



Accelerating the value of AI in government organizations

AI can deliver real value today—if you're willing to start small.

As government organizations at all levels, federal, state, and local, are asked to do more with less, artificial intelligence (AI) offers an enticing answer. AI has been put on a pedestal as something profoundly different than previous innovations, a transformative solution poised to ignite a revolution. This often has leaders of government or higher-education organizations searching for a game-changing use case that can live up to those lofty expectations.

Some may believe that realizing value from AI requires a “moon-shot” effort—a multiyear, multimillion-dollar project with sweeping capabilities involving complex agentic AI systems, custom large language models (LLMs), and extensive system integrations. Such projects often require major infrastructure upgrades, with robust data pipelines, advanced GPU clusters, increased network bandwidth, dedicated servers, edge computing capabilities, enhanced network architectures and security capabilities, and specialized AI operations (AIOps) frameworks.

Technology aside, projects of this scale typically require significant investments of time and effort to address internal legal, policy, and process requirements. Once those hurdles have been cleared, attention must then turn to the human element, with communication, training, and other change-management initiatives. The magnitude of the effort often has agency leaders looking for the perfect project that justifies the expense and attention.

Make no mistake: we're not suggesting that projects of this scale are without merit. But the perception that AI delivers value only when you “go big”—combined with top-down pressure, including the White House's “America's AI Action Plan”¹—may be

Why modern government is important

Government agencies in the US must modernize in order to keep up with changing user needs, regulations, and health and public safety requirements. Leaders of modern governments rethink business processes and service delivery models to more effectively achieve their mission. This article is one of a series that features how modernizing affects the government workforce and the user experience, improves security and public trust, and accelerates the digital journey. KPMG team members offer insights intended to help guide governments in their modernization efforts to encompass all processes, technologies, policies, and the workforce so each works together to create connected, powered, and trusted organizations.



pushing some to plan these large-scale AI projects without first considering smaller, simpler alternatives. As organization leaders either struggle to find the perfect use case or wait for it to be approved, funded, or completed, AI's potential sits untapped within the organization.

¹ “Winning the Race: America's AI Action Plan,” The White House, July 2025.





Focusing on value

While it's true that AI is revolutionizing the way our clients—and our own teams, for that matter—achieve their mission, the implementation and application of AI is no different than that of any other capability—it's just another technology tool in the transformation toolbox. If there's one thing we've learned over the last two decades, return on investment (ROI)—value—doesn't come from deploying technology, no matter how revolutionary or powerful it is. ROI comes from organizational transformations enabled by that technology that advance the mission more effectively or efficiently.

Importantly, the greatest ROI isn't always found in projects with the biggest budgets. There are dozens of opportunities for governments and educational institutions to take advantage of AI today to drive real value if they are willing to start small.

Many enterprise platforms government organizations have already adopted have extremely powerful AI capabilities built in. There are off-the-shelf solutions that can be quickly added to extend these capabilities further. There's an incredible amount of value to be found in the smaller, simpler business advancements that these AI capabilities can unlock. As they say, home runs may fill seats, but base hits win games. And if we may mix our metaphors, as organizations swing for the fences, this low-hanging fruit often sits unpicked.

Moreover, starting small isn't just about grabbing some quick wins—it's about building momentum and laying a foundation, both technologically and organizationally, that can be expanded upon. Continuing our baseball analogy, think of it as "AI spring training," helping to get the organization in shape for bigger things to come.



Four ways to AI-driven value by starting small

Organizations that are realizing value the fastest are starting with a rapid assessment of use cases with a deep dive into the functional roles that can be augmented or supported with AI to unlock capacity. You may even find opportunities to enhance processes that already incorporate traditional AI or automation by adding generative AI (GenAI) to the mix. If you prioritize these opportunities based on complexity, risk, and value, you'll likely find at the top of the list many that require no custom work at all.

1 Leverage AI capabilities in your existing enterprise platforms.

By far the fastest, easiest, and most cost-effective way to introduce AI is to take advantage of the AI capabilities embedded in the systems already in use today. Many governments and universities already license enterprise platforms (Microsoft 365, ServiceNow, Salesforce, Workday, Google Workspace, etc.) that include platform-native AI and GenAI features that can be activated incrementally with existing data access and identity controls.

