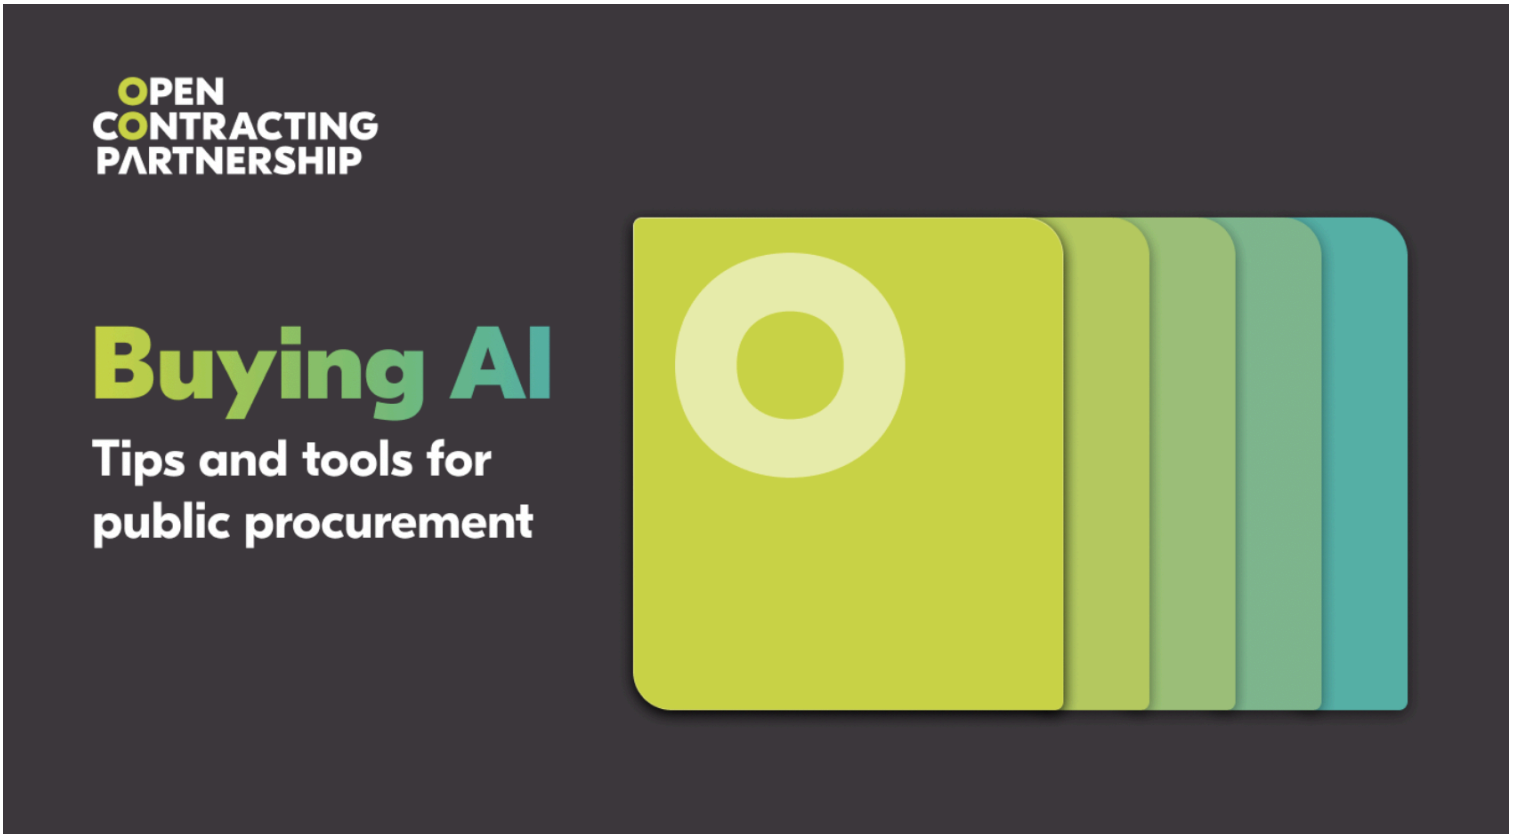


This guidance offers practical and hands-on tips to get the best results from your AI purchases.

Developed with public sector procurement, project, and technology teams, it builds on the insights of over 50 public sector practitioners and experts.

The guidance is designed to help you no matter where you are on your AI procurement journey. Each piece of content can stand alone. You are welcome to review this guidance from top to bottom, or jump around.

We provide actionable, plain-language information to help you get started with what's essential when buying AI. We hope this resource will help you filter out what you really need to know from the noise and provide procurement, project, and technology teams with a shared language and vision for working together!

Get started.

About

Open Contracting Partnership is a nonprofit organization that supports public sector organizations around the world to get better results from their procurements through leveraging data and technology, cross-sector collaboration, and goal-driven systems change.

Contact us if you would like to discuss your project. We'd love to help.

Table of Contents

1. Introduction
    1. Introduction and Insights on AI adoption
    2. Why we wrote this guidance
    3. Who this guidance is for
    4. How to use this guidance
2. AI readiness
    1. AI basics
    2. AI readiness in the public sector
    3. Your role as a procurement officer
    4. AI adoption myth-busting
3. How to buy AI
    1. Ways to buy AI
    2. Hacks for better AI procurement
    3. Planning and procurement strategy tips
4. Tips and tools
    1. Who to involve
    2. Solicitation tips
    3. Evaluation tips
    4. Contract clauses
    5. Contract management tips, roles, and responsibilities
    6. Key questions to ask
5. Resources
    1. Case studies
    2. Technical annex: RFP questions and evaluation

Introduction: *Insights on AI adoption*

Emerging technologies like GenAI and agentic AI have the potential to help public sector organizations increase efficiency through automation and strengthen decision-making, helping stretch tight human resources to better serve the public.

Procurement is the moment where AI's potential becomes practical reality. This is the strategic process by which the organization identifies their needs, what they will buy, who they will partner with, and how the solution will be implemented to deliver maximum value for the public.

Done well, the procurement process can help the public sector secure the best solution to meet their mission as quickly as possible while mitigating risks, and strengthen productivity to better achieve their policy goals. Done poorly, the procured AI solution might be underused at best, or fail to achieve outcomes at worst, resulting in expensive failures that weaken trust in government.

To make the most of AI technology, everyone involved in the procurement process must be more nimble and collaborative than ever before. AI requires specific considerations to manage the unique risks and opportunities that come from the technology itself.

Leadership must step up to create AI governance policies and support agile procurement practices. Procurement officers need a basic understanding of how AI works, what they should pay special attention to, and whom to involve. Project and technology teams need to understand how the procurement process works, and what to consider as they collaborate together with procurement organizations. Everyone needs to be aware of what questions to ask, and when.

Insights on AI adoption in the public sector

During our research process, we identified several key insights into the current state of AI adoption that informed the design of this guidance, including:

- **Hype vs reality**: While building or customizing AI tech for cutting-edge use cases is getting a lot of hype, the public sector is currently mostly focused on buying off-the-shelf licenses for productivity solutions. Therefore, we highlight specific considerations for off-the-shelf AI purchases. Some vendors also make exaggerated claims around AI. We explain how having the right people on your team and asking smart questions can help mitigate this risk.
- **Evolving attitudes towards risk**: Most organizations are taking a very risk-adverse approach to AI, but early movers are now starting to roll back burdensome processes to enable easier access to low-risk technologies. We advocate that organizations take a strategic and risk-based approach to the procurement of AI technologies.
- **AI adoption ≠ AI procurement**: Procurement plays an important and under-appreciated role in AI adoption, but its role is limited with the trend towards centralized buying of off-the-shelf licenses and integration of AI into existing products. We provide recommendations to help organizations embed procurement in their broader AI adoption strategy.

Why we wrote this guidance

Open Contracting Partnership is a nonprofit organization that supports public sector organizations around the world to get better results from their procurements through leveraging data, cross-sector collaboration, and goal-driven systems change. Given our technology expertise, our partners are now asking us questions around buying AI and looking for more resources on how to get started.

We set out to create this guidance with and for procurement directors and officers, as well as for project teams from agencies or departments that want to buy and use AI technology. Through our one-on-one conversations and group co-creation sessions with over fifty public practitioners and experts, we identified three critical challenges for successful AI procurement:

1. Procurement teams want help understanding, in plain language, what to pay attention to when buying AI technology.
2. Procurement, project, and IT/data teams seek to better collaborate and work in more agile ways to get the best products and outcomes.
3. Public sector organizations must ready their IT, data, and project teams to adopt AI technology.

This guidance seeks to address these challenges by providing actionable and plain language information to help public sector procurement officials get started with the essential information they need to know when buying AI, as well as give procurement, project, and technology teams a shared language and vision for working together. We also highlight some of the key measures that leadership should take to help make their organizations ready to buy and adopt AI technology.

Who this guidance is for

This guidance was created for **procurement officers**, who are accompanying the purchasing process and want to understand AI technology, how procuring AI is different, and what to pay attention to during the procurement process.

In addition, this guidance can also support:

- **Project teams,** from departments or agencies who want to buy AI technology and are coordinating the overall process from needs identification through implementation, and who want to understand how to get the best results from their purchase.
- **IT and data teams**, who are collaborating with procurement and project teams on AI procurements.
- **Public sector leaders,** who oversee supporting organizational readiness for AI.

How to use this guidance

This guidance can help you no matter where you are on your AI procurement journey. For example, you may refer to this guidance when you're:

- Helping buy a technology solution with AI features
- Developing your procurement strategy for a specific AI solution
- Upskilling yourself on procuring AI

Each piece of content in this guidance is designed to stand alone. This means that you are welcome to review this guidance from top to bottom, or jump around. We assume that you are familiar with the procurement process and have some knowledge of good IT procurement practices.

Several sections also feature key questions to ask, which are summarized in the **Key questions section** at the end of this guidance. You will note that these questions include both aspects for procurement and non-procurement staff to consider. This is because successful AI procurement is a team effort, and understanding generally what questions need to be asked, and when, can help you facilitate internal conversations and ensure that the right experts are involved in decision-making.

This guidance is for consulting alongside your existing government policies, and frameworks. This is not legal advice.

Applying these tools in your organization

If you are a member of leadership, we will be creating additional resources to help you deploy this guidance in your organization.

- **Buying AI training webinar:** This video training walks through the material.
- **Slides:** You are welcome to reuse our guidance in slide format with attribution.
- **Agenda template:** This customizable in-person agenda template outlines key information for an AI procurement staff training.

Acknowledgements

This guidance was made possible through the support of AWS and by the contributions of over fifty experts and practitioners from the public sector, private sector, and academia. Our deepest gratitude to you all for your personal insights, including: Albert Sanchez Graells, André Petheram, Carl Hahn, Chris Fechner, Daniel Arosa Otero, Daniel Chenok, Dmytro Solovei, Isaac Martín Delgado, Jessica Tillipman, Keasha L. Brown, Kęstutis Kazulis, LaTalya Palmer, Mariel Reed, Marion Eigner, Matt Sanderson, Milena Krogsgaard, Nadia Sava, Quinn Anex-Ries, Rebecca Piazza, Ryan Maher, Sabine Gerdon, Savina Kalanj, Waldo Jaquith, Warren Smith.

AI basics

Who: Procurement teams
What: Overview of GenAI and agentic AI technologies, common use cases, and limitations and risks

What is AI

Artificial intelligence, or AI, is a term that describes computer systems designed to perform complex human-like tasks, such as perception, communication, analysis, and content creation.

An overview of different forms of AI include:

- **Machine learning,** which maps input data (often structured or numeric) to outputs such as predictions or classifications.
- **Deep learning**, a branch of machine learning, uses multilayer neural networks to learn from complex, unstructured inputs like images, text, or video, and can produce outputs ranging from simple labels to complex generated content.
- **Generative AI** (GenAI) uses deep learning models to generate new content, such as text, images, audio, or code, by predicting what is most probable based on patterns learned from training data. Large language learning models (LLMs) are a form of GenAI.
- **Agentic AI** builds on GenAI technology like LLMs, as well as connections with other databases and tools through forms of robotic process automation (RPA), to independently perform complex tasks based on new and evolving inputs. For example, summarizing email text and then proactively scheduling calls.

This guidance is focused on the latest evolutions of AI: GenAI and Agentic.

Defining artificial intelligence

There is no one single standard way of defining artificial intelligence and its various forms. For further information on how entities define AI, check out:

- The definition provided in the John S. McCain National Defense Authorization Act for Fiscal Year 2019, a commonly referenced US federal definition
- The EU Artificial Intelligence Act: General Provisions: Definitions
- The International Organization for Standardization's What Is Artificial Intelligence?

