

Deloitte.



Maximizing the public good: How Generative AI can enhance government programs and services

The technology offers benefits for the public sector but making the most of it and avoiding potential risks takes strategy, governance, and training

By Ed Van Buren, Tasha Austin, and Brendan McElrone

Thousands of terabytes of unstructured data each day can overwhelm even the best of intelligence analysts tasked with tracking threats to national security.

Similarly, benefits case workers in state and federal agencies handle tens of millions of claims and complex situations. Millions of pages of policy documents must be interpreted and adhered to by government workers and residents alike. And billions of dollars must be disbursed accurately in support of government programs, while complying with Federal Acquisition Regulation (FAR) guidelines that run more than 2,300 pages.

These are daunting tasks in the face of open positions, retirements, and budget cuts.

Time-consuming administrative duties coupled with a large caseloads and backlogs contribute to annual turnover of 14% to 22% nationwide.¹

Fortunately, government leaders and employees don't have to go it alone anymore. Generative artificial intelligence (GenAI) provides a path to help government leaders navigate these challenges and deliver program results and mission outcomes. It's the era of government powered by humans and machines.

GenAI can cut the time child services case workers spend on administrative tasks and case management, giving them more time in the field with the clients they serve, improving job satisfaction and decreasing turnover. GenAI can parse unstructured information and pre-generate content for intelligence analysts to review, reducing the time required to draft reports and allowing analysts to focus on more complex and nuanced tasks that require human judgment, such as interpreting the strategic implications of data, verifying the accuracy of AI-generated information, and making critical decisions based on a broad range of intelligence inputs. It can answer questions about federal contractor qualifications, contracts, and bids with the latest information, making procurement processes more efficient and compliant.

GenAI can synthesize data from existing information to create original content, simulate human-like text and speech, and quickly identify patterns in data. All those characteristics have the potential not only to radically change how government personnel do their jobs, but also to improve public service delivery and the lives of constituents through healthier children, safer streets, and better jobs. In fact, [research estimates](#) that AI can contribute 14.5% to North American GDP, equivalent to US\$3.7 trillion, by the year 2030.²

Many leaders at all levels of government are eager to harness GenAI for these and other tasks to enhance services, achieve mission goals more effectively, and augment the workforce. But they must take a cautious approach to avoid potential risks known to be associated with it, such as legal issues, ethics, privacy, and security concerns. As such, this paper outlines opportunities to implement guardrails for GenAI, to help government personnel embrace this powerful new tool while upholding their commitment to their constituents.



The age of Generative AI

GenAI marks the start of a new era for AI.

Earlier iterations of machine learning algorithms analyzed data to make predictions.



GenAI, trained on vast amounts of data, is able to make predictions about what comes next in a way that creates remarkably human-sounding content and reasoning. It can generate novel information in a variety of formats and move between them effortlessly.

Companies of all sizes are adopting GenAI in their operations to improve products, workflows, experiences, and efficiencies, which can contribute to rapid market growth. By as early as 2025, more than 10% of all data is expected to be AI-generated,³ heralding a time of expanded human-machine partnerships. Government is no exception.

GenAI can be built into new and existing IT systems or used as stand-alone tools to support the workforce or customers. Many of these solutions pair GenAI and the human worker, acting as a partner to accomplish tasks more efficiently and effectively.

