models.	
AI data center infrastructure	Advanced supercomputing infrastructure, including clusters of advanced GPUs (Graphics Processing Units) with high bandwidth network connections.	License for training and deployment of powerful AI models based on security protections, export control compliance, and safety protocols to ensure human control over autonomous systems that manage critical infrastructure.

6.1.4 Promote transparency and ensure academic and non-profit access to AI

As AI continues to develop in complexity, it is vital to preserve transparency and access to AI resources without compromising security. Microsoft has committed to provide an annual AI transparency report to improve openness in relation to its AI technology. Microsoft is also committed to increasing access to AI resources for academic research and the non-profit community on the understanding that this is important for the continued development of AI.³¹ Any licensing regime would need to consider the impact on open-source models and the importance of continuing to foster an innovative open-source ecosystem.

6.1.5 Pursue new public-private partnerships to use AI as an effective tool to address the inevitable societal challenges that come with new technology

The best way to promote the safe and responsible use of AI is for the public and private sectors to work together. A whole-of-society approach is needed if AI is to be used to safeguard fundamental rights, promote inclusive growth and advance the planet's sustainability needs.³²

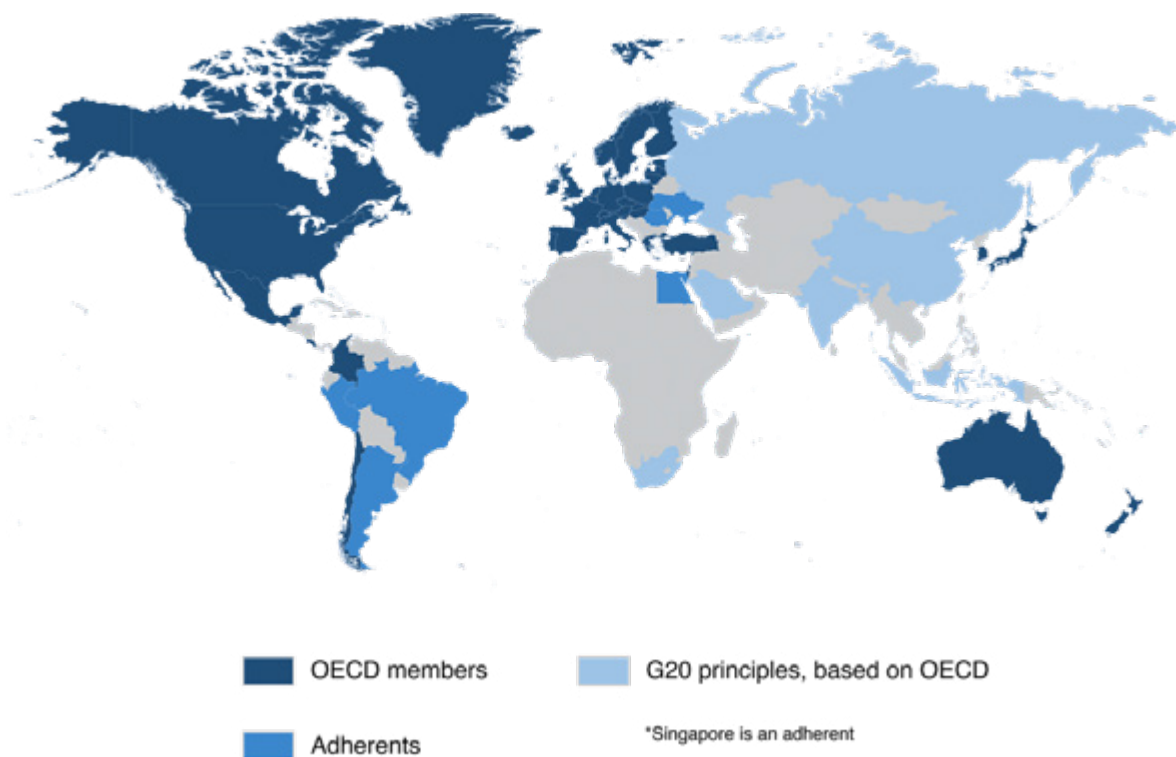


6.2 Common principles in AI regulation

While AI policy is still in its early stages of development for many countries, some key common principles are emerging. Adhering to these key principles when designing, procuring and using AI technologies in the public sector gives the best chance of remaining safe and compliant now and in the face of advancing innovation. Microsoft's own principles, covered in [Section 5.2](#) above, align closely with many of the principles proposed by the various international bodies below (most notably, those with a focus on transparency, robustness, fostering investment and accountability), further demonstrating Microsoft's commitment to ensuring its own AI policy meets the latest internationally agreed standards.