What is GenAI

Generative AI (GenAI) is a type of AI that can "generate" new content, and passively responds to queries based on recognizing and reproducing patterns from its training data. The technology has recently evolved from simple image and text generation to multi-modal models that can generate content which can be challenging to distinguish from content created by humans. GenAI use cases range from text summarization, to question answering, to digital art creation, to code generation and beyond.

What is Agentic AI

Agentic AI is an emerging technology that builds on GenAI capabilities to execute multi-step workflows. An agentic AI system can act independently to achieve directives. "Agentic" indicates the ability to act independently. Unlike other forms of AI that require prompting and step-by-step guidance, agentic AI can proactively and autonomously perform complex tasks.

AI agents can communicate with each other and other software systems to automate existing processes, as well as make independent decisions. They understand their environments and available tools, and can adapt to changing conditions, enabling them to perform a variety of sophisticated workflows. Agentic AI typically builds upon multiple hyper-specialized agents, with each focused on a narrow area of tasks. These AI-powered agents can coordinate with each other, sharing information and handing off tasks as needed. For example, agentic AI can help a developer not only write code, but automatically test and debug it, or send emails and then automatically schedule meetings based on the contents of the reply.

GenAI Stack

GenAI technology encompasses a lot of different functions. The set of technologies that work together to perform these functionalities is known as the GenAI technology stack. This stack can be thought of as three layers that build on each other and work together: the foundation layer, middle layer, and application layer.

Foundation layer: the model infrastructure

At the bottom of the stack are the foundation models (like GPT-5, Claude Sonnet, or Gemini Pro) and the infrastructure needed to run them. GenAI is powered by foundation models. Foundation models are sets of trained neural networks. They are pre-trained on tremendously large sets of data. Often these models were created based on publicly available data, such as large portions of the internet, books, and other text sources. GenAI can be customized for specific needs through a process called fine-tuning, which retrains the model on more specific data.

Middle layer: integration services

The middle layer consists of solutions and services that make the foundation models usable for specific applications. This includes the application programming interface (API) services that provide access to foundation models, development tools for customizing AI capabilities, and security and governance frameworks, as well as integration services that connect AI with existing systems.

Application layer: solutions

The top layer consists of the actual applications that people interact with. These might be AI-powered document processing systems, virtual assistants for citizen services, content generation solutions for organizational communications, or specialized analytical applications.

To understand the GenAI stack, it can be helpful to think of it like an airplane. The **foundation layer** is like the engine. Just as an engine is needed to get the airplane off the ground, the foundation layer is essential for generative capacity. The **middle layer** is like the body of the aircraft. It is the recipient of the engine's power, and is what makes the engine usable for flight. It's also customizable to meet your purposes. For GenAI, this is the layer that takes the model's power, customizes it, and makes it usable. The **application layer** is like the cockpit. Just like a pilot uses instruments and controls to direct the airplane, a user interacts with applications to use the AI model.

To take the airplane example further, **agentic AI** could be considered the autopilot: it's technology that is intended to be fully capable of flying the airplane, but you'd still want human oversight to deal with emergencies and make sure all goes well.

All the technologies in the GenAI tech stack are cloud-based. While not the focus here, this means that there are also additional hosting infrastructure considerations when buying AI.

When to consider GenAI and agentic agents

Currently, GenAI technology is often compared to an intern at work. An intern can help you do your job well and make a task more efficient. However, you still wouldn't want even the most brilliant intern making high-stakes decisions, or running crucial operations without oversight. You would also not ask an intern to carry out a task you are unable to perform yourself or, at least, you do not know what "good looks like" in relation to that task.

GenAI is great at producing results that are most *likely or plausible*. However, it does not know what is *true or correct* (it does not "know" anything). This is an important distinction. GenAI can also struggle with newer and more complex situations. Together, these factors mean that the outputs from GenAI should be checked by a human.

If GenAI is like an intern, agentic AI is more like a junior colleague to whom you have made a directive. It works independently from start to finish: from receiving a request, to making decisions and implementing solutions based on its own interpretation and analysis of the context and initial input. Since agentic AI is able to leverage LLMs with other AI agents that specialize in tasks beyond text interpretation and generation, agentic AI is capable of

dealing with more complex problems. However, agentic AI otherwise faces the same limitations as GenAI, and because it is more powerful, when things go wrong it can also be more risky. Multiple models are interconnected through agentic AI, so an inaccurate output from one of them could feed into another model that processes it inaccurately, and so on, leading to a final response that is impacted by multiple stages of errors.

Consider using GenAI and agentic AI for:

- Repetitive tasks, such as low-risk administrative activities
- Creating an initial synthesis of a lot of different types of data
- Generating first drafts or options for human review
- IT modernization and development

Avoid using GenAI and agentic AI for:

- Higher-stakes decision-making
- Decisions that require transparency and explainability
- Low-quality data

GenAI and agentic agent use cases

GenAI and agentic AI are new technologies, and the public and private sectors are learning where this AI adds value for their organizations, and how to make the most of it. This is especially true for agentic AI.

Five use cases for the public sector include:

- **Document intelligence:**
  GenAI can be programmed to automatically read, summarize, and extract key information from thousands of pages in minutes – work that could take staff weeks to complete. This means that GenAI has the potential to transform how agencies handle documents, from complex regulations to registration forms. Ultimately, this may result in faster service delivery, better compliance, and significant operational savings. For example, document intelligence for the procurement function could look like having a “super-paralegal” who instantly flips through thousands of contracts, regulations, or vendor proposals and highlights the exact clauses you need.

- **Supporting service delivery:**
  GenAI powers intelligent assistants, such as chatbots for internal or external use, which can handle thousands of inquiries simultaneously, providing information and processing routine requests automatically. These systems can integrate with existing service channels, while potentially reducing wait times and freeing up staff for complex cases. The technology can support consistent service levels in some domain areas while scaling automatically to meet demand peaks, with the potential for improving satisfaction while reducing operational costs. Imagine you're running a call center during the scheduled enrollment period for a public health insurance program. Normally, you'd need to bring in dozens of temporary staff to handle the flood of calls. GenAI is like having an infinitely scalable team of front-desk clerks who answer routine questions and process simple requests automatically. Agentic AI has the potential to build on these capabilities by developing more effective strategies to assist customers. For example, an AI agent may, after several attempts to solve a resident's issue, proceed to contact a human support agent and assign them to the case.
- **Transcription and translation:**
  GenAI can create real-time text documentation of meetings or calls, as well as provide rapid translations. This functionality can help improve internal efficiency. For example, social workers can spend less time taking notes and more time on working directly with their clients, and public-facing documents can be quickly translated into different languages. Live simultaneous translation can also enable more community members to better participate in public meetings.
- **IT modernization and development:**
  GenAI can assist developers by automating routine coding tasks, suggesting improvements, and helping modernize legacy systems. As part of this, GenAI can help technology ensure compliance with security requirements while accelerating digital transformation initiatives. Likewise, agentic AI can help with code modernization by leveraging machine learning, neural networks, LLMs, and automated reasoning, as well as potentially help with faster responses when technology breaks by expediting the incident response process. For example, agentic AI could automate the entire incident response pathway, rolling back issues, creating incident reports, and notifying any team members who need to stay informed.
- **Data analysis and insights:**
  Multi-modal GenAI can process and analyze vast amounts of information from multiple sources – including reports, sensors, and operational data – to identify

patterns and generate actionable insights that improve service delivery. These capabilities can help agencies make data-driven decisions, optimize resource allocation, and proactively address emerging challenges. This can reduce the manual effort required for complex analysis while enabling more responsive and effective public sector services. For example, transportation agencies use multi-modal GenAI to analyze traffic data and optimize signal timing, reducing commute times for schoolchildren and improving road safety for residents. Picture a highway patrol team trying to monitor every camera, road sensor, and traffic report manually – it would be impossible to catch everything. GenAI is like an always-on traffic controller that watches all feeds at once, spots patterns, and alerts you before a crash or bottleneck happens.

More on public sector GenAI use cases

- The State of California in the US offers a table with use cases
- McKinsey offers a global overview of use cases organized by archetype
- Danks Industri published a white paper on responsible use of AI assistants in the public and private sectors.
- Stanford University Human-Centered Artificial Intelligence has issued a brief on validating claims about AI for policymakers.

GenAI and agentic AI limitations and risks

GenAI and agentic AI holds strong potential to improve organizational efficiency and improve service delivery, but with these opportunities comes risk. It's critical to understand the limitations of these systems so that your organization can use them strategically, responsibly, and safely.

Common limitations include:

- **GenAI and agentic AI lack a genuine understanding of the world, and hallucinate.**
  While AI models can generate sentences, these are simply outputs – the LLM doesn't have a true human-like understanding of what the sentence means. Relatedly,

GenAI and agentic AI hallucinate, in other words, give you inaccurate results or information that sounds authoritative but is false. Hallucination rates vary by model, but no model is hallucination-free. Examples of hallucinations will be different depending on your use case, but could look like generating a realistic-sounding policy citation that doesn't actually exist, providing inaccurate or illegal advice to business owners through a chatbot, or wrongly denying access to benefits. AI's tendency to hallucinate makes it essential to have a human-in-the-loop to review and think critically about the results, especially for decision-making that directly impacts service delivery.

Agentic AI hallucinations can be particularly worrisome since they have the potential to impact workflows. For example, if the model generates false information and then relays it to the rest of the AI agents, incorrect data can rapidly spread, escalating errors in the final output.

For the public sector, both GenAI and agentic AI hallucinations could have severe real-world implications. Organizations must be confident in their use case and solution before using it extensively.

- **GenAI and agentic AI results depend on their training data.**
  GenAI models are only as good as the data they are trained on. For example, if your model is only trained on data up to October 2024, it would be unable to tell you current regulations. This is why understanding the training data sources matters a lot, as well as understanding your own data quality if you will be using it to further refine the model.

  Given their reliance on training data, GenAI and agentic AI may pose data security and privacy risks. When you input data into AI systems, that information may be stored, processed, or even used to train future models unless specific protections are in place. This is particularly important for government data that includes personally identifiable information or sensitive operational details. Organizations should ensure their AI solutions include appropriate data-handling safeguards and comply with relevant privacy regulations.

- **GenAI and agentic AI have a limited ability to explain its results.**
  GenAI solutions cannot explain in perfect detail how they arrived at their conclusions. This lack of explainability can be fine in some situations, but not ok in

others where explainability is critical and a matter of legal risk. For example, an AI chatbot that provides assistance with navigating an unemployment application can be ok, while using an AI system to make decisions around qualifying for public benefits could expose governments to legal concerns given an individual's right to due process.

Likewise, agentic AI produces results that can be hard to trace or reproduce. Since agentic AI works independently and with minimal human intervention, this can make testing, debugging, and determining where an AI model has gone wrong a challenge.

- **Agentic AI requires very strong system design.**
  The process of building a multi-agent architecture that effectively coordinates with other models, has specific information of how to carry out certain tasks, and can achieve high-level complex directives is a challenge. Moreover, agentic AI is a new area of technology that relies on successful deployment of other AI strategies. Given this complexity, public sector organizations may struggle to use agentic AI effectively even in low-risk situations where agentic AI could be a strong fit.

  For more information on limitations and how to mitigate GenAI and agentic risks during your procurement process, check out our sections on *Myth-busting* and *Key questions to ask.*

- **GenAI and agentic AI requires ongoing support, which has cost implications.**
  When implemented thoughtfully, investments in AI can deliver strong returns through improved efficiency and service quality. However, the limitations we explore here have real cost implications for implementation. The need for human oversight, data quality improvements, and ongoing monitoring means AI adoption involves more than just technology costs. Organizations should budget for additional staff time to review AI outputs, potential data preparation work, and ongoing performance monitoring.

AI readiness in the public sector

Who: Public sector leaders and procurement teams
What: How to support AI readiness in the public sector

AI readiness in the public sector

Responsible and effective use of AI in the public sector isn't just about buying the "right" product. Before making a purchase, organizations must understand how different types of AI – from traditional machine-learning models for data analysis to generative AI for content creation – all have different use cases, technical requirements, and oversight and governance needs.

AI readiness means organizations must be equipped to:

- Conduct targeted risk assessments and impact evaluations for specific use cases
- Define clear, purpose-specific use cases and enable teams to experiment, iterate, and learn
- Maintain appropriate human oversight, auditability, and accountability where it matters most

Procurement's role is centered around ensuring the right solutions are available when teams need them. Meanwhile, a lot of the work to build readiness happens where AI is actually deployed and used. For example, **leadership** sets the strategic vision for AI adoption and creates frameworks that allow teams to move more quickly on low-risk applications while ensuring appropriate oversight for higher-impact use cases. Your agency or department **project team** needs to really understand the use case or problem they are trying to solve, and think through change management to support adoption. Your **technology team** should handle the technical integration and ongoing support.

