



Unlocking Tomorrow:

**A PwC and Microsoft playbook that
explores the power of agentic AI
for the Public Sector**



Table of contents

Foreword	03
Executive summary	04
Introduction	06
01 Understanding agentic AI	08
• What AI agents can do	09
• Driving business value with intelligent automation	09
02 How agentic AI can transform the public sector	12
• Government	15
• Education	19
• Defence	22
03 Building a solid, responsible foundation for agentic AI	25
• Cloud platforms as the backbone	26
• Data quality and security	27
• Ethical concerns and building public trust	28
• People are key to the success of agentic AI	30
04 Moving towards implementation and driving business value	32
• Strategy and discovery	33
• Design and deploy	34
• Scale and collaborate	34
05 The importance of partnerships	36
Conclusion	39
Contacts	40

Foreword



Agentic AI has the potential to offer a range of new benefits for organisations across almost every industry, with the public sector set to be one of the biggest beneficiaries. By embedding agentic AI into the fabric of public service, governments can not only enhance responsiveness and citizen satisfaction but also stimulate economic growth by automating workflows, enhancing productivity and enabling more intelligent and data-driven decisions and resource allocation.

AI agents are already helping public sector organisations improve the value of the services they provide and transform how they are delivered. But agentic AI isn't just improving efficiency today; it can also help the public sector anticipate what citizens and businesses will need in the future. It enables public sector organisations to create new, innovative solutions for the people and businesses they serve. Well-designed agentic AI systems can also drive iterative and continual improvement.

When implemented with the right infrastructure and governance, agentic AI can contribute to stronger outcomes, from smoother access to needed services and proactive support for businesses, to a safer, more secure world, to open, inclusive and personalised learning experiences. If deployed across a broad range of public sector enterprises, agentic AI can empower organisations to act as dynamic agents of change and help support broader economic growth objectives.

Executive summary

The public sector faces constant pressure to do more with less, and to deliver faster and more modern digital services. Citizens and businesses have grown accustomed to always-on, personalised and responsive services in commercial settings, and they expect the same from public sector organisations. Such services will also help the public sector to combat declining trust in institutions, anticipate changing demands and adapt operations to change.

Agentic AI can help public sector organisations achieve these goals. It has the potential to reshape operations by introducing autonomous, goal-oriented systems capable of supporting complex workflows across departments with minimal human supervision. Well-crafted agents can adapt in real time, integrate with legacy systems, and assist with mission-critical decisions—capabilities that are increasingly essential in addressing today's complex societal and operational challenges.

They enable public sector organisations to not only automate tasks but optimise processes and make smarter, faster, more data-driven decisions in near-real time. Alongside other AI tools, agentic AI can help public sector organisations anticipate tomorrow's needs and develop new solutions to address them. The resulting benefits will transform how governments serve their citizens, how educational institutions teach students and how military organisations defend people and critical infrastructure against both physical and online threats. Public sector leaders recognise the vast opportunities: A recent Microsoft survey in the UK¹ found that 72% of public and private sector leaders expect agentic AI to eventually be fully integrated into their organisations, and 39% expect that to happen by 2027.



1- Microsoft: [Agents of Change](#)

But to reap the benefits of this technology, the public sector must understand the requirements and build a strong foundation for data, infrastructure and governance. This requires attention to cloud infrastructure, privacy and security regulations, effective change management and a commitment to best practices and Responsible AI. Paying attention to the people who will ensure the success of agentic AI is equally important; public sector organisations will need to focus on training and motivating current staff and recruiting employees with the right skills to implement agentic AI systems with care and get the most out of them in day-to-day settings.

Before they implement AI agents, organisations must identify clear objectives and ensure that these align with their goals. They also need to understand their existing capabilities and readiness for adoption. Deployment should begin with small pilots, with a focus on the highest-impact use cases in which agents can quickly deliver meaningful benefits. Organisations can then scale up over time, using what they've learned through monitoring and oversight mechanisms like a Centre of Excellence to continually optimise and innovate.

Working with partners such as PwC and Microsoft enables public sector organisations to build on the partners' established experience and expertise in agentic AI, and to bring about truly transformational changes in how the public sector achieves meaningful outcomes for its stakeholders. With a well-focused AI vision and a proven roadmap to AI agent deployment, governments, defence organisations, educational institutions and more can serve their constituents more cost-effectively, efficiently and intelligently than ever.



Introduction

Around the world, public sector organisations provide a wide range of essential—and in many cases vital—services to their constituents, from infrastructure and education to national defence. The roles and responsibilities vary from country to country, as do public sector size and spending.

Worldwide, the public sector
accounts for approximately

11 %
of all employment.²

The nature of the COVID-19 pandemic also required public sector organisations to adapt quickly and adopt new ways of working to maintain services while protecting public health. For many institutions, this accelerated the adoption of new digital applications and platforms³, often in collaboration with private enterprises.

2- International Labor Organization: [Who powers the public sector?](#)

3- United Nations Department of Economic and Social Affairs: [World Public Sector Report 2023](#)

With its digital transformation already under way, the public sector now has a significant opportunity to adopt another innovation with unprecedented potential to increase employee productivity, reduce costs and improve outcomes for citizens: agentic AI.

Just a few short years after the arrival of widely available GenAI tools-primarily focused on content generation and conversational interfaces-we're now seeing the emergence of an even more transformative innovation: agentic AI. AI agents are providing forward-thinking organisations with an unprecedented ability to intelligently automate repetitive workflows, extract data-driven insights for enhanced decision making, reduce costs and free up human resources to focus on more strategic and value-adding tasks.

As public sector organisations integrate AI agents into their operations, they can make noticeable improvements to their workflows. Governments will be able to deliver faster, more responsive, more efficient and more valuable services to businesses, citizens and other constituents. Defence organisations will gain a force multiplier that strengthens both offensive and defensive capabilities to keep citizens, soldiers and critical infrastructure safer. Schools can transform how they engage with students, achieving the ability to provide more personalised and effective educational support.

