

MARCH 2024

A GUIDE TO OUR DELIVERABLES

AI governance tools for
government, by government



GovAI
COALITION

TABLE OF CONTENTS

Introduction	3
AI Governance: From 0 to 90	5
Our Deliverables	6
AI FactSheet	7
AI Policy	8
AI Policy Manual	9
AI Incident Response	10
Vendor Agreement	11
Use Cases	12
Vendor Registry	13
Join the Coalition	14
About Us	15
Member Agencies	16
Timeline	17
Acknowledgements	18

Introduction

The GovAI Coalition is composed of over 400 public servants representing over 150 local, county, and state governments across the nation united in our mission to promote responsible and purposeful AI in the public sector.

In November 2023, the [GovAI Coalition](#) was established to give local government a voice in shaping the future of AI and ensure that AI development is geared toward the benefit of society. Since November, the Coalition has grown to include public agencies that represent over 100 million Americans.

As a majority of government agencies are [looking towards AI](#) to improve services, our Coalition strives towards a more just and beneficial technological future for all. From generating real-time translations of public meetings to predicting traffic patterns for improved bus routes, AI tools have the potential to make government services more responsive, efficient, and accessible to all.

AI also poses serious risks. [Racial discrimination](#), [violations of privacy](#), and [threats to safety](#) are well-documented harms of AI systems used without guardrails.

To improve the lives of all residents while protecting our communities against AI risks, it is vital that public agencies establish safeguards for the responsible use of AI.

For the past four months, hundreds have come together in the GovAI Coalition to do the work that practitioners would normally do alone. We recognize that only by working together can we produce high quality work at the speed necessary to meet the quickly-evolving AI landscape.

Introduction

(cont'd)

To this end, our members have worked diligently to create a suite of policy templates and knowledge-sharing tools that any public agency can use to jumpstart their own AI governance program. This includes:


- Template **AI Policy** and **AI Policy Manual** to present to your agency's leadership
- Template **AI Incident Response Plan** for the next generation of security risks
- **AI FactSheet** and **Vendor Agreement** template to ensure we are buying responsible AI tools
- A growing repository of guides for **AI use cases** being tested by other agencies
- A growing **Vendor Registry** of AI systems available for government use

These documents and repositories are by practitioners, for practitioners to serve our communities better through responsible AI.

We invite policymakers to borrow, edit, and adopt these templates for their own use. This guide is intended to help practitioners quickly understand, adopt, and implement the policy templates created by the Coalition. In return, we ask that you share your feedback and continue the Coalition's conversation around responsible AI governance.

We have intentionally aligned the templates with emerging national standards, including the Presidential [Executive Order on AI](#), the Office of Management and Budget's [draft memorandum on AI](#), and the National Institute of Standards and Technology's [AI Risk Management Framework](#). Our Coalition makes it easier to stay aligned with government best practices, receive support from other members, and procure new AI systems with greater transparency from vendors.

Over the next year, we will support agencies in implementing their own AI governance programs, advocate for responsible AI at all levels of government, and share effective AI use cases. Agencies interested in joining the Coalition may do so [here](#).



AI Governance: From 0 to 90

The Coalition’s plug-and-play policy deliverables can help you take your AI governance program from 0 to 90.

Question	Deliverable
How do I prepare my agency to use and govern AI?	AI Policy & AI Policy Manual principles, policy, and processes
What can I do with AI that adds value?	Use Cases starter guide and resources
Who can I work with for good, responsible AI solutions?	Vendor Registry database of AI FactSheets
How do I ensure vendors follow my expectations?	Vendor Agreement legal addendum for AI systems
How do I prepare for when things go wrong?	Incident Response AI Incident Response Plan
	AI FactSheet nutrition facts of AI system

Our Deliverables

In this section, we describe the purpose of each deliverable and the role it serves in an agency's AI governance strategy. We explain how practitioners can implement the templates in their agency and indicate specific actors that should be involved in the use of the templates in **bolded text**.

The development of these seven deliverables in just four months was made possible by the tireless dedication of our Coalition members and their commitment to responsible AI. We have intentionally aligned the templates with emerging national standards, including the Presidential Executive Order on AI, the Office of Management and Budget's draft memorandum on AI, and the National Institute of Standards and Technology's AI Risk Management Framework.

