
# Artificial Intelligence and Data Analytics (AIDA) Guidebook

ATARC Artificial Intelligence and Data Analytics Guidebook Working Group

*March 2022*



ATARC

Advanced Technology Academic Research Center

# Acknowledgements

The Artificial Intelligence and Data Analytics (AIDA) Guidebook would not have been possible without the support and contribution from the following organizations and agencies. Thank you to everyone for your immense dedication and efforts towards this project.

Advanced Technology Academic Research Center (ATARC)

Artificial Intelligence and Data Analytics Working Group (AIDA WG)

Defense Health Agency (DHA)

Department of Veteran Affairs (VA)

General Services Administration (GSA)

Internal Revenue Service (IRS)

MITRE

National Aeronautics and Space Administration (NASA)

National Institute of Standards and Technology (NIST)

National Institutes of Health (NIH)

National Nuclear Security Administration (NNSA)

Office of National Coordinator for Health Information Technology (ONC)

TekSynap

United States Department of Agriculture (USDA)

**Disclaimer:** This Guidebook was prepared by the members of the ATARC AI & Data Analytics Guidebook Development Working Group in their personal capacity. The opinions expressed do not reflect any specific individual nor any organization or agency they are affiliated with.

# Comment

The Federal Government does not endorse any private sector products or services, or attempt to transfer any intellectual property rights by way of this guidebook. The Federal Government also does not seek to lead or control the development of advice or recommendations on public policy matters that may be within the purview of any federal agencies.

# Table of Contents

# Executive Summary

The Advanced Technology Academic Research Center (ATARC) and its partners in the Artificial Intelligence and Data Analytics (AIDA) Working Group determined that there was a need for a leadership-accessible reference guide that could outline the utility and use of artificial intelligence (AI) and data analytics, as well as best practices, opportunities for implementation, and potential challenges involved. Their mission was to find answers for how government may implement AI and data analytics as well as achieve the objectives of the May 2012 Digital Government Strategy[1]. Members sought to identify, understand, and provide strategic insights to address challenges and barriers affecting the broader information technology (IT) community's data deployments and offerings to accelerate the adoption of data best practices to increase efficiency and reduce cost. It is through this working group that this document was created.

The guidebook covers essential subjects related to AI and data analytics such as Machine Learning Methodology, Privacy and Security, Data Management, Measurements of Effectiveness and Key Performance Indicators (KPIs), Human-machine interaction, Societal Impacts of Artificial Intelligence, Organizational Change Management, and Governance. The value of this guidebook is as a starting point toward understanding the AI and data analytics development process. It is not intended to be a comprehensive source for all topics nor a technical document, but should encourage the reader to seek additional knowledge and provides references to begin that journey. The guidebook should also begin to demonstrate how AI and data analytics can positively impact policy decisions and provide a return on investment for an organization.

Cooperation between government, business, and academia is critical for linking AI and data analytic policies, ethics and operations. Collaboration opportunities include crafting federal IT rules, regulations, and policies; developing best practices and exemplar use cases for data management frameworks; assuring operational security; protecting intellectual property; generating non-biased training datasets; and establishing mechanisms to verify that data being utilized is 'fit for purpose' are all tasks completed in the follow-on work. Furthermore, this structure can further develop best practices to protect data in transit and at rest, as well as propose ways to minimize administrative difficulties associated with regulatory compliance.

[1] "Digital Government Strategy - Building a 21st Century Platform to Better Serve the American People." https://obamawhitehouse.archives.gov/sites/default/files/omb/egov/digital-government/digital-government.html

# 1 Introduction

The Advanced Technology Academic Research Center (ATARC) and its partners in the Artificial Intelligence and Data Analytics (AIDA) Working Group developed the AIDA Guidebook as a reference guide for leaders to outline the utility and use of artificial intelligence (AI) and data analytics, as well as best practices, opportunities for implementation, and potential challenges involved. The AIDA Working Group's goal was to identify, understand, and provide strategic insights to address challenges and barriers affecting the broader information technology (IT) community's data deployments and offerings. Driven by the May 2012 Digital Government Strategy, they collaborated on each of the sections described below to present a holistic view for government AI and data analytics implementation. This AIDA Guidebook is designed to help organizations of all types, especially project managers, chief information officers (CIOs), senior managers, and others who wish to utilize AI and data analytics to better serve their community, stakeholders, and clients.

The guidebook structure begins with an overview of the history of the document, moving into a description of how AI models are developed, followed by key areas of consideration for leaders implementing AI and data analytics, and closing with suggested next steps to maintain and update this document moving forward. It is worth mentioning, that there is not a one-size fits all solution for establishing AI models, and this guidebook does not seek to provide that. Just as the size and scope of models may vary, so does the path to successful implementation. Sections 6 through 12 outline potential areas for consideration when planning, developing and establishing models. These are subject areas that have been explored through previous use cases and successful implementations, and the AIDA Working Group identified them up front in order to maximize the utility of this guidebook. Sections 2-4 provide background information on ATARC, the AIDA Working Group, and the public and private collaborations that preceded the creation of this document.

As described in Section 0, developers effectively build, evaluate, and manage analytic and learning systems through machine learning pipelines to codify and automate the workflows it takes to produce a machine-learning model. Machine learning pipelines consist of multiple sequential steps that do everything from data extraction and preprocessing to model training and deployment.

Section 6.4 covers privacy and security, which are key components of AI and data analytics implementation. Privacy is relevant to AI systems in maintaining the confidentiality of individuals whose data is used for model training, the stage in which developers use dataset(s) to train a machine learning algorithm, as well as the confidentiality of subjects to which the system is applied. Security refers to keeping both the AI system’s development code (i.e. the model itself, testing tools, and test results) and the data used to train it free from interference.

In Section 7, an overview of data management covers the methods, tools, procedures, and processes used to ensure the availability, confidentiality, integrity, quality, reliability, and usefulness of data sets. Data management standards include data dictionaries, data ontologies, metadata structures, quality measures, database optimization and backup standards, and data retention standards.

Section 8 provides an overview of measurement and key performance indicators (KPIs) and describes how to define observable, measurable quantities that relate to desired performance characteristics. Measuring performance in terms of cost, financial, quality, time, flexibility, delivery reliability, safety, customer satisfaction, employees' satisfaction, and social performance indicators have significant positive impact on the overall performance of organizations.

Human-machine interaction covered in Section 9 refers to the design and optimization of how users and IT systems interact collaboratively to perform a function. In the context of AI and data analytics, this approach allows human-machine teaming and to raise performance above that of humans or AI alone.

Section 10 provides an overview of the ethics and societal impact of AI and data analytics. Ethics refers to the derivation of behavioral parameters or rules that, if followed, will preserve an identified set of values. Ethical AI ensures that AI developed by the organization or entity maintains human dignity, and understands and manages the risks associated with the AI's effects with people. This Societal Impact refers to the aggregate, group, and individual outcomes generated or influenced by a system, focusing on privacy and surveillance, bias, and discrimination.

Section 11 discusses the challenges of organizational change management and navigating the changes in an organization's structures, systems, strategy, or culture as part of AI or data analytics implementation. These changes are often implemented as part of an overall plan to achieve strategic, operational, or financial goals.

Finally, Section 12 provides examples of how governance policy provides oversight of operational management, to ensure that an organization (i.e., people, processes, and technology) is being managed effectively to achieve strategic goals while complying with relevant laws, regulations, and policies. Data governance is more narrowly focused on a collection of processes, roles, policies, standards, and metrics that ensure the effective and efficient use of information in enabling an organization to achieve its goals.

Appendices A-D include references to key laws and regulations, common acronyms, key definitions and references to documents cited throughout the guidebook.

# 2 Background

## 2.1 ATARC

The Advanced Technology Academic Research Center (ATARC) was established in September 2012, initially under the name Advanced Mobility Academic Research Center (AMARC) and was formed in response to the Digital Government Strategy, described in more detail within section 2.2, which was released earlier that year. This strategy describes the need for innovation and moving the country toward 21st century technological capability to create easier access for the citizens of the United States to be able to access their data from wherever they are and whenever they need it. The Digital Government Strategy identified three distinct goals:

1. Unlock the power of government data combined with AI capabilities to spur innovations across our Nation and improve the quality of services for the American people
2. Enable the American people and an increasingly mobile workforce to access high-quality digital government information and services anywhere, anytime, on any device
3. Ensure that as industry adjusts to this new digital world, we seize the opportunity to procure and manage devices, applications, and data in smart, secure and affordable ways

AMARC founder Tom Suder played a prominent role in the formation of the Digital Government Strategy and created the organization to continue the dialogue between government, industry, and academia. AMARC hosted a pair of Federal Mobile Computing Summits before being asked by the General Services Administration (GSA) in 2013 to hold additional events on the topics of Cloud Computing and IT Networks. As AMARC's collaborative focus changed from Mobile to Federal IT, so too did its name. In early 2014, Suder renamed AMARC to ATARC, and within a year, added Cybersecurity, DevOps, and Data & Analytics to its roster.

**Figure 1** provides an overview of the concept behind ATARC and its purpose and value.

**Figure 1: ATARC Overview**

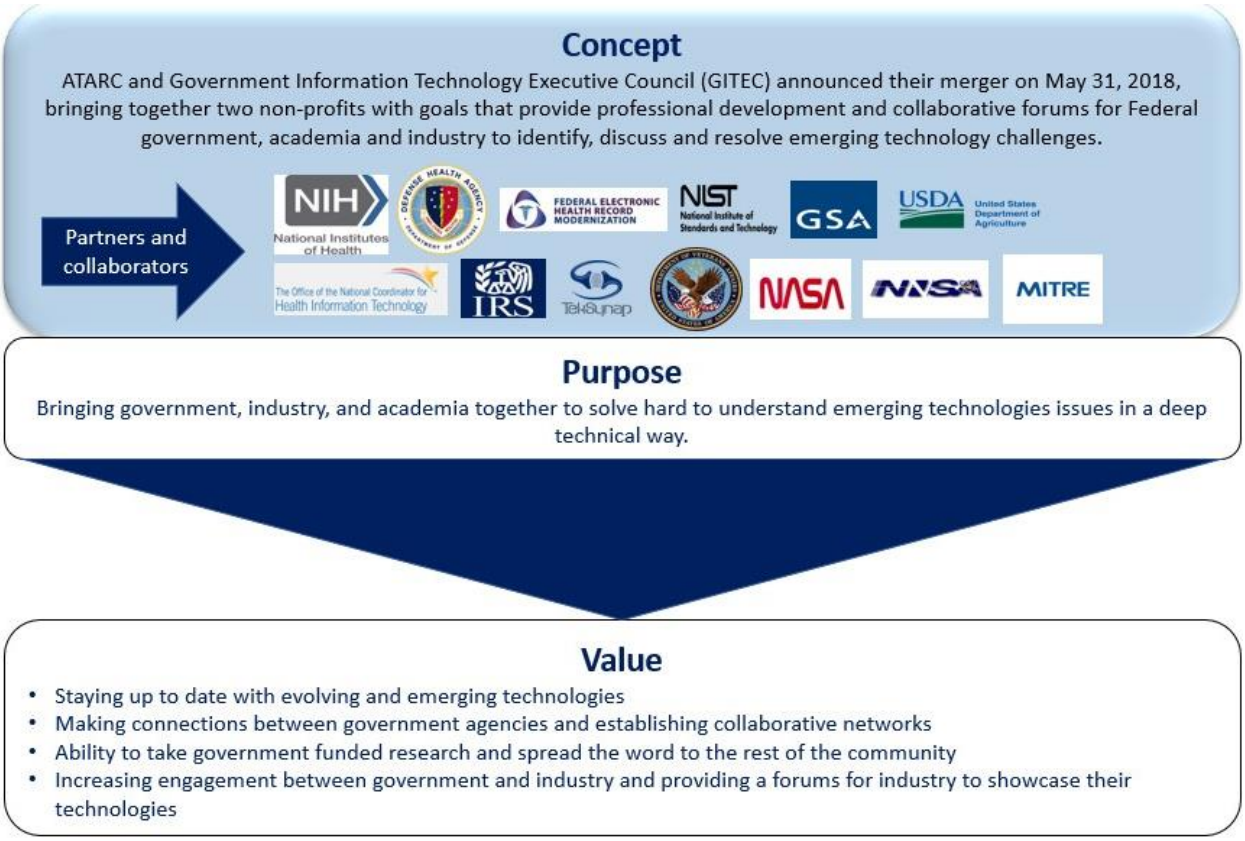


In 2015, the Interagency Program Office (IPO) and ATARC established HealthTracs, a healthcare-focused discussion during ATARC's Federal Summits as a primary method of healthcare industry trend surveillance. This partnership leveraged an existing forum to expand relationships within the government, academia, and the private sector, nurturing collaborations in order to discuss key healthcare innovation challenges relevant to the Departments and the IPO. The IPO continued to collaborate with ATARC through various Federal Summits to examine emerging technology challenges and resolutions surrounding the use of cloud computing and big data within the Federal Government.

In 2017, ATARC published a 235-page report to the American Technology Council and Federal CIO Council titled "Navigating the Future of Mobile Services."[2] Produced in the spirit of the Digital Government Strategy and in conjunction with the cross-agency Mobile Services Category Team, the report was a collaborative effort between more than 160 people representing 75 agencies, bureaus, and companies. In 2018, ATARC published its 25th white paper in conjunction with MITRE as the output from events in the ATARC Federal IT Summit Series.

[2] https://atarc.org/wp-content/uploads/2019/01/ATARC-MSCT-Report-Navigating-Future-of-Mobile-Services-2.pdf

The Government Information Technology Executive Council (GITEC) announced their merger with ATARC on May 31, 2018. The merger brought together two non-profits with similar goals that provide professional development and collaborative forums for Federal Government, academia, and industry to identify, discuss, and resolve emerging technology challenges. These engagements resulted in the establishment of the Institute of Electrical and Electronics Engineers (IEEE) P2795 Working Group in 2018, to work together in the creation of shared analytics standards. The IEE Standards for Shared analytics identified the requirements for using shared analytics over secured and unsecured networks and established consistent methods of using an overarching interoperability framework to utilize one or more disparate data systems for analytic purposes without an analytic user having explicit access to or sharing the data within these systems.

The ATARC AIDA Working Group was formed in response to these initial collaborators recognizing the need for a manual on the use of AI and data analytics tools and concepts. The purpose of this working group was to identify and provide harmonized solutions to industry challenges in implementing AI and data analytics as well as to address the presidential requirements for artificial intelligence within the Digital Government Strategy.

It is through this working group that this document was created.

**Figure 2** provides an overview of the foundational statements for this working group.

**Figure 2: AIDA WG Foundational Statements**

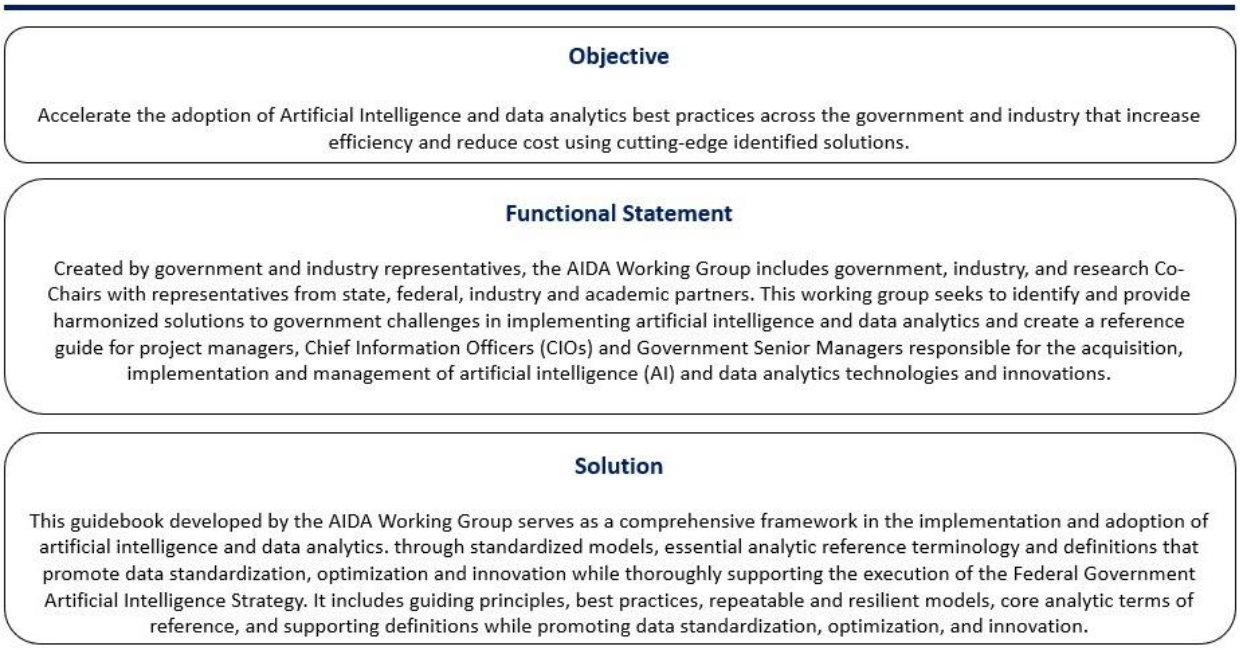

The working group agreed to develop two products to support technology leaders, this guidebook and a report assessing AI policy entitled "From Ethics to Operations: Current Federal AI Policy."[3] The AI policy framework report includes all categories of policy related to AI; a review of current AI policy, legislative, and regulatory activities; an assessment of the current federal AI policy environment; and recommendations for using the framework to promote a comprehensive, consistent, and accurate federal AI policy environment.

