Best Practices for Public Sector AI Use Case Inventories

July 2024

Authored by
Quinn Anex-Ries, *Senior Policy Analyst, Equity in Civic Technology*

Introduction

Transparency is a longstanding, bipartisan bedrock of a functioning government that best serves the people, and government use of AI is no exception. To achieve transparency around how public agencies are using AI, governments are creating and publishing AI use case inventories that detail when, how, and why public agencies develop, acquire, and use AI systems. Typically, AI use case inventories take the form of a regularly maintained, often annual, report or repository of information about the purpose and use of each AI tool an agency uses, the types of data used to train and deploy a system, how a system is tested, and other information about the development and acquisition of a tool.

Already, use case inventories are legally mandated across all federal agencies, and at least 13 states and 2 cities have implemented their own inventory requirements. Federal AI use case inventories were first required under Executive Order 13960 signed by President Trump in December 2020, and were later enacted into law through the bipartisan *Advancing American AI Act* in December 2022. At least eleven states have enacted legislation that requires public agencies to conduct AI inventories, including California, Connecticut, Delaware, Indiana, Kentucky, Maryland, New Hampshire, New York, Texas, Vermont, and West Virginia. And, several other states implemented such a requirement through executive orders, including Alabama, California, and Mississippi. Two cities — New York, NY and San Francisco, CA — passed city-wide ordinances that require public agencies to conduct AI inventories. And other cities, like San Jose, voluntarily created AI inventories through their Chief Information Officers' offices.

As government AI uses continue to grow, it is all the more urgent for policymakers to proactively communicate to members of the public and other impacted stakeholders (including researchers, academics, private companies, and other branches of government) about the use and impact of AI and

to ensure that procedures are in place to enable a robust, ongoing inventory process. Transparency is critical for advancing public trust and the efficient use of public resources, providing insight and accountability into government actions, enabling individuals to make informed decisions and to seek redress if harm occurs, and empowering stakeholders to identify potential gaps or issues in the government's assessment and oversight of these tools.

Building on past successes and failures in government transparency efforts, this brief outlines best practices in two important facets of AI inventories:

- **Structure:** AI use case inventories should have structural requirements to ensure that they are formatted, shared, and maintained in a manner that is maximally transparent and accessible to stakeholders.
- **Content:** AI use case inventories should include detailed, easily understood information about a system's purpose, use, and design so that stakeholders have access to meaningful information about how and why constituents are impacted by AI.

Structural Best Practices

For AI use case inventories to be as useful as possible for government agencies, private industry, academics and researchers, and the public, governments should consider establishing structural requirements to ensure that inventories are robust, transparent, and regularly updated.

One of the primary challenges that governments face in publicly communicating about their uses of AI is that there are multiple, distinct audiences that each have unique needs and levels of understanding. For instance, while long and detailed inventories may be useful for civil society organizations, academics, legislators, or private companies, this level of information may be overwhelming to tax payers and concerned citizens. To make this information accessible to the greatest number of stakeholders possible, governments should take steps to communicate and share AI use case inventories in different formats and mediums, including the following:

- **Provide short summaries with key points** for each tool and agency within an inventory, allowing individuals to quickly find high-level information about the general function and purpose of a specific tool or the range of use cases across an agency;
- **Label and organize** AI use case inventories **to make them easily navigable and searchable**. AI inventories should be structured such that the average layperson can quickly find the specific AI use case that they care most about (e.g., if an inventory contains 100 AI systems, an individual should be able to easily figure out and find information about any systems that specifically have to do with Medicaid benefits); and

- **Adopt innovative formats for stakeholders to interact with AI inventories** that go beyond a simple spreadsheet or database. This could include data visualizations, webinars, short-form blog posts, or other forms of interactive multi-media. For instance, the state of Mississippi plans to release a podcast to provide stakeholders with an overview of the information contained in the state's AI inventory.[1]

Currently, most existing AI inventories consistently implement two best practices:

- **Centrally collate and publish** all agency inventories to make it easy to find and navigate government-wide data. Many jurisdictions, including the federal government, Vermont, and New York City, require agency inventories to be centrally compiled;

- **Designate an agency or office** (e.g., Department of Technology) responsible for overseeing the inventory process and for compiling all agency inventories. A majority of jurisdictions appoint a central agency or office with the responsibility of overseeing the government-wide inventory process, including the federal government, California, Connecticut, Kentucky, Maryland, and Vermont.