AI FactSheet

Purpose

The AI FactSheet is a template that captures the “nutrition facts”, or essential technical details, of an AI system. Vendors are expected to complete the AI FactSheet during the procurement process. The AI FactSheet is a critical document that provides agencies the technical information they need to adequately understand, evaluate, and use third-party AI systems.

How to use it

Procurement officers require the **vendor** to complete the AI FactSheet during the procurement process. The **Chief AI Officer**, or the equivalent staff responsible for reviewing AI systems, reviews the AI FactSheet to determine if the technical capabilities of the AI system meet the agency’s needs.

Vendors may signal their commitment to responsible AI practices by proactively completing and submitting the AI FactSheet to the Vendor Registry through the [Vendor Registry Intake Form](#). **Staff involved in the use of AI systems by the agency** may view the submitted AI FactSheet in the Vendor Registry to compare vendors or determine if the vendor is a good fit for the agency.

AI FactSheet for Third Party Systems

Please provide details regarding your Artificial Intelligence (AI) product by filling out the FactSheet¹ template below. This information will be kept in the GovAI Coalition vendor registry, and may be made publicly available. You can find an example of a completed FactSheet on page 4.

AI FactSheet

Vendor Name	<hr/>
System Name	<hr/>
Overview	Brief summary of the AI system.
Purpose	What function does the AI system perform, and for what purpose? If the system performs multiple functions, list each discretely and reference below. For features that are configurable, please describe all configuration options and default settings.
Intended Domain	What domain is the AI system intended to be applied in?
Training Data	How was the AI system trained? What data was used? How often is data added to the training set? Was all training data legally obtained and its use fully licensed?
Test Data	What data was used to test system performance? Under what conditions has the system been tested?
Model Information	General description of the model(s) used (e.g., large language model, transformer, deep learning, supervised learning, built on an existing open source model, computer vision)
Update procedure	In general, how often are the models updated for users? Will the user have a choice in moving to the updated model or staying on the current model? What documentation is available for new versions of the model?
Inputs and Outputs	What are the inputs to the AI system? What are its outputs? What interfaces and integrations are supported?

¹ The FactSheet template is heavily inspired by the IBM Research [AI FactSheets 360 project](#).

GovAI

1

AI Policy

Purpose

The AI Policy is the agency's broad institutional policy that governs the use of AI systems. The AI Policy is meant to apply to all AI systems deployed by the agency and all individuals when using AI systems on behalf of the agency (e.g., staff, contractors). The AI Policy defines:

- Guiding Principles for responsible AI
- Roles and responsibilities related to the agency's use of AI systems
- AI Review process for AI procurements
- Requirement for vendors to complete the AI FactSheet and comply with the Vendor Agreement template
- AI Incident Response Plan
- Prohibited uses of AI systems
- Sunset procedures for AI systems
- Policy enforcement

How to use it

Agency staff participating in the GovAI Coalition present the AI Policy template to their highest governing body (e.g. City Council, Board of Supervisors) for approval. The **governing body** approves the adoption of the AI Policy. **Agency staff across the organization** use the AI Policy Manual to implement the AI Policy.

Artificial Intelligence (AI) Policy
[Agency]

I. Purpose

This policy establishes a comprehensive, yet flexible, governance structure for AI systems used by, or on behalf of, the [Agency]. This policy enables the [Agency] to use AI systems for the benefit of the community while safeguarding against potential harms.

The key objectives of the AI Policy are to:

- Provide guidance that is clear, easy to follow, and supports decision-making for the staff (full-time, part-time), interns, consultants, contractors, partners, and volunteers who may be purchasing, configuring, developing, operating, or maintaining the [Agency's] AI systems or leveraging AI systems to provide services to the [Agency].
- Ensure that when using AI systems, the [Agency] or those operating on its behalf, adhere to the Guiding Principles that represent values with regards to how AI systems are purchased, configured, developed, operated, or maintained.
- Define roles and responsibilities related to the [Agency's] usage of AI systems.
- Establish and maintain processes to assess and manage risks presented by AI systems used by the [Agency].
- Align the governance of AI systems with existing data governance, security, and privacy measures in accordance with the [Agency's Information and Security Policy] and [Agency's Data Policy].
- Define prohibited uses of AI systems:
- Establish "sunset" procedures to safely retire AI systems that no longer meet the needs of the [Agency].
- Define how AI systems may be used for legitimate [Agency] purposes in accordance with applicable local, state, and federal laws, and existing agency policies.