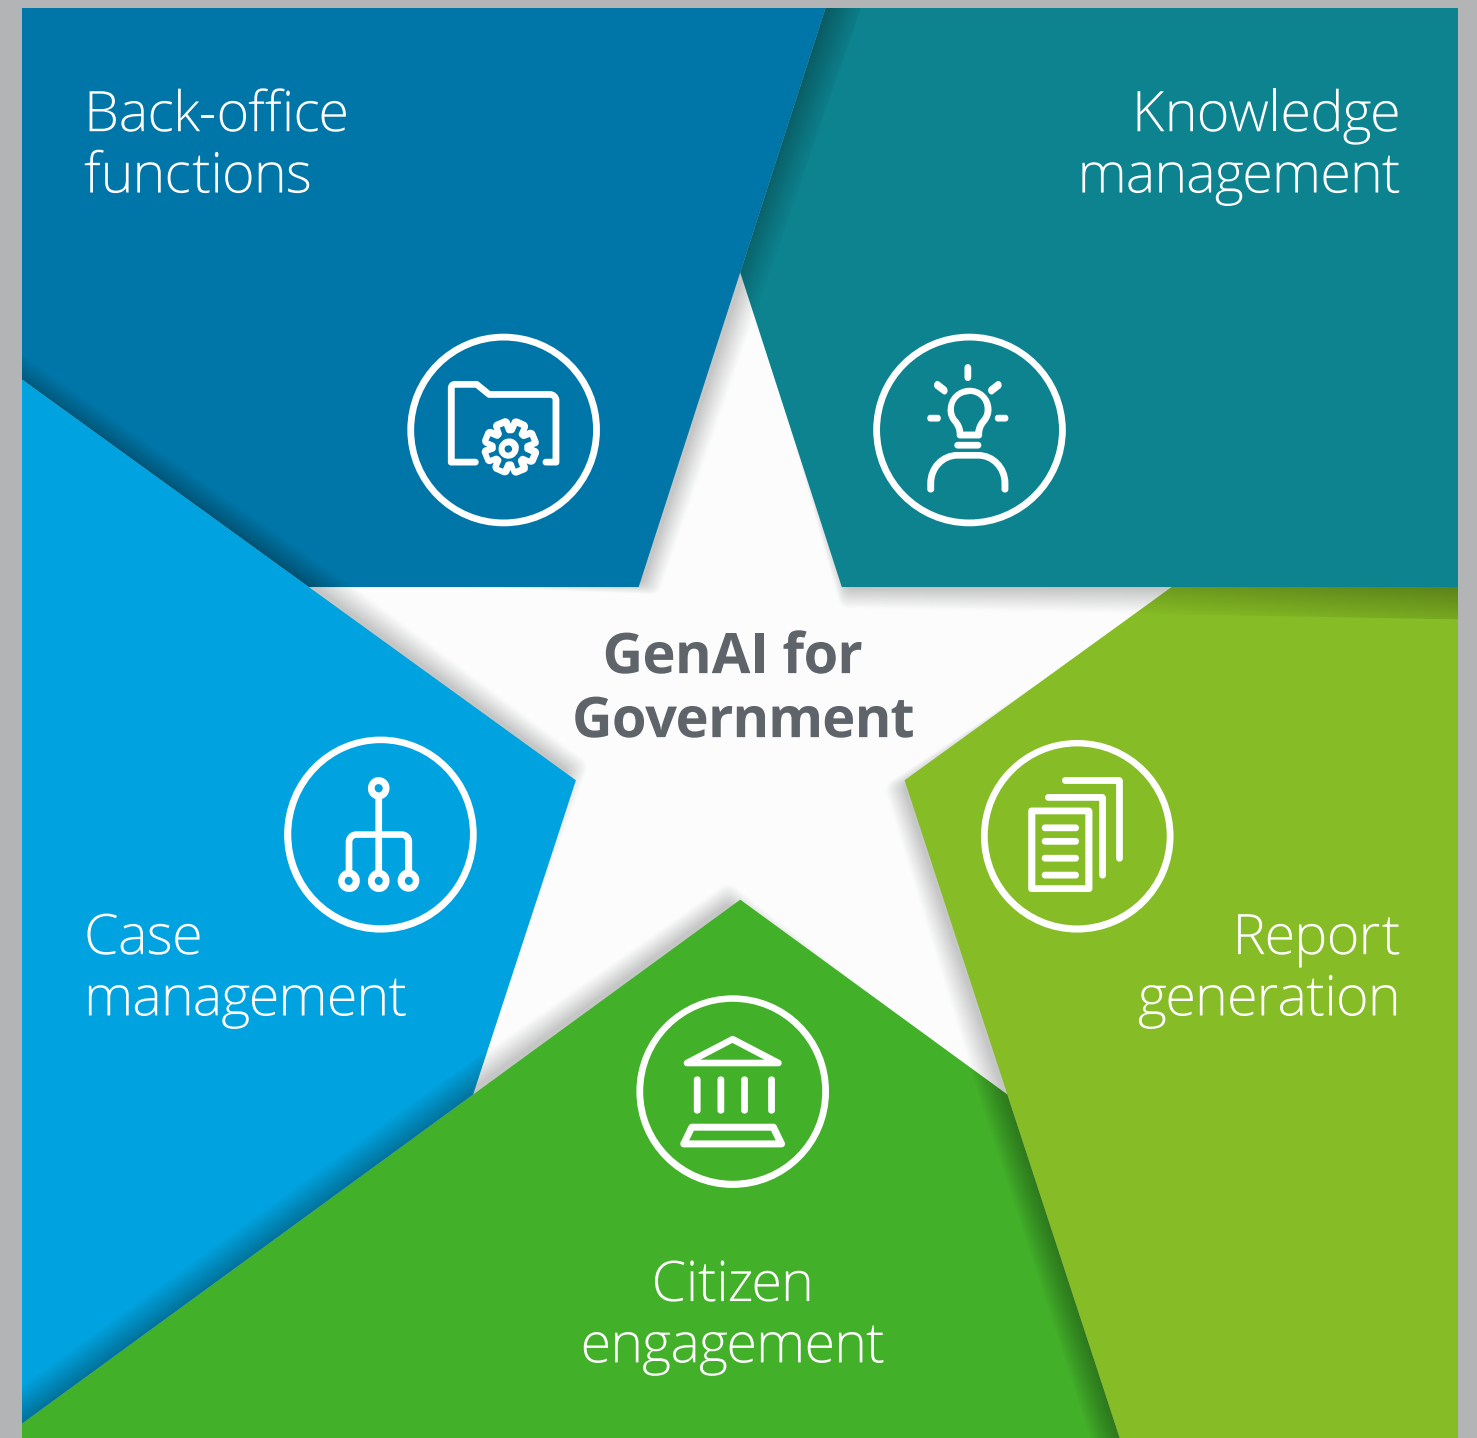


Major use cases for Generative AI for government

GenAI provides opportunities in many areas, but there are a few broad categories that have tremendous and immediate applicability.

Leaders should consider GenAI sidekicks in areas such as citizen engagement, report generation, case management, knowledge management, and back-office functions.

