

AI Guide for Government

A living and evolving guide to the application of Artificial Intelligence for the U.S. federal government.

Table of Contents

Introduction to the AI Guide for Government

0.1 — Why are we building an AI Guide for Government?

0.2 — Who should read this AI Guide for Government?

Chapter 1: Understanding AI and key terminology

1.1 — Why you should care about AI

1.2 — Key AI terminology

Chapter 2: How to structure an organization to embrace AI

2.1 — Organizing and managing AI

Chapter 3: Responsible and Trustworthy AI Implementation

3.1 — AI principles and guidelines are still evolving

3.2 — Why is DEIA essential for a responsible and trustworthy AI in practice?

3.3 — Ask questions often and repeatedly

3.4 — Moving forward

Chapter 4: Developing the AI workforce

4.1 — Start with people

4.2 — Developing and retaining AI talent

4.3 — Understanding AI job roles and career path

4.4 — Recruiting AI talent

Chapter 5: Cultivating Data and Technology

5.1 — Tools, capabilities, and services

5.2 — Data governance and management

Chapter 6: AI Capability Maturity

6.1 — AI development founded on good software practice

6.2 — Organizational maturity areas

6.3 — Operational maturity areas

Chapter 7: Solving business challenges with AI

7.1 — Understanding and managing the AI lifecycle

7.2 — Good software practice and AI development

7.3 — Identifying AI use cases in your organization

7.4 — Starting an AI project

Introduction to the AI Guide for Government

Artificial Intelligence (AI) refers to the computational techniques that simulate human cognitive capabilities. AI will transform most, if not every aspect of humanity, which presents a range of challenges and opportunities.

AI has already revolutionized the business world. Its application across the federal government is fundamentally changing the way agencies meet their mission. **The U.S. government must embrace these opportunities head-on to remain on the leading edge and stay competitive.**

This AI Guide for Government is intended to help government decision makers clearly see what AI means for their agencies and how to invest and build AI capabilities.

Because AI is such a broad term to describe new and emerging applications, we've broken the AI Guide for Government into different chapters. At this time, the Guide does not include technical sections.

The AI Guide will help leaders understand what to consider as they invest in AI and lay the foundation for its enterprise-wide use. It helps leaders understand the types of problems that are best suited for the application of AI technologies, think through the building blocks they require to take advantage of AI, and how to apply AI to use cases at the project level. It also explains how to do so responsibly.

Section 0.1: Why are we building an AI Guide for Government?

Since we started, we have engaged with leaders across the federal government to understand the excitement, challenges, and opportunities surrounding AI as an enabling technology.

The federal government need is clarity, education, and guidance around

- what AI as a set of technologies is,
- what it can and cannot do, and
- how to apply it to federal agencies' mission areas.

This is meant to be an evolving guide to the application of AI for the U.S. federal government, because AI is a rapidly evolving set of technologies.

Section 0.2: Who should read this AI Guide for Government?

This guide is made for agency senior leaders and decision makers. We specifically target senior agency leadership: chief information, data, and technology officers, chief financial and procurement officers, program directors, and agency mission leaders.

This guide gives leaders enough information to make the right AI investments. You don't need to be technically proficient in AI to use this guide. Non-technical decision makers who must determine the level of investment required, and where and how those investments will be deployed, can use this guide.

Other users could be directors and managers of AI program offices. They might delegate management of AI projects to other managers who will oversee key decisions on whether to build or acquire AI tools, techniques and algorithms, and solicit expertise from those in the agency with more advanced AI expertise, if available.

Chapter 1: Understanding AI and key terminology

Chapter 1 gives a shared understanding of the key terminology used in the AI space. Many terms are used to describe AI and the tools and capabilities associated with it. This chapter aims to help clarify where possible.

As the term Artificial Intelligence itself describes a wide range of technologies that don't have scientific or industry standard definitions, we will simply discuss these terms, not define them.

Section 1.1: Why you should care about AI

Technological advances allow both the private and public sector to use the resources needed to collect, house, and process large amounts of data, as well as to apply computational methods to it. This changes the way that humans interact with computers. AI has already changed the way that businesses interact with their customers. Entire markets have changed around this technology to provide fast, efficient, and personalized service. We should use this transformative technology to enhance and support the many ways that the government serves its people.

We are also facing increasing complexities and interdependencies in our challenges, as well as increasingly large volumes of data we need to use to understand and solve for those challenges. We are past the point where human cognitive abilities can directly process and make sense of all this information. We need AI tools and techniques to support human capabilities to process this volume of information to reveal insights that support better decisions.

Most federal agencies know that input data volumes are increasing relentlessly and not being handled thoroughly. In the past, we have blamed this kind of shortfall on lack of personnel, supplies, or equipment.

But while these factors are still true, there is no practical increase in any of those resources that would itself suffice to address the new information volumes. With thousands or millions of pages of documents, we could never even try to hire enough staff to read through them all.

The federal government needs AI. The use of AI algorithms informed by human domain expertise drive insights and inform decisions and actions. The use of AI will enable the agencies to handle millions or billions of data inputs with a feasible level of personnel and funding.

To keep up with the ever-increasing volume of data and information, we must change the way we think about our work and processes. AI capabilities will allow us to process inputs and understand reality in today and tomorrow's complex world.

Section 1.2: Key AI terminology

The current AI landscape is both exciting and confusing. Phrases like “advanced analytics” and “machine learning” are often used along with AI. You need to know what the words mean before you discuss how to adopt the technology.

One of AI’s challenges is that it’s a multi-disciplinary domain where even basic definitions are tricky. Here, we will focus on three terms and the relationship among them: AI, machine learning, and data science.

Artificial intelligence (AI)

AI combines three disciplines—math, computer science, and cognitive science—to mimic human behavior through various technologies. All of the AI in place today is task-specific, or narrow AI. This is an important distinction as many think of AI as the general ability to reason, think, and perceive. This is known as Artificial General Intelligence (AGI) which, at this point, is not technically possible.

This technology is rapidly evolving, and neither the scientific community nor industry agree on a common definition.

Some common definitions of AI include:

- A branch of computer science dealing with the simulation of intelligent behavior in computers.
- Advanced statistical and analytical methods such as machine learning and artificial neural networks, especially deep learning.
- A computer system able to perform specific tasks that normally require human intelligence, such as visual perception, speech recognition, decision-making, and language translation.

AI capabilities are rapidly evolving, and neither the scientific community nor industry agree on a common definition of these technologies. In this guide, we will use the definition of AI from the [National Defense Authorization Act for Fiscal Year 2019 \(https://www.congress.gov/bill/115th-congress/house-bill/5515/text\)](https://www.congress.gov/bill/115th-congress/house-bill/5515/text), which is also referenced in the [Executive Order on Maintaining American Leadership in Artificial Intelligence \(https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence\)](https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence).

The term “artificial intelligence” includes the following:

- Any artificial system that performs tasks under varying and unpredictable circumstances without significant human oversight, or that can learn from experience and improve

performance when exposed to data sets.

- An artificial system developed in computer software, physical hardware, or other context that solves tasks requiring human-like perception, cognition, planning, learning, communication, or physical action.
- An artificial system designed to think or act like a human, including cognitive architectures and neural networks.
- A set of techniques, including machine learning, that is designed to approximate a cognitive task.
- An artificial system designed to act rationally, including an intelligent software agent or embodied robot that achieves goals using perception, planning, reasoning, learning, communicating, decision making, and acting.

It is important to keep in mind that the definition of AI is still evolving and that achievements in the field today have been in task-specific AI or “narrow AI”, as opposed to what is commonly called artificial general intelligence that can learn a wide range of tasks—like humans.

Data Science

Data science is the practice and methodology of implementing analytics or machine learning techniques, with subject matter expertise, to provide insights that lead to actions.

Data science is a broad field that covers a broad range of analytics and computer science techniques. This field—and the various professions that perform data science—are a critical component to building AI solutions.

In practice, data science is a cross-functional discipline that combines elements of computer science, mathematics, statistics, and subject-matter expertise. The goal of data science is to produce data-driven insights and processes that can help solve business, operational, and strategic problems for different kinds of organizations. This is often, though not always, achieved using machine learning and artificial intelligence capabilities.

Throughout these chapters, we will frequently refer to data science and data science teams. These are the teams who support the many data and AI efforts underway in government agencies.

Read more about how data science fits into the broader government AI ecosystem, Integrated Product Teams (IPT), and Developing the AI Workforce in Chapter 2 of the AI Guide for Government,

Machine Learning (ML)

Machine Learning (ML) refers to the field and practice of using algorithms that are able to “learn” by extracting patterns from a large body of data. This contrasts to traditional rule-based algorithms. The process of building a machine learning model is, by nature, an iterative approach to problem solving. ML has an adaptive approach that looks over a large body of all possible outcomes and chooses the result that best satisfies its objective function.

Though different forms of ML have existed for years, recent advancements in technology provide the underlying capabilities that have enabled ML to become as promising as it is today. Increased computing capacity (especially elastic computing infrastructure in the cloud), large-scale labelled data sets, and widely distributed open-source ML software frameworks and codes propelled the development of ML models. With these advancements, the accuracy of ML prediction and the number of problems ML can address have dramatically increased in the past decade.

There are three high-level categories of machine learning: **supervised learning**, **unsupervised learning**, and **reinforcement learning**. Each has its own mathematical backbone, and each has its own unique areas of application. Occasionally in more complex workflows, they may be combined.

Supervised learning, also known as supervised machine learning, is defined by its use of labeled datasets to train algorithms to classify data or predict outcomes accurately.

- Input data is fed into the model.
- Weights are adjusted until the model has been appropriately fitted, i.e. generalizes and adequately represents the pattern.
- A training dataset is used to teach models to yield the desired output and includes inputs and outputs that are correctly categorized or “labeled”, which allow the model to learn over time. The algorithm measures its accuracy through the loss function, adjusting until the error has been sufficiently minimized.

Supervised learning models can be used to build and advance a number of important applications, such as:

- Image and object recognition are applied computer vision techniques that are used to detect instances of objects of a certain type of classification such as a car or pedestrian. For example, in health care an AI system can learn to recognize what are pre-cancerous cells and what are not, in order to assist medical professionals conduct an earlier diagnosis relative to what a medical professional could determine on their own.
- Predictive analytics is used to provide deep insights into various data points and allows for the

anticipation of results based on given output variables. Examples of predictive analytics include credit scoring to predict likelihood of paying on time based on factors including customer data and credit history.

- Customer sentiment analysis is used to extract and classify important pieces of information from large volumes of data—including context, emotion, and intent. It can be useful for gaining an understanding of customer interactions and can be used to improve customer experience.
- Spam detection is used to train databases to recognize patterns or anomalies in new data to organize spam and non-spam-related emails effectively. As the name suggests, it can be used to detect spam and help create a better user experience and reduce cyber fraud and abuse.

Unsupervised learning is often used in data exploration before a learning goal is established.

Unsupervised machine learning uses unlabeled data. From that data, it discovers patterns that help solve clustering or association problems. It's useful when subject matter experts are unsure of common properties of a data set. Unsupervised learning models are utilized for three main tasks—clustering, association, and dimensionality reduction. Clustering is a data mining technique which groups unlabeled data based on their similarities or differences. Association is used to discover interesting relationships between variables in a dataset. Dimensionality reduction is used to reduce the number of dimensions while still maintaining meaningful properties close to the original data.

Machine learning techniques have become a common method to improve a user experience.

Unsupervised learning provides an exploratory path to analyze data to identify patterns in large volumes more quickly when compared to manual observation to determine clusters or associations.

Some of the most common real-world applications of unsupervised learning are:

- News feeds: used to categorize or “cluster” articles on the same story from various online news outlets.
- Computer vision: used for visual perception tasks such as object recognition.
- Medical imaging: used in radiology and pathology to diagnose patients quickly and accurately.
- Anomaly detection: used for going through large amounts of data and discovering atypical data points within a dataset.
- Customer personas: used to understand common traits and to build better buyer persona profiles.
- Recommendation engines: uses past behavior data to discover data trends that can be used to develop tailor such recommendations.

Reinforcement learning is a behavioral machine learning model that is similar to supervised learning, but the algorithm isn't trained using sample data. This model learns as it goes by using trial and error. A sequence of successful outcomes will be reinforced to develop the best recommendation for a given problem.

Applications using reinforcement learning:

- Autonomous vehicles: used for self-driving cars, reinforcement learning improves safety and performance
- Industry Automations: used to control HVAC systems in buildings, data centers and various industrial centers, which leads to increased energy savings.
- Trading and Finance: time series models can be used for predicting future sales as well as predicting stock prices
- Language and text: used for text summarization, question and answering, and language translation using natural language processing
- Healthcare: used to find optimal policies and procedures using previous experiences of patient care without the need for previous information.

Key Messages

- Supervised learning uses labeled datasets to train algorithms to classify data or predict outcomes.
- Unsupervised learning uses unlabeled data. From that data, it discovers patterns that help solve clustering or association problems.
- Reinforcement learning sequence of successful outcomes will be reinforced to develop the best recommendation for a given problem.
- AI solutions use one, or in some cases several, of these ML techniques.

Myths about Artificial Intelligence

Though AI is a powerful technology already providing deep insight and business value, it is not magic. Understanding AI's limitations will help you choose realistic and attainable AI projects. Below are some common myths about AI and pitfalls to avoid when evaluating it as a potential tool.

Myth about AI:

“AI will replace humans in the workplace.”