The [Agency] defines "artificial intelligence" or "AI" to be a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments.¹ AI systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate

¹ Definition from [15 U.S.C. 9401\(3\)](#).

GovAI
COALITION

1

AI Policy Manual

Purpose

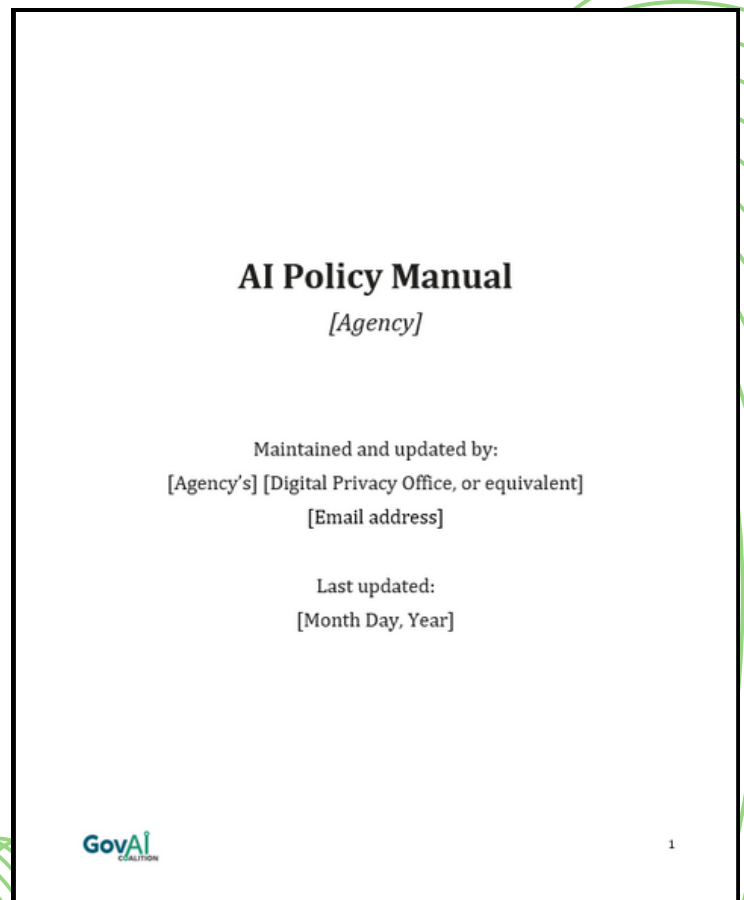
The AI Policy Manual provides guidance on how to comply with an agency's AI Policy. The AI Policy Manual details:

- AI Governance structure in the agency
- AI Review Framework
- RFP protocol for AI systems, including example questions and how to evaluate responses

The AI Review Framework is an in-depth guide for how practitioners should review procurements that involve an AI system. It consists of multiple steps, including risk analysis, impact assessment, public engagement, review, pre-launch preparation, and ongoing monitoring.

How to use it

Staff involved in the use of AI systems by the agency should refer to the AI Policy Manual to ensure they comply with the agency's AI Policy. The **Chief AI Officer**, or the position responsible for reviewing AI systems, should become well-versed in the AI Review Framework. **Procurement officers** should also become familiar with the AI Review Framework since it occurs throughout the public procurement process.



AI Incident Response

Purpose

The AI Incident Response Plan is a playbook for agencies to prepare for and address AI incidents. Examples of AI incidents include an AI system leaking confidential information or routinely producing inaccurate outputs. The AI Incident Response Plan is designed to help agencies gain a fundamental understanding of the processes, responsibilities, and actions required to prevent, mitigate, address, and resolve AI incidents.