New areas will emerge as technology evolves and leaders become more familiar with GenAI.



Use case example

Potential benefits from integrating GenAI

Streamlining citizen interactions



- Increases the efficiency and responsiveness of government offices, thereby improving citizen trust.
- Allows government personnel to focus on complex queries and issues that need human judgement, such as addressing unique cases that fall outside standard procedures.

Information aggregation



- Provides citizens with accurate, aggregated information in one location, facilitating quicker decision-making and problem-solving.
- Enables health agency staff to focus on healthcare policy-making and program management, by reducing the need for manual data compilation.

Higher education student engagement



- Enhances student engagement by providing immediate responses to queries, leading to better academic performance.
- Frees professors to focus on research, course development, and individualized student mentorship.



Citizen engagement

In late 2021, President Biden signed an executive order directing government agencies to operate more efficiently to respond faster to citizens' needs, whether those needs involve passing through a security checkpoint, claiming retirement benefits, or checking on the status of a farm loan application.⁴ According to the order, paperwork imposed by executive departments and agencies on the public runs in excess of 9 billion hours a year.⁵

How GenAI can help. AI-based chatbots can assist government personnel by streamlining knowledge management and speeding up searches. Similar chatbots and other virtual assistants could provide more options for people interacting with government agencies. The day may come when motor vehicle departments, airports, immigration offices, and other government offices are supplemented with AI-based personal assistants that can understand questions regardless of how they are posed, and answer in kind. Government departments and agencies could deploy GenAI to analyze constituent

interactions and behaviors in order to optimize how services are offered. Interacting with people when, where, and how they prefer could enhance trust in government and ultimately boost citizen engagement.

In the future, agencies could use AI-based chatbots on websites that people consult for information on government benefits and services. For example, a state health benefits agency could use an AI-based tool to search and aggregate information stored in multiple websites or databases and serve the answer in a single location.


GenAI could also improve student engagement in public higher education. AI-based chatbots could provide individualized feedback to students customized based on their performance on previous exams or papers.⁶



Use case example

Potential benefits from integrating GenAI

Report generation for government agencies



- Agency personnel can engage in analytical work such as policy development, strategic planning, and inter-departmental collaboration, thereby maximizing the utility of their specialized skills.

Reviewing reports for decision-making



- Decision-makers can focus on nuanced analysis of summarized insights, enabling them to make well-informed and quicker decisions that could be critical in time-sensitive matters.

Reviewing classified materials for public release



- Intelligence analysts and decision-makers can focus on handling sensitive, high-priority declassification tasks that require human judgement, enhancing national security by reducing the margin of error.



Report generation

Congress mandates that federal departments, agencies, the White House, and the judiciary produce annual accountings, budget summaries, or other reports.⁷ Many state and local government departments and agencies are bound by similar requirements. Compiling the data and drafting the reports takes substantial time and labor.

Agency personnel also spend substantial time reviewing reports to inform decisions or learn the status of projects or initiatives. Reviewing classified materials to determine if they can be made public is particularly tedious, and a wrong decision could affect national security.

How GenAI can help. Using GenAI can cut report production time, reduce errors, and free up agency personnel for analytical work. When GenAI models are fed relevant, conclusive background material, including text, images, or video, and given the appropriate prompts, they can compile comprehensive reports, apply plain language or style guides to them, and cite subordinate content.

For example, those intelligence analysts could feed a photo or video into a natural-language GenAI model trained on similar imagery to automatically generate a plain language report that helps them understand who or what is in the image.

GenAI also can take on tasks that require more complex decision-making, such as declassification reviews.⁸ US State Department personnel used an AI-based machine learning model trained on past declassification decisions to review more documents that didn't pass a certain confidence score threshold and achieved a baseline accuracy of about 97%.⁹



Use case example

Potential benefits from integrating GenAI

Medicaid long-term care support



- Streamlined eligibility determination for benefits, allowing more timely access to services for patients.
- Care providers can focus on improving patient care plans and coordinating additional services.

Public health grant funding management



- Scientists and committee members can devote more time to evaluate the impact and innovation of the research proposals.
- Enables quicker allocation of funds, inviting more researchers to apply and accelerating the pace of public health research.

Law enforcement content analysis



- Reduces exposure to content that can lead to burnout or mental health issues among law enforcement personnel.
- Officers can focus on investigative work that requires human intuition and ethical considerations.



Case management

Case management serves a critical role at both the federal and state level to consolidate, track, and analyze both quantitative and qualitative information to make decisions and achieve desired outcomes across a multitude of government functions. Case managers work with citizens across a variety of services, including benefits administration, employment services, long-term care, mental health services, immigration, and criminal justice, among others, to ensure citizens are receiving the right services at the right time.

However, the high administrative burden of increasing caseloads and tight budgets can sometimes leave these workers' best skills on the table as they find themselves spending more time on paperwork than they do visiting clients. The time they spend on critical documentation and administrative activities can further exacerbate these high caseloads and contributes to high turnover.¹⁰ With case workers spread thin, recipients of these services may face a lack of clarity around their case status, or potentially miss out on more personalized and timely services.