6.2.1 The OECD AI principles

In 2019, the Organisation for Economic Co-operation and Development (OECD) adopted a series of AI principles, with the aim of promoting the use and development of AI while also ensuring democratic and human rights are respected. These principles have since been adopted by 42 countries, including but not limited to the EU, the UK, the USA and Canada. When the list is expanded to include those that have adopted the G20 AI principles (which are based on the OECD AI principles), the list includes China, India and Russia.³³ It should be noted that these are legal principles and represent a formal commitment by the adopting countries to adhere to them.



The OECD principles are split into two separate categories: (i) the values-based principles; and (ii) subsequent recommendations for policy makers:

1. Values-based principles

- **Inclusive growth, sustainable development and wellbeing:** AI systems must be developed and used for the benefit and prosperity of people and the planet as a whole.³⁴
- **Human-centered values and fairness:** When developing AI systems, human rights, democratic value and rule of law must be at the forefront of the developers, with the ultimate aim of being able to contribute to a fair and just society.³⁵
- **Transparency and explainability:** Adequate disclosures and transparency must be undertaken regarding AI systems.³⁶
- **Robustness, security and safety:** Every AI system must be designed safely and in a robust manner, with adequate risk management being undertaken throughout its lifetime.³⁷
- **Accountability:** Those responsible for the development and operation of AI systems must be held accountable for their functioning in line with the above principles.³⁸

2. Recommendations for policy makers

- **Investing in AI research and development:** Governments must invest and procure the investment in research and development of AI technology.³⁹
- **Fostering a digital ecosystem for AI:** Governments must invest and develop adequate digital infrastructure to facilitate the development of AI technology.⁴⁰
- **Shaping an enabling policy environment for AI:** Governments must consider the development of AI when enacting policy and adapt policy where necessary.⁴¹
- **Building human capacity and preparing for labor market transformation:** Governments must invest in their people and workforce so they can contribute to the development of AI technology.⁴²
- **International cooperation for trustworthy AI:** Governments must work together to share information and research in order to strive towards responsible adoption and use of AI technology.⁴³

6.2.2 G7 Hiroshima Principles

On 7 September 2023, the G7 Digital and Tech Ministers and partners met virtually to discuss the opportunities and challenges of advanced artificial intelligence (AI) systems. At that meeting, the G7 ministers reiterated a commitment to encouraging “an environment where trustworthy AI systems are designed, developed and deployed for the common good worldwide including in emerging and developing economies in furtherance of democracy, human rights, the rule of law, and our shared democratic values and interests”. They also reaffirmed a commitment to work towards common international standards and tools for trustworthy AI (based on the OECD principles above) that also foster innovation.

The G7 ministers have stated an aim to develop a comprehensive framework of guiding policies and principles for responsible AI development and use by the end of the year. These principles could include:

- Take appropriate safety measures and consider societal risks before deployment, including placing on the market.
- Endeavor to identify and mitigate vulnerabilities after deployment, including placing on the market.
- Report publicly about models’ capabilities, limitations and domains of appropriate and inappropriate use, ensuring sufficient transparency.
- Work towards responsible information sharing among organizations developing AI and governments, civil society and academia.
- Develop and disclose risk management plans and mitigation measures, including privacy policies and AI governance policies.
- Invest in robust security controls, including cybersecurity and insider threat safeguards.
- Develop and deploy mechanisms such as watermarking or other techniques to enable users to identify AI-generated content.
- Prioritize research to mitigate societal, environmental and safety risks and prioritize investment in mitigations.
- Prioritize the development of advanced AI systems to address the world’s greatest challenges, notably but not limited to the climate crisis, global health and education.
- Advance the development of and alignment with internationally recognized technical standards.⁴⁴

6.3 Different Government Approaches to AI Regulation

Most governments are working hard to understand the nature of AI technology and its potential impact on the lives of their citizens so they can provide effective regulation and guidance on its development and use. Regulation in most countries is at an early stage, albeit evolving rapidly. In this section, we explore some of the approaches different countries are taking.

6.3.1 EU – AI Act

The draft EU AI Act was first published in 2021 and represents how the European Union will seek to regulate AI.⁴⁵ The purpose of the act is to ensure that the use of AI in the EU is ‘safe, transparent, traceable, non-discriminatory and environmentally friendly’.⁴⁶ The EU AI Act is founded on a belief that AI systems should ultimately be overseen by humans, rather than by the technology itself, and this is best way to ‘prevent harmful outcomes’.⁴⁷ The importance of human oversight over AI frameworks is reflected in the policy positions adopted by organizations in the private sector. (See, for example, Microsoft’s position summarized in [Section 6.1](#))

6.3.2 China – Interim Measures for the Management of Generative Artificial Intelligence

In July 2023, the Cyberspace Administration of China (CAC) issued the Interim Measures for the Management of Generative Artificial Intelligence, effective from 15 August 2023.⁴⁸ The aim of the Interim Measures is to encourage the development of AI systems while also ensuring there is adequate supervision in place.