This guidebook compliments the policy framework, establishing a solid structure that covers the key principles and best practices for AI application while still fulfilling the objectives of the Federal Government AI Strategy. The working group first announced the idea of this guidebook at the May 2021 AIDA Summit. The full ATARC timeline described above is shown in

**Figure 3.**

**Figure 3: History of ATARC**

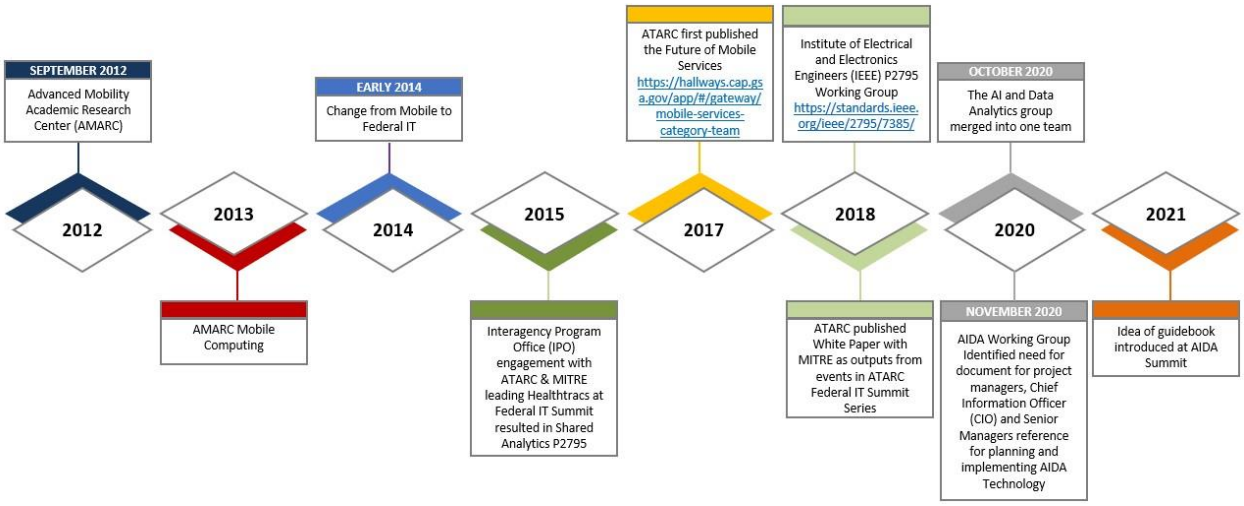


## 2.2 Digital Government Strategy

The Digital Government Strategy aims to serve the American people by combining government data and AI to fuel innovation and improve services, allow the easy access of high-quality digital government information and services at any point or place, and ensure that the government takes the opportunity to acquire and manage devices, applications, and data efficiently, securely, and affordably through the digital world. In 2011, President Barack Obama signed two Executive Orders directing both the public and private sectors to build the Digital Government

[3] https://www.researchgate.net/publication/355165955_From_Ethics_to_Operations_Current_Federal_AI_Policy

Strategy to complement several other ongoing initiatives aimed at building a 21st century government that is more efficient and effective for the American people:

- Executive Order 13571[4] Streamlining Service Delivery and Improving Customer Service aimed to improve the Federal Government service delivery through performance measurement evaluation, development, and implementation of best practices, and service improvement through technological initiative(s).
- Executive Order 13576[5] Delivering an Efficient, Effective, and Accountable Government provided the Director, Office of Management, and Budget with the authority to monitor and report on efficiencies, cost savings, and improvements implemented throughout the Federal Government.

On February 11, 2019, President Donald Trump signed Executive Order 13859[6], "Maintaining American Leadership in Artificial Intelligence," which established strategic objectives to further the Federal Government's approach to Artificial Intelligence development. The White House Office of Science and Technology Policy, which was the primary coordinating agency, also released the "American Artificial Intelligence Initiative: Year One Annual Report" in February 2020. The report provided a status update of each of the objectives set in the Executive Order and the progress realized in the previous year.

President Joe Biden signed Executive Order 13960[7] "Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government" in December 3, 2020. He also announced the formation of the National Artificial Intelligence (AI) Research Resource Task Force on June 10, 2021 to expand access to critical resources and educational tools to spur AI innovation as directed by Congress in the National AI Initiative Act of 2020. The Task Force serves as a Federal advisory committee to help create and implement a blueprint for the National AI Research Resource (NAIRR), a shared research infrastructure providing AI researchers and students across all scientific disciplines with access to computational resources, data, educational tools, and user support.

[4] https://obamawhitehouse.archives.gov/the-press-office/2011/04/27/executive-order-13571-streamlining-service-delivery-and-improving-custom

[5] https://obamawhitehouse.archives.gov/the-press-office/2011/06/13/executive-order-13576-delivering-efficient-effective-and-accountable-gov

[6] https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence

[7] https://www.federalregister.gov/documents/2020/12/08/2020-27065/promoting-the-use-of-trustworthy-artificial-intelligence-in-the-federal-government

## 2.3 Assumptions

The following assumptions are established for this document:

- The National Artificial Intelligence Initiative Act (AI-IA) establishes the national direction for AI and data analytics strategy, standards, and technologies. Implementing agencies will adhere to guidance provided by the National Artificial Intelligence Initiative Office, the National Artificial Intelligence (AI) Research and Development (R&D) Inter-Agency Working Group, and the National Institute for Standards and Technology.
- ATARC will continue to actively engage representatives from both state and Federal Governments as well as industry and academic partners to gain concurrence and define challenges and solutions.
- The current rate of investment, economic impact, and progress in AI and data analytics technologies will continue to advance at an exponential rate influencing the capabilities and possible achievements for government, industry, and academia.
- Maintaining pace with commercial AI and data analytics offerings, business practices, and R&D activities will remain a national security imperative.
- The ongoing societal discussion regarding AI and data analytics' potential impact on citizens' constitutional and human rights will benefit from informed, consistent engagement by Federal Government leadership.
- Effectively addressing and managing the cybersecurity threats related to Artificial Intelligence and data analytics technologies will remain a critical function of the Federal Government and commercial industry.
- Establishing and maintaining best practices across public and private sectors will benefit all AI and data analytics system stakeholders.
- Best practices for IT system development are generally applicable to the development of AI and data analytics systems.
- Vendor agnostic recommendations are most effective and broadly applicable.

# 3 Scope

This document focuses on practical guidance for government technology leaders interested in utilizing AI and data analytics within their operations. It is particularly oriented towards those near the beginning of their journey, though some sections will be useful for those improving their existing systems or transitioning to the next level. This document addresses key concepts needed to understand and leverage AI and data analytic technologies. There is a large and rapidly growing proliferation of guidance and strategy on AI and data analytics. This document does not try to duplicate that body of work but attempts to provide reference to relevant exemplars.

This document will do the following:

- Identify key concepts and aspects of AI and data analytics relevant to implementers
- Provide guidance on applying best practices for AI and data analytics system development
- Provide an overview and links to navigate to more detail and examples

This document does not address the following:

- Cost or effort estimates of implementing and deploying AI and data analytics
- A tutorial on AI, data analytics, statistics, machine learning, or IT system development
- Detailed AI and data analytics system risk assessment

The AIDA Guidebook aligns to other key references listed in Appendix D: References. This addendum includes the references for all key documents, organizations, and ideas mentioned throughout the guidebook.

# 4 AIDA Guidebook Fundamentals

Artificial Intelligence is defined in the National Security Commission on Artificial Intelligence's Final Report[8] as "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." This definition is appropriate for the range of artificial systems envisioned for the scope of this guidebook, as it includes essentially all of the past, present, and currently contemplated AI systems being developed, acquired, or regulated by the Federal Government. This definition also supports the distinction that AI is an engineering discipline. AI seeks to apply scientific understanding, frameworks, and techniques in order to create systems with specific behaviors and features. The suggestions and recommendations in this guidebook are provided with the intent of supporting engineering and system development efforts, rather than to influence scientific research or the discovery of new knowledge.

Developers describe AI and system development in various different terms, which can cause confusion. Understanding the history of AI development provides some context to the terms used and their origin. Given the many challenges of creating a system capable of demonstrating intelligence, historically developers took two approaches.

One approach is known as symbolic modeling and it involves the developer creating a model of intelligence. These models generally take the form of illustrations that are encoded in software, and represent sequences of taking an input, traversing the model, and generating an output. This was the approach initially taken by the attendees of the 1956 Dartmouth workshop on AI, which resulted in some very early progress examining the topics of computers, natural language processing, neural networks, and theory of computation, abstraction, and creativity[9].

Statistical modeling, which was inspired by emulating the human brain's neural network, is the second basic approach to building AI systems. Figure 4 provides an overview of symbolic modeling, statistical modeling, and how the two models compare.

[8] https://www.nscai.gov/2021-final-report/

[5] http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html

**Figure 4: Comparison of Symbolic and Statistical Modeling**

| | Symbolic Modeling | Statistical Modeling |
|---|---|---|
| Definition | Symbols, metaphors, and modeling used to facilitate positive change | A mathematical representation of observed data |
| Example applications | • planning<br>• reasoning (diagnosis or decision support)<br>• conversation generation | • pattern recognition<br>• motor skills (robots)<br>• speech generation<br>• search engines |

While these early systems were developed and altered manually, today's AI revolution is driven by machine learning algorithms – statistical models that are automatically refined by repeatedly executing thousands of data sets through an algorithm that improves itself or learns to improve its outcomes. Advances in modern AI systems are all made possible by the confluence of four converging technological trends: increases in computing speed, addressable storage capacity, network bandwidth, and advanced mathematics. We can access larger and more comprehensive data sets, as math gets more efficient, trustworthy, and robust in its representation of reality. This allows the training of more models faster and the pursuit of AI in a growing range of applications is now feasible because of advances in computational complexity.

### AI Example: Smart Assistants

Smart assistants are devices loaded with software that you use to access information, perform tasks or control other devices. You can have a smart assistant on your computer or your phone, but most people use them through a smart speaker. Well-known brands include Alexa, Siri, Cortana and Google Assistant.

In addition to symbolic modeling of AI and machine learning, there is a complementary perspective called human-machine teaming or human-machine learning that evolved from intelligence augmentation, which is the use of technology to enhance human intelligence and decision making. Although intelligence augmentation evolved around the same time as AI and machine learning, this perspective is relevant more than ever as we seek to leverage the advantages gained by human intelligence and machine learning models.

Development of modern AI systems focuses on continuously developing and deploying updates, often referred to as Agile Development. Some of the modern approaches to AI system development include:

- DevOps - the combination of software and hardware systems aimed at continuously developing and deploying software to increase operational efficiencies.
- DevSecOps - the integration of security requirements and testing within the DevOps process.
- AIOps - use of AI in supporting network operations, including anomaly and threat detection.
- MLOps - applying the continuous development and integration approach to the development of machine learning solutions, requiring additional automation step of automating training data access and preparation.

This guidebook recommends implementing MLOps to develop AI systems whenever feasible and Section 0 will provide additional detail on the Machine Learning Pipeline methodology and its current applications.

The remainder of the guidebook describes key AI and data analytics concepts and considerations relevant to the implementation of these models and analytics tools. **Figure 5** details the primary intents of this guidebook and how each section of the document relates back to those intents.

**Figure 5: Intent and Content of AIDA Guidebook**

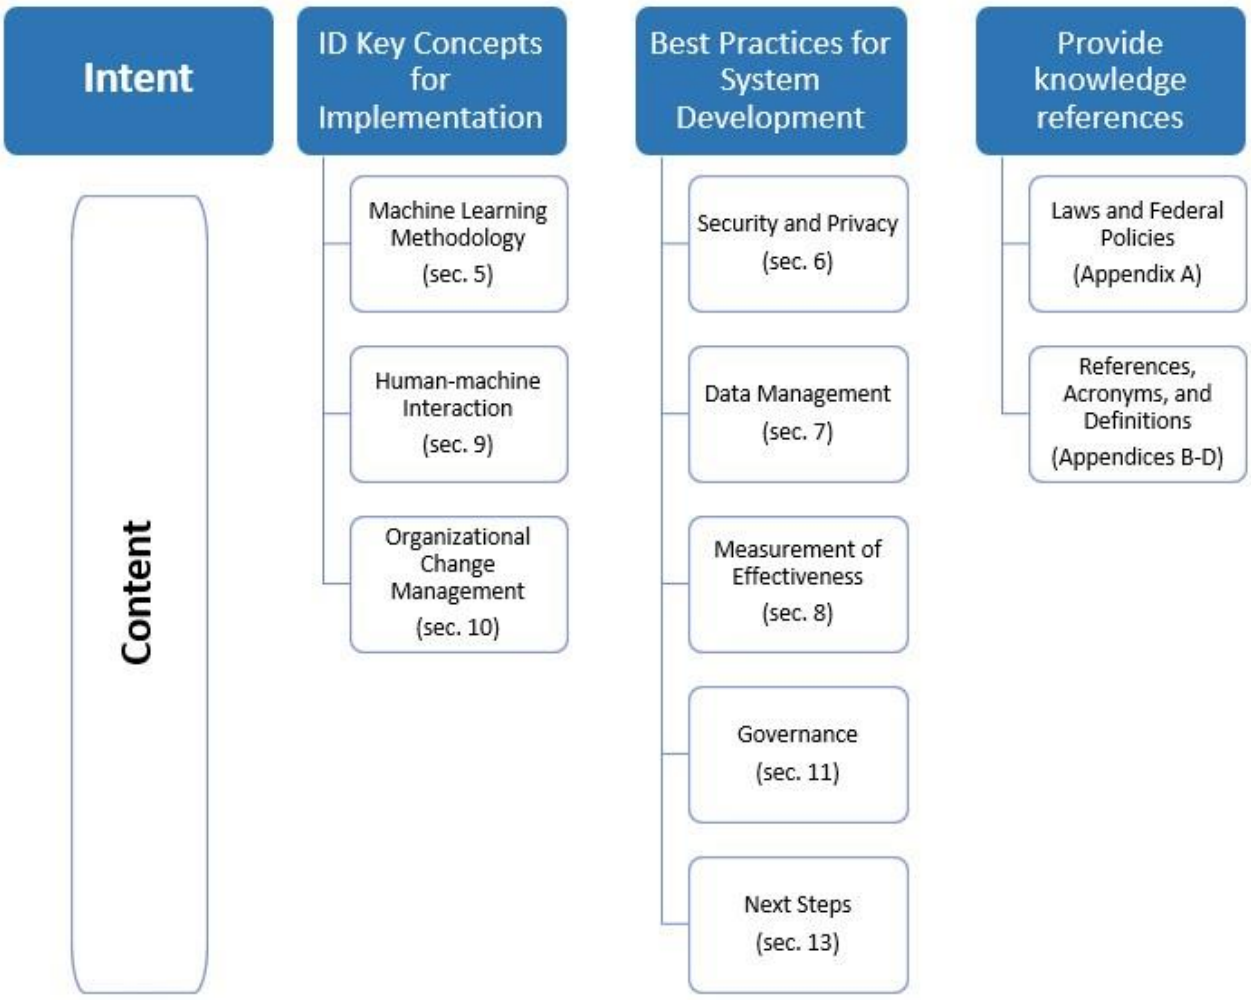

# 5 Machine Learning Methodology

This section provides an overview of how developers effectively build, evaluate, and manage analytic and learning systems through a machine-learning pipeline. A machine-learning pipeline is a way to codify and automate the workflow necessary to produce a machine-learning model. Machine learning pipelines consist of multiple sequential steps that do everything from data extraction and preprocessing to model training and deployment. **Figure 6** provides a high-level example of a machine-learning pipeline, with each step described in more detail to follow.

**Figure 6: Machine Learning Pipeline**

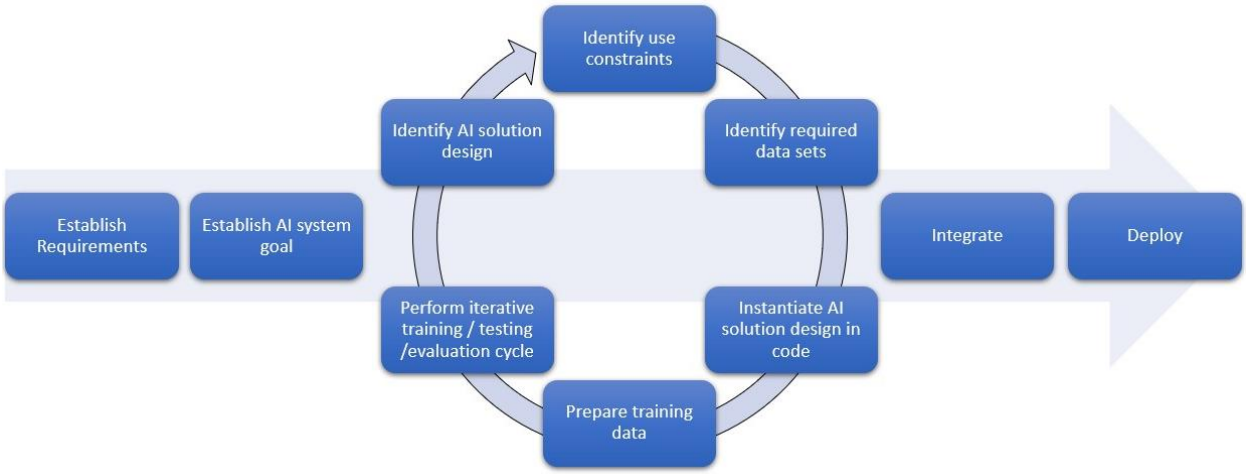


Pipeline steps:

1. Establish AI system goal – traditional goals of AI research include reasoning, knowledge representation, planning, learning, natural language processing, perception, and the ability to move and manipulate objects.
2. Establish requirements – consider desired performance, usability, integration, and statistical behavior.
3. Identify AI solution design – identify the algorithms and programming language to be used.
4. Identify use constraints – constraints enumerate the possible values a set of variables may take in a given world.
5. Identify required data sets – the more complex your model becomes, the more data you will need to determine its parameters.

6. Instantiate AI solution design in code – there are many programming languages to choose from such as C++, Java, Python, or R.
7. Prepare training data – includes cleaning the data of missing values, formatting data for consistency, making the units consistent, decomposing complex values, and aggregating simple values in the data.
8. Perform iterative training, testing, and evaluation cycle – input the data into the model in order to train it and improve model accuracy, setting a minimum acceptable accuracy threshold. If the testing and evaluation reveal that the model is not ready for deployment, return to step 3 and continue to refine.
9. Integrate – determine how machine learning will work within existing business processes.
10. Deploy – turn the model on in a real-world environment.

While the steps are sequential, the Machine Learning pipeline is often more of an iterative process, especially between the "identify AI solution design" and "perform iterative training, testing and evaluation cycle" steps. Often developers will need to revise either the training data, data preparation/augmentation, or machine learning model structure as a result of the training, testing, and evaluation step. The developer may uncover invalid assumptions which require revisiting the initial design setup and conducting continuous iteration until the overall AI system requirements are met.

A machine learning pipeline integrates both statistical behavior (i.e., statistical analysis and response requirements for the system) and use constraints (i.e., constraints on the how the system is to be deployed to support decision making – how autonomous, what timing requirements, what level of potential harm to users and/or subjects) into the AI system development process. Although powerful when implemented correctly, the machine-learning pipeline does offer challenges to a developer. One such challenge is imposing DevOps practices on a machine-learning pipeline. As previously defined, DevOps is the combination of software and hardware systems aimed at continuously developing and deploying software to increase operational efficiencies. In this process, DevOps could also be considered the code used to create the model. The machine-learning model is the combination of the data and the code, which is refined through continuous integration, continuous deployment, and continuous training of the model.