Organisations that embrace the potential of agentic AI and make transformation an urgent priority will be the ones that are best prepared to meet the challenges of the coming decades. They will gain a first-mover advantage by innovating early and building their capabilities now. This will enable them to reap great cost and efficiency benefits and meet the rising expectations of those they serve, while late adopters will find themselves struggling in the years ahead.

The public sector stands at a crossroads. The following pages provide a guide for public sector transformation through the adoption of agentic AI.

01



Understanding agentic AI

Artificial intelligence is set to revolutionise our lives and work. Past technological breakthroughs from the personal computer to the internet brought big changes, but AI promises to supercharge individual and business capabilities in exponentially greater ways—especially with the potential of agentic AI.

What AI agents can do

Agentic AI systems can act autonomously to achieve specified goals. They do this by being able to string together and automate multiple tasks, enabling an agent to execute complex processes with minimal human intervention from start to finish. AI agents are typically goal oriented—created to achieve specific objectives—and interactive, which means that they can communicate with and respond to other systems, adapting their actions accordingly.

These abilities make AI agents useful for a wide variety of often tedious and time-consuming tasks that are handled manually today. Agents can be integrated into almost any process, can be scaled quickly, and operate effectively as an extension of the workforce—whether that's in HR, procurement, inventory management, IT or another area. By streamlining and accelerating work across multiple applications, data silos and departments, they can improve efficiency, enhance customer experiences and integrate disparate business processes.

Driving business value with intelligent automation

Working together with Microsoft, PwC has developed a strategic framework for agentic AI. It's designed to serve as the foundation for enterprise transformation. This approach enables organisations to connect AI agents with collaboration tools and business applications—including CRM, ERP and other enterprise systems—while orchestrating complex logical processes through customised AI models and advanced data intelligence.

The framework is structured around four interconnected pillars that collectively drive measurable business outcomes:

- **Augmenting decision-making with intelligent insights**

AI agents provide timely, data-driven insights that support strategic planning and enable faster and better-informed decision-making. This enhances business agility and improves overall performance.

- **Empowering the workforce with AI-driven productivity**

By automating repetitive tasks and augmenting human capabilities, AI solutions increase workforce productivity, foster innovation and allow employees to focus on higher-value strategic initiatives. This also supports upskilling and talent development.

- **Reinventing customer experiences with AI agents**

AI agents enable personalised, scalable service delivery, improving customer satisfaction and engagement across channels.

- **Streamlining business processes with AI agents**

Intelligent automation accelerates workflows, reduces operational errors, and enhances efficiency across functional and industry-specific processes.

Augmenting decision-making with intelligent insights

Use AI agents to deliver actionable insights, enabling leadership to make faster, data-driven, and informed decisions.

- Faster decisions
- Aligned strategies
- Reduced manual analysis

Reinventing customer experiences with AI agents

Deliver personalised, real-time solutions powered by AI agents, transforming customer interactions and driving loyalty.

- Personalised experiences
- Scalable service delivery
- Greater customer loyalty

Empowering the workforce with AI-driven productivity

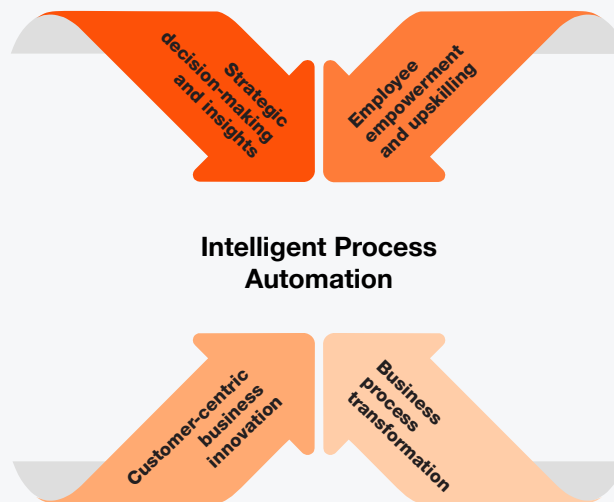
Enhance productivity and innovation with AI tools, automating repetitive tasks and enabling employees to focus on strategic work:

- Higher productivity
- Upskilled employees
- More strategic focus

Streamlining business processes with AI agents

Automate workflows with AI agents to enhance efficiency, accuracy, and speed in business operations.

- Faster workflows
- Reduced errors
- Improved efficiency



Through this integrated model, organisations can work smarter, reduce costs, extract deeper insights from data and transform how decisions are made and how value is delivered.

Well-designed AI agents are built with security and compliance as priorities. This helps organisations that deploy such tools to minimise risks and adhere to even the strictest regulatory requirements. For example, Microsoft Copilot Studio—a graphical, low-code tool for building agents and agent flows—supports security and governance controls for regulatory compliance, data residency, data loss protection and more.

Even with these built-in protections, however, organisations must always apply security best practices to the environments in which agents will operate, as well as to any internal and external interactions. They must also consider identity management, data protection and API security, among others.

Organisations will find three types of agents to be particularly valuable. These include task agents, used to automate specific actions such as scheduling meetings or processing transactions; autonomous agents, which operate with little human intervention to act on predefined goals, make decisions and return specified outcomes; and retrieval agents, designed to seek out data from a variety of sources.

By adopting agentic AI, public sector organisations can combine human and AI potential to become intelligent enterprises. This will enable them to introduce more dynamic, adaptive ways of working and gain new levels of flexibility and responsiveness.



02

How agentic AI can transform the public sector



Public sector organisations face constant pressure to deliver more and better services with fewer resources, and to make the most of every bit of funding they receive. For many public sector organisations and governments, using AI to achieve these goals is seen as a clear priority. In 2018, all EU member countries signed up to the Declaration of Cooperation on Artificial Intelligence⁴, aimed at boosting AI capacity and adoption in Europe, which commits to making AI available and beneficial to public administrations.