Case managers aren't the only public servants who must balance meeting the needs of the public with adhering to mandated reporting requirements and other stipulations. Grant-making, permitting, records administration, and government aid program personnel face similar pressures. Whenever a claim or request is made to an agency, it should be tracked from beginning to end while retaining relevant information and adhering to applicable guidelines and regulation. For example, the Freedom of Information Act (FOIA) allows citizens to submit record requests to over 100

federal agencies, which are required to respond to the requestor within 20 days.¹¹ With FOIA requests hitting a historic high of 1.1 million requests in 2023 alone, agencies are tasked with a significant administrative load in addition to their core mission functions.¹²

How GenAI can help. Personnel can use GenAI to quickly look up and summarize key file data from relevant regulations, policies, and procedures. AI models can be built to provide guidance and policy suggestions at the point of decision making and be deployed across the case management lifecycle. GenAI models can be programmed to apply plain language and work with easy-to-use interfaces, so personnel don't need significant training to use them. The technology can be integrated into other common agency applications to enhance analytical rigor and reduce staff workloads.

One state Medicaid program is investigating deploying GenAI to augment long-term care providers' efforts to support patients' eligibility for benefits. In this case, care providers could ask a chatbot backed by an AI-based system questions about state benefits policy parameters and instantly retrieve relevant information.



GenAI also has implications for case management in public health. A government health research organization gets tens of thousands of grant funding applications a year. Determining which requests to recommend for approval can take input from thousands of scientists and thousands of meeting hours. The organization is investigating GenAI-based models that review committees could use to distill data from grant applications to speed up the decision-making process, which could open the door for more researchers to apply. GenAI could be used in other aspects of grant program case management, including to predict how likely a grantee is to succeed, manage deliverables within approved grants, and detect fraud.

In addition, GenAI tools could relieve public sector workers of tasks that could negatively affect their mental health. Law enforcement agencies could use it to analyze and describe explicit content or otherwise gruesome images, relieving their personnel of a task that historically has led to high burnout and turnover.

Use case example

Potential benefits from integrating GenAI

IRS tax code navigation



- Auditors can quickly interpret tax regulations, enabling faster and more accurate audits.
- Allows for more targeted taxpayer education and compliance efforts.

Federal contracting guidelines



- Acquisition leaders and contractors can swiftly comprehend federal guidelines, reducing delays and errors in procurement.
- Federal agencies can audit contracts more efficiently.

Learning and development



- Improved training quality tailored to different roles and learning styles.
- Speeds up the development of training programs, enabling quicker employee onboarding.

Knowledge base enhancement



- Facilitates the identification of key trends or insights for strategic planning.
- Reduces the time needed for data analysis, freeing analysts to focus on interpreting results.



Knowledge management

When the US government created the modern income tax system in 1913, the entirety of the income tax code fit on a single page. In 2021, the tax code contained close to 10,000 sections across 174 pages.¹³

The tax code exemplifies the mountains of data that government departments and agencies produce on a regular basis, creating challenges for personnel who must use and manage the information.

Government data may be unstructured—not organized in a pre-defined way—including information from search engine queries, social media, real-time streaming, and internet of things (IoT) sensors, making it difficult to find and analyze, especially for untrained users. By one estimate, 80% to 90% of data growth is associated with unstructured data.¹⁴

Other government data may be structured but locked inside databases with complicated interfaces or that don't allow for ad hoc queries. Users may need to know different types of queries to search data stored in different formats, such as text, images, or video. Depending on the situation, **data may be distributed across multiple agencies**. For instance, data relevant to natural disaster response and recovery could come from the Federal Emergency Management Agency (FEMA), state and local emergency management agencies, federal and state transportation departments, state health agencies, and law enforcement. The volume of government data has increased rapidly in recent years and shows no signs of slowing, adding to search-related challenges. On top of that, agencies may **not have staff with the skills needed to interact with data** in its present form.

How GenAI can help. GenAI provides a natural language interface that simplifies performing complex searches on unstructured as well as structured data. With user friendly interfaces, agency personnel don't need to know how to code or have data science training to do searches. The technology enables common functions such as text manipulation, summarizing and headlining, image search and clustering, and tabular data filtering and visualization.

The IRS, for example, could employ AI-based tools for auditors to use to look up specific sections of the tax code.

An agency could direct a GenAI model to ingest the contents of a federal or state regulation and create a lightweight interface that lets people do natural language look ups for contract guidelines. A GenAI model tied to the FAR guidelines for government contractors could respond to natural language questions about allowable business expenses and the model could provide the limits, as well as the source of the information in the regulation.

In learning and development, agencies could use GenAI to capture the tacit knowledge of retiring employees and create standard operating procedures to pass relevant bits down to new workers on specific job tasks. For example, the agency could build a collection of training videos by recording experienced employees performing specific tasks and prompting a GenAI-based tool to write descriptions of the actions shown in the videos. An agency could also interview retiring workers to capture their knowledge of the organization and use GenAI to personalize learning materials based on those interviews to a new employee's role, experience level, and learning style.

GenAI also can be used to summarize vast amounts of information and then search for insights or trends in it, or to create new entries in knowledge databases or expand existing ones based on new information.



Use case example

Potential benefits from integrating GenAI

Finance functions



- Streamlined financial closing processes to promote accuracy and compliance.
- Enables faster financial decision-making based on analytics.
- Reduction in financial errors or discrepancies, requiring less manual auditing.

