Best Practices in Public-Sector AI Governance

A PRACTITIONER'S PLAYBOOK

Eric Hysen

Executive Fellow in Applied Technology Policy

APRIL 2026

UC BERKELEY EXECUTIVE FELLOWSHIP IN APPLIED TECHNOLOGY POLICY

UC Berkeley School of Information

UC Berkeley Goldman School of Public Policy

THE EXECUTIVE FELLOWSHIP IN APPLIED TECHNOLOGY POLICY is a groundbreaking partnership between UC Berkeley's School of Information and the Goldman School of Public Policy. This prestigious eight-month program unites the socio-technical expertise of the School of Information with the policy acumen of the Goldman School to prepare distinguished leaders for the future of technology policy. The non-residential fellowship is structured to provide a dynamic platform for policy leaders to reflect on their experiences, mentor the next generation, and document their contributions to digital transformation in government. Fellows will participate in research, teaching, and high-impact meetings and events.

The Executive Fellowship in Applied Technology Policy is grateful for the support of the John D. and Catherine T. MacArthur Foundation, the PIT Infrastructure Fund and the Berkeley School of Information Bellwether Strategic Priorities Fund.

Contents

Executive Summary

Governments across the United States are rapidly adopting artificial intelligence (AI) to improve services — from easing highway congestion and answering tax questions to supporting clinicians and detecting fraud — while facing an evolving and complex risk landscape. This paper explores how governments are harnessing the benefits of AI while managing its risks through an evaluation of the maturing process of public-sector AI governance: the systems of frameworks, rules, policies, standards, and organizational structures that ensure AI is safe, ethical, secure, and aligned with mission and values.

Our analysis is based upon a comparative review of 66 U.S. public-sector AI governance policies, as well as ten in-depth interviews of federal, state, local, and international leaders, along with the author's lived experience as Chief Information Officer and Chief AI Officer at the U.S. Department of Homeland Security, the third-largest federal agency, from 2021–2025.

We find that governance is advancing quickly but unevenly: 43 U.S. states have established some AI governance capacity, yet there remains wide variation in scope, structure, and transparency. At the federal level, the Trump Administration maintained many core governance requirements implemented by the Biden Administration, while shifting the emphasis toward speed of adoption and reducing regulatory burden.

Based on the literature and interviews, we provide a pragmatic playbook for leaders building and implementing AI governance in public-sector organizations, with best practices organized around five governance stages.

Figure 1: Best Practices in Public-Sector AI Governance

1. Policy Development
   a. Start with a transformative vision for what AI will accomplish, not just risk management.
   b. Adopt "minimum viable governance" by anchoring only core concepts in law or regulation.
   c. Refresh policies on a set cycle — but not too frequently.

   d. Leverage existing policy mechanisms as much as possible.
   e. Use standard definitions of AI, and be intentional about scope.

2. Leadership and Resourcing
   a. Put someone in charge, regardless of job title.
   b. Use governance boards for alignment, not micromanagement.
   c. Resource governance as a core function, building on and strengthening IT governance.
   d. Signal prioritization from top leadership.
   e. Build communities of practice to scale expertise.

3. Intake and Inventory
   a. Understand the key difference between use cases and systems.
   b. Intake early and often.
   c. Iterate intake to collect more information for higher-risk or later-stage uses.
   d. Make intake collaborative, not just compliance.
   e. Automate intake by embedding checks in procurement, budget, and security processes.

4. Risk Assessment and Management
   a. Define risk tiering using clear, practical signals.
   b. Develop structured, formal impact assessments and allow self-review for low-risk uses.
   c. Tie risk acceptance to non-technical business owners.
   d. Make risk management continuous, not a one-time checkpoint.
   e. Avoid framing high risk as inherently bad.

5. Publication and Engagement
   a. Invite stakeholders into the policy and implementation process.
   b. Ensure that disclosures accurately reflect the reality of AI plans.
   c. Differentiate transparency for experts versus end-users.
   d. Extend transparency even to sensitive domains like law enforcement and national security.
   e. Sustain ongoing dialogue after publication.

