

U.S. Government
Accountability Office

AI Strategy

August 2026



Table of Contents

Overview of GAO's AI Strategy	1
GAO's AI Vision	2
Our Key Enablers	2
Our Five Pillars	3
Pillar 1: Strengthen Effective Governance	4
Pillar 2: Prioritize Mission-Driven AI Use	6
Pillar 3: Develop AI-Ready, Scalable, and Secure Data and Infrastructure	8
Pillar 4: Develop an AI-Ready Workforce	10
Pillar 5: Establish Continuous Monitoring	12

Overview of GAO's AI Strategy

Artificial Intelligence (AI) is changing how government works. Federal agencies across government are increasingly using AI to help fulfill their missions from health care delivery to national defense.

Since we published our first AI use case in early 2024, we have explored how AI could deepen staff capacity to improve our mission delivery and operational efficiency. We continue to develop and pilot AI tools to accelerate and expand the depth of our analytical work. For example, we are examining how AI can help us identify patterns and trends in data that previously went uncovered. Similarly, we are actively developing and piloting tools to streamline business workflows and automate tasks.

GAO's AI strategy provides a road map for how we will advance our use of AI in mission delivery and operations. The strategy recognizes the benefits of AI and provides a pathway to realize them. It also acknowledges the potential risks and ethical considerations of using AI, including inaccurate results, data breaches, and unauthorized use (shadow AI). This strategy seeks to harness the benefits of AI while managing the risks through responsible use, accountability, and governance.

In this strategy, we begin with our AI Vision, then discuss our Key Enablers and our Five Pillars – each of which includes a Maturity Assessment.

The strategy reflects federal and industry practices, including GAO's [Enterprise Risk Management Framework](#), [Strategic Plan](#), [Standards for Internal Control in the Federal Government](#), [AI Accountability Framework](#), as well as the [Committee of Sponsoring Organizations' AI framework](#). The strategy is also built on and reflects [our core values](#).

This living strategy will evolve with emerging technologies, risks, and opportunities. It will also evolve as our use of AI increases and matures. We will review the strategy as needed, but at least once a year to ensure it reflects our current AI use.



GAO's AI Vision

The strategy seeks to integrate AI into our engagement and business processes with the singular goal of enhancing our support of Congress in carrying out its constitutional responsibilities for the benefit of the American taxpayer. We will achieve this by using AI to enhance staff productivity, accelerate mission delivery, and drive operational efficiency. When investing in, developing, or deploying AI, we will hold ourselves to the same high standards by which we hold audited agencies. AI also cannot replace first-hand evidence and direct observation gathered through site visits and stakeholder input. This strategy

will maintain the human-in-the-loop approach to make critical decisions and we will seek to learn from our and others' experiences by adopting leading practices. As we adopt these practices, we will be guided by our core values, people values, and mission.

Our reputation for objectivity, accuracy, and nonpartisanship is the foundation on which all our work rests. It is what Congress expects and relies on – they can trust what we say. Our reputation has been earned over a century of fact-based, nonpartisan work. We will only pursue AI use that enhances this reputation.

Our Key Enablers

Achieving our vision for AI use at GAO will be a journey, filled with starts and stops, successes, and setbacks. To propel us forward on our journey, the strategy will draw on some of the agency's existing enablers:




Our **strong governance structures**, including our Quality Assurance Framework, Data Governance Strategy, and Enterprise Risk Management process, to ensure accountability and transparency in our use of AI.



Our **multidisciplinary, talented workforce**, which has a proven ability to quickly master new subjects and work within established governance structures.



Our **external network of experts** across government, academia, and industry, relationships that have been forged through our reputation as a trusted partner and truth seeker.



Our **collaboration with other legislative branch agencies**, which allows us to build on each other's AI experiences to reduce duplication of effort and drive efficiency.



Our **institutional knowledge** of leading practices and agency efforts in the areas of R&D, acquisition, financial management, science and technology, and project management.