The **procurement team** supports this readiness by creating pathways for innovation while supporting responsible and good implementation. This means instituting agile procurement practices, setting up processes to streamline access to low-risk AI solutions,

and building vendor relationships to better understand the market. For individual procurements, the procurement function also drives a buying strategy that takes into account the organizational mission, program needs, and implementation of responsible best procurement practices.

Since many public sector organizations are still building their AI capabilities, this section outlines how procurement can accelerate AI-readiness rather than slow it down.

Many public sector organizations are still building these AI capacities. In this section, we provide an overview of how to get your organization AI-ready, and the role of procurement in supporting AI-readiness.

What an AI-ready organization looks like

There are several enabling conditions that can help you and your organization get the most from your AI technology. The more of these capacities you have in place, the better you will be able to procure and adopt AI technology. It's still feasible to buy AI without all of these conditions in place – and many public sector organizations do! – but the solutions are less likely to be used or used to their full potential, leading to wasted resources or possible harms to public sector mission outcomes.

If your organization is not as AI-ready as you would like, we invite you to consider how you might support building these capacities. This might look like advocating for an institutional AI adoption strategy, requesting more agile procurement practices to enable experimentation, having conversations with AI industry leaders to learn more and support market research, or dedicating staff time to upskill on AI.

Institutional AI readiness

Institutional AI readiness is established at the leadership level, often by people that work across public sector agencies, such as a Department of Technology, Department of Innovation and Technology, or Chief Information/AI Officer, Digital Transformation Agency. While procurement officials don't set this strategic direction, they play a crucial role in bringing it to life by creating purchasing pathways and building vendor relationships.

Creating and using a robust **AI adoption strategy** can help public sector organizations align on overall goals for using this new technology and accelerate institutional AI-readiness. A good AI adoption strategy will balance opportunities and risks: it should

articulate both the vision for leveraging AI to improve public services and operational efficiency, and establish governance and oversight mechanisms to mitigate risks at the point of use.

By creating frameworks that enable innovation while ensuring appropriate controls, an AI adoption strategy can help agencies move forward confidently with AI adoption while managing potential downsides. The strategy should address key risks, including technical, operational, ethical, legal, and reputational considerations, as well as set out clear requirements for data privacy, transparency and accountability, security, bias monitoring, and regular risk assessments, as well as consider integration with existing legacy technology.

The strategy should establish clear criteria for categorizing AI systems based on their risk level and intended use. As part of this, procurement should be enabled to streamline access to low-risk technologies while ensuring appropriate oversight mechanisms are in place during implementation. This approach empowers teams to innovate while ensuring governance happens where it's most effective.

Critically, an AI adoption strategy should also include provisions for safe experimentation environments where teams can trial AI solutions, learn from pilots, and build internal capabilities before full deployment. This means procurement processes should support both rapid access to low-risk AI solutions for experimentation, and structured pathways for scaling proven solutions.

Data management is another critical component of organizational AI-readiness. Data management is the practice of organizing, cleaning, and governing information so it can be trusted and used effectively. Data engineering is the behind-the-scenes process of making sure all of this data – especially from fragmented or legacy systems – is collected, standardized, and delivered in a usable format for AI. In procurement, when legacy systems lack a data management and engineering plan, AI systems may produce misleading analysis or miss critical risk indicators, introducing procurement and compliance risks.

Good data management matters for many reasons:

- **Procurement risk:** Legacy systems often store contract data, pricing details, or vendor performance records in inconsistent formats. Without careful data engineering, AI may misinterpret or ignore this data, leading to flawed risk assessments or compliance gaps.

- **Data silos:** If information is spread across multiple, unconnected legacy databases, AI cannot see the full picture – introducing blind spots in supplier evaluations or spend analyses.
- **AI dependency on structure:** Unlike humans, who can sometimes infer meaning from messy records, AI requires structured and standardized inputs. Data engineering creates that structure so AI models can function.
- **Future scalability:** Agencies that invest in strong data management and engineering now will be able to more easily modernize legacy systems and unlock advanced AI capabilities for procurement in the future.

AI strategy resources

Don't have an institutional AI strategy? Check out these resources:
- The Gov AI coalition has resources for creating AI strategies, governance frameworks and more, targeted at local-level governments.
- The OECD AI Policy Observatory has guiding principles for national and subnational governments for designing AI policies.
- The Global Index on Responsible AI provides general best practices.
- The World Economic Forum has a framework for developing national AI strategies.

Many countries already have national or subnational AI strategies. For example:
- The European appraoch to artificial intelligence (European Commission)
- Artificial Intelligence 4.0 programme (Finland)
- Department of Homeland Security Artificial Intelligence (AI) Strategy (United States)
- Artificial Intelligence (AI) Action Plan (Australia)
- Artificial Intelligence Guidelines and Action Plan (Chile)
- City of Seattle 2025-2026 AI Plan (United States)

Department or agency project team AI-readiness

The department or agency project team that is buying AI technology knows their needs best, and if and how AI can help them achieve better outcomes. The procurement team should be involved and consulted, and can support the project team through providing purchasing frameworks, conducting market research, and building relationships with vendors. To help departments and project teams become AI-ready, they can:

- **Develop leadership expertise and structures.** Supportive and knowledgeable leadership is critical for teams to meaningfully build their own AI readiness.
- **Build a strong understanding of the problem you are trying to solve.** This problem should drive the technology and purchasing decisions. AI might not be the right solution to meet your needs.
- **Know AI solutions**: Teams need to understand not just their problem space but also which type of AI solution – generative, agentic, or simpler AI technologies – best fits their needs and context.
- **Know your data.** Seek to understand your own data. Good quality data is important for good AI outputs, but may not be available for all teams at the beginning of the project. Consider your data classification requirements and ensure your AI solution will be ready to transmit and analyze the data you wish to use. Procurement can help ensure that your vendor contract addresses your specific data requirements.
- **Understand your capacity to actively manage your vendor relationships as part of your project implementation.** Proactive contract management and close vendor collaboration are important for projects, especially those with technology deployments. AI solutions require ongoing monitoring and performance assessment. By closely monitoring progress and maintaining open communication channels to address issues early, you can help ensure that your AI solution is working as intended. To ensure a successful collaboration, it's important to check that your internal team has the capacity and responsibility to manage the vendor's work following the contract award.

Help for project teams to get AI-ready

Lots of resources are available to help develop your project team's approach to implementing AI in your project, including:

- The UK government's AI playbook for agencies and project teams
- The U.S. General Service Administration's guidance for project teams to get AI ready
- 18F's De-Risking Guide has strategies for technology vendor management
- The Australian Government's AI ethics principles to design responsible AI initiatives

Procurement AI-readiness

Procurement directors and procurement officials can do a lot to build readiness to support the implementation of responsible and effective AI.

- **Understand what AI technology you can purchase already.** Depending on your jurisdiction, you may be able to acquire AI solutions through framework agreements, and/or preexisting centralized or decentralized contracts. For example, your project team might already have access to CoPilot if your organization uses Microsoft Office, or Gemini through Google Workspace, or Amazon Q if your organization uses Amazon Web Services. This existing access also provides valuable opportunities for teams to experiment and build AI literacy before making larger procurement decisions.
- **Embrace agile and innovative procurement approaches.** This might be one of the most important steps that you can support to ensure better deployment and outcomes from AI. Enable your project team to buy differently and use outcomes-based contracting, rather than rigid requirements, to promote creative solutions and maximize value. Starting small and opting for iteration when possible can help manage risks and achieve the best results for AI purchases, as well as open the door to engaging with smaller, innovative providers that don't typically participate in public procurement. You can also create umbrella contracts to hire design and delivery teams, and structure them in sprints. These approaches also enable learning from pilot projects to inform larger AI investments and strategy. (Depending on your current policies and practices, we know this can be challenging. Check out our *Hacks for better AI procurement* section for ideas on how to be more agile and iterative.)

- **Support cross-organizational collaboration.** Double-check you have the right people in the room. With technology like AI, your key stakeholders should be involved early and often. Your technology team is especially important – they should get involved as soon as a project team decides they need to make a purchase. Security and legal experts, and someone who can represent the voice of whoever will use or be impacted by the technology, also need to be consulted along the way as you develop requirements and evaluate proposals and vendor claims. Usually, the project team is responsible for pulling everyone together, but encourage them to include procurement early on in strategy meetings rather than waiting to make a purchase decision, and let them know if someone is missing.
- **Build AI procurement expertise within your team.** Consider training procurement staff on AI basics and the dynamic sector landscape. This helps procurement professionals ask better questions and provide more strategic support to project teams.

Help for public sector procurement AI-readiness

There are great resources beyond this one that support AI-readiness for procurement, and we found several helpful as we developed this guidance. Here is a selection:

- The UK procurement guidelines for AI
- The State of California's procurement journeys
- New Zealand's AI-specific guidance to its general procurement guidelines
- Microsoft Public Sector for Expertise's guide on AI procurement, targeted especially at governments in the EU
- The USDS TechFARHub for guidance on iterative procurement at the federal level
- The World Economic Forum guidance on AI procurement
- PUBLIC's[guide on Buying Generative AI in Government
- The World Bank Group's Building Trustworthy Artificial Intelligence: Frameworks, Applications, and Self-Assessment for Readiness

AI readiness: Key questions

Procurement and organizational objective
Lay the groundwork for overall good AI procurement through policies and practices

Team	Questions
Procurement	**General** • Can we use agile methods for our IT procurements? • How can our procurement processes support strong collaboration between technical teams, procurement, legal, and policy stakeholders? • What does the market for this product or service look like? How can I engage them, both to get input into the best specifications and to ensure a good, competitive process? **Procurement risks** • How do we guard against vendor lock-in and ensure that contracts include clear exit strategies and data portability? • What protections do we need in place to clarify data rights, intellectual property ownership, and the handling of sensitive information? • What is the state of our data, and how much does that need to be considered in the procurement strategy?
Agency or department buyer	**General** • How do our overall organizational mission and goals align with our AI initiatives? • What are our internal capability gaps?
IT/Data analytics	**General** • Who inside the organization can take the lead on AI and connect the dots across teams? • What governance structures and frameworks do we need for responsible and ethical AI adoption? • Do we have the necessary AI literacy and understanding across relevant stakeholders?

	● How will we manage internal change and encourage safe adoption of AI tools? ● How will we ensure inclusive design and testing of AI tools to reflect diverse public needs?

Your role as a procurement officer

Who: Procurement teams
What: The role of procurement in supporting AI adoption in the public sector

Procurement officials often play a key role in supporting AI adoption. You help define the use case, your decisions shape how AI systems are evaluated and implemented and what safeguards and monitoring are put in place, and you help ensure oversight over vendors. You play a critical role in enabling teams to move quickly on promising opportunities, test solutions before major investments, scale successful implementations, and adjust direction based on learning.

By focusing on organizational needs rather than technology hype, leveraging existing procurement best practices, and enabling agile implementation, procurement officers can help ensure AI investments deliver real value for their organizations and the public they serve.

We recognize that many of you are doing your best despite challenging circumstances, and some of the practices we suggest in this guidance might feel wildly impossible given your context. Even so, we encourage you to review the material and think about what you can do, and check out our Hacks for better AI procurement section for ideas on how to get started.

You are also not alone. You do not need to know it all or decide it all. Your project team and technology team are there to work alongside you. Here we break down what you ideally have decision-making power around, and what you can influence:

As a procurement official, you have decision-making power around:

- **The procurement method and procedure:** You can translate business requirements into effective purchasing strategies and flexible procedures that support innovation, and help experimental projects achieve better outcomes through flexible, iterative, agile approaches and pilots.
- **Timeline:** Establishing procurement timelines and milestones that encourage strong and diverse vendor participation.

- **Evaluation model and criteria:** For AI, criteria often go beyond price to include factors such as training data, risk mitigation, energy use, and much more. While procurement can advise technology and project teams on criteria, procurement can have the most decision-making power over the evaluation model.
- **Contract terms and management:** For complex and evolving technology infrastructure like AI, proactive contract management is critical for monitoring performance, providing oversight, enabling iterative learning and adaption, and ensuring the best outcomes. You can help ensure prompt payment and compliance with laws and policies, and help structure contracts to enable scaling successful solutions. IT or the buying department or agency are often responsible for monitoring performance, but sometimes this is done by procurement, too.
- **Engaging the market and the public:** You can recommend or decide to deploy strategies, such as user research with beneficiaries/end users and vendor engagement through request for information. These approaches are key to making outcome-driven projects that engage responsible vendors and deliver benefits to the public.
- **Transparency:** Sharing information and crucial documents such as requests for information, requests for proposal, awards, and contracts, with the public can not only help vendors prepare better responses but also build trust with the public. Depending on your legal context, you can decide to proactively share more information through your existing public procurement and open data portals, but you can also create project specific public websites for more complex public facing projects.