Several additional best practices in AI inventories have emerged, although they are inconsistently prioritized in existing inventories:

- **Make full agency AI use case inventories publicly available**. While some jurisdictions affirmatively require these to be made public (such as the federal government and Connecticut), others (like Texas and Indiana) do not require AI use case inventories to be made publicly available, limiting public accountability about how systems are developed and used;

- **Standardize reporting across agencies** to ensure consistency and conformity of information. Inconsistency in agency reporting is a common issue across governments.[2] To address this challenge, some jurisdictions, like Texas, have instituted requirements for a centralized government agency to create a shared submission form to encourage consistency in agency reporting. Without such standards, significant variability between agency reporting can make it more challenging for stakeholders to easily understand and use AI use case inventories, and can lead to potential errors or omissions in agency reporting;

- **Include all current and planned AI systems**, and not just high-risk, sector-specific, or generative AI systems, in use case inventories. A number of jurisdictions (such as Alabama, California, Delaware, and Maryland) only require their AI use case inventories to include information about high-risk tools or generative AI tools, meaning that a number of systems that do not meet these

[1] See, *Artificial Intelligence*, Mississippi Department of Information Technology Services, https://www.its.ms.gov/services/innovating/AI [https://perma.cc/K6BE-PVW].
[2] See, e.g.: Quinn Anex-Ries, *Exploring the 2024 Federal AI Inventories: Key Improvements, Trends, and Continued Inconsistencies*, Center for Democracy & Technology (April 15, 2025), https://cdt.org/insights/exploring-the-2024-federal-ai-inventories-key-improvements-trends-and-continued-inconsistencies/ [https://perma.cc/E8W7-3PDT].

definitions get left out of public reporting even if they can have significant consequences for the delivery of government benefits and services;[3]

- **Add new systems to use case inventories within a predetermined period** (e.g., within 60 days from first date of use) and **conduct a general update at least annually**. Most commonly, jurisdictions require that AI inventories be updated annually. However, a number of jurisdictions — including Alabama, West Virginia, and Maryland — do not establish a specific timeframe for agency inventories to be updated. Most jurisdictions lack additional requirements for new systems to be added on a more frequent basis;

- **Designate senior leadership** responsible for conducting the inventory process. While most jurisdictions direct agencies to lead their own inventory process, few explicitly require agencies to appoint dedicated leadership to do so. Appointing specific agency leadership to oversee this process helps ensure accountability and coordination across government;

- Ensure that inventory responses are **accessible, in plain language, and machine readable**. Few jurisdictions explicitly require that their inventories are accessible and understandable to the public. However, some jurisdictions require that their inventory reporting aligns with existing open government requirements that often include similar obligations. These requirements facilitate improved utility and public access to this information;

- Conduct **stakeholder engagement *before* inventories are created and published** to solicit input about what information to include and how best to communicate this information to stakeholders; and

- Conduct **datawalks or other forms of public engagement *after* inventories are published** so that stakeholders understand what inventories include and why that information is important, as well as to solicit feedback about how inventories should be updated in the future.[4] Currently no jurisdiction has an affirmative requirement for agencies to consult stakeholders about the form or contents of AI use case inventories.

[3] In California, for instance, this led to serious issues with the state's AI inventory, which included no reported uses of high-risk AI systems in spite of clear evidence that such systems are already in use. See, Khari Johnson, *State Claims There's Zero High-Risk AI in California Government–Despite Ample Evidence to the Contrary*, CalMatters (May 28, 2025), https://calmatters.org/economy/technology/2025/05/california-somehow-finds-no-ai-risks/ [https://perma.cc/432C-FJRZ].

[4] For example, the City of Long Beach launched mobile community workshops and data walks to educate stakeholders about the technologies used by the city and how those tools use data. A similar model could be used to inform stakeholders about information contained within AI use case inventories and to solicit feedback. See, *Long Beach Hosting 'Data Walks' to Increase Transparency and Public Awareness for City's Data Collection Practices*, City of Long Beach (February 20, 2024), https://longbeach.gov/press-releases/long-beach-hosting-data-walks-to-increase-transparency-and-public-awareness-for-citys-data-collection-practices/ [https://perma.cc/E7BA-ZVE7].

Content Best Practices

The true measure of an AI use case inventory's usefulness is in providing the public and other stakeholders with sufficient information to understand why and how the government is using AI. The public is entitled to know how their rights and safety are affected by AI and how taxpayer resources are being used to acquire and deploy these tools. Moreover, this understanding is critical for supporting efforts to hold governments accountable and to mitigate harms that may have been overlooked.

Drawing on best practices from legislation and public agencies across the country, ***governments should consider requiring the following 10 categories of information in AI use case inventories*** as a starting point, while recognizing that interested members of the public and other stakeholders should be consulted about what they care most about knowing. Most importantly, governments should take steps to ensure that agencies provide *detailed* and *robust* answers to each of these fields, while also maintaining easily navigable summaries as described above on page 2.

Required Categories	Why This Matters
***Description of the AI system**, including:* ● *intended purpose,* ● *anticipated benefits,* ● *how the model was trained, fine-tuned, and tested (including performance benchmarks used during testing, the legal permissions for the data used during model training, and steps taken to ensure that the model was trained and tested on representative data),* ● *outputs, and* ● *inputs (including any demographic variables).*	Providing significant detail on *why* and *how* an agency plans to use or is using an AI system as well as the specific kinds of data that the system uses and the outputs it produces is critical for providing stakeholders with a baseline understanding of the overall function and role of a system within agency operations, helping stakeholders to evaluate the veracity and thoroughness of other information provided throughout the inventory.
***Clear information about the developer of the AI system and specific details about system procurement or development**:* ● *If the system is procured from a third-party vendor, provide a description of the contracting agreement including provisions to require performance testing, privacy and cybersecurity compliance, and risk mitigation measures, and to prohibit the re-use of government data for training the AI system.*	Giving details about the developer of a system, and specific information relevant to the vendor or in-house development team, allows stakeholders to understand who was responsible for developing each system and to hold relevant actors accountable for errors or oversight.