Our **science, technology, and cybersecurity expertise**, including our strategic foresight and horizon scanning capabilities and our Innovation Lab, which pilots tools in a prototype-to-production pipeline and serves as a sandbox for testing novel AI use cases.

Our Five Pillars

GAO's AI strategy reflects five key pillars that drive our decision-making:

				
Pillar 1: Strengthen Effective Governance Establish a risk-based governance framework that ensures transparent decision-making; accountable and responsible use; and compliance with laws, regulations, and policies.	Pillar 2: Prioritize Mission-Driven AI Use Identify, prioritize, and expand (scale) AI use cases that align with the agency's mission, strategic plan, and priorities.	Pillar 3: Develop AI-Ready, Scalable, and Secure Data and Infrastructure Support data governance, quality (accuracy, completeness, and reliability) and standardization; adopt AI enabling infrastructure and scalable processes; and embed security features and practices throughout the AI ecosystem.	Pillar 4: Develop an AI-Ready Workforce Build on GAO's ethos of continuous learning and improvement by equipping GAO's staff with the knowledge, skills, and tools to responsibly use AI.	Pillar 5: Establish Continuous Monitoring Continuously monitor costs, performance, and outcomes at the use case, portfolio, and system levels to ensure the prudent use of funds and maximize the return on investment.

These pillars align with federal guidance and collectively offer a road map for GAO's responsible adoption and use of AI.




Maturity Assessment



In implementing GAO's AI strategy, we have assessed our maturity in each of the five pillars. Each pillar's maturity level is currently foundational or developing. A foundational maturity level is exploratory and in early stages. A developing maturity level is more structured and efforts are more consistent.

For each of the five pillars, we will develop a set of qualitative and quantitative metrics that measure both outputs and outcomes before the 2027 AI strategy update.




Pillar 1

Strengthen Effective Governance

GAO will continue establishing a risk-based governance framework that provides structured guidelines and practices for the agency's use of AI. The framework will provide structures and processes for accountability, transparency, and security to help ensure compliance with laws, regulations, and policies and manage risks across the agency.

Governance Structures



GAO's Executive Committee (EC) established the Technology and Data (T&D) governance board which guides and provides oversight of the agency's use of emerging technology including AI. The board is also responsible for making technology investment decisions. The board will charter subordinate boards, councils, and working groups, which will create governance processes for the development of, investment in, and use of technology, including AI. Our analysts, specialists, mission teams, and other work units are critical inputs into this governance model.

The T&D governance board chartered the AI Council to serve as the agency's focal point and key oversight body for GAO's AI use. The Council is co-chaired by GAO's Chief Technology Officer (CTO) and Chief Scientist. The Council includes GAO subject matter experts, including the Chief Information Officer (CIO), Chief Information Security Officer (CISO), Chief Data and Analytics Officer (CDAO), and the Innovation Lab Director. The Council will develop governance processes for identifying, prioritizing, and evaluating AI use cases (including their performance and return on investment). The Council will advise the T&D governance board on AI matters and make recommendations on whether AI use cases should be expanded, modified, or terminated based on piloting and ongoing monitoring. The CTO and Chief Scientist will serve as advisors to the T&D governance board on emerging technology, including AI.

Policies

GAO will develop policies to guide the use of, development of, and investment in AI. The purpose of these policies is to ensure our AI use complies with applicable laws, regulations, standards, and existing GAO policies. These policies will also promote consistency in AI decisions, oversight, and use across the agency and ensure use cases align with this strategy, the agency's strategic plan and our risk appetite.

As the T&D governance board approves AI use cases, GAO policies, orders, and other guidance, such as our internal guide for leading engagements (Electronic Assistance Guide for Leading Engagements), may need to be updated. We will establish processes to ensure policies, orders, and guidance are updated to reflect new AI use cases. We will also use these processes to ensure AI use cases comply with GAO policies.



Risk Management

GAO will manage potential enterprise-level risks, such as cybersecurity and privacy, through its enterprise risk management process. The Risk Management Council will establish the agency's risk appetite, identify and assess risks and the associated internal controls, and pinpoint needed risk-mitigation efforts for our AI use. GAO's Office of Internal Controls will periodically evaluate the effectiveness of the internal controls that are in place.