You can also often influence:

- **A coherent and cohesive organizational AI technology strategy.** In collaboration with the technology team, you can help build alignment around what technologies are purchased to avoid technology and vendor fragmentation, as well as avoid duplicate purchases across different departments or agencies.
- **Articulating the use case and what success looks like:** A strong use case grounded in real needs is critical for determining what type of solution your team should pursue. You can also help support return-on-investment (ROI) calculations, and define vendor performance metrics.
- **Who is consulted as part of the procurement process, and when**. You can help pull in experts at the right moments.

- **Requirements:** Your requirements should support the responsible use of AI, and encourage participation through balancing risk with what the market can realistically provide.
- **AI governance and risk assessments:** Your organization should have policies on AI governance and risk assessments in place, around how AI is used, and how to address risks around transparency, privacy, and bias. You can help ensure that vendors have the capacity to support responsible AI use through structuring appropriate contract terms and managing risk through procurement vehicles.
- **Understanding the market:** You can help evaluate vendor capabilities and market offerings, including their relative maturity, as well as discover new vendors and emerging solutions.
- **Vendor relationships:** Strong vendor outreach and engagement throughout the procurement cycle is critical for performance. You can also anticipate and support contract modifications as needs evolve.
- **How the AI solution is adopted and goes to scale:** You can facilitate contract expansion for successful solutions, enable cross-agency adoption where appropriate, and help manage ongoing vendor performance.

AI adoption *myth-busting*

Who: Procurement, project, and leadership teams
What: Myths around planning, technical capacity, integration, and governance, and the reality

As procurement officials play an increasingly critical role in their organizations' AI adoption journeys, it's important to separate fact from fiction. While AI offers transformative potential, successful procurement requires a clear understanding of technical realities, vendor capabilities, and implementation challenges.

This section addresses common misconceptions about AI procurement and adoption. We break down popular myths, explain the reality, and highlight key considerations when evaluating vendors. Understanding these nuances will help procurement officials make more informed decisions and set realistic expectations.

Planning myths

Myth: "Implementing AI is always cost-effective and provides immediate ROI."

Reality: AI projects can require significant upfront investment and ongoing costs. ROI may take time to realize and depend on proper implementation and use. ROI discussions can evolve quickly, given the speed that the sector introduces innovations and improvements. Some tasks could be accomplished equally well by non-AI technology. During vendor evaluation, look for solutions for cost estimation, ROI projection, and optimization to help make informed decisions and maximize the value of AI investments. Plan for pilots to demonstrate value before making large investments.

Myth: "Implementing AI is never safe or cost-effective, so let's not use it at all."

Reality: Your specific use case should inform what technology solution you pursue, and AI might be a good fit. Lower-risk, proven ROI AI solutions do exist and continue to be developed. Risk and ROI can also be managed and understood through conducting small-scale pilots before taking solutions to scale. The advantages and disadvantages of AI are nuanced, and it's worth articulating these through your organizational AI strategy.

Myth: "AI systems are set-and-forget solutions."

Reality: AI systems require ongoing attention, though the level varies significantly. Managed cloud AI services may reduce operational overhead compared to self-hosted solutions, but all AI implementations need monitoring for performance, accuracy, and bias over time. Models may need updates as data patterns change, and business requirements evolve. During vendor evaluation, consider both the vendor's managed service capabilities and your organization's capacity for ongoing AI system management. Look for solutions that match your technical resources and operational preferences.

Myth: "Agentic AI will completely replace human decision-making"

Reality: Effective AI implementation typically augments rather than replaces human judgment. Different levels of automation are appropriate for different contexts and risk levels.

Technical capacity myths

Myth: "One AI model can solve all tasks effectively."

Reality: Large language models are foundational, but fine-tuning, prompt engineering, and guardrails are critical to make them useful and safe in specific contexts. Different models may be optimal for different use cases and considerations, such as cost, latency or built-in safeguards and bias mitigation. Look for vendors that have the capability to help adapt foundation models to unique needs, embedding safeguards and improving accuracy for specific use cases.

Myth: "Complete transparency of AI systems is always possible and necessary."

Reality: While transparency is important, its meaning and implementation can vary. Transparency in AI can encompass different technological and operational aspects, such as:

- Model architecture and training data
- AI system capabilities and limitations
- Clarity for users when they are interacting with AI-enabled systems

When considering vendors, look for clear documentation on model purpose, inputs, known limitations, and responsible use cases tailored to public sector needs, commitment to user transparency, such as clear labeling of AI-generated content or interactions, regular updates on model changes and improvements and willingness to discuss the balance between proprietary technology and transparency needs. Remember: Full technical transparency of large models isn't always feasible or meaningful to non-technical users. Focus on transparency that enhances trust, supports responsible use, and aligns with your organization's needs and values, and which will allow your organization to discharge applicable duties – for example, to provide reasons for decisions.

Myth: "AI decisions can always be explained in simple terms."

Reality: Not all AI decisions can be explained like a rulebook – context matters. Seek out vendors with support for guardrails that meet the requirements for explainability /interpretability for your specific use case.

Myth: "Automated decisions are always more objective than human ones"

Reality: AI systems can perpetuate or amplify existing biases if not carefully designed and monitored. They require robust testing and ongoing oversight, especially for agentic AI making consequential decisions.

Integration myths

Myth: “AI systems from different vendors can easily work together.”

Reality: Integration complexity can vary significantly. While some AI systems, particularly within the same vendor ecosystem, can integrate smoothly, cross-vendor integration often requires additional effort. Different GenAI systems often speak different languages – from varying API structures and data formats to different security protocols

and performance standards. Making systems work together requires significant integration effort, standardization, and often custom middleware solutions. During vendor evaluation, look for vendors that support open standards, APIs, and modular architectures that allow AI solutions to integrate into existing IT environments and minimize rework when switching providers, clear data processing and storage policies and transparent terms regarding ownership of customizations and prompts. Integration challenges can also vary by AI type. While GenAI might need API alignment, agentic AI requires additional considerations around decision-making protocols, audit trails, and responsibility frameworks.

Myth: "AI models and data can be easily transferred between systems."

Reality: You can move the data, but models may need reconfiguration or even retraining when moved to new environments. Look for vendors that have support for standard model formats and container technologies, structured export options and data export/import tools, standard data format support, multi-cloud and cross-platform capabilities and assistance with planning migration strategies that minimize friction.

Myth: “Once I start with one provider, I'm stuck forever.”

Reality: Lock-in is not binary. It's about the cost and effort to switch. Look for vendors that have data available for export, such as prompts, or effort you put in initially to customize/calibrate, as well as open-source model formats, containerization options (e.g., Docker, Kubernetes), and hybrid/multicloud support to enable transitions when needed.

Governance myths

Myth: "AI implementation must be risk-free."

Reality: The goal is not zero risk – it's understood, mitigated, and monitored risk. During vendor evaluation, look for support for responsible AI frameworks, sandboxing environments, and robust logging to help assess and manage risks realistically.

Myth: “My organization owns the data that we use with an AI solution.”

Reality: Not necessarily, this depends on the terms of your contract. The vendor may want to use your data for training purposes elsewhere. Align with your vendor and be clear

in your contract on how your organization's data will be used. Regulations like the General Data Protection Regulation (GDPR) in the European Union, or the California Consumer Privacy Act in the State of California, mainly protect personal data, but AI solutions can also process many other types of information as well.

Myth: “AI governance requires new departments or roles.”

Reality: Good governance can often build on existing positions and frameworks. You don’t need to duplicate efforts or add complexity to existing frameworks. Before creating new processes or roles, think about your existing organizational capacity and where AI governance structures might be a good fit, and where you already have expertise.

Myth: "Open source AI is always safer, cheaper, or better."

Reality: Open source AI offers benefits like transparency and customization flexibility, but it's not automatically better. While you may avoid licensing fees, open source solutions can require significant internal technical expertise, ongoing maintenance, security management, and support infrastructure. Given these requirements, total cost of ownership can sometimes exceed proprietary solutions. During early market engagements and vendor evaluation, look for features like clear explanations of which components are open source versus proprietary, support for major open-source frameworks and permissive licensing, secure and scalable environments for deployment, defined processes for updates and security patches, exit and transition arrangements including knowledge transfer, and clear support and maintenance agreements.

Myth: “Open-source AI is always more risky and unsafe.”

Reality: Neither open source nor proprietary AI is inherently safer. Security and safety depend on implementation quality, ongoing maintenance, vendor support, and your organization's technical capabilities. Open source allows for greater transparency and customization of security measures, while proprietary solutions may offer enterprise-grade support and managed security features. The right choice depends on your organization's technical capacity, security requirements, and risk tolerance. Consider which components are developed using open source, what has been developed by others and is being reused, your use case, your institutional capabilities and capacities, and the extent to which you need active technical support and frequent updates.

Myth: "The same governance frameworks work for all AI systems"

Reality: Governance needs vary significantly between different types of AI solutions and different use cases. While content generation might need review processes, automated decision-making requires stronger oversight, clear accountability chains, and specific intervention protocols.

Myth: "Data must never leave the country to ensure sovereignty."

Reality: It's not just about geography. While it is important to consider geopolitical implications on where the data is held, it's probably even more important to understand who has access, how it's encrypted, and who controls the encryption keys. During vendor evaluation, consider local hosting options and sovereign environments when needed, customer-managed encryption keys, and compliance tooling to meet residency and sovereignty regulations and needs in context.

Ways to *buy AI*

Who: Procurement, project, and technology teams
What: Description of AI use cases, aligned with type of AI feature, customization level, primary purchase type, commercial model (pricing and contract type), commercial risks, AI risks, and common procurement methods.

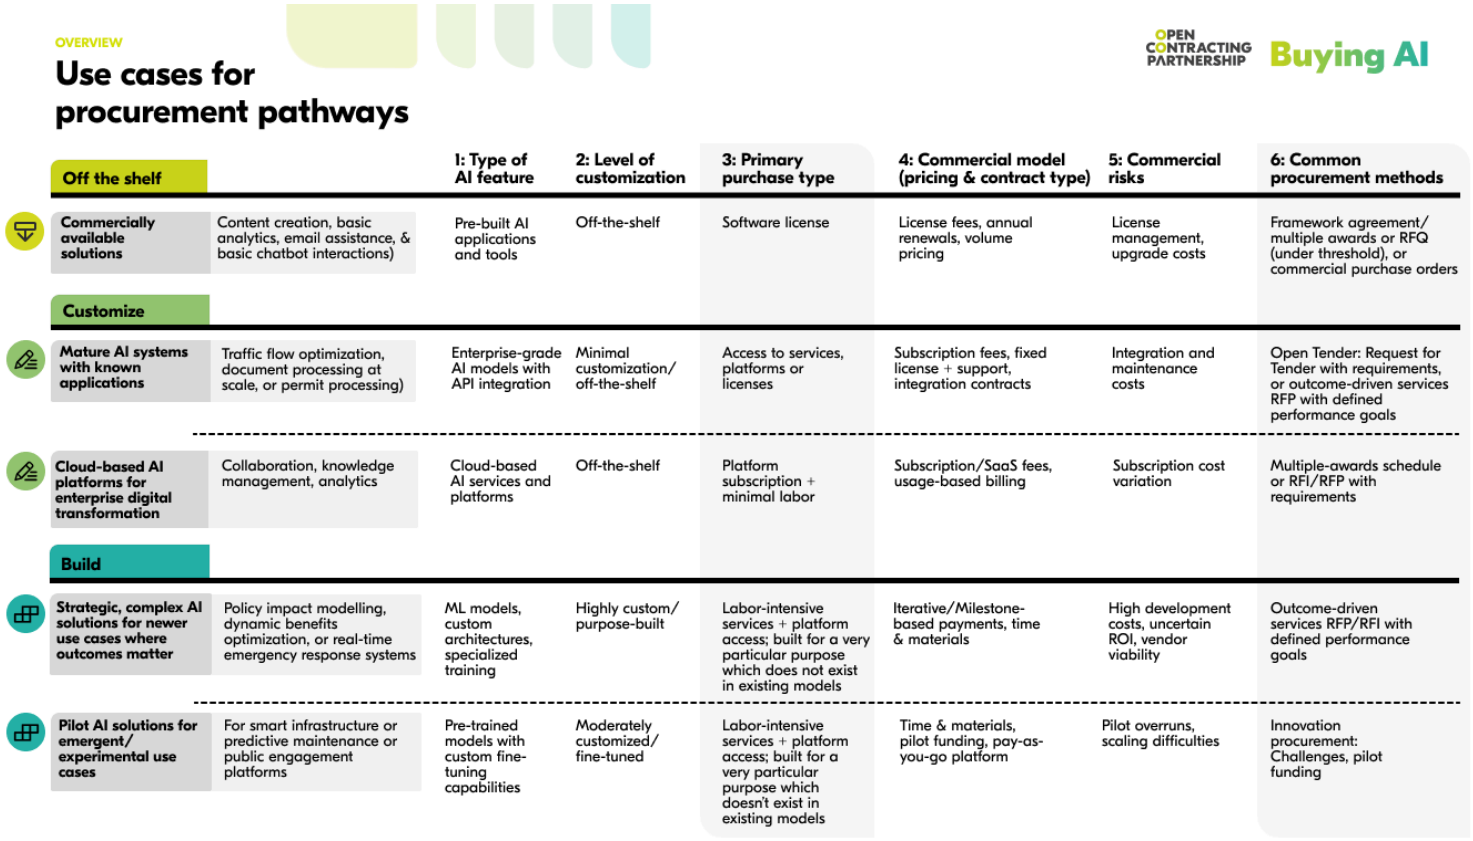
Governments often purchase AI through three pathways: **"Off-the-shelf"** for commercially available AI solutions for quick, common needs, **"Customize"** for proven AI systems to fit specific processes, or **"Build"** for new AI solutions for strategic, high-stakes, or emerging challenges. Your procurement pathway should be determined by **your use case**, and informed by other factors such as **risks.**

Right now, off-the-shelf is the most common way we hear the public sector is buying AI, and this is increasingly happening through central purchasing. This trend implies a more limited role for procurement, both in terms of the purchasing process itself and supporting adoption. Custom or build solutions mean the public sector is purchasing software development services and time, rather than simply a product. This makes good technology procurement practices critical for success.