Public-sector AI governance is still in its early days: structures vary, definitions diverge, and many programs remain thinly resourced. The experience of governments that have adopted AI successfully points to a balanced path forward: be ambitious about AI's benefits, disciplined about risk, and humble about the need to iterate. Done well, AI governance becomes an engine for responsible innovation, supporting governments in earning public trust and delivering measurable improvements in people's lives.

Introduction

Governments around the world are accelerating adoption of artificial intelligence (AI) to transform operations and improve public service delivery. AI is already being used for everything from reducing highway congestion and improving customer service for tax questions in California[1] to enhancing clinician effectiveness and better detecting payment fraud at the Department of Veterans Affairs.[2]

But AI also introduces new risks. The Cybersecurity and Infrastructure Security Agency (CISA) has identified three categories of AI risk to government services and other critical infrastructure: attacks using AI, attacks targeting AI systems, and failures in AI design and implementation.[3] These risks start with the challenges inherent to managing any information technology system (such as cybersecurity and system reliability) and expand to include unique AI impacts. When governments do not effectively manage these risks in their own uses of AI, the impacts on constituents' lives can be severe.

Moreover, AI's capabilities and risk landscape are constantly evolving. Two years ago, policy debates on AI safety focused on catastrophic risks like those involving chemical, biological, radiological, and nuclear weapons and the potential use of AI for constructing novel offensive cybersecurity capabilities. Recent reporting and debate have focused on risks of a much more personal nature, including sycophancy[4] — a failure mode that can make models "overly flattering or agreeable," leading them to validate harmful beliefs or risky impulses rather than challenge them — as well as suicide and other impacts of AI use on mental health.[5] These developments highlight that AI risk management must address not only extreme misuse scenarios but also the subtle, cumulative ways in which AI systems can shape human decision-making, wellbeing, and trust.

AI governance enables organizations to harness the enormous potential of AI while managing these ever-changing risks. In this paper, we define AI governance as the system of frameworks, rules, policies, standards, and organizational structures that ensures AI development and deployment are safe, ethical, secure, and aligned with organizational strategy and values. It spans the full lifecycle of AI systems and includes the contributions of technical, compliance, legal, and other stakeholders working to maintain accountability and trust.

1 "Governor Newsom deploys first-in-the-nation GenAI technologies to improve efficiency in state government," State of California, April 29, 2025.
2 "AI Use Case Inventory," U.S. Department of Veterans Affairs, accessed October 10, 2025.
3 "Safety and Security Guidelines for Critical Infrastructure Owners and Operators," U.S. Department of Homeland Security (DHS), April 2024.
4 "Sycophancy in GPT-4o: what happened and what we're doing about it," OpenAI, April 29, 2025.
5 Kashmir Hill, "A Teen Was Suicidal. ChatGPT Was the Friend He Confided In." August 26, 2025.

Background

Public-sector AI governance predates the current generative AI era. Canada was among the first governments to adopt a binding policy on AI use with its 2019 *Directive on Automated Decision-Making*, which requires impact assessments, transparency, and recourse for automated systems.[6] In the United States, the *AI in Government Act of 2020* required the Office of Management and Budget (OMB) to issue guidelines for federal agencies' use of AI, and in December 2020, President Trump issued *Executive Order 13960: Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government*, which required agencies to inventory and publish details about their AI use for the first time.[7]

Interest and activity in this space accelerated dramatically after the November 2022 release of ChatGPT. 2023 saw the release of two significant governance standards: the National Institute for Standards and Technology's AI Risk Management Framework and ISO/IEC *42001: AI Management Systems*, both of which are used to guide governance of AI at public and private organizations.[8] In late 2023, President Biden issued *Executive Order 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*, and in early 2024, OMB released M-24-10: *Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence*, which directed agencies to appoint Chief AI Officers and implement risk-management practices for rights- or safety-impacting uses.[9]

This paper is informed by the author's lived experience implementing these requirements as Chief Information Officer and Chief AI Officer at the U.S. Department of Homeland Security (DHS), the third-largest federal agency, from 2021-2025. DHS moved to fully harness the potential of AI while managing risks, including issuing specific governance policies for generative AI and face recognition, working across critical infrastructure sectors to develop roles and responsibilities for AI safety and security, and ultimately inventorying more than 150 AI use cases across the Department.[10] That process was complex enough to warrant publishing a lengthy explanatory blog post and a detailed public playbook to help other agencies.[11]

6 "Directive on Automated Decision-Making," Government of Canada, last modified June 24, 2025.