A number of countries across Europe, the Middle East and Africa (EMEA) have already made significant inroads into using AI in the public sector, with many taking lessons learned from successful deployments across the private sector. As of 2025, 56% of CEOs in the private sector say that GenAI has helped employees use their time more efficiently⁵ and one-third say it has increased their organisation's revenue and profitability. That represents a great opportunity for public sector enterprises to build on private sector experiences. With the addition of AI agents, organisations of all kinds now have the opportunity to gain even greater capabilities. This will help to revolutionise the breadth and depth of services that the public sector can offer.

At the same time, some segments of the public sector have struggled to adopt modern digital technologies—not only because of the investment required but because of their dependence on large and complex legacy infrastructures that are difficult to update.

By adopting cloud-based services and agentic AI, the public sector has great potential to overcome these challenges. Modern cloud services make digital transformation easier and more cost-effective than ever, while also enabling organisations to quickly deploy and scale AI agents that boost efficiency, productivity and innovation while reducing costs.

4- OECD report: [Hello, world: artificial intelligence and its use in the public sector](#), 2019


5- PwC: 28th Annual Global CEO Survey: [Reinvention on the edge of tomorrow](#)

Yasmine Rifai, Senior Strategic Account Technology Leader in the Public Sector at Microsoft Kuwait, highlights the transformative potential of AI agents:



Agentic AI stands at the intersection of innovation and necessity, empowering the public sector to accelerate services, enhance decision-making, and bridge the gap between growing demands and limited resources.

As promising as these innovations are, public sector organisations must also sharpen their focus on cybersecurity. The use of agents and AI models also brings new challenges, particularly around intellectual property rights and data protection. To mitigate these risks, public sector organisations must prioritise robust cybersecurity strategies, ensuring a secure and resilient foundation for their technological advancements.



With the right cloud and security strategies, organisations build a foundation that enables them to achieve the full potential of agentic AI. This will deliver benefits to stakeholders of all kinds, from employees to citizens and businesses to entire communities, including local, cross-regional and global—and in every context, from government to education to defence. Here, we outline the impact on these key subsectors.

Government

Around the world, governments of all kinds—from local authorities to national governments—are seeking to improve their internal ways of working and to better serve constituents, whether these are employees, citizens, businesses or others. Through the power of AI, they have a real opportunity to drive transformative change across the public sector and AI agents provide a viable solution to long term challenges. They help organisations to make more data-driven decisions, enhance and streamline citizen-facing services and improve the commercial environment by accelerating bureaucratic processes involving businesses. The public sector can also use agentic AI to streamline internal processes, making them more efficient and less complex.

Agnieszka Gajewska, PwC Global Government and Public Services Leader, Partner, PwC Poland, on how agentic AI can help improve both policy-making and citizen-facing services:



By using agentic AI governments can unlock transformative change. For so many years, we've been hearing about public sector inefficiencies and the need for more data-driven policy and innovation. While AI is not a panacea for all challenges, it offers enormous opportunities for governments—starting with tools for better and faster data analysis, predictive modelling in policy making to using AI agents to automate and improve delivery across front office and back-office functions. AI has the potential to become **a strategic foundation for a more trusted government** that delivers on its public service mandate and dramatically improves the citizen experience.

For public sector organisations that deploy agentic AI, the results can be revolutionary. Carefully-crafted agents can drive better and more citizen-focused support, faster and more personalised services and greater innovation across every area. They also unlock vast new potential to do much more with existing resources, maximising the value of every penny spent. And they can act on new data and developments in near-real time or even anticipate and respond to changing conditions.



Supporting the day-to-day work of government: faster, more efficient, more responsive services

There are many potential applications of agents in the day-to-day work of government. Municipal governments that face a high-volume of communications and requests from local citizens can develop agents to categorise correspondence according to the appropriate local services and automate responses. For example, one local council in the UK is using a Microsoft AI Builder custom prompt to streamline their response process. They expect to save more than 300 hours per annum.

Agents can also help streamline a wide variety of government services and support cross departmental coordination and better use of diverse resources. For example, civil servants can more efficiently process planning applications and more easily navigate the different departments that might be involved in making decisions. Agents can automatically check various regulations, and review government records, maps and other data sources.

Responsiveness can be greatly improved within specific workflows too. Issues with planning applications can be identified by agents earlier in the process and agents can instantly notify applicants about missing supporting documentation or other problems so planning approvals can be reached much faster. Application approvals are often slowed down because the applicant failed to submit a key piece of information or a required document. Instead of waiting months to learn this—and therefore wait even longer for a re-assessment—such issues can be resolved immediately.

These kinds of capabilities can enable planning authorities to cut through bureaucratic red tape, not only improving internal operations, but benefitting citizens and businesses directly through faster, better services. Ultimately this has the potential to support economic growth.

Giving government employees more time for key functions like providing social care

Another example of how AI agents can help reduce bureaucracy is the use of AI agents to help support the provision of social care to adults and children. Interventions in children's social care can have a lifelong impact on their wellbeing but budgets are often under intense pressure and experienced staff are scarce.

PwC has supported local authorities to refine and enhance the underlying processes to reduce bureaucracy, expedite decision-making, and allow care workers to spend more quality time with the children and their families. To achieve these advancements, PwC, in collaboration with Microsoft, has integrated AI agents into the process. These agents accelerate the drafting of comprehensive notes and initial reports following meetings and interactions and transform vast amounts of unstructured data into insightful and actionable reports.

Considering the confidential and sensitive nature of the data, these agents are carefully integrated with existing social care systems, ensuring strict data privacy is maintained. This allows case workers to manage their responsibilities more effectively and make informed short- and long-term decisions. Similar solutions could also help officials provide assistance to job seekers or process immigration requests.