Not all AI enters the public sector through these pathways. Sometimes AI isn't deliberately purchased, but is **added on to existing services under contract** that the public sector is already using. Other times, AI is provided to public sector organizations through academic partnerships, grants, or pro bono pilots with private sector companies.

Do you feel stuck? We've heard over and over again that teams don't feel they have a choice when it comes to which pathway they pursue. If this is you, check out the **Hacks for better AI procurement** section for ideas on how to get the most out of your procurement if you can't choose the method.

Use cases of procurement pathways



OVERVIEW

Use cases for procurement pathways

OPEN CONTRACTING PARTNERSHIP Buying AI

		1: Type of AI feature	2: Level of customization	3: Primary purchase type	4: Commercial model (pricing & contract type)	5: Commercial risks	6: Common procurement methods
Off the shelf							
Commercially available solutions	Content creation, basic analytics, email assistance, & basic chatbot interactions)	Pre-built AI applications and tools	Off-the-shelf	Software license	License fees, annual renewals, volume pricing	License management, upgrade costs	Framework agreement/ multiple awards or RFQ (under threshold), or commercial purchase orders
Customize							
Mature AI systems with known applications	Traffic flow optimization, document processing at scale, or permit processing)	Enterprise-grade AI models with API integration	Minimal customization/ off-the-shelf	Access to services, platforms or licenses	Subscription fees, fixed license + support, integration contracts	Integration and maintenance costs	Open Tender: Request for Tender with requirements, or outcome-driven services RFP with defined performance goals
Cloud-based AI platforms for enterprise digital transformation	Collaboration, knowledge management, analytics	Cloud-based AI services and platforms	Off-the-shelf	Platform subscription + minimal labor	Subscription/SaaS fees, usage-based billing	Subscription cost variation	Multiple-awards schedule or RFI/RFP with requirements
Build							
Strategic, complex AI solutions for newer use cases where outcomes matter	Policy impact modelling, dynamic benefits optimization, or real-time emergency response systems	ML models, custom architectures, specialized training	Highly custom/ purpose-built	Labor-intensive services + platform access; built for a very particular purpose which does not exist in existing models	Iterative/Milestone-based payments, time & materials	High development costs, uncertain ROI, vendor viability	Outcome-driven services RFP/RFI with defined performance goals
Pilot AI solutions for emergent/ experimental use cases	For smart infrastructure or predictive maintenance or public engagement platforms	Pre-trained models with custom fine-tuning capabilities	Moderately customized/ fine-tuned	Labor-intensive services + platform access; built for a very particular purpose which doesn't exist in existing models	Time & materials, pilot funding, pay-as-you-go platform	Pilot overruns, scaling difficulties	Innovation procurement: Challenges, pilot funding

Download the table: [Use cases for procurement pathways]

Buying off-the-shelf is often pursued when...

- **Use case:** Solutions are well-established and standardized
- **ROI:** Clear ROI can be more easily demonstrated
- **Capacity**: Internal technical resources are limited
- **Data security:** No custom solution required to protect data beyond the organization's standard privacy agreements
- **Timeline**: Quick implementation is needed
- **Vendors**: Established vendors exist, are pre-approved already and can deliver quickly

Buying customized or build AI is often pursued when…

- **Use case:** Unique organizational needs exist
- **ROI**: May be demonstrated or less certain for newer use cases
- **Capacity**: Internal technical capability is strong
- **Data security**: Custom solutions required
- **Timeline**: Sufficient time to accommodate agile development
- **Vendors**: Pool of vendors with known capabilities, newer vendors and collaborations are needed

Hacks for better AI procurement

Or, tips for navigating complex procurements in imperfect situations

Who: Procurement and project teams
What: Practical strategies to get the most from AI procurements, including off-the-shelf purchases

We know that not everyone lives in a world where great tech procurement feels possible. AI technology is advancing quickly, which means the best solution today may not be the best solution in a few years, making flexibility and adaptability more important than ever. But maybe you have no institutional AI strategy, feel unable to buy technology using an agile approach, or have decentralized procurement that makes it hard to ensure the right people are in the room.

In challenging environments, don't let perfection be the enemy of "good enough." Experimentation and iteration can help you better achieve your goals for a specific procurement. But this can also help you build confidence across your organization in new ways of thinking and working together, helping you shift norms around procurement practices.

Here are some strategies for those seeking to still get the most out of your AI procurements despite constraints, especially if you are buying customized or built solutions:

- **Find your allies.**
  Look for people across your organization from legal, IT, finance, and beyond who are good at getting things done, and willing to help you out. These allies can support you with your current procurement in the short term, and help with strengthening practices generally in the long term. You might also find inspiration in other cities, countries or agencies who have dealt with similar issues and might have ideas.

- **Buy just what you need to get started.** This might look like just buying research and development, a pilot, or software development services, rather than buying a final product right away with a large scope, budget, and duration.
- **Engage the market early before you buy.** There are many ways of letting potential vendors know that you are considering making a purchase. You can describe your plans publicly before you issue any Request for Proposals (RFPs) so they can get ready. If you can, start with a Request for Information (RFI) to gather market intelligence to inform your RFP.
- **Make your solicitation agile-ish, even if you can't be agile.** Go beyond technology specifications to carefully describe the problem to be solved and intended outcomes, including as much detail as possible about the current situation (such as unmet user needs, past attempts to solve the same or similar problems, organizational constraints such as legacy IT systems, reduced funding and a lack of digital, data and technology skills, etc.), and your envisioned results. You can also articulate your needs in ways that agile developers are more likely to respond to and succeed. For example, if your government blocks certain words or concepts, you can use alternatives. For example, instead of "sprints" you can ask for weekly milestones, "show and tells" can be progress reporting, "demos" can be milestone reviews, and "user research" can be a stakeholder needs assessment.
- **Don't be afraid to ask basic questions during selection.** Vendors should be able to tell you in plain language important information like what data they used to create their model, how they will use your government's data, and what success and failure looks like for their solution. You can also request information about model traceability. This can include the training data, the algorithm used, the hyperparameters, and the training environment. Even if a vendor is unable to share the training dataset itself, they should still be able to provide this general background. This information can help you understand the model outputs and identify potential issues.
- **Aim for shorter contracts (with smaller budgets).** AI technology is changing quickly, and what the best solution is now might be obsolete in two, five or ten years. Shorter contracts can help get the best technology to meet your government's current needs, and manage risks.
- **Make room for improvements (without change orders).** By incorporating guarantee periods or structuring flexibility into the contract (such as through options), buyers can better adapt the scope or request changes if the results are not meeting expectations.

- **Piggyback off existing contracts** to get services and features, this could be task orders, cooperative contracts or other pre-approved vendors.

What if I'm just buying off-the-shelf AI?

The most common way that public sector organizations are procuring AI solutions is through framework agreements or multiple awards. Although buying AI this way is more straightforward, it's still important to pay attention to several topics that we cover throughout this guidance, including:

- **Use case:** What you need should inform what you buy and how you contract.
- **Data governance and security:** You should understand how your organization's data will be used.
- **Testing and experimentation:** Before rolling out organization-wide, strongly consider pilots to help mitigate risks by creating time and space for learning what works and what does not.
- **Engaging your users, employees, and the public when possible:** Input from your end users, employees, and the public along the procurement journey can be very helpful, especially if the solution will be public-facing, or help inform decision-making for service delivery.

You can also aim to be as agile and innovative as is feasible while working inside your pre-approved structures. For example, you can:

- **Use framework agreements or multiple award schedules (MAS) as your foundation, not your final destination:** Consider if your project team can build custom prompts, models, or applications using the enterprise-licensed AI as your foundation, such as via an API.
- **Negotiate for a pilot within your framework agreement:** If this is possible, you could test different AI technologies before starting a full procurement.
- **Use procurement to add new functions:** Consider if you might…
  - Piggyback off existing contracts to get new features, such as through task orders or call-offs.
  - Partner with smaller, agile vendors that are already on the MAS. They might be able to work with you through existing contracts.
  - Leverage cooperative agreements: Some schedules allow you to use solutions procured by other agencies.

Resources on buying off-the-shelf

- A blog about platforms and AI by Mark Headd
- A blog post about "Fake COTS -commercial, off the shelf software-" and the one-day rule by Sean Boots
- This procurement of AI guide by IDC and sponsored by Microsoft offers a framework and tips on how to leverage AI offers by cloud service providers
- This toolkit by Connected by Data and Trade Unions Congress Wales provides an actionable framework for empowering worker voice throughout the procurement process for digital and AI systems

Planning and procurement strategy tips

Who: Procurement teams
What: How to develop a procurement strategy, who to consult, and what to include

Two types of planning are required for successfully adopting and purchasing an AI solution: **project planning**, to figure out the broad technology strategy, which is usually led by the agency/department project team or sometimes IT, and **procurement planning**, to determine which products/services/licenses have to be purchased as part of this project, which is led by procurement.

This section mainly focuses on procurement planning. However, at the end of this section, we also include important considerations for other involved teams during the planning of the broader project.

One valuable tool for procurement planning is a procurement strategy. A procurement strategy is a public-facing plan that explains how your state or local government will buy and deliver a specific technology to achieve measurable outcomes for residents and/or staff.

A procurement strategy should articulate the problem you are trying to address in the context of your organization's wider strategic objectives and priorities, and the outcomes you are seeking to achieve in plain language. This analysis should include a clear assessment of market solutions versus custom development, explain your chosen delivery approach, your sourcing model (including solicitation and contracts types, terms, timelines) and commercial terms up front. This helps set shared expectations for public agencies, vendors, oversight bodies, and the public, and is particularly important when procuring mission-critical technology.

Early on you'll need to decide whether to pursue off-the-shelf AI, or customize or build your AI solution. You may also consider reusing existing open-source components as a delivery

approach, or in combination with other methods. Your strategy should clearly justify this choice.

If recommending off-the-shelf,

Explain how existing enterprise platforms or applications can meet your needs through standard features and configurations. Consider how you might benefit from continuous vendor improvements, proven security frameworks, and experience from implementations in other organizations, as well as the vendor's AI services and integration capabilities. Depending on your specific use case, specialized AI applications might be a better fit.

If recommending customized solutions,

Explain why the underlying platform capabilities are strong, but standard configurations aren't enough to meet your needs. This might involve custom integrations, specialized workflows, or agency-specific data processing requirements that can be built on enterprise foundations.

If recommending building custom solutions,

Articulate why neither off-the-shelf products nor platform customization are enough to meet your needs. This typically involves unique regulatory requirements, new or unique use cases, or complex integration needs that exceed what existing platforms typically can do. Consider whether these requirements justify the additional development time, cost, and ongoing maintenance responsibilities.