The AI Incident Response Plan draws from the NIST AI Risk Management Framework and the Special Publication 800-61 Computer Security Incident Handling Guide.

How to use it

The agency's **cybersecurity, privacy, and AI staff** customize the AI Incident Response Plan template to suit their unique institutional needs. The template may be integrated and implemented with the AI Policy Manual or any of the agency's existing incident response plans, such as those for cybersecurity or privacy.

AI Incident Response Plan

Introduction

The AI Incident Response Plan (IRP) serves as the first line of defense for the [Agency] in case of an AI incident. This IRP has been created based on the NIST AI Risk Management Framework and the Special Publication 800-61 Computer Security Incident Handling Guide.

Incident response occurs in sequential phases, each one building upon the next. The following phases provide a foundation for an Incident Response (IR) Team to respond to and recover from an AI incident:

1. Preparation
2. Detection & Analysis
3. Containment
4. Eradication & Recovery
5. Post-incident Activity

Purpose

The purpose of this document is to prepare for and gain a fundamental understanding of the processes, responsibilities, and actions required to mitigate an AI incident. It is critical to identify and resolve incidents quickly before they escalate into a major incident with the potential to cause harm or damage to people, data, or the [Agency].


Scope

This IRP applies to all AI systems implemented by [Agency] staff, contractors, and any entity operating on behalf of the [Agency]. The AI IRP addresses continuity and recovery procedures to appropriately mitigate AI incidents.

Approach

The key points in development of the AI IRP include:

- **Evaluate:** Evaluate risk levels and determine the appropriate response for an AI incident, which may include obtaining senior management support.
- **Plan:** Keep the plan simple. A well-organized, systematic, and up-to-date AI IRP that is readily available will help teams get through most situations.
- **Communicate:** Communicate regularly on the incident status. Provide the relevant facts as they are available, disseminate them quickly, follow up regularly, keep relevant parties informed and resolve incorrect information.
- **Review:** Review the AI IRP at least bi-annually to ensure the documented procedures are still appropriate and that the team is equipped to respond


1

Vendor Agreement

Purpose

The Vendor Agreement is a boilerplate addendum that agencies can include in contractual agreements with AI vendors and/or in RFP solicitations. The Vendor Agreement plays a critical role in ensuring that the AI systems sold to agencies comply with the expectations for responsible AI in government.

The Vendor Agreement describes risk mitigation practices required of the vendor and outlines ten responsible AI provisions that the vendor must comply with, including requirements around performance, algorithmic bias, human oversight, incident response, and auditing. Notably, the Vendor Agreement requires that the vendor submit an AI FactSheet for their AI system and maintain an updated version.

Agencies may also benefit from example addendums regarding data privacy or cybersecurity. While outside the scope of this deliverable, agencies can find an example [privacy addendum here](#) and a [cybersecurity addendum here](#).

How to use it

The agency's **legal, procurement, and IT staff** customize the Vendor Agreement template to suit their unique institutional needs. The **procurement officer** includes the Vendor Agreement as an exhibit in relevant RFPs and other solicitations. The **procurement officer** ensures that the Vendor Agreement is attached as a legal rider to purchasing agreements for procurements that involve an AI system. The contracted **vendor** complies with the legally binding requirements outlined in the Vendor Agreement.

Addendum [X]: Requirements for AI Systems

This Addendum defines special requirements agreed to by [Agency] and Contractor regarding the AI system and/or subsystem provided as part of the Contract.

This Addendum governs over any contrary license terms and [Agency] will not agree to any terms that conflict with the Addendum. Failure of the Contractor to comply with the terms of this Addendum shall constitute a material breach of the Contract.

Contractor agrees to indemnify, defend, and hold harmless [Agency] regarding any third-party action arising out of or related to (1) any breach of any representation or warranty of Company contained in this Addendum; (2) any breach or violation of any covenant or other obligation or duty of Contractor under this Addendum or under applicable law; (3) any third party Claims which arise out of, relate to or result from any act or omission of the Contractor related to the provision of an AI system; and (4) any violations or alleged violations of intellectual property rights; in each case whether or not caused in whole or in part by the negligence of [Agency], or any other Indemnified Party, and whether or not the relevant Claim has merit.