Didier Ongena, Vice President and Global Government Leader for the Worldwide Public Sector at Microsoft, describes the technology backing up such use cases:



Summarisation can put unstructured data into different kinds of case management tools, creating new opportunities to draw valuable insights out of the vast amounts of information that governments manage. You're summarising those notes and putting that into the system as structured data. That's a huge use case. And it's where we're seeing an immediate return on investment and tangible impact.

Empowering decision-making and developing pro-active services

AI agents also have the potential to help civil servants make faster and more-informed decisions in other settings as well, for example for long-term budget planning and strategy development.



Citizen-facing services can be improved by using agentic AI to power virtual assistants, making it possible to provide 24/7 information and services. Agents can enable more proactive services, too. For instance, by analysing data to identify at-risk individuals so they can be provided with supportive services as early as possible. And agents can deliver significant value simply through the ability to summarise large volumes of unstructured data, enabling government employees to have more information at their fingertips.

All of these capabilities offer significant potential benefits for government organisations. And they offer a way in which governments can improve citizen trust in how they serve the public.

Education

The education sector had to adapt rapidly during the COVID-19 pandemic, which accelerated its adoption of digital technologies. Since then, educators and technology leaders alike have explored new ways to bring together the physical and digital worlds to make teaching and learning more effective, flexible, personalised and relevant to today's students' needs. At the same time, school and university administrators are constantly looking for ways to improve efficiency and streamline administrative processes, while still achieving strong educational outcomes. That's becoming increasingly urgent—a recent study found that European university administrators expect costs to rise over the next five years, without any corresponding increase in funding.⁶

Many are finding the answer is AI and, increasingly, agentic AI, which has tremendous potential to benefit educational organisations, teachers and learners. New solutions are emerging to help school managers automate repetitive tasks. Agents can help administrators schedule classes, generate reports and much more. Educators can also use agents to reduce the time they spend on tasks like attendance tracking, test preparation, grading or generating education plans for students with special needs. One provider of educational tools estimates that new tools can improve the efficiency of administrative work by 25%.⁷ It can also help speed up more creative activities like designing new courses. Further, agentic AI can also analyse large volumes of student data, helping staff or educators to identify areas where students might need special attention or individuals who might be at risk of dropping out.

Caitrina McCusker, Partner at PwC UK, describes the benefits of agentic AI for the education system:



We are starting to see huge gains in administrative efficiency being made in areas like grading, scheduling and managing student records. That's resulting in dramatically reduced workload for administrators, educators and staff, helping them deliver and execute better. The next challenge will be making these tools part of everyday operations, which is already happening at some forward-looking institutions.

6- European University Association: [Financially sustainable universities: State of play and strategies for future resilience](#).

7- PwC: [Agentic AI—the new frontier in GenAI](#)

Raising AI literacy to support educators, administrators and staff

To support the adoption of agentic AI the benefits of the technology need to be communicated clearly to all staff, including administrators, their staff and educators. All of these groups will also need appropriate training to raise AI literacy. Administrators and their staff may see their roles evolving from scheduling classes and faculty management to overseeing the effectiveness of AI systems and ensuring AI-driven tools support better learning outcomes. Together with educators, they will need support in developing new skills in areas like data analysis, collaboration with IT staff, and responsible AI and security awareness. Managing resistance to change and underlining how the technology will aid administrators' roles and improve working conditions for educators is vital.

The rewards of getting it right are enormous: the most profound benefits of agentic AI may come not from reducing the administrative burden, but from taking advantage of its enormous potential to enhance the educational experience for students and pupils. It can help provide learning that is custom-tailored, offer on-demand support and feedback and even increase learning accessibility and inclusion.

Exploring new frontiers of tailored and personalised learning

A growing number of schools and universities are using AI agents to provide one-on-one tutoring that's tailored to individual students' needs. These agents enable students to access the help they need 24/7—and at scale—by customising the education experience through the use of lesson plans, explanations and assignments that match the preferred learning style and pace of each student. Some universities are even exploring the application of this approach to testing, where individually tailored exams would be created to support personalised learning journeys. AI agents can also offer instant feedback on assignments and assessments, helping students understand their progress and areas for improvement. This immediate feedback loop enables students to make timely adjustments and enhances their learning outcomes.

Enhancing accessibility and inclusion

Ensuring that educational opportunities are equally available to all is another key area that is benefiting from new approaches. Agentic AI supports learning accessibility and inclusion by taking into account neurodiverse students and those with invisible disabilities. AI agents can be deployed to quickly recognise where different learning approaches might be necessary, helping ensure no learner is left behind. Speech-to-text and text-to-speech technologies tools are already providing novel ways for students to interact with content, for example those with hearing or visual impairments. AI-driven translation tools can also benefit international students and non-native speakers, a growing and important segment for higher education institutions. Agentic AI can support inclusion and diversity, which is essential in modern education systems.

Ana Rita Pereira, EMEA Public Sector Verticals Leader, explains how agentic AI can enhance the educational experience:



Intelligent agents can perform tasks autonomously, interact with students and educators, and provide personalised support. They offer unique opportunities to democratise education and implement change at scale, ensuring that all students have access to high-quality learning experiences. By embracing the power of AI agents, educational institutions can harness the full potential of AI to transform education, making it more inclusive, innovative and effective for all learners.

All these capabilities show how agentic AI can revolutionise education by enhancing efficiency, reducing the administrative burden on teachers, supporting personalised learning experiences, and improving accessibility and inclusion for all students.

Defence

AI has introduced significant new capabilities to military organisations around the world in recent years—both in physical spaces and in cyberspace. And agents promise to extend those capabilities even further by enabling AI-driven, real-time data processing, autonomous decision making and more intelligent and dynamic resource management.