Who and how to engage



Whom you gather input from as you develop and refine your strategy is as important as the content itself. Key stakeholders include:

- **Your internal experts**: You should speak to the agency or department buying the technology, and your technical, finance, and legal experts. Gathering this feedback early will help generate alignment around what you are buying, how you will buy and pay for it, and what to consider during contract implementation with your vendors.
- **Your organization's employees**: Seek to empower public servants' voices throughout the procurement process by engaging with public sector commissioners and trade union representatives.
- **Suppliers:** Strong, proactive supplier engagement is especially important for AI. AI technology and vendor capacity are changing very fast, so understanding what the market is capable of, now and in the near future, is important for many aspects of your procurement strategy. Potential engagement methods range from attending conferences or vendor fairs, to issuing RFIs and conducting one-to-many and one-to-one supplier sessions, to innovative, challenge-based procurements.
- **The public:** By engaging the public during the design of your procurement strategy, you can promote organizational transparency and accountability, and build trust with your residents. Public participation is especially important for any AI technology that the public might interact with directly as an end user, or used to influence public sector service delivery.

In addition to your key stakeholders, you may also consider gathering input from organizations that work on AI- and procurement-related topics for the public sector, such as professional associations, nonprofits, or academia.

How you engage your stakeholders matters, too. By taking a collaborative and participatory approach, you can develop a strong shared understanding of the problem you are trying to solve, your objectives, and outcomes.

Thinking through your budget

Identify your funding sources, and which team within your organization will own the product budget. Depending upon the scope of your purchase, you may have multiple teams responsible for different budget lines.

It's important to recognize that AI solutions typically use different pricing models than traditional software, which have implications for your costs. Think carefully with your

IT/data teams around what realistic costs look like, both upfront and for appropriate maintenance.

Common AI pricing models include:

- **Usage-based pricing**: Many AI services charge per API call, document processed, or query made
- **Subscription models:** Monthly/annual fees often based on number of users or processing volume
- **Hybrid approaches:** Combination of base fees plus usage charges
- **Custom enterprise pricing**: For large implementations or specialized requirements

Key considerations for procurement around pricing include:

- Costs can scale rapidly with usage, making it especially important to understand volume projections.
- Training or customizing models may involve significant upfront costs.
- Data preparation, integration, and ongoing human oversight all have cost implications.
- Consider pilot pricing vs. full-scale implementation costs when budgeting.

Procurement strategy template

Problem	The problem you are trying to solve, and how solving this would support the wider organizational strategic objectives and priorities
Objectives	The ultimate outcomes you are trying to achieve through your AI purchase
Delivery approach	The delivery approach (i.e. modular, agile etc.)
Sourcing approach	The design of the contract process (i.e. solicitation and contracts types, terms, timelines)

Budget	Your budget for the project, both for upfront costs and ongoing maintenance, as well as your funding source
Commercial terms	Relevant legal considerations that will inform contract terms
Governance and accountability	Your organization's approach to AI governance and accountability, including preferences for data management, data sharing, and data security.
Market research	Your approach to market research to understand market capability
Supplier engagement and communications plan	How you will let the market know early what is coming and encourage participation
Evaluation goals and priorities	The most important factors in evaluating solutions, for example price is more important than capability, or must be delivered by a certain date
Risk management	Identifying implementation risks and your mitigation approach.
Monitoring and optimization	The KPIs for performance and contract delivery
Public transparency and accountability	Your approach to communicating with the public around your planned use of AI and the procurement method

Building your procurement strategy: Key questions

Procurement objective
Develop a procurement strategy that fits the overall organizational strategy and the needs of the specific procurement.

Organizational objective
Identify your needs for your specific procurement

Team	Questions
Procurement	**General** • Are we engaging the right internal and external stakeholders in our strategy development process? • Is our budget realistic, given ongoing maintenance needs?
Agency or department buyer	**Model and system architecture** • What is the problem we are trying to solve? • What are our needs and use cases? • How will this project impact end users, including public sector staff, frontline workers and communities or service recipients? • What are the expected outcomes of the solution we want to acquire? • How might we safely test our assumptions, approach, and AI tool options before launching the formal procurement process? **Operational considerations** • Which role or team will be accountable for managing risk and compliance? Do we have the internal capacity for this oversight?
IT/Data analytics	**Model and system architecture** • Is AI the right tool to meet our needs? What deterministic and rule-based solutions might deliver comparable results to generative AI for our use case? • If we pursue AI, what type of model best suits our specific use case (e.g., general-purpose vs. specialized models)? • What integration capabilities do we need with existing systems? • What is our internal capacity to maintain the tool? • Where will the tool be hosted? Will we need to acquire additional processing or storage capacity? • What are environmental considerations?

Who to involve

Who: Procurement teams
What: Key stakeholders in the AI procurement process

Procuring complex technology like AI may require coordination and input across a range of stakeholders. Off-the-shelf technology can require less engagement, depending upon the use case. But for customized or built AI, or higher-risk use cases for off-the-shelf technology, it's important for stakeholders to be involved as soon as you can, and often throughout the deployment process. Bringing together people with diverse backgrounds, perspectives, skills, and experiences can help your organization reflect on the overall impact and how to make the most from AI, while also considering requirements for security, privacy, explainability, audibility, and legal compliance.

It's especially important to bring in your technology experts into your conversations as early as possible to help you understand if AI is a feasible and valuable tool to address your challenge, as well as navigate the technical components of the procurement and implementation processes. If you lack in-house expertise, consider what experts from your community and nonprofit partners or other experienced buyers you can consult, and think carefully about what solutions you can feasibly pursue given your capacity constraints.

Precisely whom you will involve will be different depending on your specific context. Your local laws, procurement system (centralized versus decentralized), and the specific purchase type will all impact your own final list of whom to engage.

Key roles include:

- **Procurement officer,** who leads the procurement process. They should help articulate the scope and strategy of the procurement, and often have final signoff on the vendor selection process, and then provide oversight during contract implementation.
- **Project team leader (buyer)**, or the agency/department representative that is seeking to make a purchase to meet a need. They usually facilitate the overall

discovery process, ensuring smooth collaboration and progress. In a best-case scenario, this person is a product owner, who also has technology expertise.

- **IT/Data science expert,** who can provide knowledgeable guidance around responsible AI use and its capacities and limitations, as well as infrastructure and security requirements, and more. They are usually in an advisory role, unless they are also in the position of being the buyer.
- **Legal officer,** who can help navigate considerations around data, liability, intellectual property, and AI use cases. They often provide sign-off on both the RFP and final contract.

Two other critical roles are the **advocate or representative of the end users, and the public**. Although they are different types of stakeholders, it's still important to consult them and gather their feedback along the way. Public participation is especially important if the AI will have a direct impact on service delivery for residents.

Public participation resources

For tips on public participation, check out Meaningful Stakeholder Engagement in Public Procurement for Artificial Intelligence: A Mission-Oriented Playbook by Public AI

Solicitation tips for AI

Who: Procurement, project, and technology teams
What: Best practices for solicitations for AI that go beyond off-the-shelf purchases

These tips can help if you are seeking to customize or build an AI solution.

As you develop your approach...

- **Support cross-functional collaboration.** Collaborate early and often with stakeholders across the organization, including AI specialists, legal, security, compliance, and end users, to ensure the solution meets diverse needs and regulatory requirements. Have an expert who understands AI solutions on your evaluation team.
- **Engage the end users and vendors early.** In line with good technology procurement, make sure that you have a deep understanding of your user or beneficiary needs through user research and testing. Equally important, the vendor market can give you helpful feedback not only on what is possible, but also what would incentivize them to participate in solicitations. Engagement efforts may include conducting a request for information (RFI), requesting demonstrations, and more.
- **Use an outcomes-based approach.** Focus on the problem and measurable features or Key Performance Indicators (KPIs) rather than prescribing specific technologies. Consider solutions that can leverage current commercial and emerging AI capabilities while meeting your requirements. Understand how the underlying technology works, both in terms of the algorithms and training data. For example, if a health agency wants to deploy GenAI to rapidly summarize information about patients, an outcomes-based RFP might say: "we seek to reduce time/cost for feedback summaries by 50%, ensure explainability, minimize bias—all while safely handling sensitive patient data."
- **Support testing and iteration.** Encourage pilot projects, sandbox environments, and modular procurement to enable safe testing of AI models. Consider proof-of-concept deployments or trial periods to evaluate capabilities before full

implementation. Look for documented evidence of real-world implementations in similar contexts.

As you develop your solicitation documents and requirements... {.collapsible}

- **Incorporate compliance and risk management.** Include requirements for relevant IT security and AI-specific certifications and regulatory alignment, while balancing flexibility to encourage vendor participation. This can look like specifying any needs for bias mitigation, stakeholder participation, explainability, audit trails, human-in-the-loop controls, and transparency about training data, ensuring these align with your organization's risk management framework. Also ensure alignment with appropriate security frameworks and regional compliance requirements set out by your jurisdiction. The NIST Risk Management Framework can be a helpful resource.
- **Address intellectual property (IP) and data rights.** Clearly define ownership of outputs, terms for using public sector data for model improvement, and data portability expectations. Consider contractual requirements for data residency, usage restrictions, and data handling protocols.
- **Ensure legal and regulatory alignment.** Reference applicable AI and data protection laws in your jurisdiction and incorporate them in solicitation and contract clauses. Consider requirements for demonstrating compliance with relevant legal frameworks.
- **Include vendor and contract management.** Plan for knowledge transfer and evaluate the long-term flexibility of your chosen solution, including data portability and integration capabilities.
- **Specify integration requirements.** This includes existing IT environments and compliance with relevant accessibility standards, ensuring inclusiveness and usability for all users. Evaluate technical compatibility with your current systems and infrastructure. Consider data flow and interoperability requirements, and solutions that support industry-standard APIs and data formats to enable integration with your existing systems and provide flexibility for future technology evolution.
- **Align on technical infrastructure needs for development.** Clarify whose technological infrastructure will be used for developing and deploying the solution.
- **Include data governance and privacy requirements.** These can include data governance, breach notification protocols, and data protection responsibilities. Consider appropriate security certifications and privacy controls for your context.

Understand clearly how the solution provider will use your data and address associated risks.

- **Consider sustainability.** You may have sustainability requirements for environmental impact reporting, such as energy usage metrics and sustainability practices, aligned with your organization's environmental goals.

As you draft your RFI/RFP documents, include... {.collapsible}

- **Core structure**: Start with your standard organizational RFI or RFP structure (e.g., timeline, instructions, scoring) and make it even stronger by focusing on the problem that you are trying to solve and the outcomes you want to achieve. Keep it short and use simple language.
- **Technical evaluation**: Incorporate sections for AI capabilities, security measures, multi-model integration, vendor support and transitions.
- **Quality and capability assessment**: Include questions around topics like model accuracy, hallucination, mitigation, benchmark validation, ethics, and compliance.
- **AI disclosure**: Consider adding a requirement for transparency and accountability depending upon your policies and use case.
- **Scoring rubric**: Assign weights to elements like security and compliance, disclosure, model performance and evaluation, support and maintenance, and pricing/value.
- **Guarantee period:** Include a guarantee period to help ensure that the vendor will fix issues early on to meet your specifications.

In line with good procurement practices, make sure to keep your RFP long enough to enable vendors to prepare thoughtful responses, and spend time and budget to spread the word to vendors whom you might want to attract.

Resources for drafting solicitations

The UK guidelines for AI procurement include a section on drafting requirements. PUBLIC has a guide on buying GenAI that has specific tips for drafting requirements.

Solicitation tips: Key questions

Procurement and organizational objectives

- Decide on your procurement procedure
- Draft your solicitation

Team	Questions
Procurement	**Project needs** • How might we reflect the project team's needs in a RFP? • Can the vendor demonstrate a roadmap and flexibility that aligns with our evolving needs? • What level of vendor support or service guarantees will we need if something breaks or goes wrong? • Will the vendor support integration with other AI tools or models we may adopt in the future? **Operational considerations** • What kind of selection criteria should be used to rank solicitation responses? • What commercial or pricing models are appropriate? • Are the right people in the room for this procurement? **Risk and compliance** • How will we ensure compliance with current and evolving AI-related legislation, procurement rules, and standards? • How will we manage and track AI-specific costs (e.g., compute, storage, API usage)? • How much risk does vendor lock-in pose in our context, and what mitigations are practical?
Agency or department buyer	**AI solution** • Do we know what "good enough" model response performance looks like for our specific use case, and what metrics will we use to measure AI success and ROI (qualitative and quantitative)? • What are our fallback strategies if the model gives low-quality or inaccurate responses? • How will users interact with the AI solution? If there is a platform or dashboard, who is responsible for developing it?

- Can we retrieve or export our data and models if we switch vendors?