Recruiting and HR



- Allows HR professionals to spend more time on employee engagement and talent development.
- Accelerates the hiring process by automating initial screening steps.
- More efficient updating of employee records, schedules, and benefit selections.

Procurement



- Enables proactive supply chain adjustments based on real-time performance metrics.
- Accelerated procurement cycles, ensuring that departments get what they need more quickly.
- Strengthens relationships with vendors through better communication and performance metrics.

IT and cybersecurity



- Enhanced capability for preemptive cybersecurity measures, minimizing breaches.
- Reduced IT downtime due to proactive maintenance, leading to increased operational effectiveness.
- Allows IT staff to focus on strategic technology planning and cybersecurity strategy.



Back-office functions

Nearly 70% of federal employees work in some type of administrative capacity, which can include managing files, communications, scheduling and similar tasks.¹⁵ Tasks in these back-office functions can often be some of the most time consuming as well. On average, Federal workers spend 10% of their time documenting and recording information. Across the entire Federal workforce, that comes out to be more than 400 million hours per year, the single most time-consuming task across the entire Federal government.¹⁶

How GenAI can help. AI could save time by automating functions that augment staff work, enabling them to focus on more strategic tasks.

In finance functions, GenAI can build on existing AI and machine-learning functions to provide end-to-end “touchless” processes and speed up decision making. The technology has applications for managing accounts payable invoices, vendor payments, collections, predictive forecasting, report analytics, and optimizing spending. It can automate monthly, quarterly, or yearly financial closing processes by pulling data from various sources and reconciling discrepancies and generating reports, as well as applying analytics to identify trends. It can also be used for data self-healing, automatically detecting and correcting inconsistencies in financial data or alerting people to them.

For recruiting and HR, governments can implement GenAI for tasks such as writing job descriptions, screening job applications, and onboarding employees. Many organizations in and outside of government have already upgraded to

cloud-based systems so employees can update their own work and benefits information online. GenAI could take that a step further by letting workers interact with AI-based chatbots to update records, request work schedules or time off.

In procurement, AI-based tools can power autonomous sourcing, analyzing offerings from existing vendors to determine the best match for a given need, and then ordering the parts or services accordingly. It could automate complex processes related to vendor relationships and supply chain operations, including automatically generating requests for proposals and analyzing potential vendors’ responses, as well as tracking vendors’ performance against contract terms, and flagging contracts that are up for renewal. It could enhance workers’ management of contracts, inventory, and purchase orders; procurement operations, including demand forecasting; and relationships with suppliers and other partners.

Security threats against government operations at all levels have escalated¹⁷, from cyberattacks on federal agencies to physical attacks on power grid substations. GenAI can help security teams enhance their IT and cybersecurity processes by proactively identifying and preventing potential issues. Additionally, the technology can enhance service desk and incident management by aiding in the prioritization of and response to reported issues.

GenAI can also reduce the time and effort of software development by performing autonomous coding, writing code based on the requirements it is given or assisting a human coder. It can help to support and improve IT operations by maintaining technology applications and infrastructure by analyzing patterns and trends in related performance data to predict potential failures before they occur. This can allow for proactive maintenance and reducing downtime.



Key considerations for implementing Generative AI

Despite its seemingly endless possibilities for improving operations, services, and engagement, **the technology is still new and comes with risks.**



The most significant are related to legal issues, ethics, privacy, and security. Like other forms of AI, GenAI raises concerns about managing bias, transparency in where the data used in AI models comes from and how it's used, and its potential impact on the workforce.

It has attracted new types of threats and security challenges. All these risks and threats exist in the private sector, but may be magnified in the public sector, where government agencies must preserve the trust of the people they serve.

At the federal level, steps are being taken to strike this balance between trust and productivity. In October 2023, the Biden Administration signed an Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence, reflecting this balance of cautious enthusiasm at even the highest level of government. The Executive Order states, "Harnessing AI for good and realizing its myriad benefits requires mitigating its substantial risks. This endeavor demands a society-wide effort that includes government, the private sector, academia, and civil society."¹⁸ The order includes a multitude of provisions across sectors, including the development of new AI industry standards, and

directives to support innovation, worker protections, and bias mitigation, among others. The Executive Order underscores the importance of instituting guardrails for AI training, governance, and support in the public sector.

However, significant pressures can make implementing this technology difficult for some government organizations. For example, GenAI can also be cost-prohibitive because it requires substantial computing power to run models and perform the calculations needed to respond to user prompts.

Still, GenAI holds so much promise for so many different aspects of public sector work that departments and agencies should find ways to mitigate risks and overcome other challenges to adopt it.

Here are some key factors to consider when implementing GenAI based on the use cases we've identified.

Using the technology

GenAI components should work in harmony for the technology to maximize its potential: large language models trained on the desired tasks, relevant data, and well-engineered prompts so the models provide the intended responses.



Partnering humans and machines.

Solutions should be designed to take advantage of the unique abilities of both the government workforce and GenAI. This will involve the core functions of the solution (e.g., models, data, interfaces) as well as the user interface that enhances the partnership experience. The interface can reinforce the role of GenAI—sidekick, assistant, partner, or position it for more direct autonomy. These are important choices for government leaders to make.



Training and tuning GenAI models.