### AI Example: Language Models

An AI model that has been trained to predict the next word or words in a text based on the preceding words, it's part of the technology that predicts the next word you want to type on your mobile phone allowing you to complete the message faster.

While a DevOps code may be relatively set once developed, machine learning's challenge is how to keep code up to date with data while they change in parallel. Model accuracy and

resulting decisions can degrade with time due to data drifts and organizational overconfidence in the model. Machine learning is not a one-and-done process, creating an algorithm that is infallible for all time, but an ongoing and indeed constant evolution, where the AI algorithms repeatedly encounter new data and modify themselves to account for it. To counter this, organizations use continuous integration (i.e., merging code changes into a central repository), continuous deployment (i.e., using automated testing to validate if changes to a codebase are correct and stable), continuous training (i.e., testing of the model's validity), and a human element in the development loop.

Conducting these additional steps is what differentiates DevOps from MLOps, democratizing and streamlining the analytics process. On the technical side, MLOps bypasses the bottlenecks in the deployment process, i.e., between machine learning design and implementation or deployment framework. Strategically, MLOps makes machine learning accessible to those with less data and coding expertise. Additionally, an organization may benefit by exposing the quantitative rigor to qualitative subject matter, and by combining strategy and tactics to work together. This is important since only 13% of machine learning projects[10] make it into production due to a lack of organizational engagement.

There are risks to MLOps in addition to the benefits stated above. MLOps may oversimplify the development process, cloaking intermediate steps, which may pose a challenge to those with less data and coding expertise. This may lead to downstream impacts if the code & data fall out of alignment. Developers often weigh the risks and rewards of MLOps, asking questions such as:

- How much additional infrastructure is required to make MLOps sustainable?
- Does the organization already have a substantial infrastructure?
- How to measure the increase in productivity vs. the increase in risk?

Each organization will then identify acceptable risk level when determining how to proceed. Additionally, organizations must often consider how tightly they link operational SMEs and data or modeling SMEs; if model accuracy monitoring includes ethical metrics (race, gender, etc.); and maintaining an organizational culture of respect for all contributors' expertise. Infrastructure SMEs manage the CI/CD technical side, working closely with other partners on CT. Data SMEs understand operational SMEs as well as manage the CI/CD data side and work closely with other partners on CT. Operational SME coordinates with data SMEs and are responsible for proactively engaging data and infrastructure SMEs on CT.

[10] https://venturebeat.com/2019/07/19/why-do-87-of-data-science-projects-never-make-it-into-production/

# 6 Privacy and Security

Privacy and Security are key components of AI and data analytics implementation. Even though privacy and security are complementary, there are key distinctions in how they both contribute to AI systems.

Privacy is relevant to AI systems in two ways – maintaining the confidentiality of individuals whose data is used for model training, the stage in which developers use dataset(s) to train a machine learning algorithm, as well as the confidentiality of subjects to which the system is applied. The term "subject" is useful to distinguish between individuals to whom the system is applied, as in a person diagnosed using AI to read their computed tomography (CT) scan, or a person whose medical risk is evaluated by a categorization system and a user applying the system.

Security refers to keeping both the AI system's development code (i.e., the model itself, testing tools, and test results) and the data used to train it free from interference. A secure AI system means both its development code and its training data will not be destroyed, altered, or made inaccessible to authorized users. Securing artificial systems is inherently more complicated than securing non-AI systems because of the need to secure training data.

Adhering to data protection policies for information in transit, at rest, and in use is very important when planning AI and data analytic systems. Data in transit, also known as data in motion, is commonly defined as data that is actively moving from one location to another such as across the internet or through a private network. Data protection in transit is the protection of this data while it is traveling from one network to another or being transferred from a local storage device to a cloud storage device. Along with data in transit, there is data at rest, which is data that is not actively moving from device to device or network to network such as data stored on a hard drive, laptop, flash drive, or stored in some other way. Additionally, data in use or in the memory space of a program running, such as a public cloud environment, must also be protected. This data is widely protected via commercial and open source tools running on a root of trust enabled by encrypted on-chip memory access. All major public clouds offer some form of these as trusted execution environments or enclaves.

Data encryption is one of the simplest technologies to implement to secure data in transit and at rest. Encryption, which entails the process of converting information or data into a code, is a key element in data security. Prior to transporting sensitive information, businesses generally choose to encrypt the data so that it may be protected during transmission. There are several methods for doing this.

There are connection-level encryption schemes that can be enforced, and the most widely used types of encryption are connections using Hypertext Transfer Protocol Secure (HTTPS), Transport Layer Security (TLS), and File Transfer Protocol Secure (FTPS). HTTPS is encrypted in

order to increase security of data transfer. This is particularly important when users transmit sensitive data, such as by logging into a bank account, email service, or health insurance provider. Any website, especially those that require login credentials, often uses HTTPS. A primary use case of TLS is encrypting the communication between web applications and servers, such as web browsers loading a website. TLS can also be used to encrypt other communications such as email, messaging, and voice over IP (VoIP). FTPS is a secure file transfer protocol that allows businesses to connect securely with their trading partners, users, and customers. Sent files are exchanged through FTPS and authenticated by FTPS supported applications such as client certificates and server identities.

When compared to the data in transit, data at rest is generally harder to access, which means that oftentimes private information, such as health records, are stored this way. Making the interception of this data more valuable to hackers and more consequential for victims of cyber-attacks. Despite the greater security, there is still a risk of this data being intercepted by hackers through cyber-attacks, potentially causing private information such as addresses and financial records to be released, putting an individual's safety at risk. Protecting all sensitive data, whether in motion or at rest, is imperative for modern enterprises as attackers find increasingly innovative ways to compromise systems and steal data.

If the data must be protected for many years, one should make sure that the encryption scheme used is quantum-safe. Current publically available quantum computers are not powerful enough to threaten current encryption methods. However, as quantum processors advance, this could change. Most current public-key encryption methods (where different keys are used for encryption and decryption) could be broken with a powerful enough quantum computer. On the other hand, most current symmetric cryptographic algorithms (where the encryption and decryption keys are the same) are not susceptible to quantum attacks, assuming the keys are sufficiently long.[11]

For applications where confidentiality of the data in use is of utmost importance, additional technologies could be used. When one wants to keep the data private even while it is being processed, there are a number of technologies that can be employed independently or, in some cases, even together. These include homomorphic encryption, differential privacy, federated computing, and synthetic data. Homomorphic encryption is a technique that allows operations to be performed on encrypted data without decrypting it.[12] This permits the confidential processing of data on a system that is untrusted. The results of the computation can only be only decrypted with the original key. The biggest barrier to widespread use of homomorphic encryption has been its poor performance. It is significantly slower than performing the

[11] http://www.pqcrypto.org/www.springer.com/cda/content/document/cda_downloaddocument/9783540887010-c1.pdf

[12] See https://eprint.iacr.org/2015/1192 for an overview of homomorphic encryption and related technologies

corresponding computation on unencrypted data. Differential privacy is a technique to remove some data from a dataset so that the identity of individuals cannot be determined, even when combined with other datasets.[13] Federated computation, such as federated learning[14] or shared analytics[15], allows machine learning or data analytics, respectively, to be performed remotely over decentralized data without transmitting that data. Google developed federated learning to update predictive typing models for Android keyboards by learning from millions of users by tuning the models on the users' devices without exposing the users' data.[16] Others have experimented with using federated learning across multiple organizations, such as healthcare systems. Synthetic data is an approach to take sensitive or private data and produce mock data that has similar statistical characteristics to the original data.[17] This is helpful for algorithm developers to test and improve the performance of their algorithms, typically, before being used on the actual private data in a more difficult to use secure enclave.

## 6.1 Data Transparency

Today, firms across the world are incorporating AI-based and data analytics systems to automate their business processes and enable their workforce to focus on customer and operational needs. Data Transparency is paramount to make data, analysis, methods, and interpretive choices underlying researcher claims visible in a way that allows others to evaluate them. However, the incorporation of AI technology is impeded by several challenges, and the biggest one is the lack of data transparency. Often you cannot manipulate or control a system if you do not know what goes into it. AI systems are often considered as a black box, which means regulators are able to see what goes into the AI and what comes out but are unable to see how the algorithms and technology actually works, and this makes it challenging to pinpoint logical errors in the underlying algorithms. Tech companies are hesitant to share their algorithms because it can lead to potential intellectual property infringement and theft. These challenges in data transparency are often at odds with efforts to enforce data privacy. Sometimes protected data is not transparent data, and in these cases, cybersecurity can introduce challenges in efforts to expand, control, and monitor AI as it is applied. In this instance, having greater security might present barriers to having more visible, transparent AI, making it imperative to develop explainable AI. It is easy to see how this can be a problem when faced with an application of this technology for government or in a public domain.

---

[13] https://www.vanderbilt.edu/jetlaw/

[14] https://arxiv.org/abs/1511.03575

[15] https://sagroups.ieee.org/2795/

[16] See https://ai.googleblog.com/2017/04/federated-learning-collaborative.html for how Google uses Federated Learning on Android devices

[17] https://ieeexplore.ieee.org/document/7796926

## 6.2 Security and Privacy Compliance

Security and Privacy Compliance refers to one of the very special and unique ways that AI can help government as well as industry. AI has a unique and efficient ability to identify, monitor, and address compliance issues in many different processes. This process was historically done manually where a compliance team would sift through emails or phone records and flag anything that appeared nefarious. While somewhat tedious, these compliance measures have helped institutions maintain industry wide standards and regulations as well as their own internal standards for employee conduct. However, over time technology has evolved and has now become a mainstay across all industries, offering support and assistance to organizations and institutions, making them more effective and efficient in their daily operations. More specifically, artificial intelligence has provided industries and organizations with comprehensive compliance tools that automate the process. There are, however, significant security and privacy issues that are raised when such a system is utilized in a corporate or government environment including concerns regarding bias, discrimination, and privacy. These concerns have led to some calling for policymakers to introduce regulations on the technology.

## 6.3 Risks, Opportunities and Mitigations

Artificial intelligence-based and data analytic technologies present several unique risks, opportunities, and corresponding mitigations. Some of these risks include increasing automation of certain jobs, gender, and racial bias issues stemming from outdated information sources or autonomous weapons that operate without human oversight. This is especially the case with autonomous systems, a collection of networks all under the management of a single organization or establishment. Autonomous intelligence refers to systems that can adapt to different situations and can act without human assistance. The significant benefit to the use of autonomous intelligence is that there are tremendous opportunities for reducing workload and generating efficiencies, such as energy efficiency, more accurate demand forecasting, predictive maintenance, optimized manufacturing processes, or automated material procurement, amongst others. Care must be taken that such systems function in an environment where the risk of failure does not result in a catastrophic result. When dealing with a high-hazard scenario, the chance of failure rises, and mitigation factors and techniques are often implemented to prevent or repair such incidents. The frequency and severity of these events, however, play a huge part in determining how safe AI is considered. The possibilities in the financial services industry are comparable. However, the risk and opportunities can be addressed differently from the mitigation standpoint. This is the case because we are not dealing with a situation where there could be a loss of life. An organization might apply an early warning system into an early learning system that prevent threats materializing for real. While AI is still developing, it can already be used to mitigate risk in some key areas.

## 6.4 Laws, Regulations, and Rules

Concerns about potential misuse or unintended consequences of AI have prompted efforts to examine and develop standards within the public and private sectors around the establishment of reliable, robust, and trustworthy AI systems. State and federal lawmakers have been evaluating AI's benefits and challenges and a growing number of measures have been introduced to oversee the impact and regulate the use of AI and data analytics. During 2016-2019 a wave of AI Ethics Guidelines were published in order to maintain social control over the technology. Regulation is considered necessary to both encourage AI and manage associated risks. There are a host of commonly referenced laws and regulations included in Appendix A: Representative Laws and  for reference and they are relevant to designers, developers, testers, and users of AI-based and data analytic systems. For data privacy, the European Union's General Data Protection Regulation (GDPR) is often cited as the strictest in the world. While it applies to European citizens, it is already affecting how US entities provide privacy controls over data they control. It also helped shape the far reaching California Consumer Privacy Act (CCPA) of 2018.

# 7 Data Management

Data analytics is referenced throughout this guidebook and refers to studying features, patterns, or characteristics represented by a data set. Common data analysis includes statistical methods for discovering and characterizing patterns in the data, such as averages, means, medians, ranges, and skew. Meta-analysis refers to assessing features of the data itself – in effect creating data of the data. Examples include sources, units, time stamps, quality, and physical location. Big Data refers to techniques used for analyzing large, heterogeneous, distributed data sets, and often involves tasks to establish data standards for representation, units, acceptable ranges, and standard data transformations. Shared and Distributed Analytics refer to techniques for defining analytic methods, including their representation and system-specific implementation, to be shared across heterogeneous environments and using these tools to perform orchestrated analytics. An example of orchestrated analytics would be a piece of software taking siloed medical data from multiple data storage locations, combining it, and then using it within a data analysis tool to evaluate various healthcare outcomes.

Data Management refers to the methods, tools, procedures, and processes used to ensure the availability, confidentiality, integrity, quality, reliability, and usefulness of data sets. Data management standards include data dictionaries, data ontologies, metadata structures, quality measures, database optimization and backup standards, and data retention standards. Data management processes include both manual (e.g., Extract, Transform, and Load (ETL) processes, penetration testing) and automated system processes (e.g., data backup, database optimization, continuous monitoring, load balancing). As organizations migrate their IT environment to the cloud, more data management is automated and/or performed by the cloud service provider.

Most organizations lack the enterprise data management framework to ensure that trusted and verified critical data is widely available to, or accessible by, mission commanders, warfighters, decision-makers, and mission partners in a real-time, useable, secure, and interlinked manner. The lack of that framework limits data-driven decisions and insights, which hinders organizations' ability to execute AI and machine learning solutions in an efficient and effective manner.

Additionally, organizations will often acquire, test, upgrade, operate, and maintain software and hardware solutions to fulfill data interoperability needs. These gaps are frequently addressed with complicated human-machine interfaces that add unnecessary complexity and risk of error. This limits an organization's capacity to operate effectively across the full spectrum of threats that may be used against you at top-machine speed across all domains.

Organizations often evaluate how to improve and encourage key skills and abilities (KSAs) in data fields, across the organization as appropriate, for effective data management implementation. Efforts to move organizations to a more data centric structure will keep the

workforce as the central driving component. This shift will require organizations to assess current talent, recruit new data experts, and retain and develop our workforce while establishing policies that ensure data talent is cultivated and maintained.

## 7.1 Guiding Principles

The Department of Defense (DOD) published the DOD Data Strategy Framework in September 2020. As shown in **Figure 7** below, the Department of Defense leverages eight guiding principles to influence the goals, objectives, and essential capabilities in this strategy. These guiding principles are foundational to all data efforts within DOD[18].

**Figure 7: DOD Data Strategy Framework**

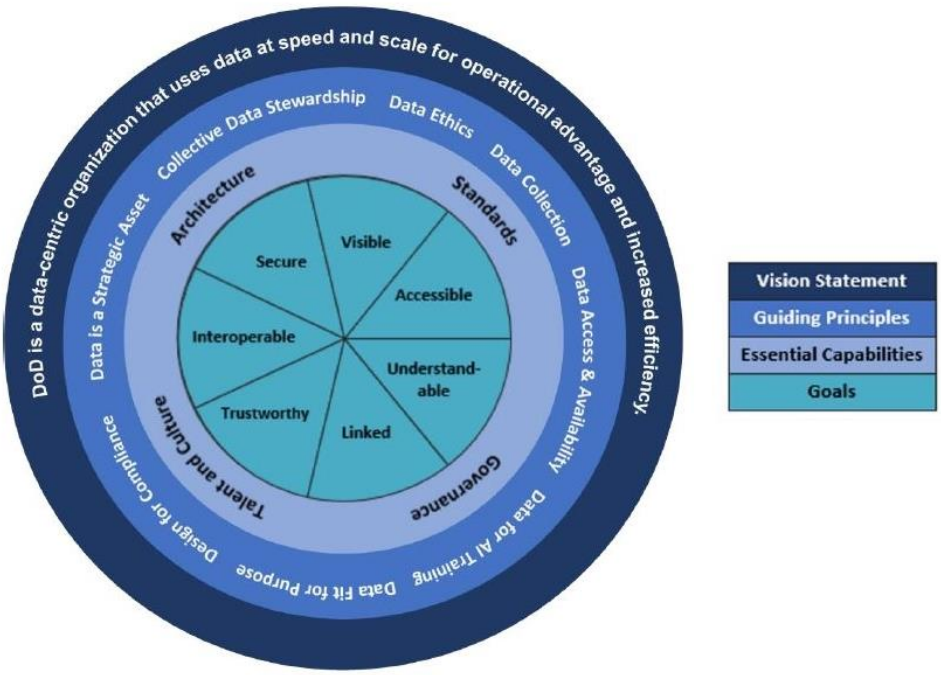


## 7.2 Data is a Strategic Asset

Data is a high-interest commodity and can be leveraged in a way that brings both immediate and lasting advantage. As an organization shifts to managing its data as a critical part of its overall mission, it gains distinct, strategic advantages over competitors and adversaries alike.

[18] Source of graphic: DoD Data Strategy: Unleashing Data to Advance the National Defense Strategy, 30 September 2020

These advantages will be reflected in more rapid, better-informed decisions using trustworthy and integrated data.

## 7.3 Collective Data Stewardship

To exploit data fully for decision-making, an organization must define roles and responsibilities for data stewardship; assign data stewards, data custodians, and a set of functional data managers to achieve accountability throughout the entire data lifecycle. Data stewards establish policies governing data access, use, protection, quality, and dissemination. Data custodians are responsible for promoting the value of data and enforcing policies, and functional data managers implement the policies and manage day-to-day quality.

**Figure 8** provides an overview of how policy, standards, guidelines, and procedures are linked in support of data stewardship.