Reality:

AI is more likely to replace tasks within a job, not the entire job itself. Almost all present-day AI systems perform specific tasks rather than entire jobs. The purpose of AI and automation is to make low-value tasks faster and easier, thus freeing up people to focus on high-value work that requires

human creativity and critical thinking.

Historically, automation has created more jobs than it replaces. **AI will mostly replace tasks, not jobs. It is more appropriate to think in terms of human-machine teams where each does the tasks for which it is best-suited. Many forecasts predict that new jobs will be created, i.e. people are and will continue to be needed for certain tasks and jobs.**

Myth about AI:

“AI can think like a human and learn on its own.”

Reality:

AI uses mathematical models and finite computing power to process information. Though some AI techniques might use “neural nets,” these algorithms only remotely resemble human biology. **Their outputs are still entirely based on data and rules prepared by humans.**

Myth about AI:

“AI is always more objective than humans.”

Reality:

AI applications are a product of data and algorithms combined into models. Data is collected, prepared, and managed by humans. Combining it with algorithms may still produce unfair and biased results. Machines and humans have different strengths and limitations. Humans are good at general tasks and big-picture thinking. Machines are good at doing specific tasks precisely. **Human plus machine combinations are almost always superior in performance to a human alone or a machine alone.**

Myth about AI:

“You can just buy AI solutions that will work across the board.”

Reality:

Identifying AI use cases and the data required for them can be specific and localized. Further, the nature of algorithms and model training can require varying degrees of customization as the data is aggregated, cleansed, assimilated, and the outcomes are generated. Barriers to consider beyond technology include organizational culture, appetite for risk, the acquisition process, and agency willingness to experiment. **Buy vs. build decisions require careful assessment.**

Myth about AI:

“Artificial General Intelligence (AGI) is just around the corner.”

Reality:

Artificial General Intelligence refers to AI that achieves general human-level intelligence. For most systems, there is a trade-off between performance and generality. An algorithm can be trained to perform one specific task really well, but not every possible task. Whether AGI takes decades or centuries to achieve, it's more complex than most imagine. **The more tasks we want a single machine to perform, the weaker its general performance becomes.**

Myth about AI:

“A large team of data scientists is required to implement an AI project.”

Reality:

Developing AI solutions might require only a couple of people a few weeks, or it could take years with a large team. It all depends on the nature of the objective, data, required technical infrastructure, and integration into the existing environment. Depending on the maturity of the AI applications related to the specific problem of interest to your agency, the level of data science involvement can vary significantly. Examples of how this may depend based on agency need are:

- Some applications, such as voice recognition, can be deployed from commercial-of-the-shelf (COTS) products.
 - Some AI applications require training of an existing algorithm using agency-specific data, needing a small data science team.
 - Some AI applications are still in the research and development stage. A relatively large data science team is needed to explore the data characteristics and identify the suited AI method to solve the problem.
-

Chapter 2: How to structure an organization to embrace AI

Many organizations seek to “bolt-on” AI into existing structures and business practices. However, upsetting existing workflows and operating outside of established structures can be tedious and counterproductive.

This chapter focuses on how to redefine business as usual and incorporate AI as a core capability, thereby transforming the approach to existing models.

Section 2.1: Organizing and managing AI

The structure for organizing and managing AI should accomplish two broad goals: 1. enhancing mission and business unit abilities to meet their objectives and 2. supporting practitioner effectiveness.

Goal 1: Enhancing mission and business unit ability to meet their objectives means including AI knowledge in mission and program offices

Though each agency is structured slightly differently, the mission units, business units and/or program offices—for the purposes of this chapter, we will call them business units—are primarily responsible for delivering on the mission of the agency. These units are staffed with subject matter experts on how the organization operates, what its primary goals are and the current operating model to deliver on that mission. As these business units look to deliver on their mission in the 21st Century, they should be looking to innovative technologies to enhance and streamline their operations. In order to achieve this, business units must consider a few key priorities:

- **Business units own the challenge.** Since the business units carry out the mission, they are primarily responsible for identifying business challenges that should be innovated through the use of technology generally, but specifically through adoption of AI. Each AI project should be directly linked to a business challenge with expected outcome and benefits identified early on.
- **Investment.** The level of investment in AI should match the level of value it adds in achieving mission and business goals. Therefore, mission or business executives should allocate funding for AI. Think about your mission and business objectives, and which tasks to support those objectives could be done better with the addition of AI techniques.
- **Team structure.** Embedding AI-focused work in the mission centers and with the customer ensures that AI is integrated into how the agency functions and achieves its mission and business goals. AI should not be approached in a silo, but rather integrated into the rest of the workforce and the agency's core workflows. Concretely, this means business units should use integrated product teams—which we discuss further down—that include AI talent as the basic unit of operations. These teams focused on implementing and running major products or services will need to be more specialized, but they should still ultimately report to and be accountable to whatever mission or business center they support.

- Don't use AI for the sake of using AI; use AI where it will be an effective tool for dealing with the current and future state of the business and missions.
- Invest in AI tools, talent, and tradecraft within the business centers that are most able to use it. The mission owners who are best able to judge AI's value to their objectives make those decisions.

Getting to Goal #1: Place AI support around the use case, not technical skills

Perhaps the most common pitfall to avoid is creating a central group of practitioners that does AI work for different mission areas or program offices upon request. One example of this is a central team of data scientists that can be loaned to a specific mission center or program office.

Even when leaders recognize that IT work should and can be done with personnel native to the various mission centers and program offices, some IT offices will discourage such encroachments on their turf. In the meantime, the IT office cannot handle the increasingly domain-specific requirements of so many mission and business centers.

AI needs to avoid some of the challenges that IT modernization efforts can experience by being directly accountable to the customers that the work is supposed to support; *AI practitioners ultimately report to mission center and program office leaders*, just as all other personnel who work on those missions and business functions do. AI talent must be placed in the organizational chart, not in a central AI group or shared service.

- Avoid centralizing AI practitioners and leaders in one unit. AI talent must be accountable to the business needs and therefore should exist across the organization.
- Avoid “data scientist” or “AI staff” for loan situations that remove accountability to the mission center or program office responsible for implementing the AI solution.

Goal 2: Support the AI practitioner's effectiveness by creating a technical AI resource with the tools needed to get the work done

Goal one establishes that the AI workforce should be spread throughout the agency—in the business units. But how do we ensure that individual AI practitioners have the tools, resources and infrastructure they need to succeed?

Addressing this requires an organizational resource that functions as the one-stop-shop to provide all of the practitioner's technical resource needs:

Technical Tool Domain Examples

- development environments,
- server space
- code libraries, etc.

Technical tools do not include AI practitioners. The people doing the work are the AI practitioners who need the tools to do the work. These practitioners across the organization can easily access centralized tools that they need to build and implement an AI system.

Agency support resources

AI practitioners will also need support and guidance in areas such as legal considerations, acquisition, and security to fully integrate an AI system into existing processes.

Institutional resources:

- security review
- legal review
- acquisition support for buying licenses, etc.
- talent/HR support

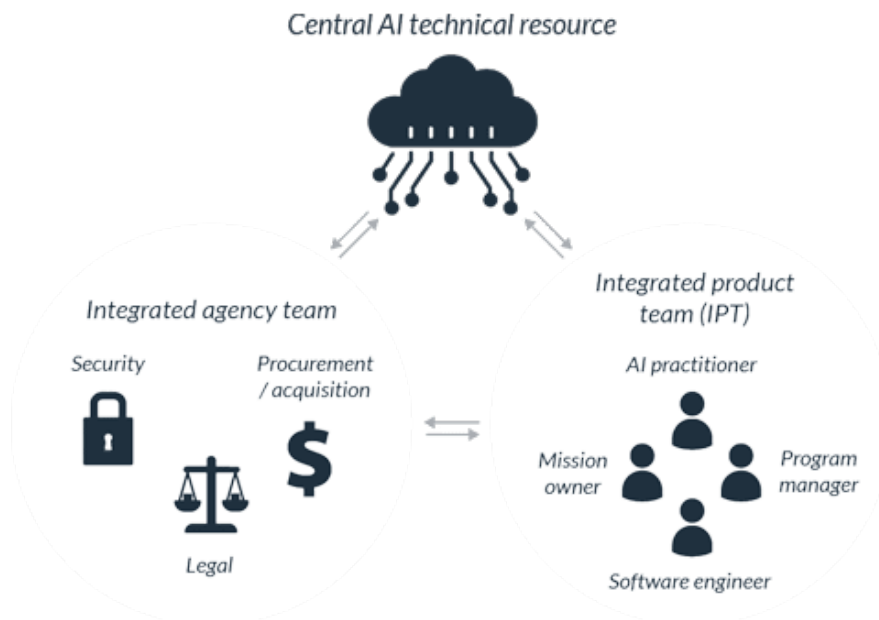
A central AI technical resource—discussed below—will serve as a hub of AI expertise for program offices to seek the support they need. These are members who help make up the Integrated Agency Team (IAT).

Chief Information Offices, Chief Information Security Offices, Chief Data Offices, acquisition and procurement offices, and privacy and legal offices must collaborate to establish this AI resource to support AI professionals in the mission and program offices.

A note on talent and HR resources

The people managing the central AI resource should also be involved in AI talent recruitment, certification, training, and career path development for AI jobs and roles.

Business units that have AI practitioners can then expect that AI talent will be of consistent quality and able to enhance mission and business effectiveness regardless of the customer's own technical understanding of AI.



- Embed AI professionals in the mission centers and program offices, but they should have access to a **one-stop-shop** for AI tools and resources in a central AI resource.
- The central AI technical resource provides technical and infrastructure support as well as access to legal, security, and acquisition support for AI professionals in the mission center and program office to succeed in AI adoption efforts.

Three key components of an AI-enabled organization



Component 1: The Integrated Product Team – doing project-based AI

Certainly, organizing an agency for successful AI development can be a challenge. With the organizational structure we've already shared in mind, let's shift our focus to the actual personnel and staffing required to implement an AI project. The central AI technical resource that provides tooling and security, legal, and acquisition support, combined with mission-based AI practitioners, make up an AI-enabled organization.

The Integrated Product Team (IPT) is responsible for implementing AI products, systems, and solutions. Given the output and work products of the IPT are generally software, structure these teams similar to a modern software team, applying agile/scrum principles. Led by a technical program manager, the bulk of the team should consist of AI practitioners and software engineers.

But in the spirit of the IPT, also consider roles like change management experts who can help develop training or processes or workflows that may be necessary, and user researchers who can generate real-time feedback about the work being done.

What makes an IPT team made up of a variety of roles so successful is that they each have the perspectives and knowledge necessary to ensure true value is delivered at all points of a project. The IPT will make many of the decisions that will impact the final deliverable(s).

Component 2: The Integrated Agency Team – Security, legal, and acquisition support for practitioners

Practitioners embedded within mission teams can leverage more standard tools, policies, and

guidance for their AI work. Therefore, they may not need as much security, legal, and acquisition support on a day-to-day basis.

On the other hand, IPTs usually deal with more novel situations that require deeper support from security, legal, and acquisition personnel. While an IPT should be structured like a typical modern software development team, consider augmenting the Integrated Agency Team (IAT) that supports the IPT with security, legal, and acquisition professionals needed to successfully launch the project.

The IAT should address these types of issues:

- Data rights
- Intellectual property provisions
- End-user licensing agreements
- Appropriations implications for different types of pricing
- The extent that software or hardware must be integrated into the existing infrastructure
- Security measures that must be complied with to start a pilot or scale a solution

The individuals who answer these questions tend to be found in the Office of General Counsel (OGC), or wherever the Chief Financial Officer (CFO), Chief Information Officer (CIO), Chief Technology Officer (CTO), or Chief Information Security Officer (CISO) are housed.

As with all other support services, you should coordinate these through the central AI resource so they will be easier to access. Some agencies' offices may lack enough specific knowledge to guide AI use. You'll need to provide training to build this knowledge.

Meeting with these individuals that make up the IAT at the beginning of your program's efforts will help establish the operating framework within the team as well as pain points that must be solved before moving into more involved phases of the project. Meeting with the IAT in the middle of planning will help validate the ideas being explored and ensure they are practical. Finally, meeting with everyone right before starting a solicitation or internal development effort in earnest will ensure that the entire organization is not only on the same page, but ready to react should a critical issue arise that has to be escalated.

Given AI's emerging and rapidly evolving nature, new and unforeseen challenges may come up at any time. Build a process that allows IPTs to ask questions or clarify issues at any point in the process, like, for example, a legal concern.

AI implementation is not only important from a technical perspective, but also from an administrative perspective. AI's technological implications will affect many people, from the public interacting with the agency or department to the employees supporting the program itself.

The value of an IAT is to support the IPT members who will actually develop and manage the day-to-day affairs, and finally, the true end-user who will actually interact with the final product. Without their insight, the IAT's input would go into vacuum, and the ability to achieve the true practical solution that results from a cross-functional collaboration is lost.

- Offices of General Counsel (OGC), Chief Financial Officer (CFO), Chief Information Officer (CIO), Chief Technology Officer (CTO), or Chief Information Security Officer (CISO) may need to provide answers or support to the Integrated Product Team (IPT). However, some offices may not be ready to support AI. Your agency may first require AI training.
- AI projects have many unseen or unique challenges and IPTs need easy access to support (technical, legal, etc.) these services as they develop.