The Interim Measures have set up four new systems for the regulation of AI. These are:

- Graded and categorized supervision: The relevant state authorities will create ‘graded and categorized supervision’ based on specific industries and fields, which are yet to be identified.
- Service agreements between providers and users: Service providers must execute service agreements with their users of generative AI in order to allocate risk between service providers and users. The Interim Measures are unclear, however, as to the full content or application of these agreements.
- Regulation on AI provided from outside China: AI services provided from outside China must comply with Chinese laws and regulations, and should they fail, the CAC has the power to impede the service in question.
- Foreign investment in generative AI services: Foreign investment in generative AI services must comply with all relevant rules and regulations imposed by the Chinese government.⁴⁹



6.3.3 Singapore – National AI Strategy

Singapore has set itself a goal of being a 'leader' in the development of AI by 2030, and its National AI Strategy 2019 sets out how it seeks to achieve this goal.⁵⁰

The strategy sets out three aims it seeks to accomplish. These are to:

- Identify areas to focus attention and resources at a national level.
- Set out how government, companies and researchers can work together to realize the positive impact of AI.
- Address areas where attention is needed to manage change and/or new forms of risks that arise when AI becomes more pervasive.⁵¹

The strategy also specifies specific sectors the Singapore government believes should be the focus of AI initiatives: Transport and Logistic; Manufacturing; Finance; Safety and Security; Cybersecurity; Smart Cities and Estates; Healthcare; Education; and Government.

Through targeting the above key sectors, which have a heightened social or economic value for Singapore, the government can maximize the financial and societal benefits of the use of AI technologies.⁵²



6.3.4 United Kingdom – Pro innovation

The UK government published its AI whitepaper in March 2023, which underlined the country's 'pro-innovation approach to AI regulation'.⁵³ The UK sees AI as one of its five 'critical technologies', acknowledging that, while AI technology is far from at its full potential, with the right strategy to help it thrive it can 'transform all areas of life'.⁵⁴

The whitepaper outlines a pro-innovation framework, which is governed by five key principles:

- Safety, security and robustness.
- Appropriate transparency and explainability.
- Fairness.
- Accountability and governance.
- Contestability and redress.⁵⁵

While these principles are not currently underpinned by statute, there is a plan to do so eventually, and the similarity of the UK's principles to the principles espoused by the OECD (see [Section 6.2.1\(a\)](#)) demonstrates a motivation on the part of the UK government to align its AI strategy with that of the international community.

In November, the UK will host the world's first summit on AI safety at Bletchley Park to consider the risks of AI, especially at the frontier of development, and discuss how they can be mitigated through internationally coordinated action. The UK's Technology Secretary, Michelle Donelan, has said, "the UK is consistently recognized as a world leader in AI and we are well placed to lead these discussions. The location of Bletchley Park as the backdrop will reaffirm our historic leadership in overseeing the development of new technologies". The summit will bring together international governments from 'like-minded' countries and experts in AI research.⁵⁶

6.3.5 United States of America

The Biden-Harris administration has been active in evaluating and responding to AI advancement with policy:

In October 2023, the Biden Administration issued an Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence, another critical step forward in the governance of AI technology. The Executive Order builds on the White House Voluntary Commitments and complements international efforts through the G7 Hiroshima Process. The Executive Order aims to ensure that AI is used responsibly and ethically, and that it benefits the public while mitigating the potential risks and harms.

Some of the main actions directed by the Executive Order are:

Requiring that developers of the most powerful AI systems share their safety test results and other critical information with the US government before making them public.

Developing standards, tools, and tests to help ensure AI systems are safe, secure, and trustworthy.

Protecting against the risks of using AI to engineer dangerous biological materials by developing strong new standards for biological synthesis screening.

Directing federal agencies to root out bias in the design, and use of new technologies, including AI, and to protect the public from algorithmic discrimination.

Establishing a new AI Safety and Security Board, an AI Advisory Committee, and an AI Interagency Policy Committee to coordinate and oversee the implementation of the Executive Order.