**Figure 8: Data-Centric Enterprise Scientific Data Stewardship Framework**

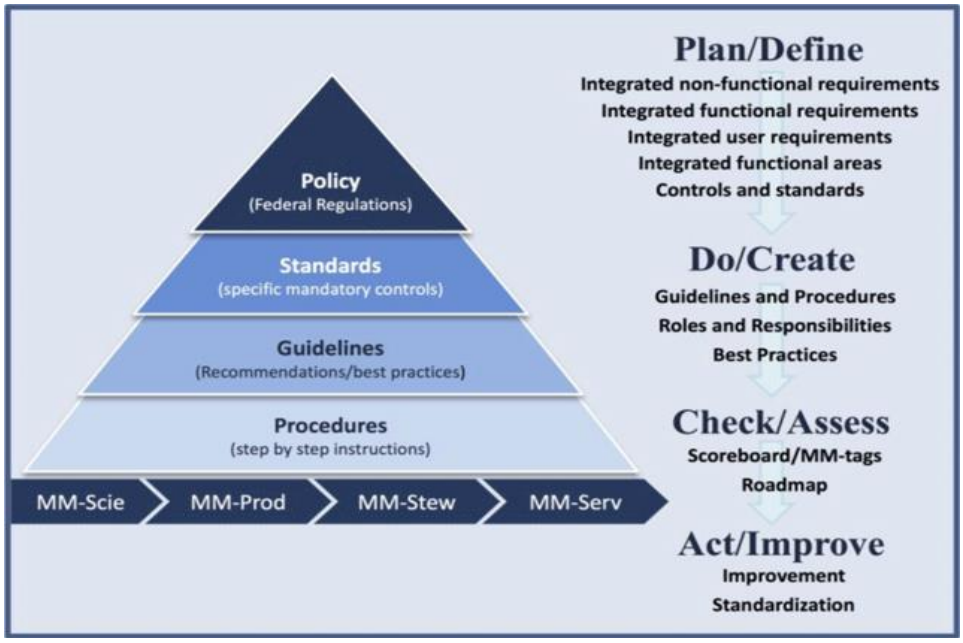


Image originally from article titled “A Conceptual Enterprise Framework for Managing Scientific Data Stewardship." Authors give permission to use under Creative Commons Attribution 4.0 License (CC BY 4.0)

## 7.4 Data Ethics

The ethical use of data is at the forefront of all plans and actions for how data is collected, used, and shared. As the Secretary of Defense stated in his guidance on AI Ethics on February 21,

2020, "*Although technology changes, the Department's commitment to the Constitution, the Law of War, and the highest standards of ethical behavior does not.*" Whether for AI or advanced analytics, ethical principles regarding the responsible use of data remain important for data and analytics leaders.

The DOD Joint Artificial Intelligence Center (JAIC) has instituted the following action to ensure military ethics and AI safety[19]:

- The Department will articulate its vision and guiding principles for AI ethics and safety in defense matters
- Investing in research and development for resilient, robust, reliable, and secure
- Continuing to fund research to understand and explain AI-driven decisions and actions
- Promoting transparency in AI research
- Advocating for a global set of military AI guidelines
- Using AI to reduce the risk of civilian casualties and other collateral damage

## 7.5 Data Collection

Regardless of the data domain, community, or use, the challenge remains the same – to discover and collect data and continuously add value to best inform the decision-maker. Consequently, organizations will enable electronic collection of data at the point of creation and maintain the pedigree of that data. When implementing best data practices, the moment data is created, tagged, stored, and cataloged. When the data is combined or integrated, the resulting product is to be immediately collected, tagged, curated, and desensitized. To expedite these processes and to minimize the risk of human error, these steps are to be automated to the maximum extent possible.

## 7.6 Data Access and Availability

Closely aligned with data stewardship and collection are the accessibility and availability of data. This is enabled by successful implementation of enterprise capabilities, such as an enterprise cloud; Identity, Credential, and Access Management (ICAM); and associated data-sharing tools. To achieve success in the adoption of these technologies, organizations will embrace new and improved technologies, processes and policies, as well as new cultural norms. An example is the cultural shift from the need to know (i.e., information withholding) to the responsibility to provide (i.e., information sharing). Making data available across business systems is essential to gaining an enterprise-wide view into the daily operations of an organization and critical to the

[19] DoD JAIC AI Strategy: https://dodcio.defense.gov/About-DoD-CIO/Organization/jaic/

success the Digital Modernization Strategy[20]. Data is often made available for use by all authorized individuals and non-person entities. Leaders that embrace information sharing, set the example by educating their organizations and enforcing data sharing best practices. Data sharing is at times restricted when required by law or organization-wide policy and where security, privacy, or ethical considerations are involved. Furthermore, we must ensure not only that data is protected, but also that it is handled properly throughout its lifecycle.

## 7.7 Data for Artificial Intelligence Training

AI is a long-term data competency grounded in training- quality datasets (TQD) that are the pieces of information and associated labels used to build algorithmic models. TQD and the algorithmic models will increasingly become a valuable digital asset. As organizations modernize and integrate AI technologies, generating organization-wide visibility of and access to these digital assets will be vital in an era of algorithmic warfare. We must also understand that our competitors gain advantage if these assets become compromised. Therefore, there is a requirement to create a modern governance framework for managing the lifecycle of the algorithm models and associated data that provides protected visibility and responsible brokerage of these digital assets.

## 7.8 Data Fit for Purpose

Fit for purpose data is quality data that is readily discoverable and understood within the context of its intended use. It includes careful consideration of any ethical concerns in data collection, sharing, use, representation of the information intended, rapid data integration, and minimization of any sources of unintended bias. Customers of an organization's data have their own requirements for accessing data, which may or may not align with the purpose or intent of the original data collection. Additionally, in some instances, legislation or a regulation may specify how data is to be used and from which source the data must be consumed. Data quality is measured in terms of accuracy, completeness, reliability, accessibility, and timeliness. Data accuracy refers to error-free records that can be used as a reliable source of information. Data completeness is the percentage of all required data currently available in a dataset. Data reliability means that data is complete and accurate, and it is a crucial foundation for building data trust across the organization. Data accessibility refers to the barriers to fully leveraging the data contained within a database. Finally, data timeliness is the degree to which data represent reality from the required point in time.

[20] https://media.defense.gov/2019/Jul/12/2002156622/-1/-1/1/DOD-DIGITAL-MODERNIZATION-STRATEGY-2019.PDF

## 7.9 Design for Compliance

Implementation of IT solutions provides an opportunity to automate the information management lifecycle fully, properly secure data, and maintain end-to-end records management. Many organizations make data management and compliance with policies a top priority. Compliance with required data policies is a critical success factor for continued funding of future solutions and will be a gate for authorizations to operate.

## 7.10 Essential Capabilities

Four essential capabilities are needed to accomplish data goals: architecture, standards, data governance, and talent and culture. These capabilities are not specific to a single goal, but they are necessary to enable achievement of all goals.

### 7.10.1 Architecture

This strategy emphasizes access to data and the capability to adjust requirements in stride with changes in technology and data sources. Architecture, enabled by enterprise cloud and other open-architecture capabilities, allows pivoting on data more rapidly than adversaries are able to adapt. The ability to create and deploy small applications rapidly and on a regular basis in support of user demands revolutionizes how businesses utilize data, putting them at a competitive advantage. **Figure 9** provides an overview of an agile architectural approach, which enables incremental value to be delivered by balancing emergent design and intentional architecture. This agile approach allows the architecture of data and systems, even a large solution, to evolve over time, while simultaneously supporting the needs of current users.

**Figure 9: The procedure of developing an integrated system in Agile Development approach.**

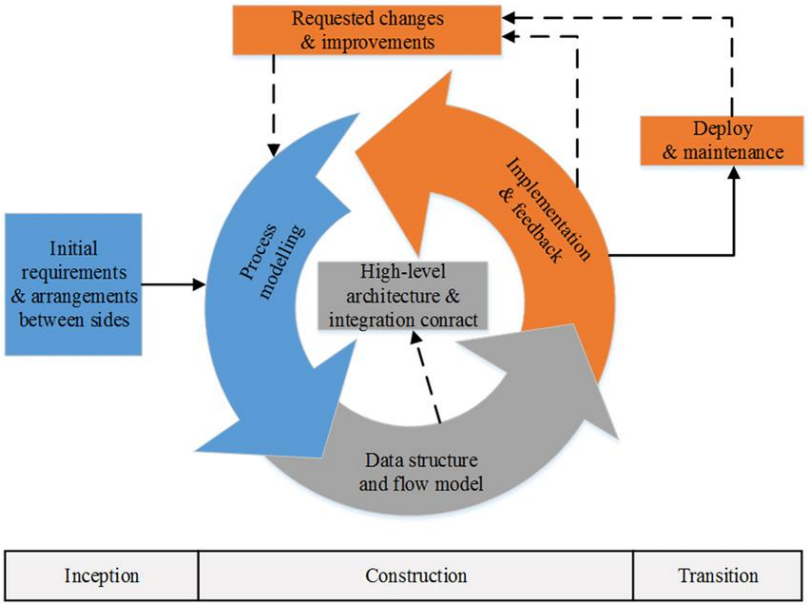

Source- Web Oriented Architectural Styles for Integrating Service e-Marketplace Systems - Scientific Figure on ResearchGate.

### 7.10.2 Standards

Organizations often employ a family of standards that include not only commonly recognized approaches for the management and utilization of data assets, but also proven and successful methods for representing and sharing data. Given the diversity of systems, these standards can be applied at the earliest practical point in the data lifecycle, such as during initial on premise data encryption, and industry standards for an open data architecture can be used wherever practical. In the instance of secure encryption, standards for maintaining data security can be applied through all stages of the data lifecycle, not exclusively during the initial backup or download of data. Standards enable data and information to be readily and securely utilized and exchanged. Without a standardized method for data creation, collection, storage, and sharing, errors in data accuracy and security can arise and jeopardize operations.

### 7.10.3 Data Testing and Evaluation

Data testing and evaluation allows organizations to make sure that the data is correct and complete. This process can determine if a database can go successfully through any needed transformations without loss, that the database can deal with specific and incorrect data in a proper way, and that all the expected data in the front end of your system has been represented correctly in the corresponding input. A test is defined as a program, procedure, or process to obtain, verify, or provide data for determining the degree to which a system (or subsystem) meets, exceeds, or fails to meet its stated requirements. To evaluate in this context means to review, analyze, and assess data obtained from testing to project system performance under operational conditions. The overall goal of testing and evaluation is to identify data errors, system shortfalls, and risk associated with go live, providing crucial information to decision makers.

### 7.10.4 Data Governance

Data governance provides the principles, policies, processes, frameworks, tools, metrics, and oversight required to effectively manage data at all levels, from creation to disposition. Data governance allows stakeholders to be heard and represented in an organized fashion. Data governance is conducted at cascading levels with all problems addressed at the lowest level possible. Data governance comprises of local system choices that impact data throughout its lifecycle. In order to support localized system decisions and resolve issues at the lowest level possible, organizations often utilize a low-code or no-code IT environment, where technology is utilized that enables non-developers to build custom apps without coding, so that easy data integration, no matter an employee's technical background, is possible.

### 7.10.5 Talent and Culture

Data-centric organizations are culturally transformed workforces. This transformation occurs over time with the development of cultural change advocates and buy in from all levels of an organization. Modern organizations continue to adapt their decision-making culture in order to reflect the demands of implementing new technologies that support artificial intelligence and data analytics. Recruiting and training contemporary, agile, information-advantaged workforce fosters a supportive ecosystem for collaboration among data experts establishes centers for data engineering excellence and equips leaders and managers to make data-informed decisions.

## 7.11 Goals and Enabling Objectives

Data is an essential and integral part of an organization's mission. Data is ubiquitous. Organizational business systems generate enormous volumes of data of which all retain and share their data for broader use. It is critical that data be of high quality, accurate, complete, timely, protected, and trustworthy. The following goals support an organizations cultural shift in the adoption of data as a strategic asset.

Data will be:

- **Visible** – Consumers can locate the needed data
- **Accessible** – Consumers can retrieve the data
- **Understandable** – Consumers can find descriptions of data to recognize the content, context, and applicability
- **Linked** – Consumers can exploit complementary data elements through innate relationships
- **Trustworthy** – Consumers can be confident in all aspects of data for decision-making
- **Interoperable** – Consumers and producers have a common representation and comprehension of data
- **Secure** – Consumers know that data is protected from unauthorized use and manipulation

# 8 Measurement of Effectiveness/Key Performance Indicators (KPI)

Measurement of effectiveness and key performance indicators (KPIs) refer to the practice of defining observable, measurable quantities that relate to desired performance characteristics or levels. These are generally defined at the level of a business process in order to provide insight into the detailed operation of an organization while performance improvements are pursued. Measuring the performance in terms of cost, financial, quality, time, flexibility, delivery reliability, safety, customer satisfaction, employees' satisfaction, and social performance indicators have significant positive impact on the overall performance of organizations. Tracked over a defined period, KPIs help an organization make better, data-driven decisions, guide behavior, productivity, and decision-making, and provide transparency and accountability.

A dashboard to track and display measures and KPIs can provide an organization with an idea of how they are progressing toward their goals. This section discusses how KPIs may be used by an organization to track development of data science skills, how well they are leveraging their training and staff development, and the demand for data science skills.

## 8.1 Progress Gaining and Applying Data Science Skills

It is important for an organization to identify and track employee progress completing their planned data science training. Developing employee data science skills is important to ensure they have the required knowledge, skills, technology, and expertise to support the organization and allows employees to stay on top of the latest trends in the domain. Depending on the employee role and the training needs, the training schedule could be six-months to two-years. For example, delays in completing training may be reviewed to understand if the delay is due to a lack of employee interest or work requirements interfering with staff development. Management needs to take steps to address training delays due to work assignments.

Another way of assessing the success of data science training is whether employees can put what they have learned into action. Organizations can evaluate time since employees finished training and whether they were able to apply their new data science abilities in their current role. If personnel are spending 90% of their time on non-data science activities after a year of training, management can take steps to improve efficiency or risk losing these talents. Corrective actions might include additional training, career counseling, change of role, or change of organizational alignment.

## 8.2 Skill Demand

Organizations often measure the ongoing demand for data science skills, such as programming, data visualization, machine learning, etc. They can also examine if the demand for these skills exceeds the current pool of available and trained staff. If yes, then the organization may choose

to continue to invest in staff training and new hires with desired skillsets as required. Depending on the particular skills in demand, the organization can focus on staff incentives to learn desired skills.

**Figure 10** provides as example, the Kanban training tracking board of staff in training, completed training, working on tasks, and completing data science tasks to provide staff and management with near-real-time information to anticipate training and task needs over time. The illustration shows that staff in training for all skills listed except for Extract, Transform, and Loading (ETL) has staff in the queue. Fully trained personnel for AI, machine learning, and Data Visualization exist. Nevertheless, ETL is a gap. In the illustration, notice that AI and machine learning only have a few staff working on a project. Machine learning and ETL have unstaffed needs according to the example. Management can take a closer look at why AI has only two trained staff working and whether there is forecast demand for AI skills in the near and mid-term. Looking at machine learning, management will notice that there is demand for four trained staff, but only one is working. Why? Is it because the trained machine learning staff are being used on non-machine learning tasks? Is this something that management can fix? Notice that there are four ETL staff trained. Yet only one is working an ETL task even though there are three open positions for ETL skills? Management can investigate the situation to see if they can make appropriate adjustments to better utilize trained staff.

**Figure 10: Kanban board of data science training**

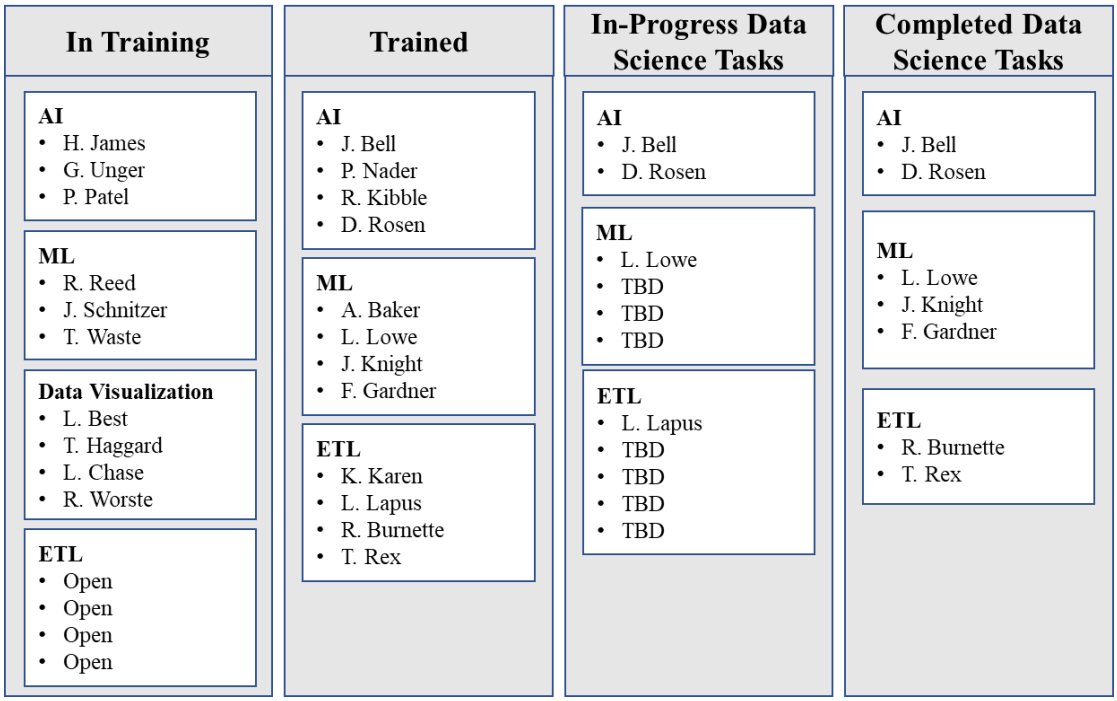

The Kanban board or other similar dashboards provide valuable insights into an organization's demand for particular data science skills such as staff availability, status of training, and staff required. Thus, informing hiring and training goals for the organization.

Additional items to include in a dashboard are a list of data science projects with the skills sought or required, the level of effort for each skill, and the anticipated start and end dates. This helps managers and staff plan availability to apply skills to task more effectively.

A scorecard of enrollment and completion rates via each training delivery method provides insight into training the effectiveness of each delivery mode (e.g., computer-based training, in-class training, and university courses). Tracking the training dropout rate provides information on how well each training mode fits the skills and interests of the staff.

An interesting measure is to determine if there is a reduction in the time to implement a data science solution as staff gain experience and proficiencies in designing and implementing data science best practices obtained through training. Measuring the time to deliver a data science product can be reviewed periodically (e.g., quarterly, annually) to see if process improvements are needed.

Charting the demand for AI and machine learning and other emerging technologies skills versus trained staff available provides useful information for making hiring and training decisions.