Component 3: The central AI Technical Resource – Infrastructure, tools & resources

The technical arm of the organization, which is often the Office of the Chief Information Officer or the IT shop, plays a key role in ensuring that AI practitioners across the organization have access to the tools they need for AI development. Centralizing platforms has a number of benefits:

- **Ease of access.** With access to a suite of tools already available, AI practitioners—or the vendors that business teams interact with—can easily experiment with datasets, create proofs of concept or even deploy at scale with ready to use AI tools without having to go through the process of individually procuring, installing and gaining security approvals individually.
- **Cost and infrastructure optimization.** A shared resource allows for optimization and automation of shared infrastructure, which can translate into significant cost savings. By sharing infrastructure across the organization, the central AI technical resource can coordinate and optimize infrastructure usage for model training, hosting, and deployment.
- **Governance.** Creating a central AI technical resource allows for greater insight into all of the AI initiatives underway in the organization. This makes it possible to create and enforce common governance policy and create a structure for monitoring AI projects.
- **Expertise.** Though the central technical resource DOES NOT provide data scientists to loan out to the rest of the organization, it is staffed by deeply technical experts who can aid in the selection of additional AI talent. HR should coordinate with these experts when hiring AI talent to be staffed in the business units.

The path to reaching goals #1 and #2: getting to a central AI technical resource

Understandably, agencies may be in the early stages of their AI integrated product team (IPT) journey; fortunately, the path to get there is a natural and methodical one.

Mission leaders and AI practitioners should identify use cases where AI could easily have a large impact. Once found, they can make the case to the executives in charge of the mission or business area that covers those use cases and start building AI capabilities in those mission/business centers.

Most likely, these early teams will need to supply their own ad hoc infrastructure, which will limit their effectiveness. Choose a use case that is so compelling—ideally to agency leadership—that the teams are still able to show great value to the mission/business center. Early wins build momentum toward more mission centers and program offices wanting to incorporate AI capability to boost their own effectiveness.

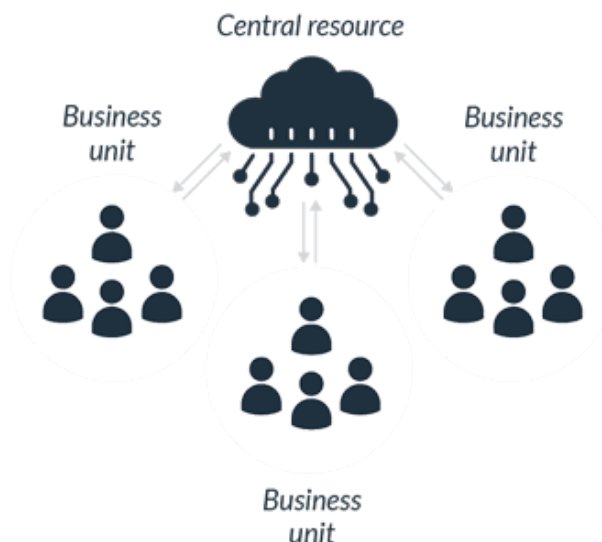
After more than a few mission centers and program offices start building AI, you should have critical mass. Move the AI infrastructure provider to a more accessible part of the organization, towards the central AI resource.

Crucially, these steps do not require immediate enterprise-wide changes. This model works at whatever scale of organic AI adoption the agency can handle at the time. Adding AI personnel to more mission centers and program offices and continuing to scale up AI practitioners' skills offers a natural and gradual path to the goal state of enabling all agency functions to use AI.

How an agency might start



Where an agency should be



- Start small with a use case that focuses on a unique mission or program challenge.
 - As more small use cases emerge, consolidate technical infrastructure to support these AI projects to avoid buying duplicate and overlapping infrastructure. Infrastructure may sit in one specific mission center or program office, but share it with others.
 - Once there is a critical mass, which will depend on your agency, move towards a central AI technical resource at the enterprise level.
-

Chapter 3: Responsible and Trustworthy AI Implementation

AI's continually evolving landscape promises to change businesses, governments, society, and the world around us. Like with development of any other technology, we must carefully consider how AI affects the people who use it and are impacted by it.

Building responsible and trustworthy AI systems is an ongoing area of study. Industry, academia, civil society groups, and government entities are working to establish norms and best practices. The National Institute of Standards [defines \(https://www.nist.gov/programs-projects/trustworthy-and-responsible-ai\)](https://www.nist.gov/programs-projects/trustworthy-and-responsible-ai) the essential building blocks of AI responsibility and trustworthiness to include accuracy, explainability and interpretability, privacy, reliability, robustness, safety, security, and importantly the mitigation of harmful bias. These building blocks raise important questions for every AI system such as “safety for whom?” or “how reliable is good enough?”

Some of the worst negative impacts are due to harmful biases in AI system outcomes, including many cases where AI further entrenches inequality in both the private and public sectors. One of the key ways the outcomes of AI systems become biased is by not carefully curating, evaluating, and monitoring the underlying data and subsequent outcomes. Without this preparation, there is a greater likelihood that the data used to train the AI is unrepresentative for the proposed use case. For many use cases, a biased dataset can contribute to discriminatory outcomes against people of color, women, people with disabilities, or other marginalized groups. Along with bias in the input datasets, the design of the AI system (such as what to optimize for) or simply how the results are interpreted also can cause biased outcomes.

An additional concern with biased outcomes is that the “black box” nature of the system obfuscates how a decision was made or the impact of certain decisions on the outcomes. Due to this, biased AI systems can easily perpetuate or even amplify existing biases and discrimination towards underserved communities. Algorithmic bias has been shown to amplify inequities in health systems and cause discriminatory harm to already marginalized groups in housing, hiring, and education. The people that build, deploy, and monitor AI systems, particularly those deploying government technology, must not ignore these harmful impacts.

Section 3.1: AI principles and guidelines are still evolving

Because what we consider AI currently is so new, there are a lot of uncertainties and nuances around how to embed responsibility into AI systems.

As discussed previously, responsibility includes: accuracy, explainability and interpretability, privacy, reliability, robustness, safety, security, the mitigation of harmful bias, and more as the field evolves. A challenge to science generally is that there are no perfect answers or approaches. Due to the speed and scale of progress in AI, practitioners in the field likely will be learning by trial and error for the foreseeable future.

Some foreign governments, international entities, and U.S. agencies have already begun to create high-level AI principles, and even some policies around AI's responsible and trustworthy use. These are important first steps, but next these principles must be translated into actionable steps that agencies can use throughout the AI development process.

When it comes to the practical implementation of AI in government — again with the fundamental requirement of responsible and trustworthy AI — researchers and practitioners are continually iterating and learning. If readers of this guide want to dive more deeply into responsible AI, there are numerous sources including within the Federal government such as the [Department of Defense Ethical Principles for AI \(https://www.defense.gov/News/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/\)](https://www.defense.gov/News/Releases/Release/Article/2091996/dod-adopts-ethical-principles-for-artificial-intelligence/) and the [Executive Order on Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government \(https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government\)](https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government).

Section 3.2: Why is DEIA essential for a responsible and trustworthy AI in practice?

To provide the AI team with the best tools for success, the principles of DEIA should be at the forefront of any technology project. Responsibility for ensuring responsible design decisions that result in equitable outcomes falls on all team members, from the practitioners to managers.

The importance of this can be illustrated by examining all of the different decisions within the AI development lifecycle. These decisions, especially ones that might be considered purely “technical,” require thorough consideration including an assessment of their impact on any outcomes. The following two examples illustrate the importance of DEIA considerations within AI system design and development:

One suggested action to encourage deep consideration of key technical questions during AI system design is to implement iterative review mechanisms, particularly to monitor for bias, and be transparent regarding tradeoffs made in decision making regarding model performance and bias. This process begins with the assumption that there are biases baked into the model, and the review’s purpose is to uncover and reduce these model and historical biases.

These may seem like technical questions that a leader, program, or project manager may not normally focus on. However, successfully managing an AI project means establishing structures that ensure responsible and trustworthy AI practices are the responsibility of the entire team, and not just left to the day-to-day developers. As demonstrated, seemingly simple day-to-day design decisions that AI teams make have implications for marginalized communities. Contributors from the entire team, which can include designers, developers, data scientists, data engineers, machine learning engineers, product owners, and project and program managers must work together to inform these decisions.

Section 3.3: Ask questions often and repeatedly

To drive this point home: A good starting point to responsibly implement AI is to ask questions, especially around key decision points. Ask them early; ask them often.

Ask the same question over and over again. (Answers might change as the team learns.) Ask different people on the team to get a collection of answers. Sometimes, you may not have an answer to a question immediately. That's ok. Plan to get the answer, and be able to explain it, as the project progresses.

As each project differs, the questions required to assess for responsibility and trustworthiness may be different. The questions outlined in this module are designed to guide teams that are building and implementing AI systems but are not official standards or policy. Rather, they are good questions to consider as your team progresses through the AI lifecycle to begin to embed responsible and trustworthy AI in the process. These questions are intended to foster discussions around broad ethics topics. Ask these questions often through the design-develop-deploy cycle and combined with testing to help reduce unintended consequences.

It's too late to start asking about responsible and trustworthy AI implementation when you're readying a complex system for production. Even if you are only one person playing with data to see if AI might be a possible solution, ask these questions. Some of the questions may not apply in the early stages of discovery. That's ok. Continue to ask them as the project evolves and document answers to these questions to track your project's progress. These answers will help identify when in the AI lifecycle these questions will become relevant and can inform future systems.

Here are some suggested questions that any team attempting to develop responsible and trustworthy AI needs to consider:

1. Focus on the root problem

Government research projects and pilots are all looking to improve the function of our government, be it via a better chatbot to help with customer service or to detect cybersecurity threats faster and more efficiently. Whatever their purpose, exploring the use of new technologies, such as AI must be done in a way that evaluates whether AI is actually the best-fit solution. Teams that are building models and systems need to clearly understand the problem to be solved, who is affected by this problem, and how AI may — or may not — be a solution.

Questions to consider include:

- Why are you considering using an AI solution in the first place?
- Is it the best option to solve this particular problem? Have you evaluated alternative solutions?
- Will it actually solve the problem? What metrics are important to assess this hypothesis and how will you measure them?
- Will it equally benefit all users or just disproportionately help some, possibly at the cost to others?

Like previously highlighted, creating a team environment where all stakeholders are educated and empowered to participate in evaluation of these types of questions is essential. For example, if the metrics don't require assessment of accessibility of the chatbot tool, the right questions were not asked.

2. Be accountable to the users

AI systems cannot exist in isolation. The outcomes produced by these systems must be able to be justified to the users who interact with them. In the case of government use, users range from government employees to recipients of benefits. This may also mean the systems must be able to demonstrate how the answer is reached, which is also critical to identifying the cause of negative outcomes. This also means that a person, or a team, needs to own the decisions that go into creating the systems.

Question to consider include:

- When something deviates from the intended output or behavior, who is responsible for noticing and correcting this?
- Is someone responsible for making sure that every step is not just done, but done correctly?

The process starts with establishing clear roles and responsibilities for data and model management. At a minimum, an aberrant outcome can be linked to its training source. This is often significantly harder than you would think, especially in the case of deep learning. Nevertheless, it should be ongoing.

3. Define and avoid harm

Researchers, advocates, and technologists on AI teams have concerns about numerous types of harms caused by ill-designed AI systems. These include risks of injury (both physical or psychological), denial of consequential services such as opportunity loss or economic loss, infringement on human rights (such as loss of dignity, liberty, or privacy), environmental impact, and even possible erosion of social and democratic structures. In looking at these harms it is important to remember that bias can impact who suffers from any of these types AI harms.

Bias can enter an AI system in many ways. While, some of the most commonly discussed bias issues

are about discriminatory opportunity loss, seen in employment, housing, healthcare, and many other fields, it's important to remember bias occurs in many forms. For example, a biased AI system for say self-driving cars could cause increased rates of physical harm to people with disabilities requiring mobility aids simply because the model data for pedestrians mostly consists of able-bodied data subjects.

Though it may be impossible to completely eliminate all bias (and that may not even be the goal) an AI team must be able to evaluate what possible harms of their system could be and how bias might cause disparate negative impacts across different populations. To reduce this possibility, the team must evaluate for bias in datasets, the model, and the design choices throughout the product life cycle. It must also evaluate for bias in the outcomes the systems produce to ensure the output does not disproportionately affect certain users.

Questions to consider include:

- What are the possible negative impacts of these systems? How do we measure this harm and what could we do to mitigate that impact?
- What is the demographics of people involved in the domain that the AI system works within? Who are directly and indirectly impacted?
- What data is required to ensure equitable outcomes across the universe of people affected?

4. Monitor the outcomes

Even after asking essential questions during system design and development, the team must rigorously monitor and evaluate AI systems. The team should create structured oversight mechanisms and policies, ideally developed throughout the design process and in place before implementation, to identify anything that is potentially causing a problem so they can intervene quickly.

Questions to consider include:

- Are there regular management reviews of changes made to the input, throughput, or output of the developed system?
- Are there clear roles and responsibilities for the management of the AI system?
- Are there automated system checks for issues such as model drift, anomalous behavior, or other potential changes?
- Are the systems auditable so that the drivers of incorrect or inequitable outcomes can be identified and fixed?
- Does the AI system provide clear notice of its use to impacted people, including what relevant factors are important to any decisions or determinations? Is there a mechanism for impacted people to contest, correct, or appeal or even opt out of the use of an AI system?