The large language models GenAI run on are time-consuming and costly to build, and the workforce with skills to create them is in short supply—both in industry and government. A more likely possibility is working with a vendor with existing AI models and training or retraining them to fit their intended use. Then organizations could integrate them into existing systems.

Depending on the application, organizations may opt for cloud-based GenAI or build on-premises systems with dedicated servers and limited access, which may offer more security and control but at the cost of greater computer power. In addition, models need to be fed the appropriate data for their appointed use to return the best possible answers.



Data lineage.

Governments should be able to trace where data for the AI models comes from and how it is used, what's referred to as data lineage or data provenance. Not only does traceability build trust and confidence in AI-based processes, but it also provides the information for the routine audits required of many organizations. Personnel should be able to explain model data's origins, how it propagated through the pipeline, and how it appears in end results.



Prompt engineering.

Prompt engineering is the technique of iteratively refining and optimizing questions or commands fed to AI models to elicit higher quality responses. Just as the information a search engine returns depends on how a question is phrased, a GenAI-based program's results depend on the prompts it is fed. The more specific and detailed the prompt, the more specific and detailed the results. Prompt engineering is a new and evolving field, but some techniques in use as we go to press include:

- **Static Prompts¹⁹:** Predefined, unchanging prompts that are used to generate a specific type of response. They can serve as building blocks for more complex prompt engineering techniques and are useful for tasks with consistent requirements.
- **Contextual Prompts²⁰:** Prompts incorporating contextual information, such as the user's role, past interactions, or specific conditions, to generate more accurate and relevant responses. This type of prompting is crucial for nuanced tasks like policy analysis.

- **Prompt Templates²¹:** Reusable and customizable text frameworks that serve as a starting place for generating responses. They help streamline the generation process and are useful for creating standardized outputs.

- **Prompt Decomposition²²:** This prompting technique involves breaking a complex query into simpler, more manageable sub-prompts. Each of these sub-prompts is processed individually, and their outputs are then combined to form a comprehensive response.

- **Prompt Pipelines²³:** A sequence of prompts and processing steps that work in tandem to produce a final output. This allows for multi-stage problem solving and can incorporate other AI or analytical techniques at different stages.


- **Prompt Chaining²⁴:** This method involves using the output of one prompt as the input for another, enabling more complex interactions. Often used in conjunction with other techniques like prompt decomposition or prompt pipelines to handle intricate tasks.

When prompting GenAI systems, it is essential to use straightforward language, provide context, and avoid ambiguous or subjective terms. Prompts produce better results if they are short, focused, spelled correctly, and grammatically correct. Simple prompts yield simple results, however. If the goal is to produce a detailed report on a relatively complex topic, the prompt should be equally as detailed.

Addressing limitations

Early Generative AI models' limitations are well chronicled. Models may confidently respond to prompts with information that is not grounded in the data on which they were trained. They may inherit biases from the training data.

Because they are built on statistical features, they lack the human capacity for logical reasoning. Agencies should continue to keep these limitations in mind as they grow their use of GenAI and should follow directives from the EOs and Office of Management and Budget (OMB), which aim to directly address many of these limitations, as they emerge.



Hallucinations

The large language models that GenAI run on create realistic-sounding language by predicting the next word in a sequence based on the body of data it has been trained on. But the models themselves do not understand the meaning of the language they process. A model is said to “hallucinate” or confabulate when it spits out illogical or counterfactual “information” in its response. Models may also produce plausible-sounding but erroneous or fictional information based on inferences, in some cases, if the prompts it is fed allow for more creativity than less. In cases where government personnel are using AI-based models to generate reports, manage case files, or write computer code, hallucinations could create results that are not entirely valid or true, undercutting the technology's potential value. Careful prompt engineering that specifies the exact sources a model should draw inferences from can help eliminate hallucination issues. So, it is important to keep people with subject matter expertise in the information creation loop to review and validate drafts.

Mitigating bias

AI models are powerful tools, but as tools, they only do as they're told, and lack a human's innate ability to understand context and values. This aspect of GenAI may impact diversity, equity, and inclusion (DE&I) because some groups may not be represented in the AI models, resulting in [potentially biased algorithms](#). One example of that appeared in early AI-based image creation applications. When prompted to generate images of doctors, models only presented pictures of doctors who were white males because of the lack of diversity within the material it was trained on.²⁵ Options for mitigating bias may include:

- **Reviewing underlying training data** to ascertain if it accurately represents the relevant population
- **Building models with clear objectives**, and researching the processes being modeled
- **Adopting data standards** to provide context for how data was sourced or modified

Asking for the perspectives of representatives from diverse populations, which could help avoid confirmation bias.

Additional considerations

Aside from technological considerations and limitations, implementing GenAI has legal, ethical, privacy, security, workforce, and systems integration implications.