## 8.3 Department, Agency, Division, Bureau Perspective

The impact of data science on the organization can be tracked in terms of improved decision-making and/or timeliness of decisions. This will provide leadership with a qualitative measure of the value of data science investments to the organizations. Items for measurement include:

- Number of impactful insights that resulted from applying data science
- Number of projects that plan to use data science compared to total projects
- Time saved because of using data science
- Cost avoidance because of using AI or machine learning solution(s)
- Human resources freed from manual repetitive work to focus on higher value tasks
- Total cost of ownership of AI or machine learning solution
- Time and cost to operate and maintain the AI or machine learning solution versus the impactful benefits – is the juice worth the squeeze?

## 8.4 Project and Task Perspective

Measures to track from a project or task perspective can provide more detailed information on the type and mix of skills needed to complete a typical project. Examples of measures include:

- Staffing requisition
- Skills required (e.g., data analyst with SAS programming skills, DBA, HDFS)
- Level of Effort (e.g., 50%, 75%)
- Duration (e.g., 3-months, one-year)
- Clearances required (e.g., SAP, TS, S)
- Location of work – work needs to be done in a SCIF in Chantilly, VA
- Time to implement a data science solution

## 8.5 Speed and Quality of Decisions Made Because of Applying Data Science Techniques

Measuring the speed and quality of decisions made because of applying data science techniques help leadership understand the impact of their investment. Examples of KPIs that leadership might want to include in the dashboard include:

- Time it takes to prepare and process data for AI or machine learning model use – Based on the size and quality of data set(s). More poor-quality data sets require more time to process than a single high quality data set
- Time to produce an actionable insight, prediction, recommendation, etc. Are we getting better at providing actionable insights faster over time?
- Time to model obsolescence – models no longer provide accurate, actionable, or useful results to the end users –models obsolete within three-months, six-months. Perhaps the environment is too volatile and data science techniques are a poor technique.
- Ethical and fair decisions made because of applying data science techniques – Are the models we build ethical and fair? Did we find out too late that the models are unethical or biased?
- Number of predictions or recommendations that were rejected – How many times did we build a model that was rejected by the end user?
- Number of predictions or recommendations that were found to be erroneous
- Cost to build, maintain, and operate AI or machine learning solution

# 9 Human-Machine Interaction

Human – Machine Interaction refers to the design and optimization of how users and IT systems interact to collaboratively perform a function. In the context of AI and data analytics, the concept of a centaur user (where human and artificial intelligence work together toward a common goal) is proving effective, whereby a user takes advantage of the capabilities of an AI and data analytics systems. This approach has been shown to allow the user to focus on solving elements of a problem, such as chess playing, in ways that AI and data analytics might not be able to, due to the user being freed from the need to track data, apply known rules, or evaluate simple scenarios.[21]

To create user-centered AI or machine learning products, there are a number of steps a team can take. First, all team members can take the time to read "AI fails and how we can learn from them."[22] Why start there? Understanding the challenges of standing up great AI projects means learning from the ones that didn't go so well. In particular, read the fails in the sections:

- The Cult of AI: Perceiving AI to be more mature than it is
- We're Not Done Yet: After developing the AI
- Failure to Launch: How people can react to AI
- Lessons Learned

If interested in additional AI Fails, the Partnership on AI now tracks what they call AI Incidents in the AI Incident Database (AIID). An initial set of over 1000 incident reports have been included in the open-source database [23]

## 9.1 AI Challenges Are Multidisciplinary, so They Require a Multidisciplinary Team

The challenges to overcome when developing or implementing AI are diverse and can be both technical and social in nature. As a result, no one person or discipline can singlehandedly "fix" AI. Those of us on the front lines of building the AI share many attributes (i.e., similar education and degrees, life experiences, and cultural backgrounds). If we do not actively work to incorporate other valid perspectives into the development process, we risk having the AI reflect our assumptions about how the product will be used and by whom, instead of being based on

[21] This chapter: ©2021 The MITRE Corporation. ALL RIGHTS RESERVED. Approved for public release. Distribution unlimited 20-03275-4

[22] J. Rotner, J. Hodge, L. Danley, "AI Fails and How We Can Learn From Them," MITRE, July 2020 https://sites.mitre.org/aifails/wp-content/uploads/sites/15/2021/02/AI-Fails-and-How-We-Can-Learn-from-Them-MITRE-2020.pdf

[23] Partnership for AI "AI Incident Database" Nov. 2020-ongoing https://incidentdatabase.ai/

research evidence and empirical data. Therefore, development teams need members with diverse demographic and professional backgrounds.Examples of members of a well-rounded team include:

> **Example: Chatbot**
>
> A computer program that simulates human conversation, either via voice or text communication. Organizations use chatbots to engage with customers alongside the classic customer service channels of phone, email, and social media.

- **Data engineers** to ensure that data is usable and relevant
- **Model developers** to help the AI achieve the project's objectives
- **Strategic decision makers** who understand the technical aspects of AI as well as broader strategic issues or business needs
- **Domain specialists** to supply context about how people in their field actually behave, existing business practices, and any historical biases. Domain experts can be scientific or non-scientific; they may be military personnel, teachers, doctors and patients, artists … any people who are actual experts in the area for which the AI is being designed.
- **Qualitative experts or social scientists** to help technologists and decision makers clarify ideas, create metrics, and objectively examine factors that would affect adoption of the AI
- **Human factors or cognitive engineers** to help ensure that AI is not just integrated into a technology or process, but is adopted willingly and with appropriately calibrated trust
- **Accident analysis experts** who can draw on a long history of post-accident insights and frameworks to improve system design and anticipate areas of concern
- **Legal and policy experts** to oversee that data use and governance are covered by relevant authorities, to identify legal implications of the deployed AI, and to ensure that the process is following established mechanisms of oversight.
- **Privacy, civil liberties, and cybersecurity experts** to help evaluate and if necessary mitigate how design choices could affect concerns in their respective areas
- **The users of the AI and the communities** that will be affected by the AI to reinforce the importance of meeting the desired outcomes of all stakeholder
- **Educators and technical writers** to document and train the users on the system, as well as how to provide feedback to the system owners.

How do we involve all of these different stakeholders in the process of standing up the AI project? Do we need all of these participants in our project?

That decision depends on the type of AI project, the stakeholders and the ultimate users of the AI system or those impacted by it.

## 9.2 Integrating Tools and Techniques From User Experience

One way to improve AI projects is to adopt a human-centered approach which includes tools and techniques that support the user experience and consider human factor engineering as identified in the AI Fails document referenced above. One commonly used approach is the Double-Diamond model as shown in **Figure 11**.

The original Double-diamond approach was developed by the British Design Council in 2003 in an effort to "promote the positive impact of adopting a strategic approach to design and the value of 'design management' as a practice. " [24]

The Double Diamond is based on a series of Design Principles:

1. Be people centered
2. Communicate visually and inclusively
3. Collaborate and co-create
4. Iterate, iterate, iterate

[24] The Design Council "The Double Diamond: a universally accepted depiction of the design process" https://www.designcouncil.org.uk/news-opinion/double-diamond-universally-accepted-depiction-design-process

**Figure 11: Design Council Double Diamond**

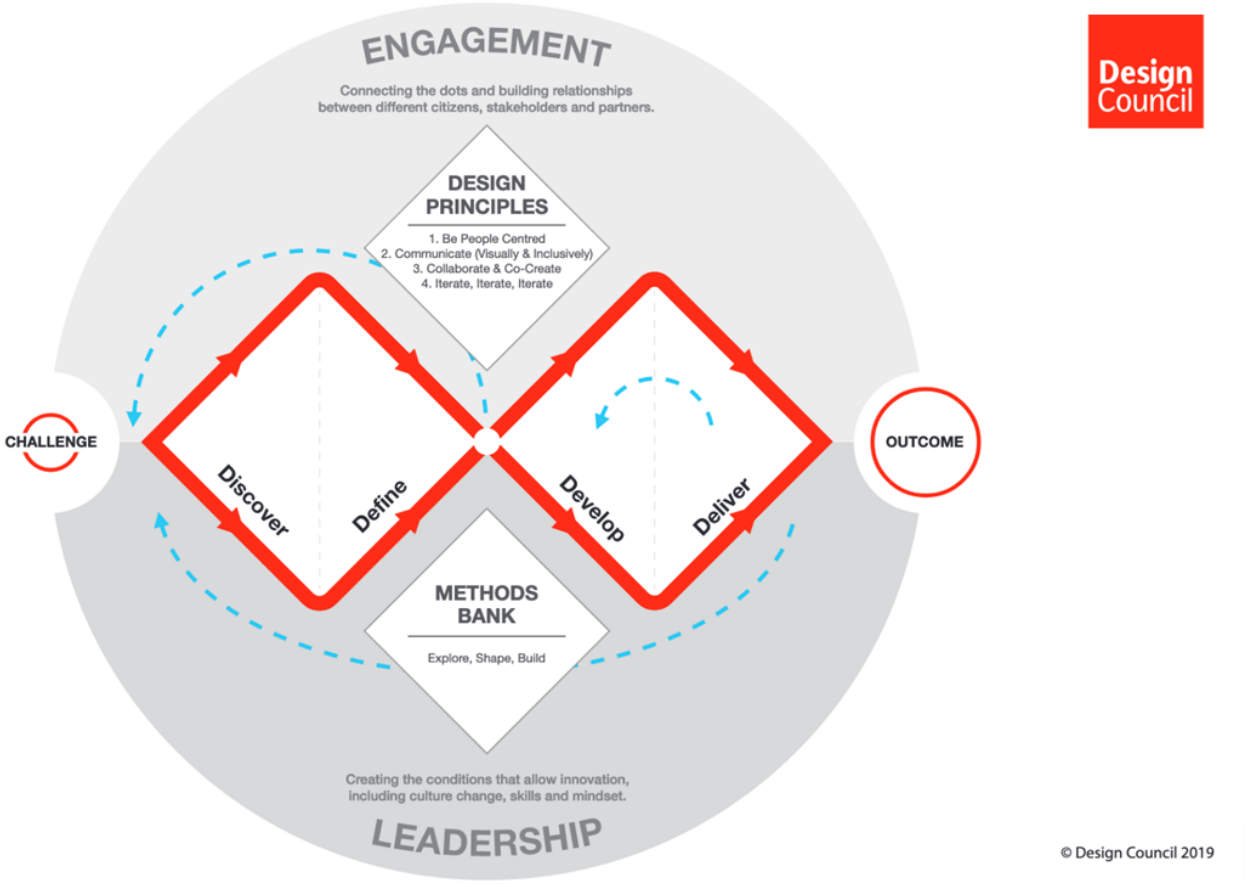


Since the inception of the Double-diamond, there have been a number of variations that add explanation and focus to the design process. One of these is **Figure 12**: Nessler's 2018 Double Diamond.

**Figure 12: Nessler's 2018 Revamped Double Diamond**

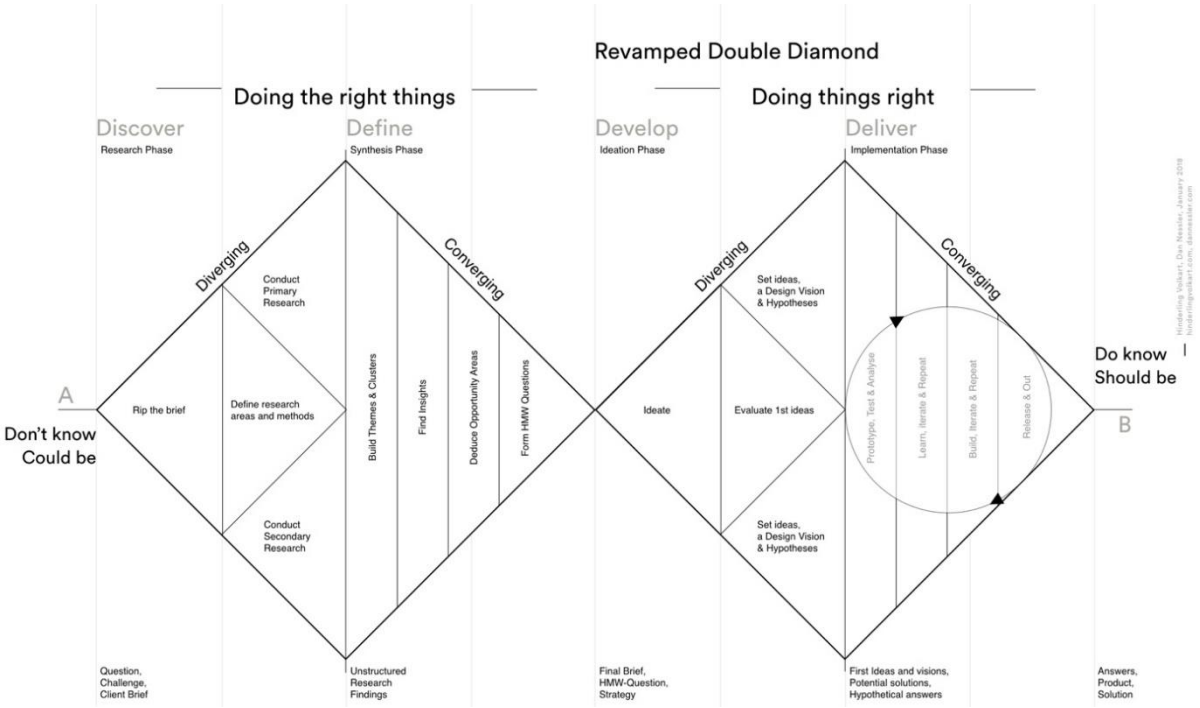


What Nessler's revamped Double-Diamond brings to the table is the concept of "Doing the Right Things" as a part of the discover process and *then* focusing on "Doing the Things Right." In other words, the discover and definition phases of any AI project must focus on defining the right problem, before considering any alternative solutions. Many of the AI Fails referenced above may have failed because the problem was not well-defined.

Regardless of version, the Double-Diamond model includes the concepts of diverging and converging, techniques often used in design processes. In the Double-Diamond model, this process occurs twice. During diverging, the team considers all of the possible options for research that could inform the problem at hand. As the reseach is sysnthesized in the definition phase, the team will converge on the real opportunities at hand.

The process repeats as the team begins to develop solutions; diverging during the ideation phase and converging for implemtation.

Nessler explains that his revamped Double-Diamond has been criticized for taking something very simple and making it complex. Nessler even agrees. What he goes on to explain is that this

model includes was a personal attempt at making sense of the various deisgn processes, tools and methodologies in the marketplace. [25]

By using the Revamped Double-Diamond as a guide, the User-Centered Engineer or Human Factors Engineer can employ the best tools and methods for the particular AI project.

The Double-Diamond processes are not linear and they are not intended to be. There is room for iteration and learning at many steps along the way.

## 9.3 Discovery and Definition

Each AI project begins with a discovery phase, focusing on the idea or need. Once the team or organization has created a 'brief' explaining what problem they are trying to solve, they can then define their research plan. One of the goals of the reseearch plan is to confirm that the team is solving the right problem.

First, the team needs to define who needs to be included in this initial research pass. As noted above, there are a variety of people who may be needed to represent different perspectives. The project team identifies who to include in the initial round of interviews and conversations.

These stakeholder interviews can be conducted by members of the core AI team. However, we recommend that at least one expert in human-centered design or user experience design, perhaps from an outside team specializing in these skills, be tasked with designing the questions and moderating the interviews. We will call this person the User-Centered Engineer, understanding that it could actually be multiple people or a team of people with these skills. Other members of the team participate by taking notes and *actively listening*.

For example, the MITRE Human-Machine Teaming (HMT) Systems Engineering Guide[26] recommends starting with subject-matter expert and stakeholder interviews. The interviews follow these guidelines:

- Paint a picture of the envisioned autonomy
- Top challenges
- Current state of automation and autonomy
- Critical decision method probes followed by an HMT audit

[25] Nessler, Dan "How to rethink the Design process, fail, reflect and iterate" (2016) https://uxdesign.cc/how-to-fuck-up-the-design-thinking-process-and-make-it-right-dc2cb7a00dca

[26] P. L. McDermott, D. C. O. Dominguez, D. N. Kasdaglis, M. H. Ryan, I. M. Trahan, and A. Nelson, "Human-Machine Teaming Systems Engineering Guide," MITRE, Dec. 2018. https://www.mitre.org/publications/technical-papers/human-machine-teaming-systems-engineering-guide

Once the initial rounds of interviews have been conducted, the project team will gather to analyze the findings from the interviews. Guided by their User-Centered Engineer, the team will then prepare for a Design Thinking Workshop. This workshop will:

- Bring together the multidisciplinary representatives
- Go through a Discover Phase:
    - Define the problem and create a problem statement (Consider Exercise 1 from Google's worksheet on User Needs + Defining Success[27]
    - Map the current process (if there is one)
    - Identify needs assumptions
    - Identify critical assumptions
    - Ask the questions from the HCAI (Human-Centered Artificial Intelligence) Framework[28]
        - What is best controlled by humans (augmentation)?
        - What is best controlled by automation?
        - This is not an either/or condition
        - Consider Exercise 2 from Google's worksheet on User Needs + Defining Success[29]
    - Map the process where the user, and the system are included—defining where the system will be augmenting or automating processes.
- Identify and Evaluate Data Sources
    - Map the existing data sources
    - Identify what data is most likely needed to solve the problem.

Once the team has analyzed the research, they converge on their approaches to the soluctions. This means synthesizing all that was learned from the stakeholders and the others impacted by the design of the AI and bring it all together.

[27] Google PAIR, "User Needs + Defining Success Chapter Worksheet" in People + AI Guidebook, May 2019. https://pair.withgoogle.com/worksheet/user-needs.pdf

[28] B. Shneiderman, Human-Centered AI, Oxford University Press, 2022

[29] Google PAIR, "User Needs + Defining Success Chapter Worksheet"

## 9.4 Implementation and Interface Design

Once the initial discovery workshop is completed, members of the project team will have an outline of all inputs for designing the AI system. The AI system comprises at least two major components:

- The front-end interface that allows the users to interact with the system
- The back-end system which houses the database, the algorithms that power the system and other supporting technology

Once the team has an idea of the problem, they can begin the solutioning process, moving into the second diamond of the Double Diamond model.

For the interface, the User-Centered Engineer can walk the stakeholder group, with at least two potential users of the system along with the rest of the user experience team, through a Design Studio Workshop. These workshops are typically two-to-three hours, with additional time if there are new folks to be brought up to speed on the decisions made during the initial workshop. (These should not be confused with Design Thinking workshops—Design studios are focused on creating the interface the users will see.) The User-Centered Engineer will then create several versions of the interface based on the outcomes of the workshop.