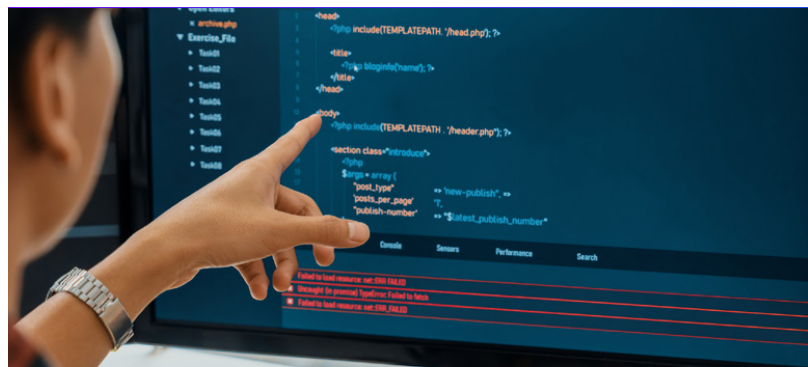
Start small with foundational pilot use cases such as automating meeting notes, writing or proofreading email drafts, or summarizing knowledge base articles. Agency publications can be checked for alignment with mission and consistency with organizational positions or regulatory requirements. These capabilities can help you streamline document creation and assist with case management, reducing manual overhead while avoiding any costly, large-scale transformation efforts. They also help your workforce become accustomed to using AI to augment their day-to-day work through productivity applications—before embedding it more deeply into their core workflows.

2 Use process automation plus AI to address bottlenecks.

Robotic process automation (RPA), which has long been used to handle repetitive manual tasks, has become much more powerful with the introduction of AI. RPA has historically been limited to highly structured tasks given its reliance on rigid rules. But AI-powered RPA agents can understand context, make decisions, and adapt automatically to new tasks or circumstances, opening up a host of new opportunities for their use.

The growing adoption of the model context protocol (MCP) is facilitating this advance by providing a standardized way to connect AI models to different data sources and software tools. As the backend systems behind the interfaces that RPA bots are designed to crawl are themselves increasingly being integrated with LLMs, AI-powered RPA agents are using MCP to tap into this knowledge to increase their capabilities further.

Easy-to-implement, off-the-shelf tools now embed AI models for smart form extraction, summarization, and email interpretation. Start by focusing on simple service requests, internal approvals, or document-handling workflows that can reduce staff burden and improve turnaround. This provides a foundation for AI-assisted process reengineering down the line—with low initial risk.





3

Deliver value through department-specific use cases.

Instead of broad AI strategies, start with narrow, high-impact use cases that resonate with departmental pain points. For example:

- In state and local government, you can automate FOIA request processing, permitting communications, or citizen FAQs.
- At the federal level, you can prototype AI agents for acquisition, HR onboarding, security, or policy drafting support—all with guardrails in place.
- In higher education institutions, you can use AI to assist in academic advising, course catalog search, or syllabus generation.

4

Introduce lightweight, domain-specific AI agents.

Even if you're already using a large-scale LLM for certain tasks, you may find that small AI agents or copilots tailored to a single domain can be easy to implement and highly effective. Domain-specific models can be less computationally intensive, and they lower hallucination risks associated with general-purpose models.²

Consider using open-source agent frameworks such as LangChain or CrewAI or vendor-specific options to build agents tied to a single system or dataset. Focus on agents that assist with basic but time-consuming tasks such as form completion, document review, or internal help desk support. These can be sandboxed for risk control, deployed iteratively, and integrated into familiar interfaces such as Microsoft Teams, Slack, or existing service portals.

² "3 Bold and Actionable Predictions for the Future of GenAI," Arun Chandrasekaran, Gartner, April 12, 2024.



Laying the groundwork

Although we're advocating for "quick-win" projects, we want to be careful not to oversell the ease of "out of the box." For some, that term may imply you simply flip a switch, change the logo at the top, and you're done. While that may be true in some cases, in many other cases, at a minimum, these software solutions will need to be properly configured to meet your needs.