Of course, oversight will not solve all potential problems that may arise with an AI system, but it does create a plan to watch for, and catch, some of the foreseeable issues before they become harmful.

Section 3.4: Moving forward

This chapter is a first step in responsible and trustworthy AI implementation, but like the iteration and innovation occurring in AI, this will be an ongoing effort. Asking these types of questions will not solve all challenges, nor does answering them ensure compliance with any standards, guidelines, or additional principles for using AI responsibly. Practitioners must be critically thinking of these key questions, considerations, and potential risks while building or implementing AI systems.

As this space is evolving rapidly, many experts are putting considerable thought into how to [implement responsible and trustworthy AI \(https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270-draft.pdf\)](https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270-draft.pdf) principles and embed them into the [day-to-day operations \(https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf\)](https://www.nist.gov/system/files/documents/2022/03/17/AI-RMF-1stdraft.pdf) of a product or project team. As authoritative literature in responsible and trustworthy AI grows and changes, it's advantageous to stay up-to-date and look to members of the AI community, both in the public and private sector for guidance. The Technology Transformation Services Centers of Excellence look forward to helping you on this journey and will provide updates as the field grows.

Chapter 4: Developing the AI workforce

Even the most advanced and technically sound AI solutions will fail to reach their full potential without a dedicated team of people that understand how to use them.

The main questions are:

- How do we do this?
- How do we get from where we are now to an AI-capable workforce?
- How do we attract and retain the right people and take care of them as employees?

This chapter will discuss what an Integrated Product Team might look like, how to build and manage AI talent, and how to develop learning programs that cultivate transformational AI capabilities.

Section 4.1: Start with people

Identifying AI talent

As much as you can, survey your organization to map existing analytics talent or teams with an analytics orientation. Though analytics and AI are not the same, there are many overlapping baseline skills. Existing analytics knowledge can grow into AI knowledge.

If there are already people in the organization with AI skills, where do they sit and who do they report to? Are they in IT, in one of the business functions, or part of the Office of the Chief Experience Officer (CXO)?

How do you know if your existing talent has AI skills? Start looking for people who exhibit some of these qualities:

- Support their decisions and arguments with data
- Are comfortable with statistics and math
- Make their own macros in excel
- Have expressed an interest in or started to learn computer programming
- Recognize that technology can make a process faster, easier or more efficient
- Know the data your organization uses well
- Follow the latest technology trends closely

An important part of assessing an organization's existing talent is acknowledging that some people may already be leveraging defined AI and ML skills. Others, however, may work in technical roles or have skills that are not directly AI related, but could easily be supplemented to become AI skills.

Your organization employs intelligent and skilled people who may already be working in AI and ML. It may also have an even broader pool of people with skills related to AI. These people may not even realize that they already have many of the skills and capabilities to help use AI and data science to advance the objectives. Your agency can train these people so they can become AI professionals.

Augment talent when needed

Certainly, many agencies want to increase the AI know-how of their internal staff. However, much of the innovation emerging in the AI field comes from private industry. Public-private partnerships are often an excellent way to get more support for AI projects.

When to bring in outside talent or vendors:

- The agency has had difficulty attracting, training, and retaining data science talent to achieve some of its objectives.
 - The use cases in question are limited and require niche skills that may not be worth hiring for and developing over the long term. These niche skills are needed for the long-term solution's maintenance, not only for the build.
 - The agency needs to quickly test the potential benefits of an AI solution before deciding whether to invest in developing internal capabilities. However, the reverse may also be true. An agency may wish to use internal talent to quickly test a new capability being pitched by a vendor before deciding to invest in the outside resources. We discuss this further in [Chapter 7, Module 4: Starting an AI project \(../starting-an-ai-project/index.html#internal-prototype-and-piloting\)](#).
-

Section 4.2: Developing and retaining AI talent

Mission and practitioner support

The most powerful tools for retaining government AI talent are ensuring that AI work is closely tied to the agency mission and ensuring that AI talent has the technical and institutional support to work effectively as AI practitioners.

This combination forms the unique value proposition for an AI career that only federal agencies can provide, and is usually the reason AI practitioners chose government over industry and academia.

If AI practitioners discover that their work isn't clearly contributing to the agency mission, they are unlikely to stay because they could do data science work outside for better money or causes they care about. If AI practitioners love the agency mission but are unable to function as AI practitioners, they are also unlikely to stay if the agency is unable to leverage their skill set. Both meaningful work and practitioner support are crucial for retaining AI talent

Retention incentives and skill development

One way to make the best use of these usually limited incentives is to ensure federal employees have full awareness and access to AI related training and skill development opportunities.

This will demonstrate that agencies are committed to the progression of employee skill and career development, and encourages AI talent to invest in their careers.

Formal education

AI and data science are fields that often require a significant technical and academic background for success. However, it's also important for people to be open-minded about who might have (most of) the relevant skills and capabilities.

They should not assume that only people with computer science or statistics education are going to be appropriate for AI-centric positions. A culture that prizes and generously supports learning not

only ensures the continued effectiveness of the AI workforce, but also serves as a powerful recruitment and retention tool. Agencies should recruit AI talent at all career stages; bringing in early-career AI talent offers a special opportunity to create a cadre of AI practitioners with deep experience with the agency. But this opportunity requires investing in formal education for these early-career practitioners in order to realize their full potential.

Many agencies already have formal education programs; for these programs to be most effective for AI practitioners, they need to be more flexible than they are now. For example, full-time degree programs should be eligible for tuition reimbursement, not just part-time programs. Agencies can make up for the higher cost of full-time degree programs by extending service agreements accordingly. Agencies shouldn't force their best AI talent to choose between continuing employment and attending the most competitive degree programs, which tend to require full-time attendance.

Training, conferences, and exchanges with industry and academia

In AI and data science, the state of the art advances by the month. AI practitioners must regularly attend industry and academic conferences to maintain their effectiveness.

AI practitioners who feel they may be falling behind in their field while working in government are more likely to leave to maintain their own competence as AI practitioners; agencies should actively prevent this situation to improve retention.

In general, interaction with industry and academia allow government AI practitioners to benefit from the vast scale of innovation happening outside the confines of government.

One of the deepest forms of government interaction with industry and academia are exchanges where government AI practitioners spend a limited time fully embedded in a partner organization. That language is now codified into law in the 2019-2020 [National Defense Authorization Act \(https://www.congress.gov/bill/116th-congress/house-bill/6395\)](https://www.congress.gov/bill/116th-congress/house-bill/6395). After completing these assignments, AI practitioners return to their agencies armed with the latest best practices and new ideas. From the other end, partner organizations can embed some of their employees in government agencies, bringing fresh perspective to the agency and offering access to a different pool of AI talent. Partner organizations benefit from these exchanges by promoting better government understanding of their industry or institution, while also developing contacts and relationships with government agencies relevant to their domains.

Section 4.3: Understanding AI job roles and career path

AI practitioner ecosystem

Chapter 2 outlines where AI practitioners should sit within mission areas and program offices. Mission areas should create a space for these emerging data science roles to become part of an Integrated Product Team (IPT) ready to take on AI implementation.

Typically, those roles include the following:

- **Data analyst:** focuses on answering routine operational questions using well-established data analysis techniques, including AI tools.
- **Data engineer:** focuses on carefully building and engineering data science and AI tools for reliability, accuracy, and scale.
- **Data scientist:** focuses on thoughtfully and rigorously designing data science/AI models, tools, and techniques. A data scientist should usually have an advanced technical degree and/or significant specialized technical experience.
- **Technical program manager:** manages software development teams, including teams building AI tools and capabilities. The job responsibilities of the role are nontechnical, as with all management roles, but a technical background greatly enhances this particular type of manager's effectiveness.

However, AI practitioners are not only doing technical work. When agencies are planning AI projects, it's important to narrow in on the sponsors and individuals required to execute key project components.

Roles that support data science teams should include:

- **AI champion:** Advocates for the AI solution's value, but ensures the clear, effective, and transparent communication of the AI solution to ensure that it is developed responsibly and produces the intended results.
- **Project sponsor:** Identifies and approves opportunities and makes go/no-go decisions. This person coordinates with the AI champion, if they are not the same person, to communicate progress up and down the chain of command.
- **Mission or program office practitioner:** Identifies opportunities and provides business and workflow understanding. This person knows the organization's mission and the day-to-day

details of the work performed. This person helps ensure that the AI solution not only performs the task intended, but can also integrate with and the existing program office team.

- **Project manager:** Ensures day-to-day progress and communicates with stakeholders and vendors.
- **Business analyst:** Provides business, financial, and data understanding.

The roles above may need to liaise among data science, IT, and the mission area's business needs. The number of most of these roles varies depending on the size of the initiative.

An AI's project success depends on the makeup of the Integrated Project Team (IPT). Though technical know-how is certainly important, without adequately understanding the challenge you are trying to address and getting buy-in from the mission and program team, the project will fail.

How is this different from any other IT project team?

Due to the iterative, data-dependent nature of AI, misguided or unsupported AI development could have serious consequences down the road.

Career path

While the most common starting point of a data science career is the data analyst role, AI-focused practitioners tend to have more of a computer science background.

They may be more likely to start as a junior data engineer or a junior data scientist. AI practitioners with a pure math background will probably start as a junior data scientist. Data engineering continues to be its own track; otherwise, with more experience and ideally an advanced technical degree, the practitioner becomes a full-fledged data scientist.

Agencies with significant AI implementation talent may also have senior technical positions such as **senior data architect** or **principal data scientist**; these expert roles usually indicate extensive technical experience and tend to have decision-making authority on technical matters, and/or advise executives. Some agencies also have academia-like groups dedicated to research and not part of mission or business centers; these groups have positions like **research scientist**, which tend to require PhDs and very specialized technical knowledge.

AI practitioners may also choose to pursue a management career path, with the most natural transition being from data engineer or data scientist to technical program manager. After that, because data science is embedded in mission and business centers, AI technical program managers

are on the same track for higher management positions as all other front-line management positions in mission and business centers.

Section 4.4: Recruiting AI talent

Competing with private industry

As AI has become prominent in recent years, the government has problems hiring AI talent. However, the government has to solve this problem if agencies want to remain relevant in the future.

The government cannot compete with private industry on salary and bonuses, but it CAN compete on offering interesting, meaningful work and recognition. Federal recruitment can use this unique advantage when AI work is closely tied to meaningful mission and business objectives that only federal agencies offer.

AI practitioners, even if they love the agency's mission, expect to actually practice AI in their jobs. That's why the supportive and powerful work environment that the central AI resource provides is just as important to the pitch as creating space for AI practitioners in mission areas and program offices.

Centralized recruitment and certification

The [central AI resource \(./organizing-managing-ai/index.html#goal-2-support-the-ai-practitioners-effectiveness-by-creating-a-technical-ai-resource-with-the-tools-needed-to-get-the-work-done\)](https://www.gao.gov/assets/330/330111/330111main.pdf), which is the place in the organization that provides all technical and institutional support to AI practitioners, knows how to actually practice AI in the agency. They are also the group most able to certify that AI talent coming into the agency are well-qualified for their roles, and suitable for the agency's particular practitioner environment.

For example, the central AI resource knows whether certain programming languages or certain hardware capabilities are prevalent. They can assess candidates' suitability accordingly. If there's a strategic decision to increase certain platforms or skill sets, the AI resource knows how to do that. While the agency's HR office is still ultimately in charge of all workforce recruitment, the AI resource works closely with HR to provide the AI domain expertise.



Placing AI talent

The central AI resource, with connection to resources like technical infrastructure, data, security, legal and human capital support, supplies a pool of well-qualified candidates for the agency. The mission and business centers looking to fill AI roles should coordinate with existing AI practitioners who know the subject to evaluate whether candidates are qualified as AI practitioners. Once the AI resource confirms candidates' AI capabilities, the mission and business centers can focus on how these AI qualified candidates can contribute to their mission and program goals. Mission centers and program offices should also coordinate closely with the AI resource to ensure that the pool of vetted candidates aligns with staffing needs.

Chapter 5: Cultivating Data and Technology

Chapter 5 covers elements of the data and technological infrastructure that goes into building AI capabilities. This includes tools, capabilities, and services as well as governance and policy to manage these tools and the data that feeds them.

Section 5.1: Tools, capabilities, and services

New AI tools, capabilities, and services are released almost daily. They promise to revolutionize the way government operates. When evaluating these AI tools and capabilities, note that there's more to AI than simply building models. This is particularly true when considering more customizable options of interactive AI platforms or building from scratch.

You'll need to evaluate development environments, infrastructure, data management, data manipulation and visualization, and computing power technologies. Some of these are offered as services through software, platform or infrastructure as a service (SaaS, PaaS and IaaS). Some are available as hardware or software installations, and some are available open source.

Though not an exhaustive list, the tools and platforms outlined below highlight what you may need to create an AI solution.

Cloud & Infrastructure

Many AI tools and solutions are tied to a cloud platform. Elastic storage, computing infrastructure, and many pre-package ML libraries help accelerate ML model development and training. Agencies with limited in-house computing resources need a cloud platform; so do ML models that require intense computing resources for training, such as for Deep Learning and GPU acceleration. A cloud platform can be more economical when the computing requirements are short-term and sporadic, depending on data security requirements.

Use orchestration tools to help manage complex tasks and workflows across the infrastructure. A variety of open source tools are available.

DevSecOps

DevSecOps is the integrated practice of bringing together software development, IT operations, and the security team. It's a critical part of successful AI delivery.