## 9.5 Testing with Users

Once designed, the User-Centered Engineer will test the interface with representative users. Usability testing can pinpoint issues early, especially when conducted early in the Develop phase of the double diamond through prototyping. Interface design is an iterative process where usability testing informs the project team of what needs to be changed or clarified. Usability testing ensures that the interface empowers users to perform critical tasks effectively and efficiently and can involve measuring task duration, error rate, and user frustration.[30]

The design team will then update the design to reflect the changes. Usability testing can be repeated as the fidelity of the prototypes approaches what the users will actually see.

## 9.6 Human-Centered Metrics

Beyond typical usability testing procedures, we recommend the project team take advantage of latest developments in human-machine teaming research to identify additional metrics for success. These metrics can act as a "north star" and drive discussions early in the process.

[30] E. Frøkjær, M. Hertzum, and K. Hornbæk, "Measuring usability: are effectiveness, efficiency, and satisfaction really correlated?," in Proceedings of the SIGCHI conference on Human Factors in Computing Systems, New York, NY, USA, Apr. 2000, pp. 345–352. doi: 10.1145/332040.332455

Shneiderman (2021) recommends thirty attributes of human-centered AI categorized into five themes as laid out in **Figure 13**.

**Figure 13: Human-Centered AI Attributes**

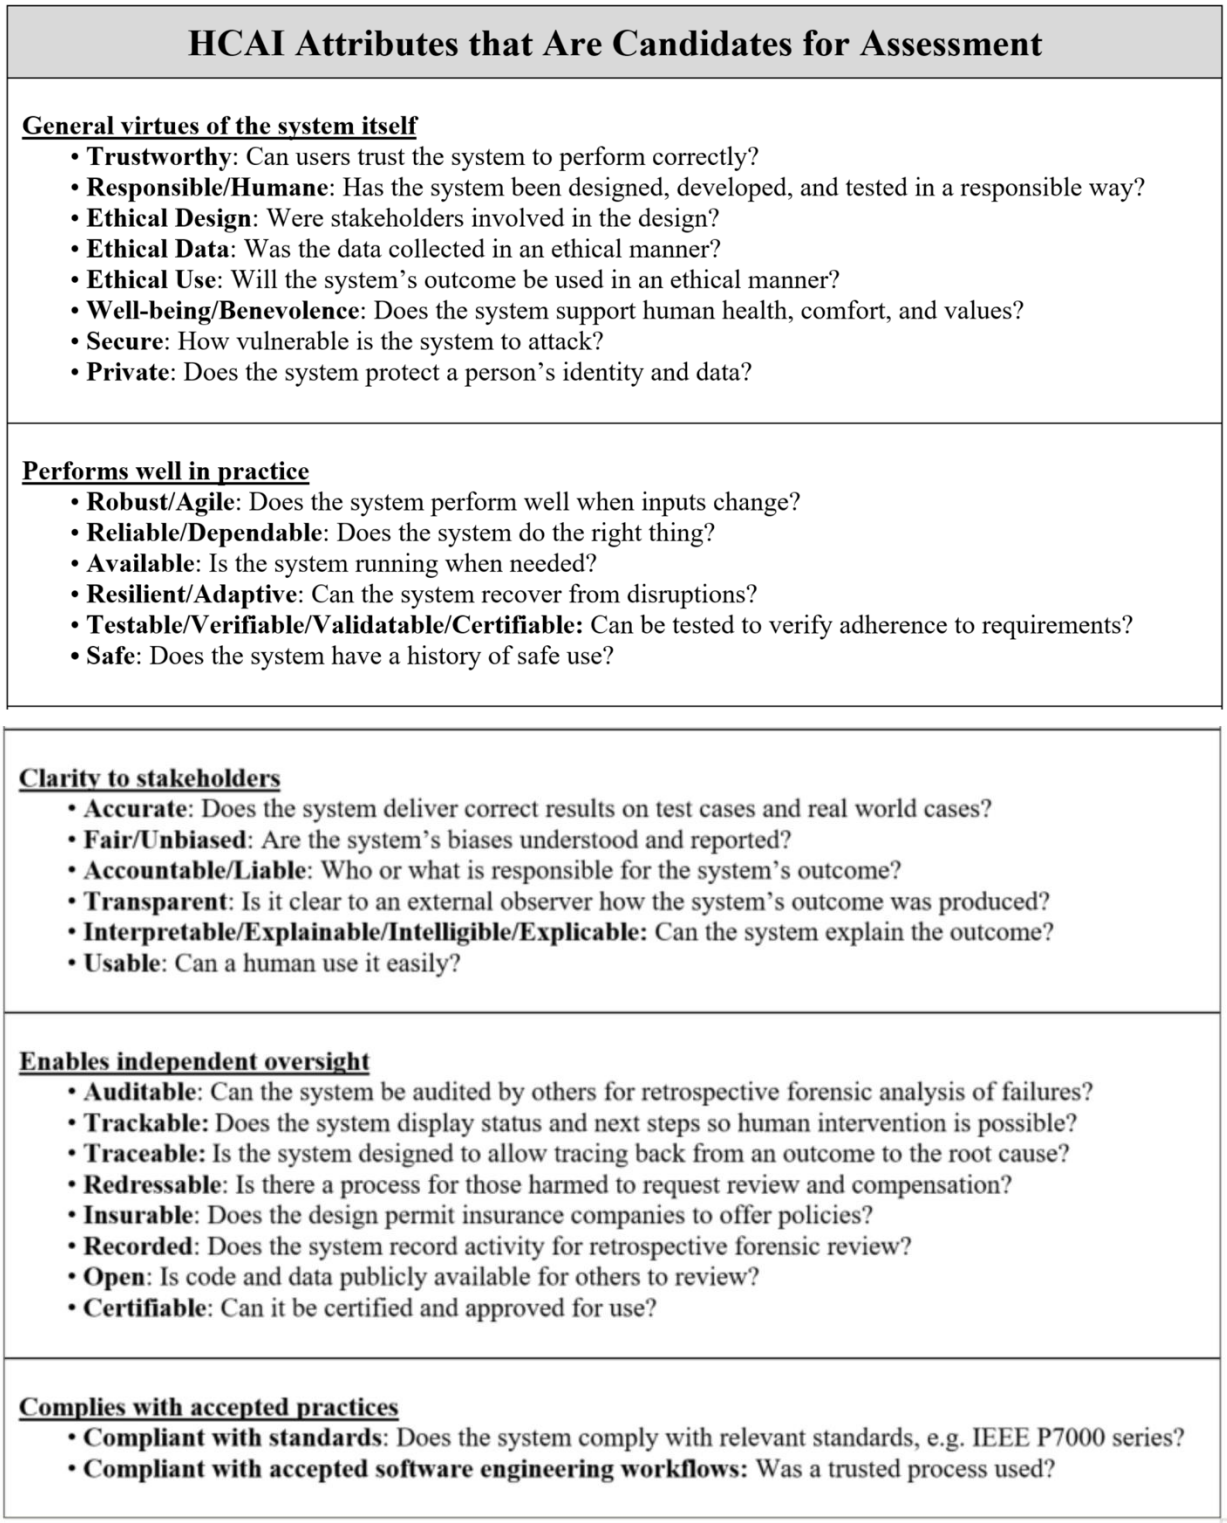

**HCAI Attributes that Are Candidates for Assessment**

**General virtues of the system itself**

- **Trustworthy**: Can users trust the system to perform correctly?
- **Responsible/Humane**: Has the system been designed, developed, and tested in a responsible way?
- **Ethical Design**: Were stakeholders involved in the design?
- **Ethical Data**: Was the data collected in an ethical manner?
- **Ethical Use**: Will the system's outcome be used in an ethical manner?
- **Well-being/Benevolence**: Does the system support human health, comfort, and values?
- **Secure**: How vulnerable is the system to attack?
- **Private**: Does the system protect a person's identity and data?

**Performs well in practice**

- **Robust/Agile**: Does the system perform well when inputs change?
- **Reliable/Dependable**: Does the system do the right thing?
- **Available**: Is the system running when needed?
- **Resilient/Adaptive**: Can the system recover from disruptions?
- **Testable/Verifiable/Validatable/Certifiable:** Can be tested to verify adherence to requirements?
- **Safe**: Does the system have a history of safe use?

**Clarity to stakeholders**

- **Accurate**: Does the system deliver correct results on test cases and real world cases?
- **Fair/Unbiased**: Are the system's biases understood and reported?
- **Accountable/Liable**: Who or what is responsible for the system's outcome?
- **Transparent**: Is it clear to an external observer how the system's outcome was produced?
- **Interpretable/Explainable/Intelligible/Explicable:** Can the system explain the outcome?
- **Usable**: Can a human use it easily?

**Enables independent oversight**

- **Auditable**: Can the system be audited by others for retrospective forensic analysis of failures?
- **Trackable:** Does the system display status and next steps so human intervention is possible?
- **Traceable:** Is the system designed to allow tracing back from an outcome to the root cause?
- **Redressable**: Is there a process for those harmed to request review and compensation?
- **Insurable**: Does the design permit insurance companies to offer policies?
- **Recorded**: Does the system record activity for retrospective forensic review?
- **Open**: Is code and data publicly available for others to review?
- **Certifiable**: Can it be certified and approved for use?

**Complies with accepted practices**

- **Compliant with standards**: Does the system comply with relevant standards, e.g. IEEE P7000 series?
- **Compliant with accepted software engineering workflows:** Was a trusted process used?

As the designers and builders of the system communicate with the stakeholders, they need to include the ongoing HCAI attributes. By addressing these considerations upfront, the stakeholders that may be resistent to change can be assured that the team is following an approach that builds trust. By demonstrating the AI system as it is built and including these considerations in the demonstrations, stakeholders can be confident that the team is covering all the bases in buidling a trustworthy AI system.

For further consideration, MITRE's HMT Systems Engineering Guide identifies the following HMT leverage points that can inform success (see the link below for definitions):[31]

- Observability
- Predictability
- Directing Attention
- Exploring the Solution Space
- Adaptability
- Directability
- Calibrated Trust

> **Example: Computer Vision**
>
> Models designed to translate visual data based on features and contextual information identified during training. This enables models to interpret images and video and apply those interpretations to predictive or decision making tasks.

Tools are available for evaluating some of these human-centered metrics. MITRE's Calibrated Trust Evaluation Toolkit (https://comm.mitre.org/calibrated-trust-toolkit/), for example, helps ensure that users' expectations match the system's actual capabilities. Value cards can be used to estimate whether models would be accepted or preferred by various stakeholders (Shen et al. 2021). There are also a variety of tools for exploring the fairness of machine learning models.

## 9.7 Adaptation Over Time

AI systems should be continuously assessed, evaluating consistency of the workflow and how well humans and systems are working together. Appropriate channels will be in place for users to easily give feedback on system performance and experience. The feedback systems have to be easy-to-use and easily available to all of the users of the system. This increases involvement as well as trust in the working AI systems. There may be a call for modifications as time goes on, based on changing circumstances, technologies, or settings. It's particularly important to check and assess adaptive or teachable systems to ensure that learning continues to function as intended. These systems can be continuously usable, dependable, safe and trustworthy; free of prejudice and ethically proper.

[31] https://www.mitre.org/sites/default/files/publications/pr-17-4208-human-machine-teaming-systems-engineering-guide.pdf, pp. 1-3.

# 10 AI Systems and Organizational Change Management

Organizational Change Management refers to changes in an organization's structures, systems, strategy, or culture. These changes are often implemented as part of an overall plan to achieve strategic, operational, or financial goals.

## 10.1 Challenges of AI Systems Driving Organizational Change Management

The most successful organizations proactively manage organizational changes, using strategy to drive all decisions and a pre-defined implementation plan that encompasses the stakeholder analysis, training, and project management needed to be successful. Implementing AI systems brings a unique set of challenges to any organization.

- Clearly understanding the present and near-term capabilities, constraints, and limitations of AI systems is the first step. Unrealistic expectations for AI system "magic" dooms many projects before the first data set is identified or model designed.
- Access to sufficient data, the ability to access many kinds of data in many different environments, operating systems, and databases
- The ability and effort needed to curate - clean, standardize, de-duplicate, and prepare large data sets
- Defining appropriate performance targets often requires reconciliation among competing views of what is important to the organization
- Discomfort with the application of AI systems to what has historically been knowledge work, especially as roles incorporate more decision-making
- Fear of staff displacement

While standard organizational change management approaches – e.g., keeping employees informed, involved, and aligned with the crucial business need for the change from the start – are undoubtedly beneficial, these challenges demand additional consideration and mitigation.

## 10.2 AI-Centric Change Management Approach

To define a change management strategy, an organization defines the goal end state, who is responsible for implementation, how implementation will occur, and how success will be determined. Organizational change driven by implementing an AI system is no different, but the end state may be less familiar to the management and staff. Organizational structures may be significantly reduced in number as AI systems and RPA bots accomplish administrative work. Existing staff can acquire new skills and be deployed in new positions, taking on more policy-oriented and strategic roles, or focusing on tasks that were previously understaffed. Clearly defining the new structure, roles, and responsibilities, while taking full advantage of the increased data flow and quality is essential to deriving all the available value from these new technologies. AI-centric organizational change management is likely to require even more time

and effort in this definition of the end state, in communication with staff, and in training staff for new roles and responsibilities.

The high-level approach for AI-centric organizational change management is to build on initial successes, advantage AI system deployment owners to expand their footprint, and create a general awareness and understanding of AI systems' success across the enterprise. Awareness expansion will be driven by governance, communication, and training. The following steps are typical for organization change management, although every implementation will require solutions tailored to a specific focus:

1. Clearly define the change and align it to business goals
    - Understand the 'Why'
    - Understand the 'How' and assess feasibility and preparedness
2. Create a roadmap to understand the current and future goal state after implementation of the change
3. Identify the leadership and implementation teams to begin assigning roles and responsibilities
4. Identify the impacts and all individuals to be affected by implementation
5. Develop a communication structure and plans for training and onboarding for the desired changes
6. Move forward with implementation – ensure consistent communication with the implementation team throughout the process to work through roadblocks as they emerge
7. Set up a structure to measure success, identify challenges, and assess the change management process overall. Utilize this structure during and after implementation to measure performance and identify key areas for improvement or future growth

## 10.3 Change Management Governance

To ensure successful organizational change, define clear leadership roles and decision-making protocols within the team responsible for implementing the change. If the groups differ, leadership develops a steady cadence of communication between the leadership team and the implementation team. Since AI systems' implementation will affect data and personal information storage, ensure that the appropriate processes for data security are in place and being maintained throughout the process. See the discussion of AI system development and data management methodologies above.

The governance function can be supported with proper documentation of processes change status, system training and testing status, compliance with all relevant procurement, development, and deployment policies, and archive steps for future reference and understanding. Key performance indicators (KPI) are used to evaluate the success of a transition and to adjust as needed throughout the change management process, which is

regarded a best practice. Examples of these metrics include employee readiness assessment results, employee engagement and participation measures, employee feedback, employee satisfaction results, training effectiveness measures, compliance reports, and qualitative observations of behavior change.

## 10.4 Communication, Socialization of Implementation

Particularly for users whose actions will be impacted directly by a change, communication of an upcoming change begins early and be reiterated on a regular basis until implementation; the level of frequency likely depends on the magnitude of effort that may be required from the end user to adapt. Communications around a change includes rhetoric that eases any concerns of apprehensions that exists in the general population of the organization around the new modification. Communications can be in-person meetings, webinars, email blasts, or newsletters. In-person meetings, live calls, or webinars can be a great way to allow for questions and open conversation about changes.

Direct managers can educate and update their teams on upcoming changes and to hold them accountable for any necessary action items. Having managers as a liaison and an intermediary for resources and questions will make the process less daunting for impacted employees. Getting direct managers on board early is crucial to this step.

## 10.5 Learning and Training

Providing end users with ample resources to aid in change adoption can be critical; examples of these resources are instructions, guidebooks, information sessions, and trainings. Over time, measure the effectiveness and the utilization of the implemented change. Leadership can continue communications if there are potential improvements to be made post-implementation, to continue to encourage adoption, and as a resource for questions, feedback, or concerns. An anonymous feedback survey can be an effective way to get honest feedback on the specific change and the management throughout the process overall.

Use appropriate KPIs to determine the success of the transition and identify any shortcomings in this process to allow for future improvements. These KPIs include employee satisfaction and feedback surveys, as well as internal help desk metrics to understand where there were difficulties, miscommunications, and shortcomings, with an implemented change, if applicable.

## 10.6 Project Management

In many larger projects, change management efforts are planned and tracked against the larger project schedule. Because some AI and DA solutions do not have a project schedule, the change management efforts will need to adapt to the evolving needs of the AI and data analytics community.

The efforts of the Change Agents, who develop a plan to engage with all organizations across DHA within a given timeframe might be an exception. The Change Agents then plan activities and track their progress against that plan.

## 10.7 Performance Management

It is possible to track and measure the efficacy of change management efforts when the goals of the change effort are clearly defined. Through end-of-deployment surveys, teams can track the effectiveness of the Change Agents, communications, and change leadership efforts. Results from these surveys provide valuable insight into what changes might be necessary to the change management efforts to have a bigger impact on AI and data analytics adoption and deployment.

These additional metrics will enable the team to understand how rapidly the community is growing, how mature it is in relation to how mature it should be, and the skills and abilities across the community.

# 11 Governance Policy

Governance refers to the function of providing oversight of operational management, to ensure that an organization (i.e., people, processes, and technology) is being managed effectively to achieve strategic goals while complying with relevant laws, regulations, and policies. Data governance is more narrowly focused on a collection of processes, roles, policies, standards, and metrics that ensure the effective and efficient use of information in enabling an organization to achieve its goals. Governance policy frameworks structure and delineate power and the governing or management roles in an organization.

IT Governance is an integral part of enterprise governance and consists of the leadership and organizational structures and processes that ensure that an organization's IT sustains and extends the organization's strategies and objectives. Ensuring related policies and management decisions are a part of this governance process provides for a high-level, organization wide strategic approach to IT and data analytics management. It also provides a structured decision-making process around investment decisions and promotes accountability, due diligence, and efficient and economic delivery of enterprise IT services.

- **Implementing Governance** – any project, investment or program that receives IT governmental funding can be aligned to one of the 19 Governance Boards. Each of these Governance Boards (GB) are assigned to one of four Executive Steering Committees (ESC).
- **Establishing a Board of Governance (BoG)** – assignees are designated by members of ESC.

**Figure 14** provides examples of governance structures that might be used within an organization at various levels and to support different functions (i.e., IT, Strategic Development, Operations, etc.). Ultimately each governance structure will be tailored to the organization's unique needs and leadership structure.