The defence sector has already been using machine learning and analytics tools for many applications—from threat detection in cybersecurity to analysis of patterns for intelligence, surveillance and reconnaissance (ISR). AI agents are expected to act as force multipliers, enabling defence teams and organisations to operate many times more efficiently and effectively with their existing resources. Because AI agents can act and react autonomously—most likely faster than a human in in structured environments or predefined scenarios—countries that adopt agentic AI will have operational and strategic advantages.

Steven De Bruyn, Defence Leader, Partner at PwC Belgium, explains some of the benefits of agentic AI for defence:



Agentic AI will play a major role in improving ISR. It will help bring all the different data points together and speed up analysis so that insights from all kinds of sensors and platforms can be interpreted more quickly. Agentic AI will also move the needle in responding to and addressing cyber threats. And it will improve interoperability and coordination across forces in joint operations too.

Improving data analysis and interoperability

With the ability to draw insights from large volumes of diverse data sources, agentic AI can help defence organisations to eliminate data silos. This could provide better analysis—and enable more-informed responses—during fast-changing events and support improved training and preparation through scenario simulations. All of these applications could advance both defensive and offensive capabilities. Pooling data across security alliances boosts situational awareness, faster decision-making and more effective threat detection through the sharing of intelligence. Interoperable AI systems can also allow for coordinated use of

autonomous assets, such as for surveillance units improving the accuracy and timing of any response.

And AI can support interoperability in less direct ways too, by enhancing knowledge exchange and transfer. PwC helped the NATO Standardization Office (NSO) develop an AI-based solution to provide easier access to about 1,400 NATO standardization documents that are used by about seven million professionals in 32 nations across the entire NATO alliance to protect a population of one billion.⁸

By using AI agents to automatically monitor social media feeds and other sources organisations will be able to rapidly identify areas where urgent support—such as aid in evacuations—is needed most. And leaders could potentially make faster, better and more data-driven decisions about deploying personnel, vehicles and other equipment by using agents that track changes in areas such as weather, traffic, flooding conditions, fuel supplies and warehouse inventories. Another potential application is for operational support, where AI agents could help to analyse and update workforce schedules in response to a wide range of changing conditions.

Defending cyberspace

Beyond defence on the ground, in the air and on the seas, nations around the world have become increasingly reliant on networked systems, so they must also defend themselves in cyberspace. AI already plays a vital role in cybersecurity, providing real-time threat detection, automated defences and other capabilities to protect IT systems against ransomware, distributed denial of service attacks, data theft and more.

As part of this, AI agents are rapidly becoming an essential part of the military's cybersecurity arsenal. Agentic AI for such purposes today is broadly focused on simplifying integration between security and service management controls, which helps to manage vulnerabilities and analyse organisational impacts of threats. AI agents can also help to automate security process workflows that previously would have been delivered by humans. And they have an important role in triaging security alerts, enabling overworked IT teams to quickly distinguish real threats from background noise and false positives. Agentic AI is now expected to play an increasingly important role of national cyber defence. Conflicts increasingly take place across multiple domains, with air, sea, land and cyber operations all launched at the same time. This complexity can further tax the ability of human cognition in being able to react to all threats effectively. Agentic AI brings the ability to respond simultaneously across these multiple domains and improve efficiency.

8- PwC: [Facilitating more effective collaboration for the NATO Standardization Office](#)

Strong and secure cloud foundations

Because defence activities often involve multiple nations working together on planning and coordination, it's also essential that any agentic AI applications are supported by highly secure and always-available cloud infrastructures. Private, hybrid, and multi-cloud approaches, such as those enabled by Microsoft Azure, make that possible.

Building on such cloud foundations, defence organisations today continue to explore new ways in which they can use AI and AI agents across a wide range of applications. For example, the European Defence Fund recently launched a 42-month Artificial Intelligence Deployable Agent (AIDA)⁹ programme with an agentic AI focus. The project seeks to develop autonomous and semi-autonomous agents for defence against cyberattacks.

Gus MacGregor-Millar, GM, Worldwide Defence & Intelligence at Microsoft, describes how the commercial cloud can help militaries securely share data:



Data sharing is an important consideration and if I am a military commander working with other nations, I want to do it securely. I want to operate in a collaborative working environment that meets anyone's standards, whether it's in Australia or Canada, an EU country or the UK. I can meet those requirements today in the commercial cloud. I can provide software-defined security within the cloud. And that allows me to securely share that data."


Agentic AI is helping defence organisations get more value from data and share it securely to improve offensive and defensive capabilities, enhance situational awareness and decision-making, improve cybersecurity, and manage conflicts across multiple domains more easily. These capabilities will be vital to keeping populations safe and secure.

9- European Commission: [Artificial Intelligence Deployable Agent](#)

03



Building a solid, responsible foundation for agentic AI



Before they can reap the benefits of agentic AI, public sector organisations must be sure to build on a solid foundation that supports their responsibility of care for constituents. That starts with ensuring that their IT infrastructure is fit for task. Organisations will also need to prioritise data privacy and security issues and fully address ethical concerns to build public trust. Finally, they will need to manage the resulting change in their organisations and make sure that their people are well-equipped to make the most of the new possibilities provided by these tools.

Cloud platforms as the backbone

Adopting and deploying AI agents requires two essential elements: a modern cloud infrastructure and strong cybersecurity.

The cloud, for example Microsoft Azure, provides the scalability and resources that any public sector organisation needs to use agentic AI, whether for internal operations, citizen services or other tasks. To build the right cloud strategy and infrastructure, every organisation must assess where it currently is on its cloud journey, understand what it wants to achieve with agentic AI and determine which cloud model will deliver on those goals.¹⁰

Achieving this begins with a review of the organisation's infrastructure and then determining how to ensure regulatory compliance, including data residency and policy alignment. Good data governance also involves practices such as updating retention policies and assessing third-party connectors to ensure that data flows are transparent, controlled, and auditable. Meeting these goals is critical for public accountability.