1. Definition of AI System

The [Agency] defines "artificial intelligence" or "AI" to be a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. ¹ AI systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.

The [Agency] defines an "AI system" to be any data system, software, hardware, application, tool, or utility that operates in whole or in part using AI. ²

2. Guiding Principles

The Contractor shall demonstrate that the AI system and its usage, deployment, and maintenance as it pertains to the services outlined in this agreement do not conflict with the [Agency]'s AI Policy and Principles therein.

¹ Definition from 15 U.S.C. 9401(3)

² Definition from United States Executive Order No. 14110 on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence

Use Cases

Purpose

The Use Case template provides agencies with a standardized format for recording and sharing their experiences using AI. The template structures AI projects in such a way that agencies can copy from each other to build stakeholder buy-in, identify necessary resources, and guide implementation. The template enables agencies to report their learnings from an AI use case they have implemented, including strategies for success and pitfalls to avoid for other agencies.

How to use it

Staff involved in the use of AI systems by the agency customize the Use Case template for each AI use case they plan to share. **Staff** may view and download the completed use cases shared by other agencies to advance best practices and gain insights into a real-world implementation.

[Project Title]
GovAI Coalition Use Case

Overview

This document serves as a template for agencies to create tailored use cases of AI. The document is meant to illustrate a strategic method for using AI technology in solving challenges in state and local governments.

When filling in the template, you do not need to complete each section. Sections can start as considerations (or blanks) and more information can be added in depth as the use case expands.


When addressing the complexity and risk of your use case, consider the depth and detail required in different sections of the template. This varies based on the sensitivity of data involved. For use cases dealing with highly sensitive data like protected health information or personal identifiers, exercise greater caution and ensure compliance with data protection laws. In contrast, projects involving less sensitive, publicly available data like code or ordinances may not necessitate the same level of detail. This template should be viewed as a dynamic document, adaptable to evolving risks and project needs. It is important to recognize that not all risks are immediately apparent, and the project will evolve over time. Keep this in mind while documenting your approach and decision-making process.

Examples can be found in the AI Use Cases Working Group folder.

Proposed Use Case

In a sentence, what is your use case? What is the question you are trying to answer or problem to solve?

In defining your proposed use case, start with a clear, one-sentence statement of the problem or question you aim to address. Consider conducting an organizational assessment or survey, possibly in collaboration with public agency leadership or CTOs/CIOs, to pinpoint business needs across various service areas like Public Safety, Health and Human Services, Public Works, etc. Recognize that some issues may be common to multiple areas, while others are unique to specific categories. Use these insights to articulate your use case, ensuring it is either broadly encompassing or precisely focused, based on the unique and common challenges identified within your organization. This approach will help in clearly defining the scope and relevance of your use case.