At a high level, DevSecOps includes an environment to manage development tools such as:

- programming languages like Python, R, Java and C++
- code repositories
- build and unit testing tools

- version control management for code and models
- tools to manage code quality
- version control management to perform security scans, and monitor and perform testing and ongoing performance

Data Management

Data collection, ingestion, management, and manipulation are AI development's more critical and challenging elements. Tools are available to handle various tasks associated with data operations, including: tools for data acquisition, data cataloging, data collection and management frameworks, data ingestion frameworks, data labeling tools, data processing tools and libraries, data sharing services, and data storage.

Not every AI solution requires every tool. The relevant tools depend on the size, complexity, structure, and location of the data being used to train AI models.

Artificial Intelligence (AI) and Machine Learning (ML)

Many new tools and products support AI development. These include data science toolkits (combined offerings for applied mathematical statistics); visualization tools to explore data, understand model performance, and present results; and machine learning frameworks that provide pre-build architectures and models to reduce the development effort.

AutoML tools—tools that automate many of the model training and deployment processes—are an emerging area of AI tool development. They further reduce the access barrier to AI technology. A number of these solutions offer low-coding ML tools.

Many of these tools are open source, but many are offered as commercial products. Selecting the right tools for the job will require careful evaluation and involve all members of the Integrated Product Team (IPT). These tools are often provided by and operated by the central AI resource.

Section 5.2: Data governance and management

Data governance and management are central to identifying AI use cases and developing AI applications. Data governance is the exercise of authority and control (planning, monitoring, and enforcement) over data assets.¹

Though Chief Data Officers (CDOs) and their counterparts have many resources available to them, we are presenting some of the key takeaways for the benefit of all those interested in data governance, especially as it relates to moving towards AI.

Legislation and guidance

In recent years, data governance and management have been codified into statute via The Foundations for Evidence-Based Policymaking Act² (Evidence Act) that requires every executive branch agency to establish a Chief Data Officer (CDO) and identifies three pillars of work for which the CDO bears responsibility: data governance; the Open, Public, Electronic, and Necessary (OPEN) Government Data Act³; and the Paperwork Reduction Act⁴ (PRA).

This legal framework assigns the CDO the area of responsibility for “lifecycle data management” among others to improve agencies’ data management.

Offering further guidance, the Federal Data Strategy⁵, which describes 10 principles, and 40 practices designed to help agencies and the federal government improve their data governance practices. Find more details about implementing the Federal Data Strategy at strategy.data.gov.

Data governance organization

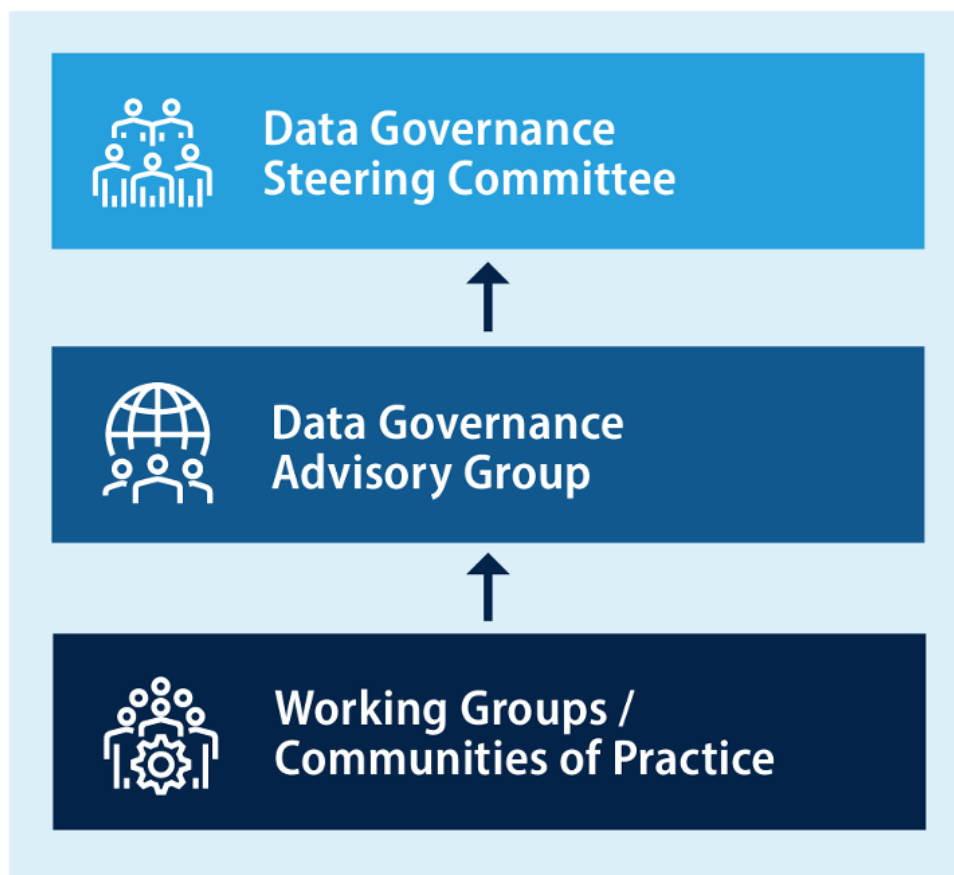
M-19-23, issued by the Office of Management and Budget (OMB) and titled “Phase I Implementation for the Foundations for Evidence-Based Policymaking Act of 2018,” provides implementation guidance pertaining to the Evidence Act. It states that each agency must establish an agency data governance body chaired by the CDO by September 2019 to support implementing Evidence Act activities⁶. The 2019 Federal Data Strategy (FDS) Year One Action Plan administers similar requirements⁷.

In executing the [Evidence Act \(https://www.whitehouse.gov/wp-content/uploads/2019/07/M-19-23.pdf\)](https://www.whitehouse.gov/wp-content/uploads/2019/07/M-19-23.pdf) and corresponding guidance, practical requirements have been found for resource allocation, technical leadership, and practitioner and user input.

Many avenues address these requirements organizationally. Combining any three of these

requirements may work for agencies depending on their specific factors. Agencies who would like to separate these requirements into separate organizational functions might do so with the organizational charts and description below:

- **Data Governance Steering Committee** - Chaired by the CDO (mandated in OMB 19-23), the Committee makes resourcing and policy decisions for enterprise and lifecycle data management.
- **Data Governance Advisory Group** - The Group collects user and practitioner needs from the enterprise, generates recommended solutions, and prepares decision documents and corresponding materials regarding data policy and resourcing for the CDO and steering committee.
- **Working Groups / Communities of Practice** -Practitioners who may or may not be involved in advisory group activities. These groups bring specific data governance and management needs to the advisory group to judge.



Establishing a multi-tier governance structure consisting of working groups, advisory boards, and decision-making bodies can distribute decision-making authority across tiers so activities and

decisions can be made quickly. Consider elevating decisions only when they cross a defined threshold like resource allocation or level of effort.

Data governance enables organizations to make decisions about data. Establishing data governance requires assigning roles and responsibilities to perform governance functions such as:

- data strategy, policy, and standards
- oversight and compliance
- sponsoring and reporting related to data management projects
- issue elevation and resolution

Example outputs of data governance activities include (but are not limited to):

- data governance framework
- data strategy and policies
- data asset inventories
- data quality plans
- data management scorecards
- business glossary
- communications plan
- data management processes
- best practices & lessons learned
- workforce skills assessment

Data lifecycle management through metadata tagging

Data lifecycle management is the development, execution, and supervision of plans, policies, programs, and practices that deliver, control, protect, and enhance the value of data and information assets throughout their lifecycles⁸. Data management in the context of this guide focuses on the data lifecycle as it moves through an AI project.

Many AI projects are iterative or involve significant monitoring components. Any given dataset is not itself static and can quickly change as the project uncovers new insights. Thus, we need ways to manage those updates, additions, and corrections.

Data is an asset that has a value in business and economic terms. Information on data sources and their subsequent use can be captured, measured, and prioritized much like business decisions of physical inventory assets.

Secondly, data management is multidisciplinary. This activity requires a broad range of perspectives and interactions across many different classes of “users” that make up the IPT, including data scientists, engineers, mission owners, legal professionals, security experts, and more. Using existing data governance processes to engage those stakeholders is essential to effectively managing data.

Data management activities start with identifying and selecting data sources, framed in the context of business goals, mission-defined use cases, or project objectives. As you identify data sources, have engineering teams integrate them into the overall data flow, either through data ingestion or remote access methods. Include a clear description of the data through relevant metadata with datasets as they are published.

Effective data governance will influence resourcing decisions based on the overall business value of datasets. In order to influence these resourcing decisions, get usage metrics and business intelligence across your data inventory.

User research and use case development help governance bodies understand high-impact datasets across different members of the IPT or mission areas.

One example of data lifecycle management is standardizing metadata captured for new data sources by populating a data card used to describe the data.⁹ Each dataset should contain common interoperable metadata elements that include, but are not limited to, the following:

- data source origin
- originating collection authority and organization
- format and type of data (e.g. jpeg, text, wav, etc)
- size of data
- periodicity of data transfer (one-time batch, every two weeks, real-time streaming)
- requisite security controls
- pre-processing or transformations applied to the data
- points of contact for data suppliers, owners, and stewards
- data makeup and features, including obvious and non-obvious constraints, such as variable names and meanings

While minimum tagging requirements vary across different organizations, the list above is provided as a general guideline. For operations that are highly specific or deal with high-impact or sensitive data, the receiving organization may need to capture more metadata fields earlier in the data lifecycle.

One example of this is access rights and handling, need-to-know, and archiving procedures

associated with classified data, which requires highly restricted and governed data handling procedures.¹⁰

Other descriptive dimensions include the mission or use case context. For example, metadata can be applied to datasets that detail information on the Five Vs, listed below:

- Volume
- Velocity
- Veracity
- Variety
- Variability

In addition to the Five Vs, apply use case specific metadata to datasets, which further supports future data curation and discovery efforts. Some examples of use case specific metadata, as noted in the NIST Big Data Working Group Use Case publication¹¹, include:

- use case title
- description
- goals
- current solutions (hardware, software, etc.)
- additional stakeholders/contributors

Once data is accurately described, it can be cataloged, indexed, and published. Data consumers can then easily discover datasets related to their specific development effort. This metadata also sets the basis for implementing security procedures that govern data access and use.

Data should be securely accessed and monitored throughout its useful lifespan, and properly archived according to handling procedures set forth by the original data owner.

Footnotes

1. Earley, Susan, and Deborah Henderson. 2017. *DAMA-DMBOK: data management body of knowledge*. ↩
2. “The Foundations for Evidence-Based Policymaking Act of 2018.” P.L.115-435. Jan. 14, 2019. <https://www.congress.gov/115/plaws/publ435/PLAW-115publ435.pdf> (<https://www.congress.gov/115/plaws/publ435/PLAW-115publ435.pdf>) ↩
3. The OPEN Data Act is part of P.L.-115-435. ↩
4. The Federal Data Strategy <https://strategy.data.gov/> (<https://strategy.data.gov/>) ↩

5. In addition to Evidence Act, administration policy published by the Office of Management and Budget titled “Managing Information as an Asset” (M-13-13) declares that executive departments and agencies “must manage information as an asset throughout its life cycle to promote openness and interoperability, and properly safeguard systems and information.” Corresponding policy requirements detailed in M-13-13 include the adoption of data standards, the development of common core metadata, and the creation and maintenance of an enterprise data inventory. ↩

6. Ibid. Pg. 20. ↩

7. “Draft 2019-2020 Federal Data Strategy Action Plan.” Federal Data Strategy Development Team. June 2019. Pg. 11. <https://strategy.data.gov/assets/docs/draft-2019-2020-federal-data-strategy-action-plan.pdf> (<https://strategy.data.gov/assets/docs/draft-2019-2020-federal-data-strategy-action-plan.pdf>) ↩

8. Earley, Susan, and Deborah Henderson. 2017. *DAMA-DMBOK: data management body of knowledge*. ↩

9. Gebru, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. “Datasheets for Datasets.” ArXiv:1803.09010 [Cs], March 19, 2020. <http://arxiv.org/abs/1803.09010> (<http://arxiv.org/abs/1803.09010>)

DISCLAIMER: The chapter includes hypertext links, or pointers, to information created and maintained by other public and/or private organizations. We provide these links and pointers only for your information and convenience. When you select a link to an outside website, you are leaving the GSA.gov site and are subject to the privacy and security policies of the owners/sponsors of the outside website. GSA does not control or guarantee the accuracy, relevance, timeliness, or completeness of information contained on a linked website. We do not endorse the organizations sponsoring linked websites and we do not endorse the views they express or the products/services they offer. The content of external, non-Federal websites is not subject to Federal information quality, privacy, security, and related guidelines. We cannot authorize the use of copyrighted materials contained in linked websites. Users must request such authorization from the sponsor of the linked website. We are not responsible for transmissions users receive from linked websites. We do not guarantee that outside websites comply with Section 508 (accessibility requirements) of the Rehabilitation Act. ↩

10. “Intelligence Community Enterprise Data Header” Office of the Director of National Intelligence. <https://www.dni.gov/index.php/who-we-are/organizations/ic-cio/ic-cio-related-menus/ic-cio-related-links/ic-technical-specifications/enterprise-data-header> (<https://www.dni.gov/index.php/who-we-are/organizations/ic-cio/ic-cio-related-menus/ic-cio-related-links/ic-technical-specifications/enterprise-data-header>) ↩

11. NIST Big Data Interoperability Framework: Volume 3, Big Data Use Cases and General Requirements [Version 2]. June 26, 2018 <https://www.nist.gov/publications/nist-big-data-interoperability-framework-volume-3-big-data-use-cases-and-general> (<https://www.nist.gov/publications/nist-big-data-interoperability-framework-volume-3-big-data-use-cases-and-general>)

[general](#) ↩

Chapter 6: AI Capability Maturity

This AI Capability Maturity Model (AI CMM), developed by the Artificial Intelligence Center of Excellence within the GSA IT Modernization Centers of Excellence (CoE), provides a common framework for federal agencies to evaluate organizational and operational maturity levels against stated objectives.