Key risk mitigation strategies include ongoing monitoring of costs, impacts, and risks as described in Pillar 5: Establish Continuous Monitoring. We will also mitigate risk by allowing GAO's Innovation Lab to continue innovating in a secure environment. The Lab provides a separate, access-controlled space for safe experimentation, enabling innovation while mitigating risk.



Key Policies and Documents

- Third-party use guidance
- AI use case review process
- GAO Enterprise Risk Management process
- GAO Strategic Plan
- GAO AI Accountability Framework
- GAO Quality Assurance Framework
- GAO Policy Manual
- Data Governance Strategy (draft)
- GAO Cybersecurity Manual



Responsible Entities and Leaders

- Executive Committee
- T&D governance board
- AI Council
- CIO, CTO, CISO, CDAO, and Chief Scientist
- Chief Quality Officer
- Risk Management Council
- Innovation Lab Director
- Chief Risk Officers




Maturity Assessment



Key Efforts

- The T&D governance board updated the AI Council's charter to reflect this strategy and an expanded scope of Council responsibilities.
- The EC reestablished the CTO position and expanded the CTO's scope of responsibilities.
- The AI Council developed some key guidance documents (e.g., a policy on third-party use), but it will need to develop others, such as standard operating procedures and templates for tracking AI use case spending and return on investment. Existing tools, templates, and dashboards, such as those currently used by the Innovation Lab, will inform these efforts.
- The AI Council has a process and cadence for reviewing and recommending AI use cases to the T&D governance board. As part of that process, we are establishing additional controls and monitoring.
- AI-related risks are considered as part of the agency's enterprise risk management process, and a risk assessment is performed each year.



Pillar 2

Prioritize Mission-Driven AI Use

GAO will identify, prioritize, and scale AI use cases that align with the agency's core mission, strategic plan, and priorities. As prudent stewards of the federal resources entrusted to us, we will prioritize AI use cases that tie directly to our mission and deliver measurable results.

Identifying AI Use Cases

GAO will take a top-down and a bottom-up approach to identify potential AI use cases for our engagement work and business operations. The Executive Committee, T&D governance board, and AI Council will suggest AI use cases and set priorities. We will also leverage the Innovation Lab's existing process to solicit and collect staff's suggested AI use cases.

Approving AI Use Cases Through a Standardized Process

GAO uses a standardized review process to evaluate and approve all AI use cases. The AI Council, in collaboration with the Innovation Lab, maintains a review process, which will culminate with an AI Council recommendation to the T&D governance board for approval. The AI Council has screening criteria (e.g., readiness, return on investment, mission alignment, buy vs. build), templates, guidance, and other tools to help ensure a consistent and efficient review process. Projects that advance the agency's core mission and priorities and deliver measurable results will be prioritized.

Piloting and Evaluating AI Use Cases

To inform significant or long-term investment decisions, we will pilot all broad AI use cases. Pilots will be of proportionate scale and duration, and we will use [GAO's pilot design and evaluation criteria](#). The AI Council is responsible for developing the piloting process – selecting, evaluating, and recommending AI use cases to the T&D governance board for full deployment across the agency. Narrower AI use cases that fall under broadly approved applications, particularly for GAO engagements, will be subject to rigorous, ongoing quality assurance requirements.

To support the scaling of AI use cases, we will standardize AI deployments by, for example, reusing components and processes that are easily repeatable. These deployments will be continuously monitored to ensure performance, integrity, and compliance with agency policies, as described in Pillar 5: Establish Continuous Monitoring.



Centralized AI Use Case Inventory

The AI Council will develop and maintain an authoritative, centralized inventory of the agency's AI use cases. The inventory will be posted on GAO's intranet. This centralized approach will help harness the benefits of decentralized experimentation and innovation while promoting transparency and agency-wide learning about effective and responsible uses of AI. This approach will help reduce redundant or duplicative efforts and promote transparency. As we update our AI Use Case Inventory, we will also update the inventory on our public [website](#).