 1

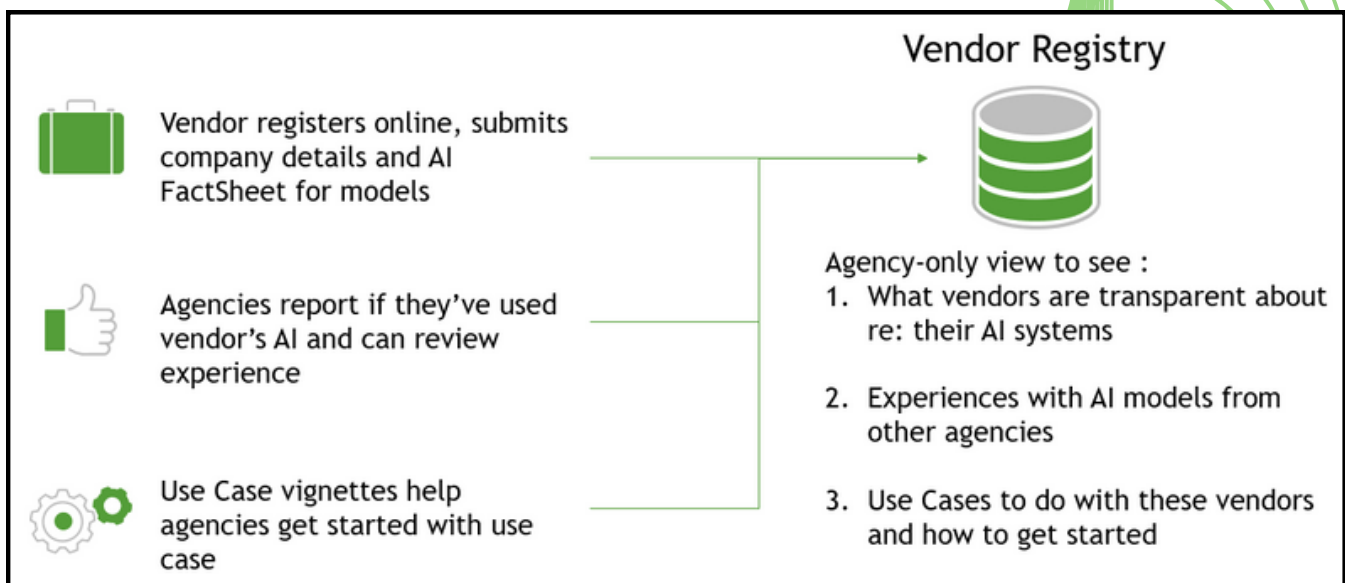
Vendor Registry

Purpose

The Vendor Registry is a platform intended to promote greater transparency from vendors around AI systems. Vendors provide information about their AI system (e.g., technical capabilities, performance metrics, optimal deployment conditions) via the AI FactSheet to the Registry, which can then be viewed by agencies in the GovAI Coalition. Agencies can also share their experience working with certain vendors to other agencies on the platform.

How to use it

Vendors complete and submit the AI FactSheet via the [Vendor Registry Intake Form](#) to be included on the Registry. **Agency staff looking to use AI tools** can browse the Registry for potential AI solutions and vendors. The agency's **procurement officer** views vendor submissions to gain more information on potential vendor systems. **Staff involved in the use of AI systems by the agency** share their experiences directly in the Registry to help others understand the effectiveness of those AI systems or vendors.



Join the Coalition

The GovAI Coalition welcomes public agencies, non-government organizations, academic institutions, and companies to join our effort. We recognize that effective AI governance takes place in a larger economic and regulatory ecosystem, and all sectors of society must be involved in shaping a brighter and more equitable technological future.

Coalition membership will look different for each type of organization. The table below illustrates how membership varies for the four organization types.

	Government	Civil Society	Academia	Private sector
How to join	Agency Application Form	Civil Society & Academia Application Form	Civil Society & Academia Application Form	Vendor Registry Intake Form
Voting rights*	Yes	No	No	No
Board member eligibility*	Yes	No	No	No
May attend	General Coalition meetings Working group meetings	General Coalition meetings Working group meetings	General Coalition meetings Working group meetings	General Coalition meetings Limited working group meetings
Additional activities	All Coalition activities	Research projects Special projects	Research projects	Cross-agency AI applications

*Only government agencies will have the ability to vote in the election of GovAI Coalition board members. Only government agencies may serve as board members.

About Us

The GovAI Coalition is composed of over 400 public servants representing over 150 local, county, and state governments across the nation united in our mission to promote responsible and purposeful AI in the public sector. In November 2023, the GovAI Coalition was established to give local government a voice in shaping the future of AI and ensure that AI development is geared toward the benefit of society. Since November, the Coalition has grown to include public agencies that represent 100 million Americans.

As a majority of government agencies are looking towards AI to improve services, our Coalition strives towards a more just and beneficial technological future for all. From generating real-time translations of public meetings to predicting traffic patterns for improved bus routes, AI tools have the potential to make government services more responsive, efficient, and accessible to all.

AI also poses serious risks. Racial discrimination, violations of privacy, and threats to safety are well-documented harms of AI systems used without guardrails. To improve the lives of all residents while protecting our communities against AI risks, it is vital that public agencies establish safeguards for the responsible use of AI.

Over the next year, we will focus on supporting agencies as they implement their own AI governance programs, advocating for responsible AI at all levels of government, and sharing effective AI use cases.



The Coalition serves as a link between global AI actors and local governments.