The AI CMM is not meant to normatively assess capabilities, but to highlight milestones that indicate maturity levels throughout the AI journey.

The AI CMM is a tool that helps organizations develop their unique AI roadmap and investment plan. The results of AI CMM analysis enable decision makers to identify investment areas that meet both near-term goals for quick AI adoption and broader enterprise business goals over the long term.

Section 6.1: AI development founded on good software practice

Much of what is discussed in this and other chapters relies on a foundation of good software development practice.

Though this chapter is not intended to explain agile development methodology, Dev(Sec)Ops, or cloud and infrastructure strategies in detail, these are fundamental to successfully developing AI solutions. As such, the AI CMM elaborates on how this solid IT infrastructure leads to the most successful development of an organization's AI practice.

Section 6.2: Organizational maturity areas

Organizational maturity areas represent the capacity to embed AI capabilities across the organization.

The first approach is top down, generally used by senior leaders and executives. It's influenced by the Capability Maturity Model Integration (CMMI) system¹ of evaluating organizational maturity. This particular top-down method helps you evaluate organizations. The top-down view in the table below is based on CMMI levels 1-5.

The second approach is user-centric, generally used by program and mission managers, who focus on delivery. It uses the CMMI system as an industry standard for evaluating organizational maturity. This bottom-up view is synthesized from CoE engagements with federal agency partners, market research, case study evaluation, and lessons learned.

AI investment decision makers can analyze activities from a top-down, bottom up, or dual-view perspective.

Top-Down, Organizational View

LEVEL 1 Initial / Ad Hoc

Process is unpredictable, poorly controlled, and unmanaged.

LEVEL 2 Repeatable

Process is documented and steps are repeatable with consistent results. Often reactionary.

LEVEL 3 Defined

Documented, standardized processes that are improved over time.

LEVEL 4 Managed

Activities are effectively measured, and meeting objectives are based on active management and

evidence-based outcomes.

LEVEL 5 Optimized

Continuous performance improvement of activities through incremental and innovative improvements.

Bottom-Up, User-centric View

LEVEL 1 Individual

No organizational team structure around activities. Self-organized and executed.

LEVEL 2 Team Project

A dedicated, functional team organized around process activities. Organized by skill set or mission.

LEVEL 3 Program

Group of related projects managed in a coordinated way. Introduction of structured program management.

LEVEL 4 Portfolio

Collection of projects, programs, and other operations that achieve strategic objectives.

LEVEL 5 Enterprise

Operational activities and resources are organized around enterprise strategy and business goals.

Footnotes

1. “Draft 2019-2020 Federal Data Strategy Action Plan.” Federal Data Strategy Development Team. June 2019. Pg. 11. <https://strategy.data.gov/assets/docs/draft-2019-2020-federal-data-strategy-action-plan.pdf> (<https://strategy.data.gov/assets/docs/draft-2019-2020-federal-data-strategy-action-plan.pdf>) [↵](#)
-

Section 6.3: Operational maturity areas

Operational maturity areas represent organizational functions that impact the implementation of AI capabilities.

While each area is treated as a discrete capability for maturity evaluation, they generally depend on one another. The operational maturity areas are:

- **PeopleOps**: Recruit, develop, retain, and organize an AI-ready workforce.
- **CloudOps**: Provide and allocate storage, compute, and other resources in the cloud.
- **SecOps**: Ensure secure deployment of code, access to systems and data, and identity resolution across storage, compute, and data assets.
- **DevOps**: Deploy and manage software throughout development, test, and production environments.
- **DataOps**: Maximize data discovery, access, and use throughout its lifecycle.
- **MLOps**: Test, experiment, and deploy AI or ML models.
- **AIOps**: Identify and resource AI initiatives within the organization.

Each operational maturity area below is supported by related key activities and key questions to focus the conversation on the current organizational state.

PeopleOps

People are at the core of every organization. PeopleOps covers not only the defining skills required to build and deploy AI capabilities, but also creating a culture that promotes the behavior necessary to successfully create intelligent systems.

Organizations must be able to explicitly identify skill sets and traits that make a strong AI team. We focus intentionally on continuously developing skill sets, technical acumen, and dynamic team management. We also identify talent gaps and provide training. Strong PeopleOps practices include the structure and program support to organize high-functioning teams and manage them throughout the entire talent lifecycle.

Organizational Maturity for PeopleOps

LEVEL 1 Initial / Ad Hoc Individual Project

- ✓ Self-directed identification and acquisition of AI related skillsets

LEVEL 2 Repeatable Team Project

- ✓ Identify types of AI talent for Lifecycle
- ✓ Employee journey created
- ✓ KSA's identified
- ✓ Active talent acquisition for AI related skills begin

LEVEL 3 Defined Program

- ✓ PeopleOps Program established
- ✓ Conduct talent gaps analysis
- ✓ Needs documented
- ✓ Talent map created
- ✓ Training identified

LEVEL 4 Managed Portfolio

- ✓ PeopleOps program linked to measurable performance objectives and active efforts to create a culture around AI is a central organizational theme

LEVEL 5 Optimized Enterprise

- ✓ Uses engagement tools for employee growth
- ✓ Integrates an ownership culture throughout organization

Key Activities for PeopleOps

- Understand the AI talent requirements for the organization (needs, skills, and training)

assessments, use case identification)

- Develop the AI talent lifecycle talent pipeline, staffing plans (FTE + contractors), learning paths, project team composition
- Enhance the AI talent lifecycle (automate the talent pipeline, develop internal best practices, expand employee engagement throughout the organization)

Key Questions for PeopleOps

- Can you identify the barriers that prevent your agency from building AI talent, or the precursors necessary to do so?
 - Are you able to identify AI talent within your organization?
 - How do you align culture with mission objectives to create an engaging, creative workplace?
 - Do you understand staff's training and resource needs?
 - How can you identify ideal AI practitioner candidates?
 - How can you ensure that AI practitioners within your organization are able to succeed?
-

CloudOps

Modern infrastructure and platform management best practices have consolidated in the cloud, so the CoE assumes that mature AI organizations will be cloud-capable.

CloudOps functionally manages development teams' virtual onboarding, and provides compute and storage resources, as well as services and tools to enable development teams. Organizations must have a repeatable process to safely and securely stand up and scale development environments for a wide range of AI activities.

Organizational Maturity for CloudOps

LEVEL 1 Initial / Ad Hoc Individual Project

- ✓ Minimal Cloud Resources or Individual User Account

LEVEL 2 Repeatable Team Project

- ✓ Innovation Sandbox in the cloud

LEVEL 3 Defined Program

- ✓ Dev, Test, and Prod environments available, but manual resources allocation

LEVEL 4 Managed Portfolio

- ✓ Self-service or templated cloud resource allocation

LEVEL 5 Optimized Enterprise

- ✓ Balanced cloud resource sharing across the organization with robust cost/benefit/usage metrics

Key Activities for CloudOps

- Assess existing cloud capabilities in the organization (accounts, environments, service providers, etc.)
- Provide environments for development, test, and production-level activities, ensuring secure access controls within and between environments.
- Enhance the provided environments' use through secured and automated pipelines.
- Continually monitor and optimize cloud resources to manage usage rates and compliance.

Key Questions for CloudOps

- Which cloud resources or platforms are available in your environment?
- Is there a systematic, enterprise-wide process by which to allocate cloud resources?
- How do AI practitioners get access to cloud resources?
- Are you able to track resource utilization across the enterprise?

DevOps

A mature AI organization must be able to securely move software into a production environment and provide continuous integration and delivery of

software updates.

DevSecOps best practices enable organizations to shrink the time to deliver new software while maintaining the security of highly reliable and available AI-enabled capabilities. DevSecOps allow organizations to produce microservices in hardened containers; move them throughout development, test, and production environments; and scale services based on user demand.

Organizational Maturity for DevOps

LEVEL 1 Initial / Ad Hoc Individual Project

- ✓ Development on local workstation

LEVEL 2 Repeatable Team Project

- ✓ Process exists for moving locally developed tools into production, but some parts are manual
- ✓ Test Driven Development (TDD)

LEVEL 3 Defined Program

- ✓ Established secure process for containerizing tools and moving into production environment

LEVEL 4 Managed Portfolio

- ✓ Increasingly automated process for deploying secure software with emphasis on reducing iteration and delivery timelines

LEVEL 5 Optimized Enterprise

- ✓ Fully-managed secure software container orchestration
- ✓ CI/CD/CATO

Key Activities for DevOps

- Understand where tools are currently developed within the organization.
- Create and secure development pipelines from local to cloud environments.
- Contain and isolate dependencies for tools to effectively use and scale resources.
- Reduce delivery and iteration timelines through development pipeline automation and full containerization.

Key Questions for DevOps

- What processes exist for moving new tools into production?
 - When are security concerns resolved during development?
 - How quickly can any new tool be deployed in a containerized environment?
-

SecOps

Secure access to code repositories, infrastructure, and platform resources, and data assets depends on understanding how person and non-person entities (NPEs) operate.

SecOps unifies best practices in security for software development, system administration, data access, and protection of AI techniques from attacks. This function also supports the ability to audit and account for user activity, while defending against security threats across the system.

Organizational Maturity for SecOps

LEVEL 1 Initial / Ad Hoc Individual Project

- ✓ Code security/validation is manually accomplished in the Test environment
- ✓ Container security is default to orchestration

LEVEL 2 Repeatable Team Project

- ✓ Code security/validation is manually accomplished within the pipeline
- ✓ Container security is baselined at orchestration

LEVEL 3 Defined Program

- ✓ Code/Container security and validation is automated within the pipeline and manually approved

LEVEL 4 Managed Portfolio

- ✓ Code/Container security and validation software is automated and automatically approved
- ✓ Software rollouts are “trusted”

LEVEL 5 Optimized Enterprise

- ✓ Pipeline security software feeds central Security Data Lake
- ✓ Automation embedded at the Pipeline Orchestration layers
- ✓ Automated code rollbacks
- ✓ Ongoing Authorization

Key Activities for SecOps

- Threat detection: mean time to detect
- Threat resolution: mean time to respond
- Threat impact: mean failure cost

Key Questions for SecOps

- Are corrective actions fully integrated into the overarching security framework?
- Is mean time to detection (MTTD) near real-time?
- Does automated resolution implementation occur throughout pipelines, including automatic rolling back suspect code/version control?
- Is mean time to response (MTTR) near real-time?
- Are impact analyses to emerging threats automated and integrated into organizational strategic decision making?
- Are threat impact: mean failure costs projected?

DataOps

Effective AI capabilities require highly tailored data to train and develop models. To deliver data to development teams, DataOps enables data discovery, access, and use for AI development.

This includes support for batch and streaming data sources, the pre- and post-processing of data, and managing data at all lifecycle phases. Effective DataOps includes a thorough inventory asset catalog, dynamic methods of accessing data, and necessary tools to manipulate data to conduct documented AI experiments. Datasets are described and subjected to version control, creating a history of how the data was sourced and transformed throughout the development process.

Organizational Maturity for DataOps

LEVEL 1 Initial / Ad Hoc Individual Project

- ✓ “Shoeboxes” of data stored locally, not discoverable
- ✓ Copied from one machine to another

LEVEL 2 Repeatable Team Project

- ✓ Routine data sources available and well documented but new data discovery is ad hoc
- ✓ Exploratory data analysis (EDA) initiated

LEVEL 3 Defined Program

- ✓ Engineering support for data management activities is explicit
- ✓ Data pipeline exists

LEVEL 4 Managed Portfolio

- ✓ Self-service for adding new data sources, preparing datasets, and curating data for ML projects

LEVEL 5 Optimized Enterprise

- ✓ Intelligent, secure data discovery, access, and use across all organizations with metrics on business usage and compliance

Key Activities for DataOps

- Determine the locations and method of access for data sources within the organization
- Make tools and resources available to actively manage data
- Align data sources within a data lifecycle management framework
- Optimize use of data sources within the data management framework through automation and self-service data curation
- Continually monitor creation and usage of data to ensure business alignment and compliance

Key Questions for DataOps

- How is data transferred within the organization?
 - Are there enterprise-wide data governance policies in place?
 - What resources are available to ensure high-quality data management?
 - Who is able to access data sources?
-

MLOps

MLOps is the selection, application, interpretation, deployment, and maintenance of machine learning models within an AI-enabled system.

Through these functions, AI development teams conduct data analysis and feature engineering to create a model that achieves a threshold level of performance. Over time, MLOps will enable deployed models to become more accurate as long as they are properly maintained. Another key function in MLOps is to create an audit trail of experiments to document design decisions throughout experimentation and create transparent, explainable AI-enabled capabilities.