Promoting innovation and competition in the AI sector, and advancing American leadership around the world.⁵⁸

In January 2023, the National Institute of Standards and Technology, an agency of the US Department of Commerce, published the 'Artificial Intelligence Risk Management Framework'. The framework recognizes that, while AI technologies have the potential to advance global society for the better, there are significant risks posed by these technologies.⁵⁷ The aim of the framework is to ensure parties that use and rely on these technologies understand the risks posed by AI and how best to manage these risks.⁵⁹

It splits risk management into four different 'functions':

Govern

This function seeks to promote and encourage a culture of risk awareness and management at all levels of an organization through the implementation of different policies, processes and procedures.⁶⁰

Map

Through the better understanding of the technology behind AI systems, it will be easier to understand and find the risks posed by them.⁶¹

Measure

By using a mixture of quantitative and qualitative methods to analyze the risks posed by AI systems, the framework suggests it will be possible to better track the risks identified by the Map function and better inform by better understanding which risks to prioritize for the Manage function.⁶²

Manage

The Manage function takes the risks identified in the Map function, and the risks prioritized using the Measure function, to adequately allocate risk resources in developing plans to respond to incidents.⁶³



In February, President Biden signed an Executive Order that instructs federal agencies to root out bias in their design and use of new technologies, including AI, and to protect the public from bias and/or discrimination in algorithmic outputs.⁶⁴

Between July and September 2023, the Biden-Harris Administration has managed to secure voluntary commitments from 15 providers of AI solutions, including Microsoft, to 'help drive safe, secure and trustworthy development of AI technology'.

These commitments are:

Ensuring products are safe before introducing them to the public.

This includes committing to:

- Internal and external security testing of their AI systems (with independent experts) before their release.
- Sharing information across the industry and with governments, civil society and academia on managing AI risks.

Building systems that put security first. This includes committing to:

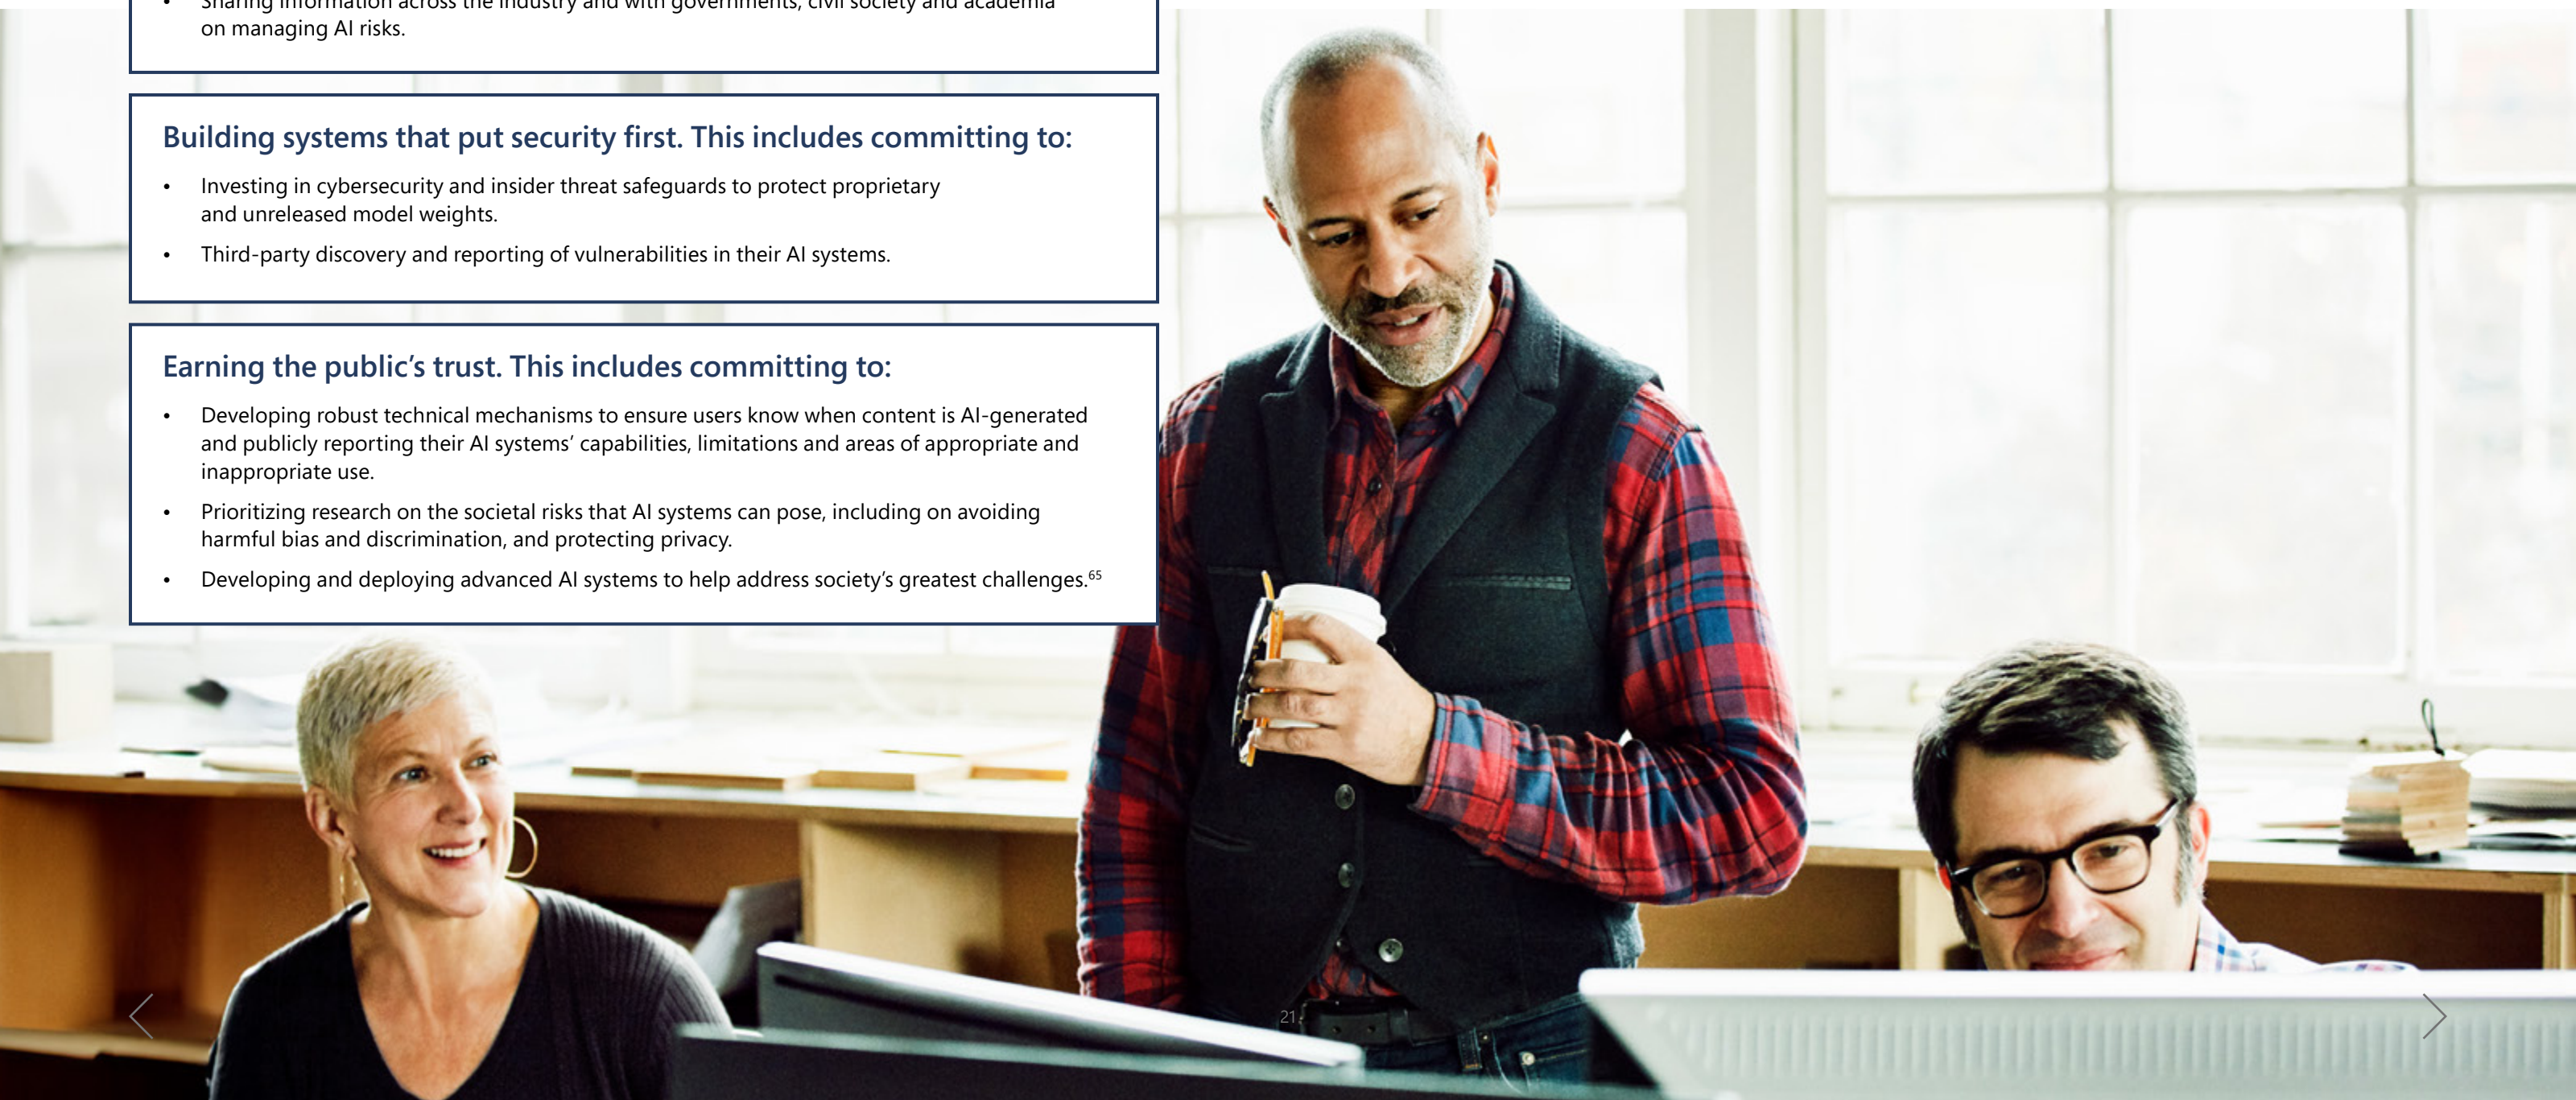
- Investing in cybersecurity and insider threat safeguards to protect proprietary and unreleased model weights.
- Third-party discovery and reporting of vulnerabilities in their AI systems.