## Figure 14: Governance Structures

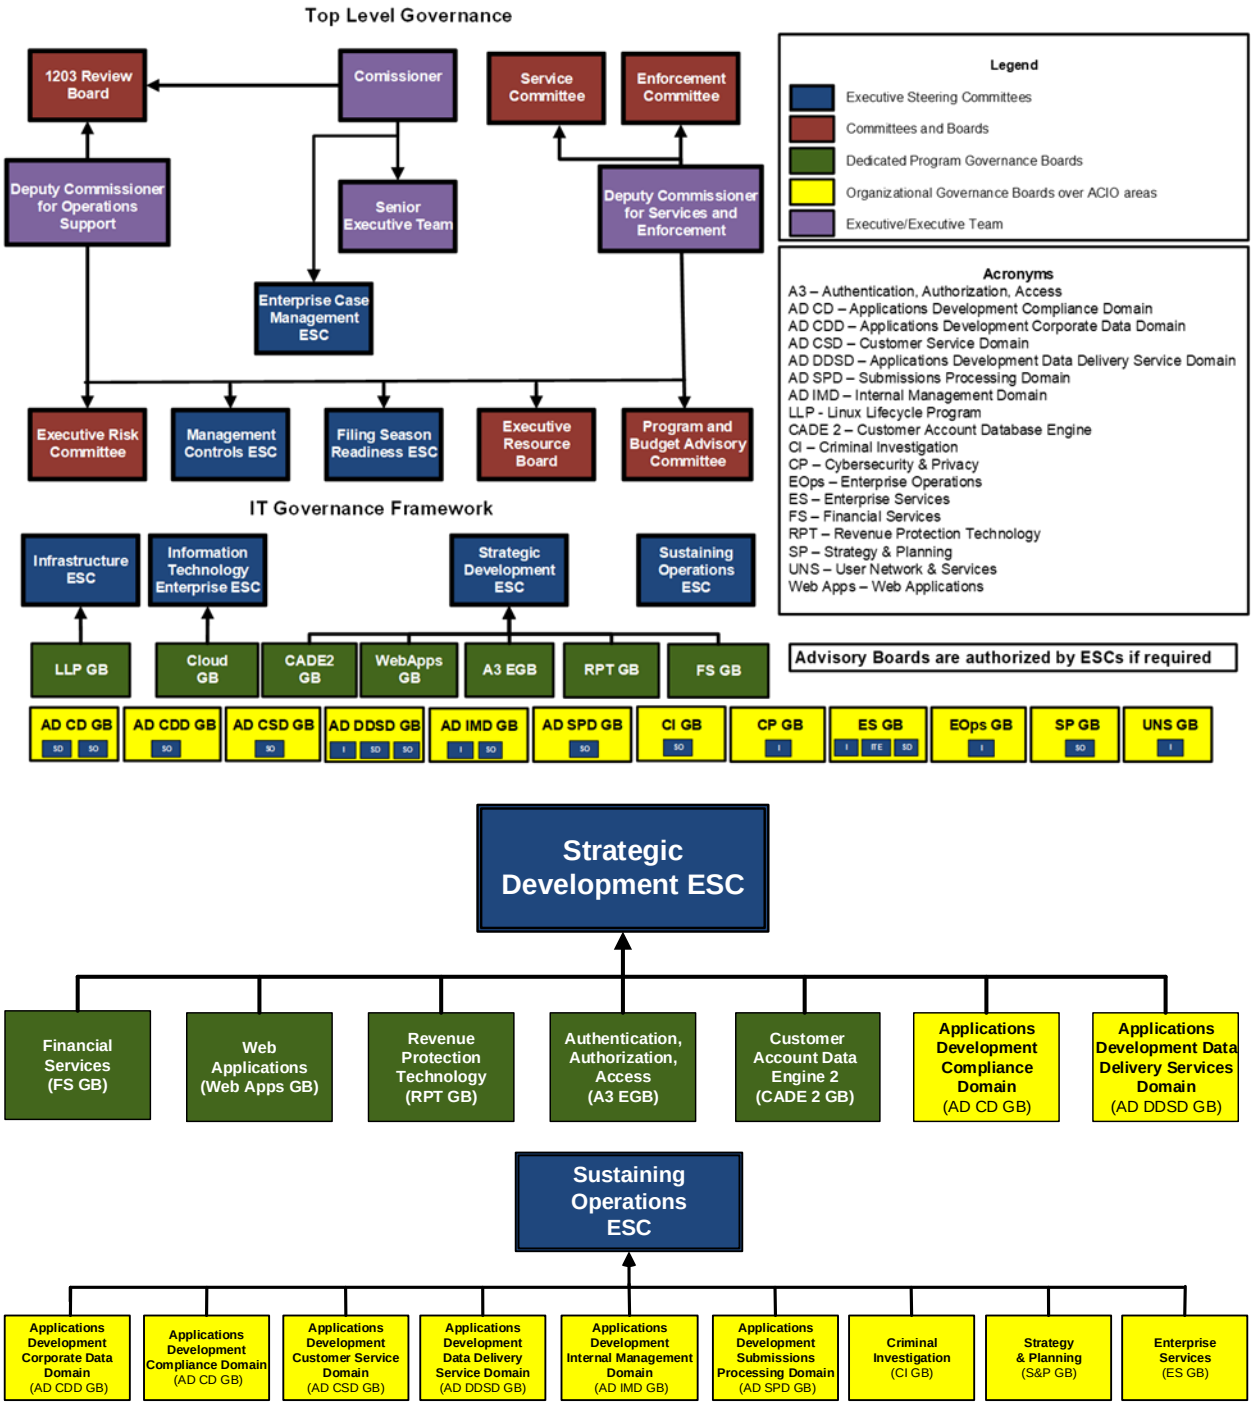

# 12 Summary

Technology is fundamentally transforming how government interacts with the public. Integration of the strategies discussed in this guidebook promote a Federal Government that is more efficient, effective, and better equipped to deliver services to the American people. Exponential advances in computing power, the rise of novel information networks, and unleashed innovation have created new platforms that are enabling the development of a 21st century digital government.

This guidebook provides tools, tips, and strategies to link data in the information layer, the platform layer, and the presentation layer using secure interoperable cloud-based platforms. It also provides a comprehensive framework for implementing shared analytics, machine learning, AI, and other emerging technologies. Finally, the AIDA Guidebook outlines when and how to leverage shared analytics throughout government agencies and other organizations using repeatable and resilient models, core analytic terms of reference, and associated definitions while promoting data standardization, optimization, and innovation – supporting the implementation of the Federal Government AI Strategy.

# 13 Next Steps

In order to ensure usability it will be necessary to link this framework to federal IT laws, regulations, policies, and procedures underpinned by Federal Information Security Modernization Act (FISMA), Federal Acquisition Reform Act (FARA), Information Technology Management Reform Act (ITMRA), Paperwork Reduction Act (PRA), Federal Financial Management Improvement Act (FFMIA), Federal Managers Financial Integrity Act (FMFIA), and Government Performance and Results Act Modernization Act (GPRA-MA). Furthermore, this framework needs to expand upon best practices to ensure data protection in transit and at rest, and best practices to reduce administrative burdens associated with regulatory compliance such as the Authorization to Operate (ATO) process. Additional activities include building best practices and exemplar use cases for data management frameworks, operational security, protecting intellectual property, building non-biased training datasets, and ensuring there are mechanisms to ensure that data being used is 'fit for purpose'.

This guidebook is intended to become a living document, updated as new information and guidance becomes available. It is also recommended that ATARC and its partners continue to identify projects, use cases, and lessons learned for community reference and identify potential venues to share those findings in partnership with academia.

# Appendix A: Representative Laws and Federal Policies

**The Federal Trade Commission Act, 1914**

The Federal Trade Commission (FTC) is one of the few agencies solely focused on consumer data. Their mission is to protect consumers by enforcing actions to safeguard consumers against unfair or deceptive practices and to enforce federal privacy and data protection regulations. Deceptive practices include a company or organization's failure to comply with its published privacy policies and its failure to provide adequate security of personal identifiable information (PII) or use of deceptive advertising or marketing methods. Many other federal laws around consumer data have been slotted under the FTC.

**Title 26, 1939**

Title 26 pertains to the statistical work carried out by the US Census Bureau for the collection of Internal Revenue Service (IRS) data regarding households and businesses. It states the conditions in which the IRS may communicate Federal Tax Returns and Return Information (FTI) with other agencies, including the Census Bureau. Title 26 specifically allows the IRS to send FTI to the Census Bureau for the purpose of building censuses and national economic accounts in addition to conducting other federally authorized statistical tasks.

**Title 13 - Covers US Census Bureau Data, 1954**

Both individuals and businesses are covered by Title 13. These protections include never disclosing or publishing personal information such as names, addresses, phone numbers, or Social Security Numbers, and the data collected by the Census Bureau cannot be used against respondents in a court of law. Census Bureau workers are also held to a higher standard of confidentiality; they must safeguard respondents' information for life and, if it is broken, face the following penalties: a federal prison sentence up to five years and a $250,000 fine.

**The Fair Credit Reporting Act (FCRA), 1970**

The Fair Credit Reporting Act (FCRA) protects consumers from consumer reporting agencies. Amended by the Fair and Accurate Credit Transactions Act (FACTA), it further restricts the use of information with a bearing on an individual's credit worthiness, credit standing, credit capacity, character, general reputation, personal characteristics, and mode of living to determine eligibility for credit, employment, and insurance among other restrictions on how credit card information can be viewed and seen. The Consumer Financial Protection Bureau and FTC provide additional authority to operate.

**Family Educational Rights and Privacy Act (FERPA), 1974**

The Family Educational Rights and Privacy Act (FERPA) is a federal law, enforced by the Department of Education (ED), that protects the privacy of student education records and applies to all schools which receive funds under an applicable program under ED. FERPA gives parents certain rights pertaining to their children's education records and the rights are transferred to the child at the legal age of 18. FERPA prohibits improper disclosure of PII derived from education records. Violations of FERPA result in a withdrawal of federal funding. ED has changed their enforcement tactics of FERPA with now just focusing on the highest-risk issues for investigations.

**Privacy Act of 1974**

This Act establishes a code of fair information practices that governs the collection, maintenance, use, and dissemination of information about individuals maintained in systems of records by federal agencies. Agencies must give public notice of their systems of records by publication in the Federal Register. US citizens are given the right to examine and edit their records and are protected against unwarranted invasion of their privacy resulting from the collection, maintenance, use, and disclosure of their personal information.

**42 CFR Part 2 regulations, pertaining to the Confidentiality of Substance Use Disorder Patient Records, 1975**

Similar to the protection of certain patient information under HIPAA and its implementing regulations, the confidentiality of alcohol and drug abuse patient records is protected by federal law under 42 U.S.C. § 290dd-2 and its implementing regulations under 42 C.F.R. Part 2. Specifically, Part 2 protects the confidentiality of patient records maintained in connection with the provision of substance abuse education, prevention, rehabilitation, treatment, training, or research by, or as part of, a federally assisted program. Part 2 is enforced by the federal Substance Abuse and Mental Health Services Administration (SAMHSA). 42 C.F.R. Part 2 was most recently updated in 2020.

**Federal Managers Financial Integrity Act of 1982 (FMFIA) (Public Law 97-255)**

The purpose of the Federal Financial Integrity Act of 1982 (FMFIA) is to update the Accounting and Auditing Act of 1950 to require Federal agencies to create internal accounting and administrative controls. These controls are created to prevent the waste or misuse of both agency funds and property as well as confirm the accountability of assets.

**Cable Communications Policy Act (CCPA), 1984**

The Cable Communications Policy Act (CCPA) is an amendment to the original Communications Act of 1934, which aligned regulations of telephone, telegraph, and radio communications

under the Federal Communications Commission (FCC). This Act gave the FCC jurisdiction and authority over the cable television industry and extended the protection of subscriber privacy. The FCC has grown into a large independent government agency that regulates all interstate communications.

**Chief Financial Officers (CFO) Act of 1990 (Public Law 101-576)**

The Chief Financial Officers (CFO) Act passed with the intention to improve the general and financial management practices of the Federal Government by outlining standards for financial performance and disclosure. The OMB was also given an increased role of management over federal financial management in addition to all twenty-four departments and agencies given a new position of chief financial officer.

**Paperwork Reduction Act (PRA) of 1995 (Public Law 104-13)**

The Paperwork Reduction Act (PRA) was enacted with the goal of reducing the paperwork load for individuals, small businesses, education and nonprofit entities, federal contractors, state, local, and tribal governments, and all other persons impacted from the collection of information for or by the Federal Government. It requires that every federal agency receive approval from the OMB before using identical questions to collect information from ten or more people.

**The Health Information Portability and Accountability Act (HIPAA), 1996 (Public Law 104-191)**

The Health Information Portability and Accountability Act (HIPAA), enforced by the Department of Health and Human Services (HHS), is a federal law that protects sensitive patient health information from being disclosed without the consent or knowledge of a patient through national standards. HIPAA included Administrative Simplification provisions that required HHS to adopt national standards for electronic health care transactions and code sets, unique health identifiers, and security. At the same time, Congress recognized that advances in electronic technology could erode the privacy of health information. Consequently, Congress incorporated into HIPAA provisions that mandated the adoption of federal privacy protections for individually identifiable health information. HHS published a final Privacy Rule in December 2000, which was later modified in August 2002. This Rule set national standards for the protection of individually identifiable health information by three types of covered entities: health plans, health care clearinghouses, and health care providers who conduct the standard health care transactions electronically. HHS published a final Security Rule in February 2003. This Rule sets national standards for protecting the confidentiality, integrity, and availability of electronic protected health information. In November 2019, HHS updated its regulations to reflect required annual inflation-related increases to civil monetary penalties, including those violations of HIPAA's "administrative simplification" provisions. Administrative simplification generally includes HIPAA's privacy and security requirements, including rules as to how health

plan data are exchanged, and the affected penalties are included in the Code of Federal Regulations at 45 C.F.R. § 160.404(b).

**Federal Financial Management Improvement Act (FFMIA), 1996**

The purpose of the Federal Financial Management Improvement Act of 1996 (FFMIA) is to advance Federal financial management by ensuring that Federal financial management systems provide accurate, reliable, and timely financial management information to the government's managers. The intent and the requirements of this Act go well beyond the directives of the CFO Act and the Government Management Reform Act of 1994 (GMRA) to publish audited financial reports. Compliance with the FFMIA will provide the basis for the continuing use of reliable financial management information by program managers, and by the President, the Congress and the public.

**Clinger-Cohen Act of 1996**

Divisions D - Federal Acquisition Reform Act (FARA) and Division E – Information Technology Management Reform Act (ITMRA) of the National Defense Authorization Act of 1996 are collectively referred to as the Clinger-Cohen Act. The Clinger-Cohen Act eliminates the General Services Administration's (GSA) single authority to acquire technology and permitted individual federal agencies to accept that role.

**Federal Financial Management Improvement Act of 1996 (FFMIA) (Public Law 104-208)**

The purpose of the Federal Financial Management Improvement Act of 1996 (FFMIA) is to improve Federal financial management by certifying that federal financial management systems provide correct, reliable, and prompt financial management information to government managers.

**Children's Online Privacy Protection Act (COPPA), 1998**

The Children's Online Privacy Protection Act (COPPA) prohibits the collection of any information from a child under the age of 13 online and from digitally connected devices. The Act requires publications of privacy notices and collection of verifiable parental consent when information from children is being collected. COPPA has not been updated since 2013 and continuously is competing against the growing rate of technology. COPPA potentially has a much larger role to play as digital content develops for younger audiences.

**Gramm Leach Bliley Act (GLBA), 1999**

The Gramm Leach Bliley Act (GLBA) governs the protection of personal information in the hands of banks, insurance companies, and more within the finance industry. This statute specifically addresses Nonpublic Personal Information (NPI), which includes any information that a financial service company collects from its customers in connection with the provision of its services.

Organizations not only have to protect and safeguard personally identifiable information (PII) and NPI, but also must explain their information-sharing practices to their customers. The FTC is charged with enforcement of GLBA.

**Government Paperwork Elimination Act (GPEA) (Public Law 105-277), 1999**

The Government Paperwork Elimination Act (GPEA) requires that federal agencies publish electronic forms, electronic filing, and electronic signatures when conducting official business with the public by 2003.

**Homeland Security Presidential Directive (HSPD-12), Policy for a Common Identification, 2004**

The Homeland Security Presidential Directive (HSPD-12) implemented the standardization of the badging process to improve security, decrease identity fraud, and protect privacy of those with issued government identification.

**Government Performance and Results Act Modernization Act (GPRA-M.A.), 2010**

The GPRA Modernization Act modernizes the Federal Government's performance management framework, retaining and amplifying some aspects of the Government Performance and Results Act of 1993 while also addressing some of its weaknesses. GPRA 1993 established strategic planning, performance planning, and performance reporting as a framework for agencies to communicate progress in achieving their missions. The GPRA Modernization Act establishes some important changes to existing requirements. It was intended to systematically hold Federal agencies accountable for achieving program results; improve program performance by requiring agencies to set goals, measure performance against those goals and report publicly on progress; improve Federal program effectiveness and public accountability by promoting a focus on results, service quality and customer satisfaction; help federal managers improve service delivery, by requiring that they plan for meeting program goals and by providing them with information about program results and service quality; and improve congressional decision-making by providing more information on achieving statutory objectives and on the relative effectiveness and efficiency of Federal programs and spending.

**32 Code of Federal Regulations 2002, Controlled Unclassified Information (CUI), 2010**

Executive Order 13556 established CUI on November 4, 2010 and Part 2002 of 32 Code of Federal Regulations prescribed Government-wide implementation standards on September 14, 2016. This order acknowledged that certain types of UNCLASSIFIED information are extremely sensitive, valuable to the United States, sought after by strategic competitors and adversaries, and often have legal safeguarding requirements. Unlike classified national security information, DOD personnel at all levels of responsibility and across all mission areas receive, handle, create, and disseminate CUI.

**Federal Information Security Modernization Act (FISMA), Public Law 113-283, 2014**

The Federal Information Security Modernization Act (FISMA) gives the Department of Homeland Security (DHS) authority to oversee the implementation of information security mandates for non-national security, specifically within the federal Executive Branch systems. The OMB is also given the authority to supervise federal agency information security practices along with the OMB being required to amend OMB A-130 to "eliminate inefficient and wasteful reporting."