Organizational Maturity for MLOps

LEVEL 1 Initial / Ad Hoc Individual Project

- ✓ Models and methods are selected ad hoc and not documented

LEVEL 2 Repeatable Team Project

- ✓ Implement methods for documenting experiments and model selection

- ✓ Utilization of GPUs initiated
-

LEVEL 3 Defined Program

- ✓ Model and methods catalog exists
- ✓ Model to use case matching leverages previous historical knowledge
- ✓ Model measures model accuracy and speed prediction

LEVEL 4 Managed Portfolio

- ✓ Increased use of infrastructure/Server based GPU acceleration for model development
- ✓ Automated selection, testing, and evaluating ML models using AutoML

LEVEL 5 Optimized Enterprise

- ✓ Use of Hyper-scale GPU acceleration
 - ✓ Feedback from ML tools captured and models are continuously updated and improved
-

Key Activities for MLOps

- Build resources to document machine learning model and method selection
- Use historical documentation to develop and enhance machine learning models
- Monitor tool usage and feedback to continuously enhance tools

Key Questions for MLOps

- How are machine learning use cases identified and documented?
 - Are practitioners able to access an enterprise knowledge-base of existing machine learning models?
 - Are models tested, evaluated, and optimized?
-

AI Ops

AIOps support AI capability development as a discipline of systems engineering and product management.

Initially, organizations may identify processes to create a funnel of AI pilots and move them to operations. Over time, effective AIOps will allow organizations to put institutional structure around AI activities through project management, product development, and systems integration.

AIOps focus squarely on integrating AI into the larger enterprise from a technical point of view as well as planning technology roadmaps for future investment. Program evaluation and support functions are also included in AIOps to show the operational impact of AI investments. AIOps include user-centered activities that organize tasks into iterative atomic units of work, manage allocating resources and personnel to work efforts, and set long-term objectives for AI-enabled capabilities.

Organizational Maturity for AIOps

LEVEL 1 Initial / Ad Hoc Individual Project

- ✓ Reactionary and ad hoc AI capability identification

LEVEL 2 Repeatable Team Project

- ✓ Established process to capture AI product requirements and user workflows

LEVEL 3 Defined Program

- ✓ Formal AI Product Management team, Strategy, and roadmap, including test and evaluation plan

LEVEL 4 Managed Portfolio

- ✓ AI Product Management goals are linked to organizational performance objectives
- ✓ T&E synthesis and evaluation included for each AI product

LEVEL 5 Optimized Enterprise

- ✓ AI Capability dependencies are mapped across org boundaries and linked to an enterprise strategy with measurable objectives

-
- ✓ Retrospective analysis of past efforts to continuously modify business objectives related to AI investment
 - ✓ Dedicated T&E efforts to optimize AI cap

Key Activities for AIOps

- Establish processes by which to determine AI use cases
- Develop AI capabilities through AI product teams
- Align AI products with organizational objectives and enterprise-wide dependencies
- Continuously evaluate AI products to optimize business usage, resource allocation, and compliance

Key Questions for AIOps

- How are AI products identified within the organization?
- Who manages AI products?
- Are product dependencies mapped to both organization objectives and data dependencies?

[Download Operational Maturity table \(../images/operational-maturity-matrix.png\)](#)

Chapter 7: Solving business challenges with AI

Innovating with AI

AI is a great tool for improving operations and providing deeper insights and more informed decision making, but AI can provide even greater benefits when viewed as an innovation pathway. The most successful AI-driven organizations don't just use AI and analytics to improve existing products and services, but also to identify new opportunities.

Building a culture of data literacy and collaboration across boundaries is a foundational step for innovation and digital transformation. Use these ideas to educate and empower teams to embrace an innovation mindset.

Learn from others in government and share use cases

You don't need to reinvent the wheel for every AI solution. No matter the size of a project or method of development, you can learn critical lessons. Through sharing AI project stories at both intra-organizational and inter-agency levels, government teams can learn which approaches to business challenges with AI succeed and which don't. This means evaluating the technical journey, but also the business journey. What was the business problem? What tools did you use to solve it? How did you write the solicitation? What clauses did you include to ensure successful delivery?

Courses and certifications

The learning and training space for data science, machine learning (ML), and AI is overwhelming. As we discussed in Chapter 4, investing time and money in a platform to target your workforce's educational needs enhances their ability to apply new skills and become educated members of an AI team

Encourage experimentation and learn by doing

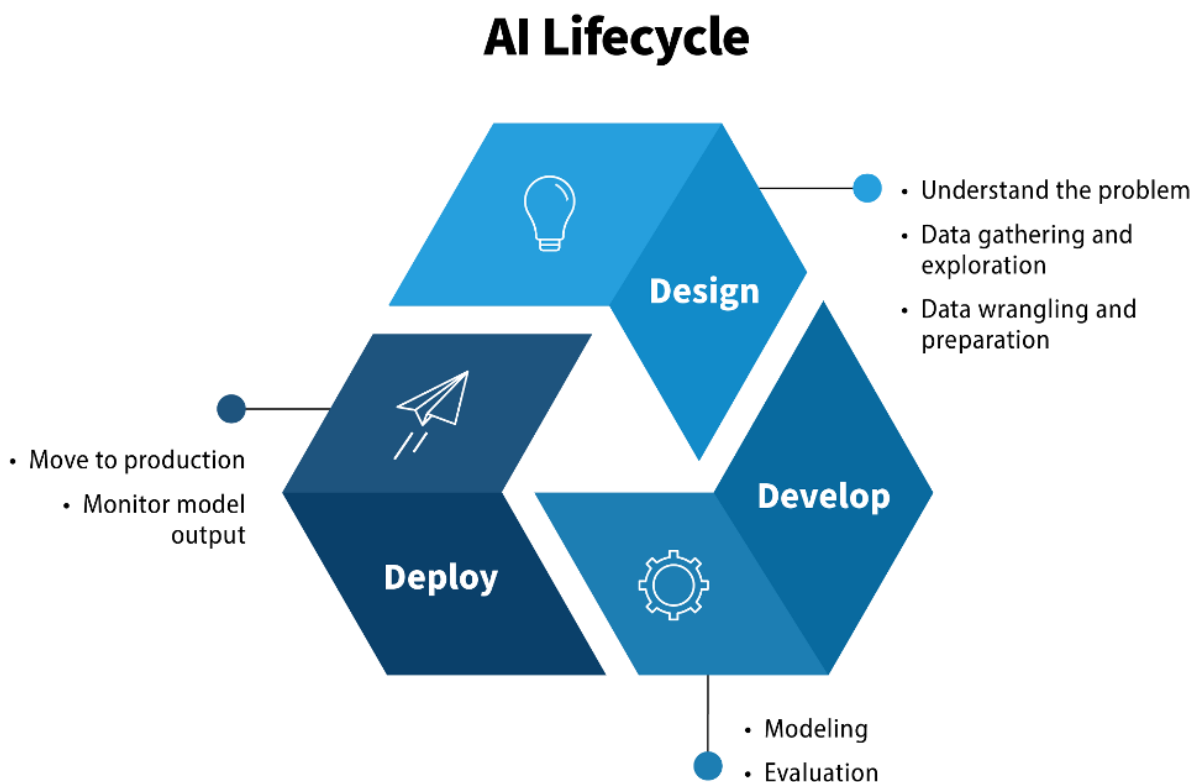
Creating a culture and space for teams to explore, discover, and experiment empowers them to be informed AI builders and buyers. It also means increasing access to data science tools and data sets across the organization. If senior leaders have concerns about experimenting, start small or in a lower-risk area. This may be an entirely internal effort or could involve a contractor or vendor.

Hackathons, challenges, and citizen scientists

Government has recently embraced the technical expertise of individuals and small organizations including students, citizen scientists, or coding groups. You can explore and experiment with your idea for an AI use case. For government lead challenge opportunities, check out [challenge.gov](https://www.challenge.gov) (<https://www.challenge.gov/>) and the Opportunity Project.

Section 7.1: Understanding and managing the AI lifecycle

The AI lifecycle is the iterative process of moving from a business problem to an AI solution that solves that problem. Each of the steps in the life cycle is revisited many times throughout the design, development, and deployment phases.



Design

- 1. Understand the problem:** To share your team's understanding of their mission challenge, you first have to identify the key project objectives and requirements. Then define the desired outcome from a business perspective. Finally, determining AI will solve this problem. Learn more about this step in the framing AI problems section.
- 2. Data gathering and exploration:** This step deals with collecting and evaluating the data required to build the AI solution. This includes discovering available data sets, identifying data quality problems, and deriving initial insights into the data and perspectives on a data plan.

3. **Data wrangling and preparation:** This phase covers all activities to construct the working data set from the initial raw data into a format that the model can use. This step can be time consuming and tedious, but is critically important to develop a model that achieves the goals established in step 1.

- No AI solution will succeed without clearly and precisely understand the business challenge being solved and the desired outcome.
- Data is the foundation of any AI solution. Without clearly understanding of the data required and the make up of that data, a model cannot use it.
- Data preparation is often the hardest and most time-consuming phase of the AI lifecycle.

Develop

4. **Modeling:** This step focuses on experimenting with data to determine the right model. Often during this phase, the team trains, tests, evaluates, and retrains many different models to determine the best model and settings to achieve the desired outcome.

The model training and selection process is interactive. No model achieves best performance the first time it is trained. It is only through iterative fine-tuning that the model is honed to produce the desired outcome. Learn more about types of machine learning and models in [Chapter 1 \(../what-is-ai-key-terminology\)](#).

Depending on the amount and type of data being used, this training process may be very computationally expensive—meaning it requires special equipment to provide enough computing power and cannot be performed on a normal laptop. See [Chapter 5 \(../cultivating-data-technology\)](#) to learn more about infrastructure to support AI development.

5. **Evaluation:** Once one or more models have been built that appear to perform well based on relevant evaluation metrics, test the models on new data to ensure they generalize well and meet the business goals.

As this step is particularly critical, we discuss it in much greater detail in the test and evaluation section.

Deploy

6. **Move to production:** Once a model has been developed to meet the expected outcome and

performs at a level determined ready for use on live data, deploy it into a production environment. In this case, the model will take in new data that was not a part of the training cycle.

7. **Monitor model output:** Monitor the model as it processes this live data to ensure that it is adequately able to produce the intended outcome—a process known as generalization, or the model’s ability to adapt properly to new, previously unseen data. In production, models can “drift,” meaning that the performance will change over time. Careful monitoring of drift is important and may require continuous updating of the model.

As with any logic and software development, use an agile approach to continually retain and refresh the model. However, AI systems require **extra attention**. They must undergo rigorous and continuous monitoring and maintenance to ensure they continue to perform as trained, meet the desired outcome, and solve the business challenges.

Section 7.2: Good software practice and AI development

Much of what is discussed in this and other chapters relies on a foundation of good software practice.

Though this chapter is not intended to explain agile development methodology and Dev(Sec)Ops, both of these are foundationally important to developing successful AI solutions. As such, part of the process of building out AI capabilities in an organization may involve refining or restructuring the current software infrastructure to support the AI development lifecycle.

Section 7.3: Identifying AI use cases in your organization

Making an organization AI-enabled is certainly challenging. As you build the processes and structures outlined in [Chapter 5: Cultivating Data and Technology \(../cultivating-data-technology\)](#), parallel efforts should focus on applying these structures to new use cases.

A use case is a specific challenge or opportunity that AI may solve. Start small with a single use case limited in scope. Eventually, as business needs and agency infrastructure grow, the number of AI use cases will grow.

Your goal should be to create an evergreen collection of AI use cases where business units organically identify, evaluate, and select the best to proceed to proof of concept, pilot, and production. This ideal state may be far in the future for some agencies, but all agencies can start to explore AI applications to current business challenges right now.

How to select which AI use cases to pursue?

To identify use cases, consider the following:

Focus on agency mission

- Problems directly tied to operational or strategic priorities
- Problems connected to KPIs with significant gaps to their targets

Find the right data

- Areas that are rich with accessible data
- Areas with under-explored data

Identify a champion

- AI needs executive sponsorship to be successful
- Align mission, data, IT, and end-user needs

Framing the problem for an AI project

As with non-AI projects, framing the problem is critical. To ensure success, follow these steps:

User interviews

Interview users for an AI project. Too often, leadership, program managers, and algorithm developers create solutions for problems that don't exist, or solutions that are simply not aligned on the right problem. By interviewing an application's actual users, AI practitioners and implementers can understand the needs and intricacies of the exact audience they're working to help.

User research ought to be iterative and consistent throughout development. Gathering user feedback is essential not just for the experience the applications would provide, but also how AI ends up being applied.

Market research

Whether you develop a potential AI solution internally as a prototype or procure it through the private-sector, you have to know what's currently available in the market to design and implement an AI solution successfully.

To find the right AI solution or vendor, leaders must know the field from meeting with companies in person, calling people, and doing market research. This vital systematic approach is improved with professional company assessors. Ultimately, the decision to buy or build AI is informed by what is available in the market and the team's internal capabilities.

Prioritizing projects

Not every identified use case is a great fit to go forward as a project. Technical, data, organizational, and staffing challenges may impede its progress.