Earning the public's trust. This includes committing to:

- Developing robust technical mechanisms to ensure users know when content is AI-generated and publicly reporting their AI systems' capabilities, limitations and areas of appropriate and inappropriate use.
- Prioritizing research on the societal risks that AI systems can pose, including on avoiding harmful bias and discrimination, and protecting privacy.
- Developing and deploying advanced AI systems to help address society's greatest challenges.⁶⁵

On 12 September 2023, two US Senators, Richard Blumenthal and Josh Hawley, released a bipartisan framework for artificial intelligence (AI) legislation. The Blumenthal-Hawley framework includes proposals for a licensing regime under an independent oversight body with a risk-based approach for AI models and uses. In his testimony to the US Senate Judiciary Committee, Microsoft's Vice Chair and President, Brad Smith, set out why Microsoft supports this approach and believes it strikes a sensible balance that can both protect the public and advance innovation. He noted that, "by also incorporating transparency and security requirements, the Blumenthal-Hawley framework puts Congress on a path to provide the public with the safety standards it deserves".⁶⁶

In addition, different US states have enacted, or are planning to enact, different legislation affecting the development and use of AI. Connecticut is the first US state to regulate the use of AI. SB 1103 was signed on 7 June 2023, paving the way to a clear path to AI use in public services.⁶⁷



7 Conclusion

The increasing use of AI technologies to drive economic efficiencies, improve lives and help solve global challenges is inevitable.

Governments wanting to take advantage of the benefits of AI in a safe and responsible manner require a route to procurement that will necessarily include:

- (i) understanding how a particular AI solution can solve their specific requirements;
- (ii) preempting and preparing for the skills and competencies that will be required to maximize the opportunities afforded by AI;
- (iii) working with trusted providers of AI solutions who adhere to globally accepted responsible AI principles; and
- (iv) adopting policies and procurement practices that encourage innovation while protecting the rights of private citizens.

Technology is a proven pillar of competitiveness and growth for governments worldwide. Countries 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 whilst better engaging with and supporting their citizens.

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.

Appendix 1

Examples of Public Sector Use Cases

1. [The UK Government allocates £13 million to revolutionize healthcare research through AI. The funding supports a raft of new projects including transformations to brain tumour surgeries, new approaches to treating chronic nerve pain, and a system to predict a patient's risk of developing future health problems based on existing conditions](#) – August 2023

The UK Government has announced that it plans to invest £13 million into research that aims to increase developments in the healthcare sector through the use of AI.

2. [Mexico's Supreme Court of Justice of the Nation provides the public with digital access to the court's historical documents with the help of Azure](#) – July 2023

Mexico's Supreme Court has used language models created using Microsoft Azure AI to build a search engine for a library of digitized files. The search engine allows the members of the public to search for nearly 100,000 judgments and 300,000 theses, granting those with no legal knowledge unprecedented access to the law.

3. [With the adoption of AI, AGU seeks to improve efficiency in legal proceedings](#) – June 2023

The Attorney General's Office of São Paulo has adopted GPT-4 technology, which is a part of Microsoft Azure OpenAI, to increase the efficiency of their lawyers and the legal system in general. The technology is helping to speed up the review of lawsuits and subpoenas by producing texts of court cases from analysis of previous cases.

4. [The City of Kelowna increases and speeds access to its services with Azure AI](#) – June 2023

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. Citizens' concerns and queries are now dealt with using a combined human and AI service that has improved the time it takes to resolve issues.