Connecting Efforts Top-Down:

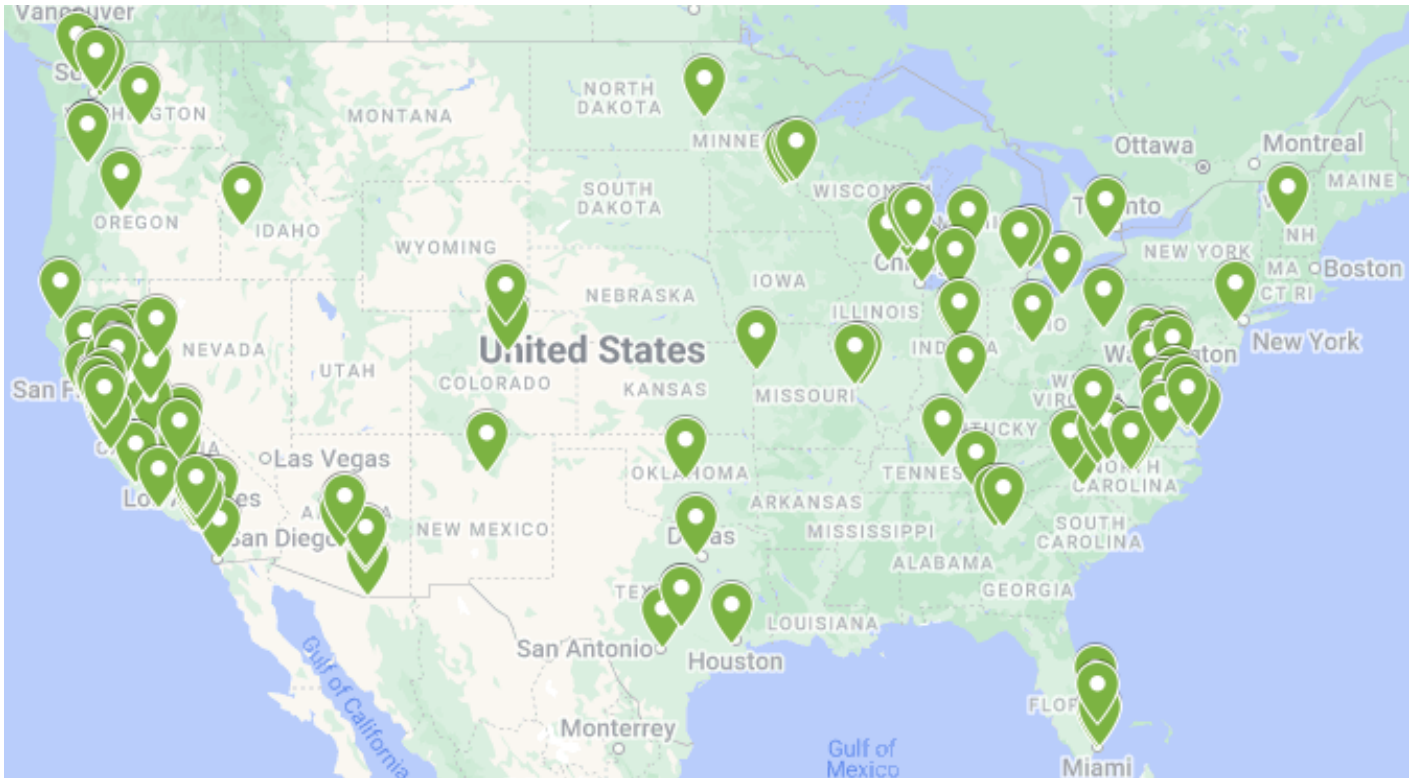
- The Coalition translates the AI policies and frameworks from regulators (e.g., White House, NIST, CPPA) into tools, “fill-in-the-blank” templates, and guides enabling even the smallest towns to use AI responsibly.
- The Coalition standardizes vendor expectations so companies can develop a common solution that meets the needs of hundreds of local agencies.

Empowering Agencies Bottom-Up:

- The Coalition is on-the-ground with local agencies implementing AI policies and solutions, pooling resources, identifying common issues, and raising collective needs to the global conversation.
- The Coalition gives local agencies a united voice to influence AI development towards trustworthy and equitable solutions that benefit our communities.