Ultimately, prioritization comes down to three factors:

Impact

- Problem size
- Impact of model output on the organization's mission priority and business metrics

Effort

- Level of effort required to acquire, ingest, and wrangle data
- Data's quality and quantity
- Complexity of analytics and model required

- Cost of development and/or implementation

Fit

- Alignment with team/unit mission and priorities
 - Capacity to build, buy, and/or adopt into existing environment
 - Ability to maintain the requisite budget and staff
-

Section 7.4: Starting an AI project

Once you've identified a project, assemble the Integrated Product Team (IPT) outlined in [Chapter 4: Developing the AI workforce \(../developing-ai-workforce\)](#) to ensure the necessary parties are engaged and dedicated to delivering success. Whether the project is a small pilot or a full-scale engagement, moving through the AI life cycle to go from business problem to AI solution can be hard to manage.

Internal prototype and piloting

Internal or organic prototypes (exploration without vendor support) provide a great way to show value without having to commit the resources required for a production deployment. It allows you to take a subset of the overall integration and tackle one aspect of it to show how wider adoption can happen. This requires technical skills, but can rapidly increase AI's adoption rate as senior leaders see the value before committing resources.

Prototyping internally can help identify where in the life cycle to seek a vendor. Not all steps require vendor support (single or many). It can also show when and how to best engage a vendor. It could reveal either that you should engage early to turn an internal prototype into a pilot, or that you should develop a pilot before engaging a vendor for full-scale production.

From pilot to production

Once you've completed a successful pilot, look to evaluate its effectiveness towards your objectives. If you determine that the pilot proved enough value—with clearly defined and quantified KPIs—that your agency wants a longer term solution, then you should seek to move the pilot to production.

If you need vendor support for scaling, take the key findings from the pilot and translate them into a procurement requirement so a private vendor can take over providing the service. The pilot's success allows the work already done to serve as the starting point for the requirements document.

Prototypes/pilots intentionally scale the problem down to primarily focus on the data gathering and implementation. When moving to production, you have to consider the entire pipeline. The requirements must feature the ways in which the results from the model will be evaluated, used, and updated.

The three most important items to consider when moving to production are these:

- **Project ownership**

What part of the organization will assume responsibility for the product's daily continuation?

- **Implementation plan**

Since the pilot addressed only a small part of the overarching problems, how will you roll out the solution to the whole organization?

- **Sunset evaluations**

At what point will the organization no longer need the results coming from the AI solution? Who and how will this be evaluated?

These are important questions to consider, which is why the test and evaluation process is critical for AI projects.

Start building AI capabilities

If there are existing data, analytics, or even AI teams, align them to the identified use cases and the objective of demonstrating mission or business value. If there are no existing teams, your agency may still have personnel with relevant skill sets who haven't yet been identified. Survey the workforce to find this talent in-house to begin building AI institutional capability.

To complement government personnel, consider bringing in contractor AI talent and/or AI products and services. Especially in the early stages, government AI teams can lean on private-sector AI capabilities to ramp up government AI capability quickly. But you must approach this very carefully to ensure sustainability of robust institutional AI capabilities. For early teams, focus on bringing in contractors with a clear mandate to train government personnel or provide infrastructure services when the AI support element has not yet been stood up.

Buy or build

The commercial marketplace offers a vast array of AI products and services, including some of the most advanced capabilities available. Use the research and innovation of industry and academia to boost government AI capabilities. This can help speed the adoption of AI and also help to train your team on the specific workflows and pipelines needed in the creation of AI capabilities.

Agencies should focus also on building their own sustainable institutional AI capability. This capability shouldn't overly rely on external parties such as vendors/contractors, who have different incentives due to commercial firms' profit-seeking nature. Especially with limited AI talent, agencies should strategically acquire the right skills for the right tasks to scale AI within the agency.

Agencies' ultimate goal should be to create self-service models and shareable capabilities rather than pursuing contractors' made-to-order solutions. Agencies must weigh the benefits and

limitations of building or acquiring the skills and tools needed for an AI implementation. Answering the “buy or build” question depends on the program office’s function and the nature of the commercial offering.

External AI products and services can be broadly grouped into these categories:

- **Software tools that use AI as part of providing another product or service**

Examples include email providers that use AI in their spam filters, search engines that use AI to provide more relevant search results, language translation APIs that use AI for natural language understanding, and business intelligence tools that use AI to provide quick analytics.

- **Software tools that help AI practitioners be more effective**

Examples include tools for automating data pipelines, labeling data, and analyzing model errors.

- **Open source**

Open-source software is used throughout industry and heavily relied on for machine learning, deep learning, AI research, development, testing and ultimately operation. Note, that many of these frameworks and libraries are integrated into many top “proprietary software applications”.

Mission centers and program offices, the heart of why agencies exist in the first place, need to ensure a sustainable institutional AI resource by focusing on building. A team of government AI practitioners dedicated to the mission’s long-term success is necessary for building robust core capabilities for the agency.

Commercial tools that enhance these practitioners’ effectiveness may be worth the cost in the short-term. However, mission centers sometimes have unique requirements that do not exist in a commercial context, which makes commercial-off-the-shelf (COTS) standalone offerings less likely to fit. Even if the agency could find an adequate COTS product, using it would be a major operating risk for an agency’s core functions to rely so much on external parties.

On the other hand, business centers are likely to benefit from AI-enhanced standalone tools. Business centers often focus on efficiency while having requirements most likely to match commercial requirements, so COTS products are more likely to fit. Business centers still need government AI talent, who can evaluate and select the most appropriate commercial AI offerings.

Besides products and services and their associated support personnel, contractors may also offer standalone AI expertise as contractor AI practitioners. These kinds of services are well-suited to limited or temporary use cases, where developing long term or institutional capability is not needed.

In the early stages of implementing AI in an agency, contractor AI personnel can help train and supplement government AI personnel. But as with commercial AI tools, wholesale outsourcing core capabilities to contractor personnel teams creates major operating risk to an agency's ability to perform core functions. Think in terms of Infrastructure as Code (IaC), the ability to rapidly provision and manage infrastructure, to design and build out your AI Platform, creating automations and agile pipelines that are conducive for PeopleOps, CloudOps, DevOps, SecOps, DataOps, MLOps and AIOps.

Infrastructure as code (IaC) brings the repeatability, transparency, and testing of modern software development to the management of infrastructure such as networks, load balancers, virtual machines, Kubernetes clusters, and monitoring. The primary goal of IaC is to reduce error, configuration deltas and increase automations, while allowing engineers to spend time on higher value workflows. Another goal would be shareable IaC across the federal government. IaC defines what the end state of your infrastructure needs to be, then builds the resources necessary to achieve and self-heal . Using infrastructure as code also helps standardize cluster configuration and manage add-ons like network policy, maintenance windows, and Identity and Access Management (IAM) for cluster nodes and workloads in support of AI, ML, Data workloads and pipelines.

Acquisition journey

After your agency decides to acquire commercial products or services, consider these practices to increase your odds of success:

1. Use a Statement of Objectives (SOO) when you are less certain of a solution's path and want to consider innovative or unorthodox methods. Consider using a Performance Work Statement (PWS) when you have clear specifications on what the product or service needs to do. A PWS, outside of a SOO, is written incorporating measurable standards that inform the contractor of the government's desired outcomes. How the contractor achieves those outcomes is up to them. The contractor is thus empowered to use the best commercial practices and its own innovative ideas to achieve the desired results.
2. Include technical tests in your solicitation as evaluation criteria. These tests should allow your technical subject-matter experts on your evaluation panel to verify the ability of any suggested approaches to apply to your program's specific circumstances.
3. Data rights and intellectual property clauses aren't the only ways to ensure a project can move from one team to another. You'll want to include deliverables like product backlogs and open source repositories with the entire source code along with all necessary artifacts to create technical and process-agnostic solutions. To minimize taxpayer exposure to repetitive buys, ensure at least government usage rights to balance private-sector concerns while maximizing the government's investments.
4. Use retrospectives on the acquisition process to identify key clauses and language that worked and those that caused problems, both in terms of the solicitation and post-award management. Document lessons learned to allow new and inexperienced members of the

team to ramp up quickly.

5. Share the results of your experiences with your federal colleagues. There is no better way to gain knowledge and improve the experience with implementing AI in a department and agency than working with others who have similar projects. You can join the [Federal AI Community of Practice \(https://coe.gsa.gov/communities/ai.html\)](https://coe.gsa.gov/communities/ai.html) to connect with other government agencies working in AI.

Test and evaluation process

Some agencies in the defense and intelligence community already emphasize testing and evaluating software. Due to the nature of AI development and deployment, all AI projects should be stress tested and evaluated. Very public examples of AI gone wrong show why responsibility principles are a necessary and critical part of the AI landscape. You can address many of these challenges with a dedicated test and evaluation process.


The basic purpose of Test & Evaluation (T&E) is to provide knowledge to manage risk that's involved in developing, producing, operating, and sustaining systems and capabilities. T&E reveals system capabilities and limitations to improve the system performance, and optimize system use, and sustain operations. T&E provides information on limitations (technical or operational) and Critical Operational Issues (COI) of the system under development to resolve them before production and deployment.

Traditional systems usually undergo two distinct stages of test and evaluation. First is developmental test and evaluation (DT&E), which verifies that a system meets technical performance specifications. DT&E often uses models, simulations, test beds, and prototypes to test components and subsystems, hardware and software integration, and production qualification. Usually, the system developer performs this type of testing.

DT&E usually identifies a number of issues that need fixing. In time, operational test and evaluation (OT&E) follows. At this stage, the system is usually tested under realistic operational conditions and with the operator. This is where we learn about the system's mission effectiveness, suitability, and survivability.

Some aspects of T&E for AI-enabled systems are quite similar to their analogues in other software-intensive systems. However, there are also several changes in the science and practice of T&E that AI has introduced. Some AI-enabled systems present challenges in what to test, and how and where to test it; all of those are, of course, dependent on the project.

At a high level, however, T&E of AI-enabled systems is part of the continuous DevSecOps cycle and Agile development process. Regardless of the process, the goal of T&E is to provide timely feedback to developers from various levels of the product: on the code level (unit testing), at the integration level (system testing, security and adversarial testing), and at the operator level (user testing).



These assessments include defining requirements and metrics by talking with various stakeholders, designing experiments and tests, and doing analysis and making actionable recommendations to the leadership on overall system performance across its operational envelope.

T&E for Projects

On a project, there are various levels of T&E; each reveals important information about the system under test:

Model T&E

This is the first and simplest part of the test process. In this step, the model is run on the test data and its performance is scored based on metrics identified in the test planning process. Frequently, these metrics are compared against a predetermined benchmark, between developers, or with previous model performance metrics.

The biggest challenge here is that the models tend to arrive to the government as black boxes due to IP considerations. For the most part, when we test physical systems, we know exactly how each part of that system works. However, we don't know how the models work; that's why we test. If you want testable models, you have to make that a requirement at the contracting step.

Measuring model performance is not entirely straightforward; here are some questions you might want your metrics to answer:

- How does the model perform on the data? How often does it get things right / wrong? How extreme are the mistakes the model makes?
- What kind of mistakes does the model make? Is there any evidence of model bias?

- Ultimately, does the model tend to do what it's supposed to (find or identify objects, translate, etc.?)

Integrated System T&E

In this step, you evaluate the model not by itself but as part of the system in which it will operate. In addition to the metrics from model T&E, we look for answers the following questions:

- Does the model performance change on an operationally realistic system? Does the model introduce additional latency or errors?
- What is the model's computing burden on the system? Is this burden acceptable?

Operational T&E

In this step, we collect more information on how the AI model ultimately affects operations. We do this through measuring:

- Effectiveness (mission accomplishment)
- Suitability (reliability, compatibility, interoperability, human factors)
- Resilience (ability to operate in the presence of threats, ability to recover from threat effects, cyber, adversarial)

Ethical T&E

Depending on the AI solution's purpose, this step ensures the system does only what it's supposed to do and doesn't do what it's not supposed to do.

Integration of acquired AI tools

To ensure that AI tools, capabilities, and services are not only acquired, but also properly integrated into the wider program's business goals, consider these practices to increase your chances of success:

- Use a Statement of Objectives (SOO) under the Simplified Acquisition Threshold (SAT) to create a prototype or a proof of concept. Begin to standardize as you move into a pilot. Be prepared to use a Performance Work Statement (PWS) when you have clear specifications on what the product or service needs to do and when you are ready to scale across the agency. As mentioned above in the Acquisition Journey sub-section, consider using a Performance Work Statement (PWS) when you have clear specifications on what the product or service needs to do, as it incorporates measurable standards aligned with desired outcomes. Different phases of AI projects require different types of solicitations and cost structures.
- Include technical tests in your solicitation evaluation criteria. These tests should allow your technical subject-matter experts on your evaluation panel to verify that suggested approaches

apply to your program's specific circumstances.

- Data rights and intellectual property clauses aren't the only ways to ensure a project can move from one team to another. Deliverables like product backlogs and open source repositories with the complete source code and all necessary artifacts are essential for creating technical and process-agnostic solutions. To minimize taxpayer exposure to repetitive buys, ensure the government maintains usage rights to balance private sector concerns while maximizing the government's investments.
- Use retrospectives on the acquisition process to identify key clauses and language that worked and language that caused problems, both in terms of the solicitation and post-award management. Lessons learned should be standardized through documentation that will allow new and inexperienced members of the team to ramp up quickly.
- Share the results of your experiences with your federal colleagues. There is no better way to gain knowledge and improve the experience of implementing AI in a department and agency than collaborating with others doing similar projects.

Thank you for printing the AI Guide for Government, a resource maintained by the Artificial Intelligence Center of Excellence (AI CoE). You can return to print a new version and see recent updates at <https://coe.gsa.gov/ai-guide-for-government/>.