• *If the system was developed in-house, the name(s) of the department(s) who developed the system and a description of the technical components and data used to develop the system.*	
Description of how different communities are anticipated to be impacted *by the AI system, including:* • *the specific communities (e.g., neighborhoods, demographics, government workers, beneficiaries) that the agency has considered and expects to be impacted by the system; and* • *the specific impacts (e.g., provision or denial of benefits, increased or decreased access to information, assignment of monetary or other penalties) that the agency anticipates such communities will face upon use of the system.*	This information is important both for improving general public awareness and understanding about the on-the-ground impact of an AI system, but also allows stakeholders to evaluate how they may be affected by a particular tool as well as the overall scope and potential risks that a tool may pose.
Determined level of risk *(low, medium, high) of the AI system and a description of how the agency made this risk determination, including a description of the specific considerations that informed an agency's determination such as:* • *scope of impact,* • *context of use,* • *level of human oversight, and* • *the significance of the system's impact on rights and safety.*	These designations are the most important factor in an agency's process for determining the level of due diligence and confidence in mitigation measures that are required for each system, providing stakeholders with critical insight into how agencies are ensuring that higher risk systems are subjected to heightened safeguards and a more rigorous review process.
Description of the measures taken to protect data privacy and security*, including:* • *how data entered into and output by the system is securely stored, processed, and shared; and* • *how the agency is ensuring that the system complies with existing legal and regulatory data privacy and cybersecurity requirements.*	Because AI systems introduce new privacy and security risks to agency operations, this offers stakeholders insight into the steps that agencies are taking to proactively address such risks during the development, acquisition, and use of an AI system.

Description of the risk mitigation measures *the agency has adopted during the development, acquisition, and use of an AI system to address any foreseeable risks posed by the system, including:* ● *pre-deployment testing,* ● *ongoing monitoring, and* ● *procedural safeguards (such as notice and appeal procedures).*	For high-risk systems especially, this enables stakeholders to assess the sufficiency of the measures that agencies are adopting to ensure that AI systems function safely and effectively and to better understand the available avenues for individuals to contest adverse, incorrect, or harmful outcomes if they occur.
Name(s) and contact information of senior official(s) *within the agency responsible for the approval and oversight of the system.*	Establishing dedicated personnel or teams responsible for a system is critical for accountability and provides the public and other stakeholders with a specific point of contact for addressing any questions or potential problems.
Description of human oversight *including:* ● *cadence (e.g., ongoing, periodic, annual),* ● *type (e.g., real-time monitoring, human approval of all outputs, randomized review), and* ● *the agency's justification for this level of oversight.*	Human oversight is an especially critical tool for agencies to identify and rectify potential errors. This information allows stakeholders to evaluate whether and how agencies are proactively implementing such safeguards.
Description of the agency's cost-benefit analysis *including:* ● *specific metrics, quantitative or qualitative, that the agency used to assess the costs and benefits of a system; and* ● *the agency's long-term plan and anticipated budget (including staffing, oversight, expected upgrades, etc.) for maintaining the system.*	AI systems can pose significant personnel and financial costs to public resources throughout a system's lifecycle, spanning system acquisition, onboarding, and monitoring. Because system performance can degrade over time due to system drift, agencies will need to dedicate significant resources to maintain connections between multiple systems and to conduct routine upgrades. This information enables greater public accountability about how agencies analyze and justify the costs of a system against the potential benefits that a system may offer to the public.
Description of the agency's plan to decommission or pause *the use of the system if errors, failures, changes, or harms are detected that lead to an unacceptable level of risk or that contradict the agency's prior cost-benefit analysis, including:* ● *the agency's incident response plan, and*	Without a plan in place, agencies may not be well equipped to respond to unforeseen errors without facing significant risk to agency operations or the public. Stakeholders should have access to clear information about how an agency plans to respond to such incidents and prevent potential harms.

• *the specific systems or backup processes that the agency has to replace the AI tool if it is paused or decommissioned.*	

Conclusion

AI use case inventories are a promising tool for governments to understand their own uses of AI across agencies, for private industry to better address public sector needs, and to build public trust about the government's use of AI. In particular, these requirements are an important forcing function to ensure that agencies have gone through a robust process to fully evaluate the benefits and risks of a system in the first place. But the utility of this governance and transparency mechanism is only as good as the quality of information documented and disclosed to the public. As public sector AI adoption grows, it is even more important for governments at all levels to create and adopt comprehensive requirements for AI use case inventories.

*****

Find out more about CDT's work as part of the Equity and Civic Technology team cdt.org/area-of-focus/equity-in-civic-tech/.

The Center for Democracy & Technology (CDT) is the leading nonpartisan, nonprofit organization fighting to advance civil rights and civil liberties in the digital age. We shape technology policy, governance, and design with a focus on equity and democratic values. Established in 1994, CDT has been a trusted advocate for digital rights since the earliest days of the internet.