5. [Broward College harnesses the power of Azure Machine Learning and responsible AI to increase student retention rate](#) – May 2023

Broward College, a higher education institution in the United States, is using Azure Machine Learning and responsible AI to create models to better understand why students drop out before finishing their studies.

6. [Brazilian state embraces explainable AI for a fair and sustainable tax environment](#) – May 2022

The state of Ceará in Brazil adopted the HMX Tax Intelligence System, which uses Microsoft Azure AI to monitor and identify any compliance issues regarding the taxation of retail vendors. The new system has led to a 21% increase in tax revenues and an 84% increase in audit efficiency.

7. [Ville de Laval speeds up community response systems with an Azure AI solution](#) – September 2021

Ville de Laval in Quebec, Canada, has integrated AI into its non-emergency hotline, which receives almost 200,000 calls a year. The initiative is using AI to help the system respond to simpler requests, reducing the need for physical agents.

References

- 1 Microsoft, [Meeting the AI moment: advancing the future through responsible AI](#) (2 February 2023)
- 2 McKinsey: [AI, automation, and the future of work: Ten things to solve for](#), Section 1 (1 June 2018)
- 3 US Department of Homeland Security and Governmental Affairs, [Peters Convenes Hearing To Examine How Federal Government Purchases Artificial Intelligence Technologies](#) (14 September 2023)
- 4 UK Information Commissioner's Office, [Registration Inbox AI](#) (6 July 2022)
- 5 [Brazilian state embraces explainable AI for a fair and sustainable tax environment](#) (May 2022)
- 6 [The City of Kelowna increases and speeds access to its services with Azure AI](#) (June 2023)
- 7 Australian Government, [Interim guidance for agencies on government use of generative artificial intelligence platforms](#) (accessed on 1 October 2023)
- 8 McKinsey: [AI, automation, and the future of work: Ten things to solve for](#), Section 1 (1 June 2018)
- 9 Gartner, [3 Barriers to AI Adoption](#) (18 September 2019)
- 10 McKinsey: [AI, automation, and the future of work: Ten things to solve for](#), Section 1 (1 June 2018).
- 11 The World Bank (International Finance Corporation), [The Role of Artificial Intelligence in Supporting Development in Emerging Markets](#) (July 2019)
- 12 Microsoft, [Announcing Microsoft's AI Customer Commitments](#) (8 June 2023)
- 13 Microsoft, [Advancing AI governance in Europe and internationally](#) (29 June 2023)
- 14 Sections 243 and 244 of the [Copyright Act 2021](#), Republic of Singapore
- 15 European Commission, [Europe fit for the Digital Age: new online rules for platforms](#) (accessed on 24 July 2023)
- 16 UK Government, [A guide to the Online Safety Bill](#) (accessed on 1 September 2023)
- 17 Ministry of Home Affairs, [Introduction of the Online Criminal Harms Bill](#) (8 May 2023)
- 18 Microsoft, [Public Sector Procurement Fit for the Digital Age](#) and [Building Blocks](#) papers
- 19 McKinsey [Getting ahead in the cloud](#) by Kreg Nichols and Kara Sprague (Autumn 2011)
- 20 [Canada's List of Interested Artificial Intelligence \(AI\) Suppliers](#) (accessed on 1 October 2023)
- 21 Microsoft, [Meeting the AI moment: advancing the future through responsible AI](#), [Microsoft On the Issues](#) (2 February 2023)
- 22 Microsoft, [Meeting the AI moment: advancing the future through responsible AI](#), [Microsoft On the Issues](#) (2 February 2023)
- 23 Microsoft, [Announcing Microsoft's AI Customer Commitments](#) (8 June 2023).
- 24 Microsoft, [Announcing Microsoft's AI Customer Commitments](#) (8 June 2023).
- 25 Microsoft, [Responsible AI Standard, V2](#) (June 2022)
- 26 National Institute of Standards and Technology, [Artificial Intelligence Risk Management Framework \(AI RMF 1.0\)](#), January 2023
- 27 LinkedIn, [Microsoft's Brad Smith post](#), 30 October 2023
- 28 Microsoft, [Governing AI: A Blueprint for the Future](#), page 10 (25 May 2023)
- 29 Microsoft, [Governing AI: A Blueprint for the Future](#), page 12 (25 May 2023)
- 30 Microsoft, [Governing AI: A Blueprint for the Future](#), page 15 (25 May 2023)
- 31 Microsoft, [Governing AI: A Blueprint for the Future](#), page 22 (25 May 2023)
- 32 Microsoft, [Governing AI: A Blueprint for the Future](#), page 26 (25 May 2023)
- 33 OECD.AI, [OECD AI Principles overview](#) (accessed 22 August 2023)
- 34 OECD.AI, [Inclusive growth, sustainable development and well-being \(Principle 1.1\)](#) (accessed 22 August 2023)
- 35 OECD.AI, [Human-centred values and fairness \(Principle 1.2\)](#) (accessed 22 August 2023)
- 36 OECD.AI, [Transparency and explainability \(Principle 1.3\)](#) (accessed 22 August 2023)
- 37 OECD.AI, [Robustness, security and safety \(Principle 1.4\)](#) (accessed 22 August 2023)
- 38 OECD.AI, [Accountability \(Principle 1.5\)](#) (accessed 22 August 2023)
- 39 OECD.AI, [Investing in AI research and development](#) (accessed 22 August 2023)
- 40 OECD.AI, [Fostering a digital ecosystem for AI](#) (accessed 22 August 2023)
- 41 OECD.AI, [Shaping an enabling policy environment for AI](#) (accessed 22 August 2023)
- 42 OECD.AI, [Building human capacity and preparing for labour market transformation](#) (accessed 22 August 2023)
- 43 OECD.AI, [International co-operation for trustworthy AI](#) (accessed 22 August 2023)
- 44 G7 Hiroshima AI Process, [G7 Digital & Tech Ministers' Statement](#) (accessed 1 October 2023)
- 45 European Union, [AI Act](#) (accessed 22 August 2023)
- 46 European Parliament, [EU AI Act: first regulation on artificial intelligence](#) (accessed 22 August 2023)
- 47 European Parliament, [EU AI Act: first regulation on artificial intelligence](#) (accessed 22 August 2023)
- 48 PwC, [Regulatory and legislation: Interim Measures for the Management of Generative Artificial Intelligence Services officially implemented](#) (accessed 22 August 2023)
- 49 PwC, [Regulatory and legislation: Interim Measures for the Management of Generative Artificial Intelligence Services officially implemented](#) (accessed 22 August 2023)
- 50 Smart Nation Singapore, [National AI Strategy](#) (accessed 22 August 2023)
- 51 Singapore, [National AI Strategy - Summary](#), at page 5 (accessed 22 August 2023)
- 52 Singapore, [National AI Strategy](#), page 18 (accessed 22 August 2023)
- 53 Gov.UK, [A pro-innovation approach to AI regulation](#) (accessed 22 August 2023)
- 54 Gov.UK, [A pro-innovation approach to AI regulation](#), executive summary (accessed 22 August 2023)
- 55 Gov.UK, [A pro-innovation approach to AI regulation](#), executive summary (accessed 22 August 2023)
- 56 UK Government, [Iconic Bletchley Park to host UK AI Safety Summit](#) (24 August 2023)
- 57 NIST, [Artificial Intelligence Risk Management Framework](#), page 1 (accessed 22 August 2023)
- 58 The White House fact sheet: [President Biden Issues Executive Order on Safe, Secure and Trustworthy Artificial Intelligence](#) (30 October 2023)
- 59 NIST, [Artificial Intelligence Risk Management Framework](#), page 2 (accessed 22 August 2023)
- 60 NIST, [Artificial Intelligence Risk Management Framework](#), page 21 (accessed 22 August 2023)
- 61 NIST, [Artificial Intelligence Risk Management Framework](#), page 24 (accessed 22 August 2023)
- 62 NIST, [Artificial Intelligence Risk Management Framework](#), page 28 (accessed 22 August 2023)
- 63 NIST, [Artificial Intelligence Risk Management Framework](#), page 31 (accessed 22 August 2023)
- 64 The White House, [Executive Order on Further Advancing Racial Equity and Support for Underserved Communities Through The Federal Government](#) (16 February 2023)
- 65 The White House fact sheet: [Biden-Harris Administration Secures Voluntary Commitments from Eight Additional Artificial Intelligence Companies to Manage the Risks Posed by AI](#) (12 September 2023)
- 66 Written Testimony of Brad Smith Vice Chair and President, Microsoft Corporation U.S. Senate Judiciary Committee on [Developing and Deploying AI Responsibly: Elements of an Effective Legislative Framework to Regulate AI](#) (12 September 2023)
- 67 LegiScan, Connecticut General Assembly, [Senate Bill 1103](#) (accessed on 1 October 2023)



© Copyright 2023. Microsoft Corporation. All rights reserved.

Microsoft, Excel, Office 365, Power BI, Teams and Viva are trademarks of the Microsoft Group of Companies. All other names are the trademarks of their respective owners. Microsoft makes no warranties, express or implied in this publication. Views expressed in this publication are not necessarily those of Microsoft. Information and views expressed in this document including URL and other internet references may change without notice.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product.
You may copy and use this document for your internal, reference purposes. 21273 12/23