Given AI's reliance on data, some level of data modernization or system integrations may be required. Existing AI capabilities in a customer relationship management (CRM) system, for example, might not truly deliver value until customer data scattered across multiple, siloed systems can be centralized and consolidated. But projects such as data consolidation are relatively minor, especially compared to the development of large-scale, bespoke AI solutions. Importantly, once completed, they should require little if any ongoing maintenance.

Beyond any such configurations or integrations, there's a certain amount of prep work that's required for even the simplest of AI projects to be safe and successful. These efforts lay the foundation for any larger AI projects that may follow. They can be used to inform your broader AI strategies by helping to build and maintain an agile roadmap that enables you to pivot quickly in the face of rapidly changing policies and technologies.

Specifically, we see three key objectives:

1. Develop governance, usage policies, and stakeholder alignment.

In the rush to implement AI, organizations may overlook critical enablers of success—not only the potential risks, but also the need for sound governance, defined roles, skilled talent, and measurable value frameworks. The Hoover Institution noted that while use of AI in the public sector is quickly gaining momentum, "adoption at this point outpaces institutional safeguards and available training options." It observed that "nine percent of respondents said their organization had no formal policies on the use of generative AI, and another 51 percent said they were unaware whether any such policies existed."³

Clear governance, usage policies, and stakeholder alignment are essential from day one. Start by measuring your AI readiness. Analyze your strategies, trust and governance, workforce skills, project management, and value tracking capabilities. Identify gaps and develop a plan to address them. NIST is actively working to help in this area with AI overlays for its 800-53 controls, which will give chief information security officers (CISOs) much needed clarity on how to treat AI systems—including when AI is added to systems that have already received their authorization to operate (ATO).



³ "Generative AI use in US public sector on the rise, survey by Hoover fellow suggests," Oliver Giesecke, The Hoover Institution, December 2, 2024.



2. Overcome skepticism and empower an army of AI strategists.

It's not unusual for government employees to be both fearful of AI and skeptical that AI can add value to their specific job. AI upskilling is essential to overcome these obstacles.

Anticipate the new capabilities you'll need from your workforce and develop reskilling or upskilling plans for employees. Train staff not only on tools but also on prompt design, responsible usage, and AI literacy, embedding change management into each initiative.

Upskilling can offer more value than just giving people a practical understanding of how to use the technology—it also can help open their eyes to the possibilities of its use. InnovateUS, a project of Northeastern University, for example, is a program to train government employees on how to leverage AI in their jobs. Students are surveyed both before and after completing the program. In a precourse survey, more than half said they had no use for AI in their jobs; but in the postcourse survey, 84 percent said they now saw its applicability and value.⁴

A key takeaway here is that you don't necessarily have to have all the answers for how AI can be used within your agency or department before you adopt it. Empowering employees through training and involving them in AI planning and implementation can create an army of imaginative people who may come up with answers to that question you may never have dreamed of.

3. Prepare data, identity, and security for responsible AI use.

A common mistake many organizations make is jumping directly to the use or development of LLMs or retrieval-augmented generation (RAG) before they've confirmed they can safely and effectively use AI with their data. It's understandable; we all want to get to the fun part as quickly as possible, but it's the equivalent of launching a new aircraft before running the safety checks.

You need to ensure that when AI tools are added, they inherit the correct entitlements and adhere to the proper security boundaries so they access only appropriate and accurate data sources. You may need to enhance data classification, update data governance controls, or modernize identity and access management systems or role-based access control models to achieve this assurance. Many of the AI agents embedded in the enterprise solutions you may already be using—ServiceNow Now Assist, Microsoft Copilot, Google Gemini, or Salesforce Einstein, for example—offer native AI integrations that can be activated incrementally with existing data access and identity controls. AI can also be used to speed the preparation process. Cleaning and inventorying data, for example, can be a tedious and time-consuming process, but one that AI can help dramatically accelerate.