Key Policies and Documents

- GAO's Strategic Plan
- AI use case submission guidance
- AI use case review process
- AI Use Case Inventory
- GAO's leading practices for pilot design and evaluation



Responsible Entities and Leaders

- Executive Committee
- CTO and Chief Scientist
- AI Council
- T&D governance board




Maturity Assessment



Key Efforts

- The agency maintains a centralized AI use case inventory and added it to its intranet and public-facing internet.
- The AI Council is updating guidance for staff on submitting AI use case suggestions.
- The AI Council is updating and expanding a standardized review process for approving and evaluating AI use cases, to include the process for designing and evaluating AI use case pilots.



Pillar 3

Develop AI-Ready, Scalable, and Secure Data and Infrastructure

To ensure responsible AI use, we will improve data quality, embed security features and practices throughout the AI ecosystem, and focus on scalable and reusable code, models, and infrastructure.

High-Quality Data

Developing and maintaining high-quality data for use in our engagement work and business operations is critical to the responsible use of AI. GAO's CDAO will lead, develop, and implement a data governance process that seeks to improve the accuracy, completeness, reliability, standardization, and integrity of our internal agency data. This will include processes to identify, clean, tag, validate, and monitor internal data and datasets. Because much of GAO's institutional knowledge resides in engagement documents rather than in structured data, the CDAO, CIO, CTO, CISO, agency records officer, and AI Council will work together to establish structured requirements for how these data are managed.

The CDAO will develop a plan to assess data readiness. AI-ready data must be machine-readable, machine-understandable, and structured for AI consumption and reuse.

The CDAO will also build systems to monitor where AI may be changing our data or impacting our data quality.

Data Provenance and Lineage

The CDAO, in coordination with the AI Council, will develop a data provenance and lineage framework that requires AI use cases to document data sources, transformation steps, and any third-party or vendor systems involved in processing. This documentation will be reviewed as part of the AI use case approval process described under Pillar 2: Prioritize Mission-Driven AI Use and will be a required component of any AI system that informs published products, findings, or recommendations. The framework will adhere to applicable National Archives and Records Administration's standards, such as those related to recordkeeping, metadata tagging, and lifecycle management.

AI Guardrails

As GAO expands its use of generative AI, we will establish clear guardrails to ensure AI-assisted outputs remain accurate, defensible, and appropriately scoped. Given that GAO's mission depends entirely on the credibility of its findings, AI-generated content that informs engagement work must be human-reviewed before use, grounded in verifiable source material, and disclosed – not simply produced and accepted at face value. These guardrails apply along with the data provenance and lineage framework described above.

Model and Infrastructure Security

AI use expands or introduces new data security risks that must be mitigated. These risks include data breaches, unauthorized access, data transfer, and cybersecurity risks. To mitigate these risks, the CISO will work to ensure that AI models and infrastructure comply with the agency's information security policies. The CISO will also collaborate with the Learning Center to develop cybersecurity training for all staff, update incident response and recovery plans for AI use, strengthen access management and authentication mechanisms for AI models,



and maintain continuous monitoring systems to detect and respond to security threats. The CISO will also update GAO's cybersecurity strategy to include AI risks.

Working with the AI Council, the CISO will embed security features and controls in the systems, pipelines, and infrastructure in which AI models are deployed, and will work to ensure these systems include built-in security, compliance, and traceability from the outset. Any AI-enabled systems developed in house will adhere to software development leading practices, including those set out in [GAO's Agile Assessment Guide](#), and will be documented, tested, and reviewed accordingly. Where GAO uses commercially provided models, security will be established at the system and interface layers and through contract terms and vendor assurance. We will develop standard contract language to ensure that all contracts for AI tools include language to protect GAO's data.