Members

The Coalition's members represent agencies across the whole country.



The Coalition is composed of a diverse range of over 400 practitioners representing over 150 agencies, from small rural towns to the largest urban centers in the country. Altogether, the Coalition's member agencies represent 100 million Americans.

Regardless of an agency's size, the Coalition provides versatile governance tools that empower practitioners to leverage the benefits of AI while safeguarding against risks.

Timeline

The GovAI Coalition enters Phase 2 during March 2024. Phase 1 of the Coalition focused on developing the policies, practices, and tools that were released alongside this Guide to Deliverables.

During Phase 2, which will last one year, the Coalition's activities are focused on **adopting** those resources, **implementing** civic AI solutions that leverage those resources, and **partnering** with non-government organizations to use AI in the public sector responsibly.

During Phase 2, our core goals are to:

- **Support** agencies in their adoption of the Coalition's resources.
- **Iterate** on existing resources, and develop additional resources as needed.
- **Connect** agencies with responsible AI solutions via the Vendor Registry, and streamline procurement via cooperative purchasing.
- **Provide** a toolkit for deploying and evaluating common use cases.
- **Partner** with national governments and international organizations on the development of international AI policy and AI implementation.

Milestones

01

Vendor Registry

100 AI FactSheets submitted to Vendor Registry by Sept 2024, and 200 by March 2025

02

Use Cases

30 AI use case "vignettes" created by agencies by Sept 2024, and 45 by March 2025

03

Policy

Release "Version 2" of templates by March 2025

04

Adoption

100 public agencies have adopted policies or practices based on the templates by March 2025

05

Platform

Beta platform for new Vendor Registry and use case resource sharing is live by March 2025

06

Community of Practice

Host a convening on government use of AI by March 2025

Acknowledgements

This work would not have been possible without the tireless dedication of our Coalition members. The past four months have been a historic collaboration between over 150 government agencies, who all banded together under the common mission of ensuring responsible and purposeful use of AI in the public sector.

We extend a heartfelt thanks to the many exceptional people who have been involved in this effort. We are particularly grateful to the members who participated in our working groups, and the exceptional working group chairs whom without none of this would have been possible. The working group members and co-chairs behind these deliverables devoted hundreds of hours outside of their normal jobs to deliver the policy templates and tools now at your disposal.

Thank you to the working group chairs below and all the Coalition members we are unable to name here who believed in our vision and helped make it a reality.

Policy

Emily Royall - City of San Antonio, TX
Anthony Fisher - State Department of Revenue, CO

Vendor Agreement

Ethan Benatan - TriMet, OR
Ryan Kurtzman - City of Long Beach, CA

Incident Response

Jonathan Behnke - City of San Diego, CA

Vendor Registry

Roy Fernando - City of Cleveland, OH

Use Cases

Jaime Wascalus - City of St. Paul, MN
Jiri Rutner - County of San Diego, CA
Omar Naseef - City of Austin, TX

Special thanks to the small but mighty team at the City of San José.

- **Khaled Tawfik**
- **Albert Gehami**
- **Leila Doty**
- **Matthew Jacquez**
- **Angela Zhao**

This team of five facilitated the amazing work of the Coalition over the past four months.

Acknowledgements (cont.)

Phase 1 Working group members (not exhaustive)

Andrew Ngui - City of Kansas City, MO

Brendan M. Daly - City of San Diego, CA

Chris Pahl - County of Santa Clara, CA

Danny Thomas Vang - Committee on Information
Technology, City and County of San Francisco, CA

Deborah Erb - County of San Luis Obispo, CA

Elizabeth Buehler - Salt Lake City Corporation, UT

Falisha (Fali) Field - County of San Diego, CA

Gene Leynes - City of Chicago, IL

Gregg Turnbull - Larimer County, CO

John Matelski - DeKalb County, GA

Lea Eriksen - City of Long Beach, CA

Levy Sun - City of Norwalk

Omar Moncayo - City of Long Beach, CA

Simon Huang - City of St. Louis, MO

Tas Jalali - AC Transit, CA



GovAI
COALITION