7 "H.R.2575 - AI in Government Act of 2020," introduced May 8, 2019; "Executive Order 13960: Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government," December 3, 2020.

8 "AI Risk Management Framework," National Institute of Standards and Technology, January 26, 2023; "ISO/IEC 42001: AI Management Systems," International Organization for Standardization, accessed October 10, 2025.

9 "Executive Order 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," October 30, 2023; "M-24-10: Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," U.S. Office of Management and Budget, March 28, 2024.

10 "Roles and Responsibilities Framework for Artificial Intelligence in Critical Infrastructure," U.S. DHS, November 14, 2024.

11 "AI at DHS: A Deep Dive into our Use Case Inventory," U.S. DHS, December 16, 2024; "DHS Playbook for Public Sector Generative Artificial Intelligence Deployment," U.S. DHS, January 2025.

In 2025, AI governance remains a priority under the second Trump Administration. While President Trump rescinded EO 14110 and OMB M-24-10, in April 2025, OMB issued *M-25-21: Accelerating Federal Use of AI through Innovation, Governance, and Public Trust*, which maintains and streamlines core governance requirements while directing increased focus on AI innovation.[12] State and local governments have also moved quickly to ramp up their AI governance. Code for America's July 2025 *Government AI Landscape Assessment* shows that 43 of 50 states have developed or established AI leadership and governance capabilities.[13] The United States Conference of Mayors' June 2025 report, *Building Your City's AI Strategy: A Roadmap for America's Mayors*, highlights 13 cities across the country that have taken steps to develop AI strategies and governance plans.

Despite its importance, AI governance remains a nascent field. The International Association of Privacy Professionals' *AI Governance Profession Report 2025*, based on a survey of 670 largely private-sector practitioners, notes "there is no clear best practice for how to build and organize an AI governance team.... AI governance is still evolving, and mature AI governance programs continue to find room to innovate."[14]

Research on public-sector practices has largely focused on reviewing the outputs of governance: policies, documentation, and inventories. Code for America's assessment is based on a review of publicly available information. The Center for Democracy & Technology has outlined best practices for AI use case inventories, and the CLeAR framework from the Harvard Shorenstein Center proposes practical documentation for datasets, models, and systems to support transparency and accountability.[15] There have been fewer accounts of the end-to-end processes governments use to design and implement AI governance.

12 "M-25-21: Accelerating Federal Use of AI through Innovation, Governance, and Public Trust," U.S. Office of Management and Budget, April 3, 2025.

13 "Government AI Landscape Assessment," Code for America, July 2025.

14 "AI Governance Profession Report 2025," International Association of Privacy Professionals.

15 Quinn Anex-Ries, "Best Practices for Public Sector AI Use Case Inventories," Center for Democracy and Technology, July 21, 2025; Kasia Chmielinski et al, "The CLeAR Documentation Framework for AI Transparency," May 2024.

Research Objectives and Methodology

This research sought to advance the practice of public-sector AI governance through:

1. Examining existing AI governance practices across multiple public-sector jurisdictions to understand current implementation approaches and their effectiveness;
2. Identifying successful governance strategies that have demonstrated measurable outcomes in enabling responsible AI deployment while meeting operational objectives;
3. Analyzing implementations to understand common pitfalls and systemic challenges that impede effective governance;
4. Developing evidence-based recommendations that provide actionable guidance for public-sector AI leaders seeking to mature their organizational governance capabilities; and
5. Contributing to the broader understanding of how governance frameworks can serve as strategic enablers, rather than administrative burdens, in public-sector AI initiatives.

Following our review of the literature and released assessments of the current state of public- and private-sector AI governance, risk management, and compliance, our research proceeded on two tracks: (1) a comparative policy analysis across jurisdictions and (2) in-depth qualitative interviews with government AI leaders to understand how policies are developed, implemented, and matured in practice.[16]

We collected and analyzed publicly-available AI governance policies from 66 U.S. government jurisdictions, spanning federal agencies, states, and selected cities. Because we relied on publicly available documents, our research may understate informal governance practices or internal guidance not published online.