Reusable, Scalable Infrastructure

This strategy will take a portfolio-level approach to our AI use, including making decisions about whether to build or buy AI. We will also seek to coordinate the sequencing of AI and non-AI projects. We will prioritize reusable code, models, and tools designed for enterprise-wide reuse. The goal is to reduce redundancies and lower costs by focusing on scalable, shared, reusable platforms rather than one-off, siloed systems.

We will also keep pace with the rapidly evolving nature of AI, continue our work on agent-based platforms, and respond with purpose-built agentic tools for building and scaling across GAO.



Key Policies and Documents

- Cybersecurity strategy
- Data governance framework
- Data and analytics strategy
- Data provenance and lineage framework
- AI use case review process
- Human-in-the-loop review policy
- AI output grounding and fact-verification standard
- GAO Agile Assessment Guide



Responsible Entities and Leaders

- Executive Committee
- CTO, CIO, CISO, and CDAO
- AI Council
- General Counsel
- Chief Quality Officer
- Chief Records Officer
- Mission team leadership



Maturity Assessment

Foundational



Key Efforts

- The CDAO will inventory internal data and data systems using a phased approach that prioritizes the data required for the agency's highest value AI use cases, and assess the quality, completeness, and reliability of those data and systems.
- The CDAO, in collaboration with stakeholders, will establish data quality framework, including data standards, and strategy.
- The CISO will update information security strategy to include AI risks.
- The AI Council will develop and implement processes for AI use case review and approval that ensure that issues of human-in-the loop, data security, scalability, and reusability are considered.
- The CDAO will create a data provenance and lineage framework.
- The Managing Director for Applied Research and Methods will update the data reliability checks that we use as part of our standard engagement process to include validating AI-generated data and information.



Pillar 4

Develop an AI-Ready Workforce

Training Curriculum

AI Community of Practice and Intranet Website

A culture of continuous learning and development is a hallmark of GAO, and we will leverage this culture to prepare our staff to responsibly use AI. We will provide staff with the knowledge, skills, and tools to responsibly use, manage, and innovate with AI by both developing a robust AI literacy training curriculum and attracting and retaining leaders and staff with AI expertise throughout the agency.

Consistent with our people values, we seek to foster a culture where staff, regardless of roles or technical experience, can learn and contribute to our adoption of AI. To achieve this goal, GAO's Learning Center will create a robust training curriculum so that staff can deepen their AI knowledge and skills and increase their AI literacy (i.e., skills that enable staff to understand, interact with, and critically evaluate AI technologies). The curriculum will be adaptive. The Learning Center, working with GAO's subject matter experts, will review and refresh course content on a recurring basis so that training keeps pace with changes in AI capabilities and the skills needed to use them responsibly. The core curriculum will consist of a combination of mandatory and elective courses, covering a range of topics, from policy-related to technical, and will be offered in a variety of formats (e.g., on demand, in person).

Individual teams and units may also develop their own AI training by coordinating with the Learning Center to avoid duplication of efforts or conflicting guidance or instruction. We will emphasize GAO-wide and mission-specific AI training so that staff can develop their knowledge, expand their skills, and better collaborate.

These efforts will help ensure that all staff across the agency have a basic level of AI literacy and are prepared to responsibly use AI in their daily work. This includes understanding how to apply the latest methods, tools, and techniques to audit federal agencies' evolving use of AI, including assessing the resulting financial, security, and safety implications. AI does not replace our professional judgment and evidence-based work.

GAO's AI Community of Practice (CoP) was established in 2023. The CoP – comprised of staff from across GAO – assists in developing GAO's AI expertise and enhancing mission team collaboration. In doing so, it holds training sessions on AI, shares information and best practices, and sponsors presentations from outside experts. This CoP also has working groups that collaborate on audits involving AI as well as GAO audits of agency's use of AI. Senior leadership actively supports this CoP, including having multiple senior executives serve as its sponsors.

We also established a web page on our intranet site dedicated to AI use at GAO. The page serves as a central repository for policies, guidance, terminology, and knowledge-sharing and as a means to promote transparency as our use of AI expands.