10- PwC and Microsoft: [How to deploy AI at scale--A PwC and Microsoft playbook that explores the critical role of cloud and cybersecurity](#)

Data quality and security

Using agentic AI can deliver significant benefits to the public sector, but it's also vital for organisations to take special care that they are deployed as securely as possible. Data must be protected per organisational policies, whether it is stored on-site or in the cloud. Organisations also need to track data storage, movement, and interactions. Public sector organisations must prioritise security, privacy, and data integrity, whether they handle data of dozens, hundreds, tens of millions or hundreds of millions of citizens. Data breaches in areas like public health, infrastructure, education, or defence have the potential to cause real harm, so extreme care is a must at all levels.

Usage of AI by public sector organisations in the European Union is generally classified as high-risk under the terms of EU AI Act. High-risk AI systems will be subject to detailed requirements to ensure data quality, security, and transparency once the Act goes into effect in August 2027. Public sector organisations need to ensure that their use of AI can comply with these strict requirements and with existing legislation on data privacy. Audit reviews and access remediation can help ensure that citizen data is well-protected.

Understanding which aspects of data protection are their responsibility and which will be handled by cloud service providers is another key task for public sector organisations looking to implement agentic AI. Under a Shared Responsibility Model, organisations using agentic AI will need to implement their own cybersecurity controls and processes to protect their data and other assets, while a cloud service provider such as Microsoft will generally oversee the safety of the underlying infrastructure. Responsibilities vary between applications deployed as infrastructure as a service (IaaS) and those deployed as software as a service (SaaS).



Frederik Blachetta, Global Public Sector Data and AI leader, Partner at PwC Germany, highlights the need for a strong data foundation:



Organisations aiming to implement advanced AI systems—including agentic AI—must recognise that success hinges on early and sustained investment in robust data foundations. This includes ensuring data quality, secure and lawful access, interoperability, and organisational readiness. Based on our experience, three out of four public-sector use cases are delayed due to missing or inadequate data. Without proactive leadership, many AI initiatives risk stalling before delivering meaningful impact.”

Ethical concerns and building public trust

Responsible AI helps confirm that AI systems are developed and deployed in ways that uphold transparency, fairness, and trust, fostering confidence among stakeholders while mitigating potential harm.

Public sector organisations have a special obligation to implement agentic AI and other AI tools responsibly.

The responsible implementation and usage of AI systems require careful identification of potential risks, which may include bias in decision-making algorithms, questions of fairness, and broader ethical dilemmas. This involves evaluating whether the AI system aligns with societal values and norms and confirming that the models are free from discriminatory practices or unintended consequences that could marginalise certain groups.

Data transparency is vital, both to comply with the EU AI Act when it comes into force and to confirm that the actions taken by agents adhere to these ethical considerations. Public sector organisations must ensure that the origins, handling, and processing of data are documented and accessible. Data sources and metadata summaries need to be clearly attributed, which can help clarify the rationale behind AI-generated outputs or decisions.

Public sector organisations have real and significant impacts on people’s lives. Having a human in the loop ensures that any AI outputs are verified and validated

by real people before they are used to take any consequential actions. AI models remain prone to generating false or misleading outputs—a result commonly called hallucination. Proper oversight by real people is key.

Shadow AI, which often refers to the use of unauthorised or insecure consumer-based AI applications within an organisation, can present an additional risk. Such tools can bypass established governance frameworks and lead to significant security vulnerabilities or unethical practices. To mitigate these risks, it is essential to implement robust ethical guidelines and promote a culture of accountability.

To ensure the solutions they deploy work as intended, public sector organisations will need to evaluate their AI systems regularly. They will also need thoroughly review enterprise policies and standards, such as Acceptable Use and Information Security policies, to be sure that they address emerging risks associated with agentic AI and other AI tools and that comprehensive measures are in place for their safe and responsible use. Regular policy reviews help maintain compliance with regulatory frameworks and reinforce a culture of accountability within the organisation.

Supporting digital public infrastructure around the world

Public sector organisations around the world are at different stages in their digital journeys. All of them will need to be able to rely on secure digital public infrastructure (DPI) as the foundation for agentic AI systems. To help foster international cooperation, the United Nations (UN) has launched several programmes, including the DPI Safeguards initiative¹¹. This multi-stakeholder effort is facilitating the creation of standards and knowledge exchange across the public and private sectors and civil society.

The UN is also focused on efforts to promote open source and open data, which will be critical in ensuring access to agentic AI systems for everyone, not just the world's richest countries. Open source, by enabling the reuse of solutions, accelerates the broad availability of applications, helps prevent redundant development, and allows new trends to be adopted more quickly. It can foster interoperability, transparency, and new forms of collaboration.¹² The open source programme at Microsoft supports the integration of open source solutions alongside custom-developed tools.¹³ That enables maximum flexibility for public (and private) customers around the world.

11- [United Nations Office for Digital and Emerging Technologies \(ODET\)](#)

12- [United Nations Office for Digital and Emerging Technologies \(ODET\), United Nations Development Programme \(UNDP\)](#)

13- Microsoft: [About Microsoft's Open Source Program](#)

Open-source AI frameworks can also be integrated into digital infrastructure to support digital sovereignty efforts, a strategic priority for many governments that are looking to retain control over their digital infrastructure. The support from Microsoft for open standards and sovereign cloud offerings, such as Azure Confidential Computing¹⁴ and Microsoft Cloud for Sovereignty¹⁵, align with these goals.

People are key to the success of agentic AI

While agents have enormous potential to streamline and automate tasks, implementing agentic AI depends first and foremost on the people in your organisation.