Legal and ethical issues

As we go to press, GenAI is so new and changing so rapidly, it's still a technology Wild West. Laws and regulations mandating AI policies and restrictions haven't caught up with the technology's growing popularity. But regulations are likely coming. When they do, government agencies must be prepared, especially those that deal with classified information, protected personal information, personal health information, and other highly sensitive data. To bear the fruits of what GenAI has to offer, jurisdictions must develop AI models tailored to the type of data they handle. They should institute policies and stipulations specific to AI models' functions, so as not to inadvertently introduce bias or cross other ethical boundaries. In the private sector, corporate GenAI users already have begun to apply guardrails to the technology to mitigate hallucinations.²⁶ NIST's AI Risk Management Framework²⁷ provides a comprehensive and adaptable set of guidelines and processes which organizations can use. Additionally, over years of helping organizations implement AI, Deloitte has developed our [Trustworthy AI Framework™](#). This framework [aligns well](#) with NIST's, and both are a way to provide structure and guide responsible decision-making in AI's design, operation, and governance. Both frameworks can [extend to GenAI](#).

Privacy and security

Safeguarding people's privacy is critical for public health departments and other agencies that use and store personal information and personal health information. Instituting rigorous guardrails and proper data pipelines and governance is one way to help accomplish that. A state health agency chatbot may require people to enter their name and social security number in order to return relevant benefits information. Although the model may capture the query to aggregate into future analytics, guardrails built into the application could automatically mask personal information. Entities can also use synthetic data, or highly representative data that simulates the intended environment without relying on actual users' private data.

Protecting GenAI models and data also means warding off cyberattacks and preventing AI-based systems from being infiltrated by adversaries. One example of [adversarial AI](#) is data poisoning, where bad actors compromise the integrity of an AI system by introducing malicious inputs to manipulate the model's outputs. AI-based machine learning models trained on open-source data or production data [can be especially vulnerable](#) to such attacks. As federal agencies upgrade systems to include AI-based technologies, they are encountering more adversarial AI challenges and working on defenses.

The Defense Advanced Research Projects Agency, for example, is developing countermeasures, including using improved metrics that make AI more robust and therefore more impervious to attack.²⁸ Agencies can cross-train their non-security workforce on defensive actions, set standards and secure models while they're being developed, and bring in specialists to evaluate models and suggest additional risk mitigation.



Workforce and culture

As the technology becomes more commoditized and more agencies adopt it for multiple uses, it's likely that a larger portion of the government workforce will interact with GenAI. As usage increases, agencies should upskill employees on applications related to their work, for instance, for roles as prompt engineers and model tuners.


There's no lack of enthusiasm in the workplace to try GenAI tools and other tools for writing, searching, creating, and summarizing content. For government, the bigger challenge is channeling employees' enthusiasm into work on applicable use cases, while ensuring they are aware of the technology's implications and the guardrails in place to maximize privacy and security while minimizing legal, ethical, and other risks.

Integration with existing systems and processes

GenAI may be the dawn of an era of closer human-machine integration. But whatever GenAI-based use cases they adopt, agencies should determine what the effect will be on the workforce and existing systems and processes and follow a pre-determined change management plan to implement them. The tools should augment rather than replace people's jobs, enabling workers to spend more time on critical tasks and enhance the quality of the services they provide.



Conclusion



Based on our observations and experience, we recommend government leaders take a measured approach to GenAI:

Issue interim guidance

As the legislative landscape continues to evolve, agencies should spell out what the workforce can and cannot use GenAI for in broad terms while working to incorporate guidance from the White House and OMB that defines how, and by when, agencies must govern, innovate with, and manage the potential risk from AI. State and local governments can take a similar approach. For example, the city of Seattle issued interim GenAI guidance that is intended to minimize issues that may arise from using the technology while officials conduct additional research and analysis. The policy directs employees to get permission from the city's IT department before using GenAI products. Employees also must validate that the output of those products is "accurate, properly attributed, free of someone else's intellectual property, and free of unintended or undesirable instances of bias and potentially offensive or harmful material."²⁹

Develop a strategy to implement and scale GenAI

GenAI adoption should be part of a coherent AI strategy which itself dovetails with the broader organizational strategy. As outlined in our 2019 article [Crafting an AI strategy for government leaders](#), an integrated AI strategy answers these five questions:

- 01. What is our level** of AI ambition?
- 02. Where should we concentrate** our AI investments?
- 03. How will AI deployment** create value?
- 04. What do we need** to execute our AI strategy?
- 05. What system** will implement and manage AI?

To thoroughly articulate answers to these questions, leaders will need to engage in holistic planning including what operating models, governance, and change management will be required, how the tech stack, data, and scaling will be managed, how the technology will be integrated into existing systems, what technology and skillsets will need to be bolstered or acquired, and how to foster a culture which will adopt and adapt to AI.

Identify the most compelling use cases

A broader AI strategy can facilitate prioritizing GenAI use cases as the value of potential use cases can be understood in context. Pick one or more use cases to pilot the technology; while the specific prioritization framework will be context-dependent, potential aspects could include level of impact, expected investment, opportunity for learning across the organization, visibility of results, mission relevance, and possibility that building for the use case would support expansion to other use cases. Assemble a cross-functional team and provide them with resources and training to produce a minimal viable product to test in a controlled environment before adding features or rolling it out on a broader scale.

However, don't get stuck in pilot purgatory. [Scaling AI](#) isn't achieved solely through pilots. While GenAI pilots will build organizational knowledge around the tools and techniques required and help create the guardrails necessary to foster trust in the technology, another benefit of an AI strategy that includes GenAI is the path it identifies for large scale implementations.