**OMB M-17-12, 2017**

This Office of Management and Budget (OMB) Memorandum sets federal policy for agencies to prepare for and respond to a potential breach of PII to include a framework for assessing and mitigating the risk of harm, as well as guidance to notify individuals whose PII was exposed. Additionally, this Memorandum provides agencies flexibility to tailor their specific responses to data breaches based on facts and circumstances per each breach and to gauge the levels of harm to affected individuals.

**National Archives and Records Administration (NARA) 1608, 2017**

This policy is part of OMB Circular A-130 "Managing Information as a Strategic Resource" to provide guidance to NARA staff for the protection of PII from unauthorized disclosure as required under the Freedom of Information and Privacy Acts. NARA 1608 emphasized the role of NARA users in ensuring appropriate physical and technical safeguards are put in place to protect all NARA systems, both text and electronic, which may contain PII.

**California Consumer Privacy Act (CCPA) 2018**

The California Consumer Privacy Act of 2018 (CCPA) gives consumers more control over the personal information that businesses collect about them and the CCPA regulations provide guidance on how to implement the law. This landmark law secures new privacy rights for California consumers, including:

- The right to know about the personal information a business collects about them and how it is used and shared;
- The right to delete personal information collected from them (with some exceptions);
- The right to opt-out of the sale of their personal information; and
- The right to non-discrimination for exercising their CCPA rights.

Businesses are required to give consumers certain notices explaining their privacy practices. The CCPA applies to many businesses, including data brokers.

**General Data Protection Regulation (GDPR) 2016**

The General Data Protection Regulation (GDPR) is the toughest privacy and security law in the world. Though it was drafted and passed by the European Union (EU), it imposes obligations onto organizations anywhere, so long as they target or collect data related to people in the EU. The regulation was put into effect on May 25, 2018. The GDPR will levy harsh fines against those who violate its privacy and security standards, with penalties reaching into the tens of millions of euros. With the GDPR, Europe is signaling its firm stance on data privacy and security at a time when more people are entrusting their personal data with cloud services and breaches are a daily occurrence. The regulation itself is large, far-reaching, and fairly light on specifics, making GDPR compliance a daunting prospect, particularly for small and medium-sized enterprises (SMEs).

# Appendix B: Acronym List

| | |
|---|---|
| **AIDA** | Artificial Intelligence and Data Analytics |
| **AI** | Artificial Intelligence |
| **AI-IA** | National Artificial Intelligence Initiative Act |
| **AMARC** | Advanced Mobility Academic Research Center |
| **ATARC** | Advanced Technology Academic Research Center |
| **BoG** | Authorization to Operate (ATO) Board of Governance |
| **CCPA** | Cable Communications Policy Act |
| **CFO** | Chief Financial Officer |
| **CIO** | Chief Information Officer |
| **COPPA** | Children's Online Privacy Protect Act |
| **CUI** | Controlled Unclassified Information |
| **DHA** | Defense Health Agency |
| **ETL** | Extract, Transform, and Load |
| **FARA** | Federal Acquisition Reform Act |
| **FCRA** | Fair Credit Reporting Act |
| **FDA** | Food and Drug Administration |
| **FERPA** | Family Educational Rights and Privacy Act |
| **FFMIA** | Federal Financial Management Improvement Act of 1996 |
| **FISMA** | Federal Information Security Modernization Act |
| **FMFIA** | Federal Managers Financial Integrity Act of 1982 |
| **FTC** | Federal Trade Commission |
| **FTPS** | File Transfer Protocol Secure |
| **GB** | Governance Board |
| **GLBA** | Gramm Leach Bliley Act |
| **GITEC** | Government Information Technology Executive Council |

| | |
|---|---|
| **GPEA** | Government Paperwork Elimination Act |
| **GSA** | General Services Administration |
| **HIPAA** | Health Information Portability and Accountability Act |
| **HMT** | Human-Machine Teaming |
| **HSPD-12** | Homeland Security Presidential Directive |
| **HTTPS** | Hypertext Transfer Protocol Secure |
| **ICAM** | Identity, Credential, and Access Management |
| **IPO** | Interagency Program Office |
| **IRS** | Internal Revenue Service |
| **ITMRA** | Information Technology Management Reform Act |
| **JAIC** | Joint Artificial Intelligence Center |
| **KPI** | Key Performance Indicator |
| **NASA** | National Aeronautics and Space Administration |
| **NIH** | National Institutes of Health |
| **NIST** | National Institute of Standards and Technology |
| **NNSA** | National Nuclear Security Administration |
| **OMB** | Office of Management and Budget |
| **PII** | Personal identifiable information |
| **PRA** | Paperwork Reduction Act |
| **SME** | Subject matter expert |
| **TLS** | Transport Layer Security |
| **TQD** | Training- quality dataset |
| **USDA** | United States Department of Agriculture |
| **VA** | Department of Veterans Affairs |

# Appendix C: Key Definitions

**Addressable storage capacity** – measure of how much data a computer system may contain.

**Advanced mathematics** - advanced portions of mathematics, customarily considered as embracing all beyond ordinary arithmetic, geometry, algebra, and trigonometry.

**Architecture** - models, policies, rules, or standards that govern which data is collected, and how it is stored, arranged, integrated, and put to use in data systems and in organizations.

**Artificial intelligence** - the development of computer systems that can perform tasks typically requiring human intelligence, such as language translation, decision-making, and visual perception.

**Autonomous intelligence** - the most advanced form of AI in which processes are automated to generate the intelligence that allows machines, bots, and systems to act on their own, independent of human intervention.

**Big data** - extremely large data sets that may be analyzed computationally to reveal patterns, trends, and associations, especially relating to human behavior and interactions.

**Cataloged** - data cataloging is the process of making an organized inventory of your data.

**Change management** - the management of change and development within a business or similar organization.

**Cloud computing** - the practice of using a network of remote servers hosted on the internet to store, manage, and process data, rather than a local server or a personal computer.

**Computing speed** - how quickly a processor can perform tasks.

**Curated** - data curation is the organization and integration of data collected from various sources.

**Data Analytics** - the science of analyzing raw data to reveal trends and draw conclusions, informing and supporting decision making and planning.

**Data ethics** - the moral obligations of gathering, protecting, and using personally identifiable information and how it affects individuals.

**Data governance** - setting internal standards and data policies that apply to how data is gathered, stored, processed, and disposed of.

**Data management** - process of ingesting, storing, organizing, and maintaining the data created and collected by an organization.

**Data science** - an interdisciplinary field that uses scientific methods, processes, algorithms, and systems to extract knowledge and insights from noisy, structured and unstructured data, and apply knowledge and actionable insights from data across a broad range of application domains.

**Data stewardship** - oversight or data governance role within an organization and is responsible for ensuring the quality and fitness for purpose of the organization's data assets, including the metadata for those data assets.

**Distributed analytics** - spreads data analysis workloads over multiple nodes in a cluster of servers, rather than asking a single node to tackle a big problem.

**Human-machine interaction** - the communication and interaction between a human and a machine via a user interface.

**Intelligence augmentation** - the deliberate enhancement of human intelligence using some technological means, such as eugenics, gene therapy, brain-computer interfaces, nootropics (smart drugs), neuroengineering, or some other means that hasn't been invented yet.

**Key Performance Indicator** - a quantifiable measure used to evaluate the success of an organization, employee, etc. in meeting objectives for performance.

**Machine learning** - the study of computer algorithms that can improve automatically through experience and by the use of data.

**Neural network** - computing systems inspired by the biological neural networks that constitute animal brains.

**Network bandwidth** - maximum amount of data transmitted over an internet connection in a given amount of time.

**Records management** - an organizational function devoted to the management of information in an organization throughout its life cycle, from the time of creation or receipt to its eventual disposition.

**Statistical modeling** - a mathematical model that embodies a set of statistical assumptions concerning the generation of sample data (and similar data from a larger population).

**Standards** - a technical specification that describes how data should be stored or exchanged for the consistent collection and interoperability of that data across different systems, sources, and users.

**Stored** - stored data means data, which although not a tangible visual depiction, may be used to render a tangible visual depiction through the use of an appropriate device or process.

**Symbolic modeling** - a therapeutic and coaching process developed by psychotherapists Penny Tompkins and James Lawley, based on the work of counselling psychologist David Grove. Using Grove's clean language, a progressive questioning technique using clients' exact words, the facilitator works with a client's self-generating metaphors to clarify personal beliefs, goals, and conflicts, and to bring about meaningful change.

**Tagged** - a tag is a keyword or term assigned to a piece of information, which helps describe an item and allows it to be found again by browsing or searching.

**Usability testing** - a technique used in user-centered interaction design to evaluate a product by testing it on users.

**Use case** - a specific situation in which a product or service could potentially be used.

**Use constraints** - how data may or may not be used to assure the protection of privacy or intellectual property.

# Appendix D: References


101st United States Congress. "PUBLIC LAW 101-576—NOV. 15, 1990." *congress.gov.* November 15, 1990. https://www.congress.gov/101/statute/STATUTE-104/STATUTE-104-Pg2838.pdf (accessed July 14, 2021).

104th United States Congress. "Clinger-Cohen Act." *fismacenter.com.* February 10, 1996. https://www.fismacenter.com/Clinger%20Cohen.pdf (accessed July 14, 2021).

105th United States' Congress. *Children's Online Privacy Protection Act of 1998.* October 21, 1998. https://uscode.house.gov/view.xhtml?req=granuleid%3AUSC-prelim-title15-section6501&edition=prelim (accessed July 14, 2021).

106th United States Congress. *Gramm-Leach-Bliley Act .* November 1999. https://www.ftc.gov/tips-advice/business-center/privacy-and-security/gramm-leach-bliley-act (accessed July 14, 2021).

113th United States Congress. "PUBLIC LAW 113–283—DEC. 18, 2014." *congress.gov.* December 18, 2014. https://www.congress.gov/113/plaws/publ283/PLAW-113publ283.pdf (accessed July 14, 2021).

91st United States Congress. *Fair Credit Reporting Act.* 1970. https://www.ftc.gov/enforcement/statutes/fair-credit-reporting-act (accessed July 14, 2021).

97th United States Congress. "Federal Managers Financial Integrity Act of 1982." *treasury.gov.* September 08, 1982. https://www.treasury.gov/about/organizational-structure/offices/Mgt/Documents/fmfia-legislation.pdf (accessed July 14, 2021).

Biden Jr., Joseph R. "Executive Order on Improving the Nation's Cybersecurity ." *The White House.* May 12, 2021. https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/ (accessed July 14, 2021).

B. Shneiderman, *Human-Centered AI*, Oxford University Press, 2022, forthcoming.

*California Consumer Privacy Act (CCPA).* https://oag.ca.gov/privacy/ccpa (accessed July 14, 2021)

*Change Management Statistics You Need to Know.* March 10, 2018. https://capacity4health.org/change-management-statistics/ (accessed July 14, 2021).

Crevier, Daniel. AI: The Tumultuous Search for Artificial Intelligence. New York, NY: BasicBooks, 1993. pp. 44–46.

Department of Defense: Office of Prepublication and Security Review. *DOD Information Resource Management Strategic Plan FY 19-23.* Modernization Strategy, Washington, D.C.: United States Department of Defense, 2019.

Department of Homeland Security. "Homeland Security Presidential Directive-12." *opm.gov.* August 27, 2004. https://www.opm.gov/news/reports-publications/management-budget-reports/homeland-security-presidential-directive-hspd-12.pdf (accessed July 14, 2021).

D. Nessler, "How to apply a design thinking, HCD, UX or any creative process from scratch," Medium, Feb. 07, 2018. https://medium.com/digital-experience-design/how-to-apply-a-design-thinking-hcd-ux-or-any-creative-process-from-scratch-b8786efbf812.

E. Frøkjær, M. Hertzum, and K. Hornbæk, "Measuring usability: are effectiveness, efficiency, and satisfaction really correlated?," in Proceedings of the SIGCHI conference on Human Factors in Computing Systems, New York, NY, USA, Apr. 2000, pp. 345–352. doi: 10.1145/332040.332455.

*Executive Order 13576--Delivering an Efficient, Effective, and Accountable Government.* June 13, 2011. https://obamawhitehouse.archives.gov/the-press-office/2011/06/13/executive-order-13576-delivering-efficient-effective-and-accountable-gov (accessed July 14, 2021).

*Executive Order on Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government.* December 03, 2020. https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-promoting-use-trustworthy-artificial-intelligence-federal-government/ (accessed July 14, 2021).

Federal Trade Commission. *Federal Trade Commission .* July 14, 2021. https://www.ftc.gov (accessed July 14, 2021).

*FTPS.* June 14, 2021. https://en.wikipedia.org/wiki/FTPS (accessed July 14, 2021).

Ge, Peng, et al. " A Conceptual Enterprise Framework for Managing Scientific Data Stewardship." *Data Science Journal*, 2018: 1-17.

Google PAIR, "User Needs + Defining Success Chapter Worksheet" in *People + AI Guidebook*, May 2019. https://pair.withgoogle.com/worksheet/user-needs.pdf

*Health Insurance Portability and Accountability Act of 1996.* 1996. https://aspe.hhs.gov/report/health-insurance-portability-and-accountability-act-1996 (accessed July 14, 2021).

H. Shen, W. H. Deng, A. Chattopadhyay, Z. S. Wu, X. Wang, and H. Zhu, "Value Cards: An Educational Toolkit for Teaching Social Impacts of Machine Learning through Deliberation," in Proceedings of the 2021 ACM Conference on Fairness, Accountability,

and Transparency, New York, NY, USA, Mar. 2021, pp. 850–861. doi: 10.1145/3442188.3445971.

J. Rotner, J. Hodge, L. Danley, "AI Fails and How We Can Learn From Them," MITRE, July 2020. https://sites.mitre.org/aifails/wp-content/uploads/sites/15/2021/02/AI-Fails-and-How-We-Can-Learn-from-Them-MITRE-2020.pdf.

"Market Guide for AIOps Platforms". Gartner. https://www.gartner.com/en/documents/3772124 (accessed July 14, 2021).

National Archives and Record Administration. *Digital Government: Building a 21st Century Platform to Better Serve the American People.* May 23, 2012. https://obamawhitehouse.archives.gov/sites/default/files/omb/egov/digital-government/digital-government.html (accessed July 14, 2021).

National Archives and Records Administration. "NARA 1608, NARA's Privacy Program." *archives.gov.* July 07, 2017. https://www.archives.gov/files/foia/directives/nara1608.pdf (accessed July 14, 2021).

National Institute of Standards and Technology (U.S. Department of Commerce). "U.S. LEADERSHIP IN AI: A Plan for Federal Engagement in Developing Technical Standards and Related Tools ." *nist.gov.* August 9, 2019. https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf (accessed July 14, 2021).

National Security Commission on Artificial Intelligence. *2021 Final Report.* Final Report for Fiscal Year 2021, Washington, D.C., United States of America: National Security Commission on Artificial Intelligence, 2021.

Office of the Press Secretary. *Executive Order 13571--Streamlining Service Delivery and Improving Customer Service.* April 27, 2011. https://obamawhitehouse.archives.gov/the-press-office/2011/04/27/executive-order-13571-streamlining-service-delivery-and-improving-custom (accessed July 14, 2021).

P. L. McDermott, D. C. O. Dominguez, D. N. Kasdaglis, M. H. Ryan, I. M. Trahan, and A. Nelson, "Human-Machine Teaming Systems Engineering Guide," MITRE, Dec. 2018. https://www.mitre.org/publications/technical-papers/human-machine-teaming-systems-engineering-guide

P. McDermott, "Calibrated Trust Evaluation Toolkit," MITRE, August 2020. https://comm.mitre.org/calibrated-trust-toolkit/

PUBLIC LAW 104–13—MAY 22, 1995. *congress.gov.* May 22, 1995. https://www.congress.gov/104/plaws/publ13/PLAW-104publ13.pdf (accessed July 14, 2021).

PUBLIC LAW 104–208—SEPT. 30, 1996. *congress.gov.* September 30, 1996. https://www.congress.gov/104/plaws/publ208/PLAW-104publ208.pdf (accessed July 14, 2021).

PUBLIC LAW 105–277—OCT. 21, 1998. *govinfo.gov.* October 21, 1998. https://www.govinfo.gov/content/pkg/PLAW-105publ277/pdf/PLAW-105publ277.pdf (accessed July 14, 2021).

Review, United States Department of Defense: Office of Prepublication and Security. *Executive Summary: DoD Data Strategy Unleashing Data to Advance the National Defense Strategy* . Executive Summary, Washington, D.C.: United States Department of Defense, 2020.

ROSS, RON, VICTORIA PILLITTERI, KELLEY DEMPSEY, MARK RIDDLE, and GARY GUISSANIE. "Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations." *nvlpubs.nist.gov.* February 2020. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-171r2.pdf (accessed July 14, 2021).

Sanger, David E., and Julian E. Barnes. "Biden Signs Executive Order to Bolster Federal Government's Cybersecurity." *New York Times.* May 12, 2021. https://www.nytimes.com/2021/05/12/us/politics/biden-cybersecurity-executive-order.html (accessed July 14, 2021).

Scarfone, Karen, Murugiah Souppaya, and Matt Sexton. "Guide to Storage Encryption Technologies for End User Devices." *nvlpubs.nist.gov.* November 2007. https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-111.pdf (accessed July 2021, 2021).

Shaun Donovan, United States Secretary of Housing and Urban Development. "OMB Memorandum." *Obama Whitehouse Archives.* January 03, 2017. https://obamawhitehouse.archives.gov/sites/default/files/omb/memoranda/2017/m-17-12_0.pdf (accessed July 14, 2021).

The Advanced Technology Academic Research Center (ATARC). *About ATARC.* July 2021. https://atarc.org/about/ (accessed July 14, 2021).

The Office of the Law Revision Counsel of the House of Representatives. August 16, 1954. https://www.census.gov/history/www/reference/privacy_confidentiality/title_26_us_code_1.html (accessed July 14, 2021).

*Title 13, U.S. Code.* August 31, 1954. https://www.census.gov/history/www/reference/privacy_confidentiality/title_13_us_code.html (accessed July 2021, 2021).

U.S. General Services Administration . *GSA IT Security Policies* . April 29, 2021. https://www.gsa.gov/policy-regulations/policy/information-integrity-and-access/gsa-it-security-policies (accessed July 14, 2021).

United States Department of Justice. "UNITED STATES DEPARTMENT OF JUSTICE OVERVIEW OF THE PRIVACY ACT OF 1974 2020 Edition." *justice.gov.* 2020. https://www.justice.gov/Overview_2020/download (accessed July 14, 2021).

*What is GDPR, the EU's new data protection law?* https://gdpr.eu/what-is-gdpr/ (accessed July 14, 2021).