Data strategy and management

- Do we and our vendor partners understand the quality, completeness, and relevance of our data?
- What role will the vendor play in helping us prepare or transform data for effective use?
- How will we manage data ownership, licensing, and usage rights?
- What safeguards will the vendor provide to ensure that sensitive, personal, or classified data is handled in line with legal and policy requirements?
- Will the AI system retain or learn from our data inputs? Under what terms, and with what controls?
- Is the algorithm specialized for our organization? If so, will we have input into how the algorithm is fine-tuned or retrained using our data? What additional data sources do we need to get the best results?

Operational considerations

- How might we approach AI implementation in stages, such as through a pilot or sandbox?
- How will we measure and monitor model performance over time?
- What policies will govern human-in-the-loop oversight for critical decision-making?
- What mechanisms will we put in place for continuous performance monitoring and system health checks?
- What accident response mechanisms will we put in place? What processes will we follow to learn from accidents and improve their prevention and the management of associated risks and harms?
- If we are involved in retraining or fine-tuning the model, what internal processes will we establish for prompt engineering and optimization?
- Who owns operational responsibility for uptime, error handling, and incident response?
- What logging and audit mechanisms will be used for AI inputs and outputs?
- How will we ensure traceability and version control for prompts and model configurations?

Risk and compliance

- What processes will we establish for bias detection, mitigation, and fairness audits?
- What controls will be in place for data privacy, consent, and security (at

	rest and in transit)? ● How will we meet model explainability requirements, especially for high-impact or regulated use cases? ● How will we manage IP and copyright issues relating to AI-generated content? ● How will we assess and audit third-party AI providers for security, ethical alignment, and accountability? ● What safeguards will we implement against hallucination, misuse, or reputational risk in public-facing services? ● How will we address bias detection and adversarial robustness?
IT/Data analytics	**Project needs** ● What are our requirements for model transparency and explainability? ● How will we measure and monitor model performance over time? ● What are our requirements for model updates and maintenance? ● What level of customization do we require? ● What are the latency, reliability, and availability requirements for our applications? ● Are there specific regulatory cybersecurity requirements, such as around data residency or geographic deployment? **Integration** ● How will the AI solution integrate with our existing systems? ● Does the vendor support open standards and interoperability (e.g., prompt formats, model interchange)? ● How do open-source and proprietary models compare in terms of cost, performance, and vendor support and accountability? ● Are there clear exit strategies or migration paths if we outgrow a particular vendor or need to switch platforms? **Data strategy and management** ● How will we jointly manage data governance – including storage, audit trails, and compliance with data retention and deletion policies? ● Are there opportunities to collaborate with vendors on domain-specific dataset creation or fine-tuning? ● What processes are in place for anonymizing or redacting data before it enters AI systems? ● How does the vendor support discoverability and metadata management for both inputs and outputs? ● How will data flows be documented, monitored, and audited across the lifecycle?

Tips for *evaluating AI proposals*

Who: Procurement teams
What: Best practices for evaluation

Your evaluation methods will vary depending on what and how you are buying. As you evaluate proposals, keep your end user and/or beneficiary in mind based on your organization's user research. Make sure that your evaluation criteria and selection process speak to their needs. Here are some tips to consider.

General good practices

Lifecycle and vendor management

- Plan for ongoing model maintenance, updates, and version control
- Confirm vendor commitments to transparency about changes and potential impacts
- Include rollback and incident response procedures in contracts

Risk and compliance

- Ensure adherence to applicable AI regulations and standards in your jurisdiction
- Establish explainability requirements aligned with use case needs
- Define data governance policies covering data sharing, retention, and deletion
- Assess risks related to third-party dependencies and supply chains

Documentation and governance

- Maintain detailed records of evaluation criteria, scoring, and decisions
- Require vendors to disclose model limitations and ethical considerations
- Embed responsible AI requirements into procurement contracts

Additional considerations

- Compatibility with open source AI models or platforms, where appropriate
- Support for data interoperability and integration with existing public sector systems, tools, and data repositories
- Transparency and communication plans for public-facing AI applications

- Preparedness for public records requests or transparency laws impacting AI outputs
- Strategies to mitigate risks related to hallucinations, misinformation, or misuse
- Clearly define support and escalation pathways for incidents
- If appropriate and legally feasible given your context, consider involving end-users/beneficiaries themselves in the evaluation process.

Best practices by contract method

Method	Best practices
Open solicitation	**Evaluate AI-specific technical criteria** • Suitability of the type of model for your use case (general-purpose vs. domain-specific) • Model accuracy, reliability, and robustness, including known failure modes • Scalability, latency, and availability under expected workloads • Security features aligned with your jurisdiction's regulations (e.g., GDPR, HIPAA, FedRAMP) • Data residency, sovereignty, and cloud region options to meet compliance needs **Understand technology and vendor transparency** • Clarity on training data provenance and governance • Vendor openness about model development, fine-tuning methods, and update cadence • Infrastructure dependencies and integration capabilities with your existing IT ecosystem • Support for open standards and interoperability to avoid lock-in (e.g., standard APIs, model interchange formats) **Review vendor and solution track record** • Demonstrated enterprise or public sector deployments • References relevant to your region and sector • Certifications and compliance attestations applicable to your regulatory environment • Support for hybrid and multi-cloud strategies or on-premises deployments **Pricing and commercial terms** • Transparent usage-based or subscription pricing models

	• Availability of cost management and optimization tools • Clear delineation of charges for pilot vs. production environments • Contract flexibility for scaling usage or changing terms as needed **Security, privacy, and compliance** • Strong data protection measures, including encryption at rest and in transit • Capabilities to comply with data protection regulations (e.g., GDPR for EU, CCPA for US) • Auditability and logging features for traceability • Support for sovereign cloud or localized data centers where required **Responsible AI and ethical considerations** • Alignment with your organization's AI ethics framework and regulatory requirements • Bias detection and mitigation strategies embedded in the solution • Explainability features for AI outputs, especially where decisions impact citizens • Human oversight mechanisms and escalation procedures
Innovation procurement	**Include all above points, plus:** • Access to sandbox or trial environments for iterative testing • Flexible contracting to support phased delivery and learning cycles • Clear frameworks for intellectual property rights and data ownership • Opportunities for co-development or joint innovation with providers • Availability of professional services and ecosystem partnerships
Framework agreement (multiple awards) or under-threshold purchase	**Focus on essential due diligence:** • Standardized compliance and security documentation • Service Level Agreements (SLAs) with clear uptime and support commitments • Transparent cost structures and usage monitoring

Evaluation criteria

As you create your evaluation criteria, align around your metrics, what the terms really mean, and the implications. For example, "accuracy" often comes up during evaluation. But other metrics matter as well, such as precision (which is also sometimes confused with accuracy by non-experts) and recall. In some use cases, precision and recall are trade-offs as parts of the overall accuracy of the solution. Speak to technical experts to understand the implications of these metrics in terms of false positives/negatives, and discuss relevant operational and legal implications.

Your evaluation criteria will be case-dependent. Some evaluation criteria for open solicitations you may consider include:

- **Problem fit,** including demonstrating a strong understanding of your need and the appropriateness of the solution.
- **AI approach and explainability,** such as model algorithm and training data, model transparency, algorithmic clarity, audit readiness.
- **Environmental impact,** such as energy use and source.
- **Data governance and bias handling,** including strategies for fairness, data quality, and addressing known model limitations.
- **Integration, training and operations,** such as fit with existing infrastructure, plans for training and upskilling internal staff, and ongoing support services.
- **Risk and accountability,** like liability mechanisms, governance structures, and human-in-the-loop system design.
- **Financial value,** such as cost over the lifecycle and real ROI expectations.
- **Vendor experience,** including relevant experience with similar AI solutions.

Contract clauses

Who: Procurement teams
What: Contract principles

Your contract clauses will depend on your specific situation, use case, and needs, and if you are buying off-the-shelf, customized, or building AI. If you are buying an off-the-shelf solution, your team can conduct a gap analysis to understand to what extent the vendor's standard license terms comply with your requirements. If you are buying something that is customized or build, then your team has the chance to put your requirements into the contract.

Here, we provide an overview of some types of commercial and technical requirements to help jumpstart your thinking on what clauses to consider. These are not intended to be comprehensive lists, or serve as legal advice and should be complemented by your organization's standard technology procurement clauses.

It's also important to remember that contracts are just one aspect of managing risk. A good contract is short and clear, and should be a tool for trust building and collaboration, versus a long legal document all about transferring risk onto the other party. Even the most stringent contract is not a failsafe for managing risk. Be aware of what is feasible for vendors to agree to around issues like transparency and accuracy. Unrealistic expectations can lead to prolonged contract negotiations or failures.

Model contract clauses and additional resources

Several governments non-governmental entities have created model contract clauses, or are in the process of drafting them, including:

- The updated EU AI model contractual clauses
- The Australian government's AI model clauses

- The Society for Computers and Law's Artificial Intelligence Contractual Clauses
- The GovAI Coalition's template vendor agreement
- The Center for Inclusive Change's AI Procurement: Essential Considerations in Contracting

Additional resources include:

- The Comprehensive AI Contract Guide from World Commerce and Contracting (WCC) with key contractual obligations, negotiation areas, and success factors specific to AI procurement in government contexts
- WCC's Most Negotiated Terms research which covers the key terms that get negotiated and disputed in contract from both a buyer and supplier perspective.
- The AI Procurement Lab's Procurement Tools for Buying AI

Potential commercial requirements include:

- **Licensing and use,** depending on the scope of your contract, to confirm licensing model and renewal terms.
- **Data use**, to align on how the organization's data will be used and allow for evolutions in security standards.
- **Intellectual property (IP) rights,** to outline who owns and will retain ownership over what.

Potential technical requirements include:

- **Performance,** to enable regular assessment of how the solution is performing in terms of accuracy.
- **Testing and validation,** to enable milestone-based payments tied to testing and acceptance.
- **Human oversight,** to provide risk mitigation and promote human-in-the-loop practices.
- **Data quality and provenance**, to document training data and sources and who is responsible for data quality, improvements, and due diligence.
- **Ongoing maintenance,** to ensure the ongoing maintenance of the solution.

- **Security compliance,** to ensure the organization can configure and run error reports, and set security standards
- **Transparency and auditability,** to provide compliance reports and feasible auditability depending on the technology.
- **Responsible AI compliance**, to ensure compliance with key AI policies as agreed upon by the organization and the vendor.
- **Exit and transition**, to ensure clear data export and transition terms.

If compliant with your legal context, you may also consider a requirement around **contract transparency** to enable you to make the full or part of the contract publicly available. This transparency can also support knowledge sharing across public entities and other jurisdictions.

Contract management for AI technology

Who: Procurement, project, and technology teams
What: Tips for contract management

While an effective procurement process and clear, simple contract can help set your internal and vendor teams up for success, the work continues in contract management.

Good contract management is critical to get the most from your AI contracts, secure good value for money, and to keep the project within scope. A lot of value can be gained or lost during the contract management phase, which is why close alignment between the procurement, project, and technology teams matters.

Contract management tips

- **Focus on communication with your internal technology and project teams.** AI technology requires ongoing maintenance. This has cost, audit, and security implications. For example, with flexible services, it can be easy for an internal team member to spin up a report or monitoring service and unknowingly go over cost. Sometimes technology teams can also accidentally set off vendor security alerts by not going through proper channels with the vendor when conducting vulnerability assessments, leading to service interruptions. Open and clear internal communication and a shared understanding of boundaries with internal teams can help you keep your project within scope, maintain good vendor relationships, and maintain smoothly operating services.
- **Fit your vendor management approach to the size and scope of the contract.** How often you check in with a vendor will vary. Buying development services is different from buying access to a service, and a large, high-profile project is different from a simpler and lower-risk project. No matter the scope, your check-in meetings should help gauge if you are continuing to get good value from the contract.

Contract management: Key questions

Procurement objectives

- Manage contract scope
- Ensure swift payment

Organizational objectives

- Implement and refine the AI solution
- Conduct change management and train staff to support solution adoption and scaling
- Prepare for decommissioning

Team	Questions
Procurement	**Operational considerations** • Is the vendor meeting our milestones? • Are we paying the vendor in a timely manner? • Do we need to issue a change order? **Innovation and scalability** • How will we support scaling what works, while allowing flexibility for departments to innovate? • What structures will support knowledge sharing and solution reuse across teams or regions? • How will we stay informed about AI capabilities and vendor offerings?
Agency or department buyer	**Operational considerations** • What training and support will we offer to end users, especially in high-risk or public-facing roles? • How will we test and evaluate deliverables during the development or implementation phase? Who will be responsible for these activities? • How satisfied are users? • How will we document and share AI learnings and best practices across the organization? • Is the payment model working out as expected?