A recent PwC survey found that

75%

of executives agree or strongly agree that AI agents will reshape the workplace more than the internet did.¹⁶

That reflects enormous interest in agentic AI, but there's still a long way to go to reach this level of impact. Public sector organisations need to develop a comprehensive communication plan along with detailed change management materials to ensure smooth adoption of agentic AI among end users. This process involves educating users about the capabilities, limitations, and ethical considerations of AI systems. By fostering clarity and understanding, organisations can build confidence in the technology while promoting responsible usage and minimising resistance to change.

14- Microsoft: [Azure confidential computing](#)

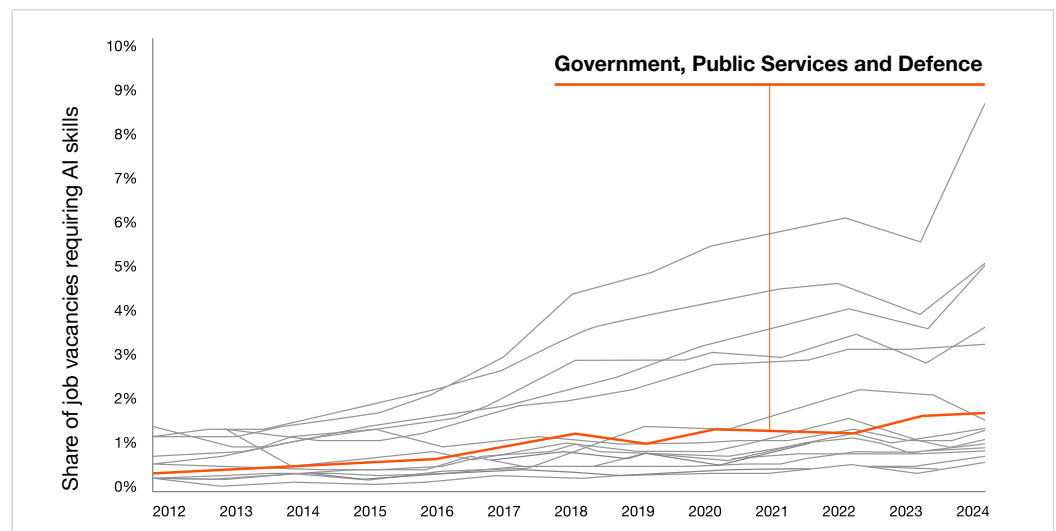
15- Microsoft: [Microsoft Sovereign Cloud](#)

16- PwC: [PwC's AI Agent Survey](#)

PwC research highlights some of the challenges public sector organisations may face. In our 2024 Global Hopes and Fears survey¹⁶, government and public sector respondents showed a cautious approach compared to their private sector counterparts regarding job changes and the adoption of new skills and tools. Embracing skills development and fostering excitement and optimism about its potential will be absolutely critical to make the most of agentic AI. Strong change management and learning programmes are a must.

This process won't be “one and done”—ongoing training and education to encourage a mindset of continuous improvement and innovation will be vital. Public sector organisations must build a culture of trust and responsible innovation and continually embed safe AI practices across their workforce. The AI Skills Initiative from Microsoft, including the AI Learning Hub and certifications through Microsoft Learn¹⁷, offers public sector organisations a scalable way to upskill their workforce. These resources can be integrated into change management plans to build AI literacy and reduce resistance to adoption.

Public sector organisations will also need to attract new workers with the right skills to fill positions as they open. PwC's recent AI Jobs Barometer¹⁸ shows a steady increase in recent years in public sector job postings explicitly requiring AI skills.



Source: PwC: AI Jobs Barometer | Government, Public Services and Defence

The results also show that workers with AI skills are able to command significantly higher wages in many industry sectors. This “wage premium” is much lower in the public sector however, which may make it more difficult to recruit urgently needed talent going forward.

16- PwC: [Global Hopes and Fears 2024](#), Government and public services: Embrace digitisation and transformation—Randa Bahsoun

17- Microsoft: [Microsoft Learn—Spark Possibilities](#)

18- PwC: [2024 AI Jobs Barometer](#) | Government, Public Services and Defence

04



Moving towards implementation and driving business value: a strategic roadmap

Successfully deploying agentic AI in the public sector requires a structured, value-driven approach. PwC's Agentic AI Value Chain provides a comprehensive framework that guides organisations through three critical stages—strategy and discovery, design and deploy, and scale and collaborate—which are aimed at ensuring readiness, delivering measurable outcomes, and achieving sustainable impact.

Strategy and discovery

This stage aligns AI goals with organisational objectives by assessing AI readiness, diagnosing current maturity levels, and identifying priorities. Organisations should evaluate their data landscape, regulatory environment, and internal capabilities as the first steps towards AI adoption. They'll also need to consider the ethical issues noted in the previous section to make sure that trust is designed into systems from the start.

Key activities in this phase include conducting an AI readiness assessment to ensure alignment of AI strategy with mission objectives and evaluate organisational maturity. Opportunity mapping assists in identifying high-impact use cases through business process and qualitative analysis, while workforce transformation planning defines training needs and prepares teams for effective human-AI collaboration. PwC supports this phase with accelerators such as the agentic Workflow Assessment, enabling public sector entities to prioritise initiatives, define key performance indicators (KPIs), and develop a compelling case for transformation.

Design and deploy

At this stage, organisations implement their AI strategy by designing AI agents for specific workflows that deliver value. This includes creating new operating models, defining roles, and equipping teams with AI skills. Key activities involve developing AI agents to automate, assist, and manage processes, establishing KPI frameworks, and validating models for reliable and responsible performance. PwC focuses on deploying AI agents with strong governance, security, and compliance controls to minimise risks and increase trust.



Scale and collaborate

With initial deployments in place, organisations can scale agentic AI across departments. This phase embeds AI agents into core business functions, supported by governance and continuous improvement. Key activities include expanding AI adoption through cross-line-of-business integration and establishing centres of excellence with monitoring systems for sustained performance and compliance. PwC provides ongoing support through managed services, model fine-tuning, and strategic co-sell initiatives, enabling public sector organisations to scale responsibly and collaboratively.