Finally, trust but verify. You must implement logging and audit mechanisms to monitor how AI systems are actually interacting with sensitive government or institutional data to detect any incorrect assumptions or misconfigurations. Simply keeping track of your use cases is essential to help verify responsible usage—and may even be a requirement. OMB M-25-21, for example, requires CFO Act agencies to annually inventory their AI use cases.



⁴ Amanda Welsh, PhD, The Burnes Center for Social Change at Northeastern University.



You are not as unique as you may think

One thing organizations at federal, state, and local levels have in common is that they often believe they are so distinct from any other organization that commercially available software designed primarily for the private sector cannot handle their unique requirements or specialized business processes. They often believe that bespoke solutions are required to handle these idiosyncrasies. They may believe, therefore, that the “low-hanging-fruit” opportunities afforded by off-the-shelf software solutions may not exist—for them.

In our experience helping government and higher education institutions with digital transformation efforts, however, we’ve found that such cases are more the exception than the rule. The enterprise platforms commonly used by these organizations have been designed with leading practices baked in.

If any customization is required, it’s typically just for the “last mile”—to optimize for the specific business process that’s involved, which is true across nearly every organization in every industry. Up to that point, the underlying technology and the fundamental business processes it’s been developed to support, remain consistent regardless of sector or industry. In most cases, we find it’s not only easier to change an organization’s business processes to match what the software does rather than the other way around, but it’s also far better for the organization’s efficiency and productivity to adopt those leading practices.

Many of the business processes these organizations have come to rely on were developed based on the technology that was available at the time, which in many cases was no technology at all. This “process debt” may have accumulated over decades—multiple, manual, and often tedious processes that continue to this day for no other reason than “because that’s how it’s done.” Taking advantage of the AI capabilities of existing off-the-shelf systems is a perfect opportunity to reduce some of this process debt.

Standardization can also facilitate the sharing of knowledge and solutions across agencies or departments. As many organizations face pressure to do more with less, function sharing or consolidation—human resources, for example—becomes an increasingly attractive or necessary option. Simply keeping track of your use cases is essential to help verify responsible usage—and may even be a requirement. OMB M-25-21, for example, requires CFO Act agencies to annually inventory their AI use cases.



How KPMG can help

KPMG LLP has worked with federal, state, and local governments for more than a century, so we know how public sector agencies, education institutions, and healthcare organizations work—and we’re helping them design and implement real AI uses cases today.

We’re experienced, nimble, and flexible. We understand the unique issues, pressures, and challenges government organizations face on the journey to AI adoption. We’ll meet you where you are on that journey and help advance your progress with no agenda other than to see you succeed. We’ll help you leverage the investments you’ve already made to help maximize their value—not try to sell you something new.

We offer clarity and insight. As a trusted advisor, we can help you make sense of everything going on in the highly dynamic world of AI that can impact your mission, from regulatory mandates and governance to emerging technologies. We can help align your efforts with leading practices from both the private and public sectors, and help keep you moving forward quickly with confidence and conviction.

We see the big picture. We can help you anticipate and adapt to the wide-ranging impacts AI can have on your organization, including budgets and financial controls, business processes and operating models, and employee growth and retention. We can help you understand your data—where it comes from, what controls are required, how to maximize value locked in it, and how to share that value across organizations. We can help you harness the power of AI ethically and responsibly with trusted AI principles and governance models for managing risk.



We can help you from strategy through implementation.

Unlike business-only consultancies, our more than 15,000 technology professionals have the resources, the skills and experience, the battle-tested tools and methodologies, and the close alignment with leading AI technology providers to help achieve your vision, quickly, efficiently, and reliably. And unlike technology-only firms, we have the business credentials, subject matter professionals, and public sector experience to help you deliver measurable results, not just blinking lights.

Contact us



Rob Breakiron
Managing Director, Advisory
KPMG LLP
773-710-4521
rbreakiron@kpmg.com



Tom Frame
Principal, Federal Advisory
KPMG LLP
703-286-6888
tframe@kpmg.com

read.kpmg.us/modgov

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us:



kpmg.com

The views and opinions expressed herein are those of the authors and do not necessarily represent the views and opinions of KPMG LLP.

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

© 2025 KPMG LLP, a Delaware limited liability partnership and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.