	Innovation and scalability ● How will we enable safe experimentation (e.g., pilots, sandboxes) while maintaining appropriate controls? ● How will we engage with the wider public sector, academic, or civic tech community to share best practices, create spaces for reflection and learning, or co-develop tools?
IT/Data analytics	**Innovation and scalability** ● What frameworks will we use to evaluate, iterate on, and scale successful AI pilots? ● How can we support interoperability between different AI platforms or ecosystems over time? ● Is there room for improvement in the solution? ● What could have been done differently?

Key questions to ask: Overview

Who: Procurement, project, and technology teams, as well as leadership
What: Questions to ask during the AI procurement processes

This is an overview of key questions featured in this guidance to help your organization make informed decisions and implement AI solutions responsibly and effectively. These questions are drawn from the **AI readiness**, **Building your procurement strategy**, **Solicitation tips**, and **Contract management** sections.

Successful AI procurement is a team effort, so as a procurement official, you are unlikely to have all the answers. However, understanding generally what questions need to be asked, and when, can help you facilitate internal conversations and ensure that the right experts are involved in decision-making.

The questions have been organized by the team member who may be best positioned to lead the inquiry. As you review these questions, consider your own unique organizational setup and who in your organization might be best equipped to provide insights on each topic. At different moments, you may need to consult IT, legal, data analytics, project management, or business units to get the answers you need, as well as consult outside stakeholders who advocate or represent the end users.

Not all questions will be relevant for your context. These questions are most relevant for custom or build purchases. For example, if you are buying off-the-shelf AI, procurement will mostly focus on aspects like licensing and access management, while if you are customizing or building software, the procurement team will work more closely with the project leaders and IT to tackle more complex issues relating to the vendor relationship and technical solution.

AI readiness: Key questions

Procurement and organizational objective
Lay the groundwork for overall good AI procurement through policies and practices

Team	Questions
Procurement	**General** • Can we use agile methods for our IT procurements? • How can our procurement processes support strong collaboration between technical teams, procurement, legal, and policy stakeholders? • What does the market for this product or service look like? How can I engage them, both to get input into the best specifications and to ensure a good, competitive process? **Procurement risks** • How do we guard against vendor lock-in and ensure that contracts include clear exit strategies and data portability? • What protections do we need in place to clarify data rights, intellectual property ownership, and the handling of sensitive information? • What is the state of our data, and how much does that need to be considered in the procurement strategy?
Agency or department buyer	**General** • How do our overall organizational mission and goals align with our AI initiatives? • What are our internal capability gaps?
IT/Data analytics	**General** • Who inside the organization can take the lead on AI and connect the dots across teams? • What governance structures and frameworks do we need for responsible and ethical AI adoption? • Do we have the necessary AI literacy and understanding across relevant stakeholders?

	• How will we manage internal change and encourage safe adoption of AI tools? • How will we ensure inclusive design and testing of AI tools to reflect diverse public needs?

Building your procurement strategy: Key questions

Procurement objective
Develop a procurement strategy that fits the overall organizational strategy and the needs of the specific procurement.

Organizational objective
Identify your needs for your specific procurement

Team	Questions
Procurement	**General** • Are we engaging the right internal and external stakeholders in our strategy development process? • Is our budget realistic, given ongoing maintenance needs?
Agency or department buyer	**Model and system architecture** • What is the problem we are trying to solve? • What are our needs and use cases? • How will this project impact end users, including public sector staff, frontline workers and communities or service recipients? • What are the expected outcomes of the solution we want to acquire? • How might we safely test our assumptions, approach, and AI tool options before launching the formal procurement process? **Operational considerations** • Which role or team will be accountable for managing risk and compliance? Do we have the internal capacity for this oversight?
IT/Data analytics	**Model and system architecture** • Is AI the right tool to meet our needs? What deterministic and rule-based solutions might deliver comparable results to generative

	AI for our use case? • If we pursue AI, what type of model best suits our specific use case (e.g., general-purpose vs. specialized models)? • What integration capabilities do we need with existing systems? • What is our internal capacity to maintain the tool? • Where will the tool be hosted? Will we need to acquire additional processing or storage capacity? • What are environmental considerations?

Solicitation tips: Key questions

Procurement and organizational objectives

- Decide on your procurement procedure
- Draft your solicitation

Team	Questions
Procurement	**Project needs** • How might we reflect the project team's needs in a RFP? • Can the vendor demonstrate a roadmap and flexibility that aligns with our evolving needs? • What level of vendor support or service guarantees will we need if something breaks or goes wrong? • Will the vendor support integration with other AI tools or models we may adopt in the future? **Operational considerations** • What kind of selection criteria should be used to rank solicitation responses? • What commercial or pricing models are appropriate? • Are the right people in the room for this procurement? **Risk and compliance** • How will we ensure compliance with current and evolving AI-related legislation, procurement rules, and standards? • How will we manage and track AI-specific costs (e.g., compute, storage, API usage)?

	• How much risk does vendor lock-in pose in our context, and what mitigations are practical?
Agency or department buyer	**AI solution** • Do we know what "good enough" model response performance looks like for our specific use case, and what metrics will we use to measure AI success and ROI (qualitative and quantitative)? • What are our fallback strategies if the model gives low-quality or inaccurate responses? • How will users interact with the AI solution? If there is a platform or dashboard, who is responsible for developing it? • Can we retrieve or export our data and models if we switch vendors? **Data strategy and management** • Do we and our vendor partners understand the quality, completeness, and relevance of our data? • What role will the vendor play in helping us prepare or transform data for effective use? • How will we manage data ownership, licensing, and usage rights? • What safeguards will the vendor provide to ensure that sensitive, personal, or classified data is handled in line with legal and policy requirements? • Will the AI system retain or learn from our data inputs? Under what terms, and with what controls? • Is the algorithm specialized for our organization? If so, will we have input into how the algorithm is fine-tuned or retrained using our data? What additional data sources do we need to get the best results? **Operational considerations** • How might we approach AI implementation in stages, such as through a pilot or sandbox? • How will we measure and monitor model performance over time? • What policies will govern human-in-the-loop oversight for critical decision-making? • What mechanisms will we put in place for continuous performance monitoring and system health checks? • What accident response mechanisms will we put in place? What processes will we follow to learn from accidents and improve their prevention and the management of associated risks and harms? • If we are involved in retraining or fine-tuning the model, what internal processes will we establish for prompt engineering and optimization?

	• Who owns operational responsibility for uptime, error handling, and incident response? • What logging and audit mechanisms will be used for AI inputs and outputs? • How will we ensure traceability and version control for prompts and model configurations? **Risk and compliance** • What processes will we establish for bias detection, mitigation, and fairness audits? • What controls will be in place for data privacy, consent, and security (at rest and in transit)? • How will we meet model explainability requirements, especially for high-impact or regulated use cases? • How will we manage IP and copyright issues relating to AI-generated content? • How will we assess and audit third-party AI providers for security, ethical alignment, and accountability? • What safeguards will we implement against hallucination, misuse, or reputational risk in public-facing services? • How will we address bias detection and adversarial robustness?
IT/Data analytics	**Project needs** • What are our requirements for model transparency and explainability? • How will we measure and monitor model performance over time? • What are our requirements for model updates and maintenance? • What level of customization do we require? • What are the latency, reliability, and availability requirements for our applications? • Are there specific regulatory cybersecurity requirements, such as around data residency or geographic deployment? **Integration** • How will the AI solution integrate with our existing systems? • Does the vendor support open standards and interoperability (e.g., prompt formats, model interchange)? • How do open-source and proprietary models compare in terms of cost, performance, and vendor support and accountability? • Are there clear exit strategies or migration paths if we outgrow a particular vendor or need to switch platforms?

	Data strategy and management • How will we jointly manage data governance – including storage, audit trails, and compliance with data retention and deletion policies? • Are there opportunities to collaborate with vendors on domain-specific dataset creation or fine-tuning? • What processes are in place for anonymizing or redacting data before it enters AI systems? • How does the vendor support discoverability and metadata management for both inputs and outputs? • How will data flows be documented, monitored, and audited across the lifecycle?

Contract management: Key questions

Procurement objectives
- Manage contract scope
- Ensure swift payment

Organizational objectives
- Implement and refine the AI solution
- Conduct change management and train staff to support solution adoption and scaling
- Prepare for decommissioning

Team	Questions
Procurement	**Operational considerations** • Is the vendor meeting our milestones? • Are we paying the vendor in a timely manner? • Do we need to issue a change order? **Innovation and scalability** • How will we support scaling what works, while allowing flexibility for

	departments to innovate? ● What structures will support knowledge sharing and solution reuse across teams or regions? ● How will we stay informed about AI capabilities and vendor offerings?
Agency or department buyer	**Operational considerations** ● What training and support will we offer to end users, especially in high-risk or public-facing roles? ● How will we test and evaluate deliverables during the development or implementation phase? Who will be responsible for these activities? ● How satisfied are users? ● How will we document and share AI learnings and best practices across the organization? ● Is the payment model working out as expected? **Innovation and scalability** ● How will we enable safe experimentation (e.g., pilots, sandboxes) while maintaining appropriate controls? ● How will we engage with the wider public sector, academic, or civic tech community to share best practices, create spaces for reflection and learning, or co-develop tools?
IT/Data analytics	**Innovation and scalability** ● What frameworks will we use to evaluate, iterate on, and scale successful AI pilots? ● How can we support interoperability between different AI platforms or ecosystems over time? ● Is there room for improvement in the solution? ● What could have been done differently?

Case studies

Supporting AI-readiness in Australia

Who: Australian Government's Digital Transformation Agency (DTA)

What they did: Develop policies and tools to support effective and safe use of AI technology, including:

- Policy for the responsible use of AI in government
- Technical standard for government's use of artificial intelligence
- AI model clauses

What worked well: The DTA meets staff where they are by directly embedding their guidance on good procurement practices for AI technology in go-to resources like their Digital Marketplace, 2 panel agreement, and the BuyICT request for quote procurement workflow. Although the DTA also promotes its materials in other ways, such as through blog posts, they've seen that the tactic of integrating the guidance in something that agencies use on a regular basis to be a successful way to support adoption.

DTA also keeps their policy documents live, so that they are able to update them on a regular basis and stay current with the latest policy and technology developments.

What was harder: DTA encourages bringing in procurement and technical teams as early as possible in the procurement process. However, sometimes DTA sees situations where agencies are bringing in these stakeholders late and are in a rush to get to market. This can result in extended contract negotiations that could have been avoided with cross-departmental engagement.

Lesson learned: Applying AI technology in low-risk internal use cases can deliver valuable ROI, while generating useful learning and experience with AI technology.

Outcome-driven contracting success in the United States

Who: Centers for Medicare and Medicaid Services (CMS)

What they bought: What they bought: AI/ML services to streamline operations, as well as some maintenance of existing legacy tools, from Skyward IT Solutions, LLC; the resulting tools include a chatbot for the CMS' Office of Information Technology team that has reduced the time it takes to draft tickets with a user story from 15-20 minutes down to one minute, as well as a procurement tool for the Office of Acquisition and Grants Management that enables faster labor category analysis, which is now saving the CMS $3 million a month.

Procurement procedure: Challenge-based procurement through an open and competitive process

What worked well: The high-level, open-ended, outcome-driven contract focused on problems rather than dictating solutions. Given the rapid pace of AI development, there are now AI capabilities that were not available when the RFP was drafted or the contract was signed. The flexible contract has allowed the vendor to deliver the most cutting-edge solutions possible to CMS.

What was harder: Institutions are still learning about AI, how it works, and what is required to get the most from this technology, and sometimes building this awareness falls on the vendor. For example, there is still a lack of understanding of the value of a specialized versus generic AI solutions. Generic models can be useful, but are not a one-size fits all solution. Narrow government use cases can require highly specific knowledge management, and be a better fit for specialized solutions.

Lesson learned: For a contracting approach like this to be successful, the government must be willing to trust their vendor partner.

Technical annex: GenAI requirements and evaluation criteria for cloud service providers

Who: Procurement teams
What: Best practice solicitations

Our guide aims to help you make your AI procurement a success – without learning every technical detail. But, if you're ready for the deep end, or if you have a colleague with IT expertise, this technical annex provides sample requirements, evaluation criteria and scoring tables to reuse or adapt for a more detailed and technical procurement.

Organizations often rely on Cloud Service Providers (CSPs) when implementing GenAI solutions. This is because CSPs offer unique advantages, including access to the massive computational resources needed for GenAI that would be cost-prohibitive for individual organizations to invest in and maintain on their own. CSPs are also responsible for ensuring that their infrastructure is safe, secure, and up-to-date, which can free up organizations to focus on achieving their core goals, without managing complex infrastructure upgrades.

Download the full technical annex with sample requirements, evaluation criteria and scoring tables.