Pick the right technology for the job

GenAI should fit into an organization's broader technology strategy and integrate with existing solutions. GenAI may not be the right type of AI for a specific situation. The key is understanding the problem and understanding the strengths of different means of solving it. Often the best results come from combining multiple tools, each suited to specific purposes. Technology alone isn't the answer. Bolting new technology onto existing processes doesn't work if they are riddled with problems. Fix process problems before applying new solutions.

For example, while in principle GenAI could be used to predict future demand, more traditional predictive analytic techniques might be more appropriate. GenAI and RPA could be used in combination to automate a process to send questionnaire emails to users—or this entire process could potentially be replaced with a web-based form requesting the same information.




GenAI provides the government leaders an immediate opportunity to improve the delivery of their programs and missions, support the government workforce, and engage residents and citizens more effectively and authentically.

Combining government workers with GenAI solutions, GenAI can deliver mission outcomes rapidly and efficiently.

While not without risks, GenAI solutions can be put in operation safely and securely. Like any new innovation, organizations should include this in their strategies so that they can align governance, policy, technology and workforce considerations.




Reach out for a conversation.



Brendan McElrone
Consulting Managing Director
Deloitte Consulting LLP
bmcelrone@deloitte.com



Tasha Austin
Advisory Principal
Deloitte & Touche LLP
laustin@deloitte.com



Edward Van Buren
AI Strategic Growth
Offering Leader
Deloitte Consulting LLP
emvanburen@deloitte.com

Endnotes

1. Child Welfare Information Gateway, "[Caseload and Workload Management](#)," United States Children's Bureau, September 2022.
2. Dr. Anand S. Rao, Gerard Verwaj, and Euan Cameron, "[Sizing the price: What's the real value of AI for your business and how can you capitalise?](#)," PricewaterhouseCoopers, accessed September 2023.
3. Gartner, "[Gartner Identifies the Top Strategic Technology Trends for 2022](#)," October 18, 2021.
4. Performance.gov, "[Transforming Federal Customer Experience and Service Delivery to Rebuild Trust in Government](#)," accessed September 2023.
5. Joseph R. Biden Jr., "[Executive Order on Transforming Federal Customer Experience and Service Delivery to Rebuild Trust in Government](#)," The White House, December 13, 2021.
6. Scott Kozak, "[NeuroTracker Collaboration with US Air Force & NASA](#)," LinkedIn, October 28, 2022.
7. Law Librarians' Society of Washington, DC, "[Sources for Mandated Congressional Reports](#)," September 13, 2018.
8. Alexandra Kelley, "[Where Agencies Are Turning to Automation to Augment Human Work](#)," Government Executive, July 13, 2023.
9. Jory Heckman, "[State Department Looks to AI for Streamlined FOIA Workloads](#)," Federal News Network, June 21, 2023.
10. Child Welfare Information Gateway, "[Caseload and Workload Management](#)," United States Children's Bureau, September 2022.
11. U.S. Department of Labor, "[Guide to Submitting Requests Under the Freedom of Information Act](#)," accessed March 2023.
12. U.S. Department of Justice, "[Agencies Processed Over One Million FOIA Requests in Fiscal Year 2023](#)," March 4, 2024.
13. eFile, "[Tax History and the Tax Code](#)," October 23, 2023.
14. Phil Goodwin, "[Tape and Cloud: Solving Storage Problems in the Zettabyte Era](#)," LTO Consortium, June 2019.
15. Partnership for Public Service, "[Fed Figures: COVID-19 and the Federal Workforce](#)," accessed September 2023.
16. Peter Viechnicki and William D. Eggers, "[How Much Time and Money can AI Save Government?](#)" April 26, 2017.
17. Apurva Venkat, "[Cyberattacks Against Governments Jumped 95% in Last Half of 2022, CloudSek Says](#)," CSO Online, January 4, 2023.
18. Joseph R. Biden Jr., "[Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#)," The White House, October 30, 2023.
19. Cobus Greyling, "[Eight Prompt Engineering Implementations \[Updated\]](#)," Medium, July 2023.
20. Priyanka Vergadia and Kaliah Williams, "[Tips to enhance your prompt-engineering abilities](#)," Google Cloud Blog, August 14, 2023.
21. Langchain, "[Prompt Templates](#)," accessed September 2023.
22. Tushar Khot et al., "[\[2210.02406\] Decomposed Prompting: A Modular Approach for Solving Complex Tasks](#)," arXiv, last modified April 11, 2023.
23. Cobus Greyling, "[The Evolution Of Prompt Engineering](#)," Substack, April 17, 2023.
24. Anthropic, "[Prompt Chaining](#)," accessed September 2023.
25. Leonardo Nicoletti and Dina Bass, "[GenAI Takes Stereotypes and Bias From Bad to Worse](#)," Bloomberg, June 9, 2023.
26. Belle Lin, "[Generative AI Pilots Have Companies Reaching for the Guardrails](#)," The Wall Street Journal, May 19, 2023.
27. Institute of Standards and Technology, "[AI Risk Management Framework](#)," accessed September 2023.
28. GARD Project, "[Holistic Evaluation of Adversarial Defenses](#)," accessed April 19, 2024.
29. Seattle IT, "[City of Seattle Adopts Generative AI Policy](#)," TechTalk, City of Seattle, May 31, 2023.



This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.