Working through these three phases will help public sector organisations move from great ideas to sustained outcomes.

Strategy and discovery

Strategic discovery and readiness



Diagnose maturity and align strategic priorities

Anchor client strategy to outcome-based transformation pillars leveraging the Intelligent Enterprise model

Business process qualification and opportunity mapping



From concept to business justification

Use accelerators like the Agentic Workflow Assessment to identify and prioritise high-impact processes, tied to business outcomes and KPIs

Design and deploy

Workforce transformation and operating model redesign



Lay the foundational for human + AI teaming model

Define new roles, skills, and team structures to support human-agent collaboration and organisational enablement

Agent design, development and deployment



Build and deploy agents that create real business value

Design, build and deploy agents across defined business areas to automate, augment, and orchestrate targeted workflows

Scale and collaborate

Operating and KPI monitoring model



Operationalise and govern agents at scale

Set up Managed Services and a Center of Excellence to monitor, govern, and continuously improve AI agent ecosystems

Cross business areas



Embed AI agents into critical business domains


Drive scaled transformation with intelligent automation - unlocking value across all business areas (Finance, HR, Operations, Legal, Sales, etc.)



05



The importance of partnerships



For organisations wanting to successfully develop secure AI solutions, it is imperative to bring in a strong technology partner with the right capabilities, industry knowledge and functional expertise. It is equally essential to tap into proven industry expertise and knowledge to overcome barriers to implementation. This helps to smooth the journey to AI and jump-start its adoption.

PwC's strategic collaboration with Microsoft¹⁹ makes it easier for organisations to take full advantage of AI agents. Our global alliance helps organisations to use AI agents as critical assets across the business to support innovation and increase business efficiencies. Together, we support industries in all sectors at all stages of the AI journey, from facilitating cloud migrations to offering tailored AI use cases to full AI agent adoption. This partnership builds on a shared commitment to making use of cutting-edge technology responsibly, to facilitate that AI solutions are not only innovative but also secure, transparent and accountable.

PwC's Agentic AI Value Chain, in collaboration with technology solutions like Microsoft Copilot Studio, focuses on empowering public sector organisations to automate workflows, enhance productivity, support intelligent decision-making, and transform their operations and citizen services. By taking full advantage of scalable agentic AI solutions, the partnership ensures seamless integration across public sector functions and platforms. Built from the ground up, these solutions incorporate responsible AI practices, robust security, effective data management, transparency and accountability—all essential elements for successfully harnessing AI in government, defence, and education. Together, we help public sector organisations to stay at the forefront of technological innovation and serve constituents efficiently.

19- PwC: [PwC and Microsoft announce strategic collaboration to transform industries with AI agents](#)

And both PwC and Microsoft have demonstrated these capabilities first-hand by successfully deploying these solutions in their own organisations. As one of the largest Microsoft 365 Copilot customers globally, PwC has Copilot licenses enabled and deployed to more than 40 countries. This achievement is underlined by our recognition as the Microsoft 2024 Global Partner of the Year for Building with AI.²⁰

Mauro Xavier, EMEA Microsoft Alliance Leader and Partner at PwC Spain explains how collaboration drives successful use of agentic AI:



Successfully deploying AI agents also requires collaboration with customers who recognise the technology's potential. To truly unlock the transformative power of agentic AI, you need the best—a cutting-edge technology partner like Microsoft, one of the best advisory firms with the ability to execute like PwC, and a visionary customer with a real business problem and ready to embrace change. When these three elements converge, they don't just work—they create groundbreaking innovations that reshape organisations and redefine possibilities. That is where the magic happens.”



20- PwC: [PwC and Microsoft deliver award-winning advisory services](#)

Conclusion

Demands on public sector organisations are changing rapidly. Now is the time for these organisations to lead with a vision. By investing in agentic AI today, they can build more resilient, citizen-centric institutions and transform themselves into intelligent enterprises that are equipped to meet the challenges of tomorrow.

Agentic AI will enable them to provide faster and more responsive services to citizens and businesses, innovate and gain unprecedented capabilities to act on data in real time—helping them to build public confidence and trust. The potential benefits to organisations, employees, citizens, students, businesses, communities and critical infrastructure are enormous.

- Governments can decrease red tape and deliver services to citizens and businesses more effectively.
- Educational institutions can achieve better learning outcomes and make education more accessible to all.
- Defence organisations can improve offensive and defensive capabilities and enhance cybersecurity to keep populations safer.

Using agentic AI, governments, educational institutions, defence organisations and more will be able to gain—and act on—insights from large volumes of complex, and fast-changing data in near-real time. Bureaucratic obstacles can be reduced, roadblocks caused by data siloes eliminated, and faster and more responsive services tailored to individual needs provided.

But successful deployment requires the right foundation, both in technology and in practices. PwC and Microsoft are prepared to assist public sector organisations in harnessing the significant benefits that agentic AI and AI can offer and let the magic begin. Get in touch with us to start piloting your most relevant use cases.

Contacts

Microsoft



Didier Ongena
VP, Global Government Leader,
Worldwide Public Sector,
Microsoft



Angus MacGregor-Millar
GM,
Worldwide Defense & Intelligence,
Microsoft



Ana Rita Pereira
EMEA Public Sector Verticals Leader,
Microsoft



Yasmine Rifai
Senior Strategic Account Technology
Leader in Public Sector,
Microsoft

PwC



Mauro Xavier
EMEA Microsoft Alliance Leader,
Partner,
PwC Spain



Agnieszka Gajewska
Global Government and Public Services
Leader,
Partner,
PwC Poland



Frederik Blachetta
Global Public Sector Data & AI Leader,
Partner,
PwC Germany



Steven De Bruyn
Defence Leader Belgium,
Partner,
PwC Belgium



Caitroina McCusker
Partner,
PwC UK