Attract and Retain AI Expertise

We will continue to attract and retain individuals with AI expertise and place them in impactful positions throughout the agency. By embedding AI literacy at every level, aligning talent with mission needs, and fostering a culture of ethical innovation, we are building the capacity to adopt AI and use it responsibly.

For multiple years, we have prioritized attracting and retaining individuals with science and technology expertise in our budget and work plans. We continued this focus with our [Fiscal Year 2027 budget request](#). We also re-established a CTO position in June 2026 and expanded the position's scope of responsibilities. The CTO is responsible for developing the agency's long-term enterprise technology road map, evaluating emerging capabilities and technologies, and fostering a culture of experimentation.



Key Policies and Documents

- GAO Strategic Plan
- Annual budget request
- Annual workforce plan
- Revised training curriculum
- AI intranet webpage
- AI CoP charter



Responsible Entities and Leaders

- Executive Committee
- Chief Learning Officer
- CTO
- CISO
- Chief Scientist
- Chief Human Capital Officer
- Office of the Chief Administrative Officer (Workforce Planning unit)



Maturity Assessment



Key Efforts

- The Executive Committee reestablished and appointed the agency's CTO position.
- Management chartered an AI CoP and established an intranet page dedicated to AI matters.
- The Executive Committee directed the Learning Center to develop a suite of AI-related courses; the Learning Center is in the process of developing these courses.
- The CTO, CIO, CISO, CDAO and the Chief Scientist, in collaboration with the agency's workforce planning office and other stakeholders, will help to identify AI workforce needs.



Pillar 5

Establish Continuous Monitoring

GAO will continuously monitor the agency's use of AI to ensure that our use is responsible; complies with applicable laws, regulations, policies; reflects leading practices; and is a prudent use of taxpayer dollars. We will continue to develop structured processes to plan, track, and evaluate all AI use and investments. The AI Council and T&D governance board will consider suspending or terminating AI solutions that do not meet established cost, schedule, or performance thresholds.

Track Benefits and Costs

The AI Council will establish a process and associated thresholds to monitor the life cycle costs, performance, and outcomes of AI use cases and investments. We will monitor these at the use case, system, and portfolio levels, to ensure the prudent use of funds and maximize the return on investment. The evaluation criteria and methods should be identified as AI pilot use cases are approved so that the needed data are collected throughout the pilot.

Investments will be guided by clear criteria, including cost, to ensure AI systems and procurement practices deliver results and promote transparency, accountability, and an efficient use of resources. To help ensure strategic investments, we will evaluate investments at the use case and portfolio levels. This will help ensure mission-driven investments.

Types of Reviews

We will evaluate the use of AI and related investments in multiple ways, including performance evaluations, compliance reviews, privacy assessments, internal controls audits, and security assessments. This will require close collaboration among multiple GAO teams and units. Continuous improvement will be driven by performance indicators (e.g., process efficiency, risk-mitigation outcomes, user experience), defined life cycle checkpoints, incident tracking, root-cause analysis, and implementation of any lessons learned.

Transparency

To help promote transparency in our AI use and investments, the CTO will prepare an annual state of AI use report for the T&D governance board. This report will summarize what we are achieving and learning through our AI use and investments.



Key Policies and Documents

- AI use case review process
- Annual State of AI Use report
- Enterprise Risk Management process
- GAO's leading practices for pilot design and evaluation



Responsible Entities and Leaders

- | | |
|---------------------------|--|
| • Executive Committee | • Chief Economist |
| • Controller | • Managing Director for Applied Research and Methods |
| • AI Council | • Financial Management and Business Operations |
| • CTO | • Office of Internal Controls |
| • CIO | |
| • Risk Management Council | |



Maturity Assessment



Key Efforts

- The agency risk register was updated to include AI use; we will establish our risk appetite and complete a risk assessment.
- We will continue to develop criteria for evaluating use cases.
- Processes need to be developed to track costs and benefits of AI use cases and to link those costs and benefits to budget and accounting systems.
- A process and thresholds for evaluating use cases at a portfolio level need to be developed.