We found two federal agencies, DHS and the General Services Administration, had issued public formal, mandatory governance policies in addition to OMB. We also found that 43 states had issued some form of governance policy (matching Code for America's number). Of these policies, 11 were enshrined in law, 24 in formal administrative policies, and eight through informal guidance or strategy memos.

City-level data was drawn from municipalities highlighted in the U.S. Conference of Mayors roadmap and supplemented by desk research; the city-level data is meant to provide examples, not a complete picture.

16 Unless otherwise cited, all quotes from current and former government officials come from these qualitative interviews.

This desk research provided a baseline to inform a series of ten in-depth, semi-structured interviews with current and former AI leaders across three federal agencies, five states, one foreign country, and one city. We selected a cross-section of federal agencies and states to ensure the analysis included a wide range of sizes, policy priorities, and technology maturity levels. These interviews provided insight into how policies are actually developed, resourced, and implemented on the ground.

By synthesizing the policy review and interviews, we identified five common stages in the maturation of public-sector AI governance.

Figure 2: Stages of AI Governance

The following sections explore each of these stages. While the process is presented sequentially, effective governance is iterative: insights from publication and engagement feed back into policy development, risk assessments reshape intake processes, and the cycle continues. Governments should expect to revisit earlier stages as their use of AI matures.

Stage 1: Policy Development

AI governance begins with policy frameworks that establish boundaries for responsible use while encouraging innovation. Effective policies articulate broad principles such as fairness, accountability, and transparency, and avoid locking in overly prescriptive requirements that can become outdated. Because AI technologies evolve rapidly, policies are often designed to be iterative, with mechanisms for regular review and refinement. This ensures that rules remain relevant and adaptable as new risks and opportunities emerge.

A second layer of policy development is how governments integrate AI oversight into existing regulatory and administrative structures. Some treat AI policies as standalone instruments, while others fold them into broader data, technology, or service frameworks. Striking the right balance between general coverage and technology-specific rules is an ongoing challenge. Overly narrow policies risk missing important applications, while overly broad ones may slow adoption or confuse agencies. Governments are also increasingly recognizing the importance of defining AI in clear, consistent terms to prevent gaps and overlaps in regulation.

Best Practices

1. **Start with not just risk management, but a transformative vision for what AI will accomplish.** Policies are more effective when they articulate how AI can improve services and outcomes, rather than focusing solely on constraints. Setting a positive vision, including highlighting mission-critical AI use cases and fostering a culture of learning and adoption, helps agencies see governance as an enabler, not a blocker.

   Maryland Governor Wes Moore combined the announcement of his AI Executive Order with other actions to transform the State's digital services, saying, "Together, we will improve the experience that Marylanders have on state websites to make it easier to access state resources; together, we will ensure that all state services are accessible to the public — including individuals with disabilities."[17] Nishant Shah, Former Maryland Senior Advisor for Responsible AI, described how this ethos informed their governance process, saying "We don't want to just be the office of 'no.' The whole point of governance is to help agencies responsibly get to 'yes' on AI projects that can actually make a difference for residents."

2. **Adopt "minimum viable governance" by anchoring only core concepts in law or regulation.** Because technology evolves quickly, binding statutes or executive orders

17 "Governor Moore Announces Action to Transform Maryland Executive Branch Digital Services," State of Maryland, January 8, 2024.

should establish high-level principles, while detailed requirements should live in flexible guidance that can be updated as needed. This avoids locking governments into rules that age poorly.

As Vanitha Zacharias, Chief AI Strategist and Administrator of Technology Governance for the State of Ohio, explained in an interview, "We are using a minimum viable governance approach to enable quicker adoption. Our process is multifaceted: low-impact projects are self-governed and additional layers of governance are applied for high-risk or high-impact projects."

3. **Refresh policies on a set cycle — but not too frequently.** AI governance benefits from predictability and continuous improvement. Regular updates keep frameworks current, but overly short cycles can overwhelm agencies with constant updates.

   Stephen Burt, Canada's Chief Data Officer, noted, "In the digital space … it doesn't work very well [to never change policy], so building in expiry dates, or some kind of a cycle of refresh, is really important." However, Canadian government leaders found their AI policy's original six-month refresh cycle to be excessive in practice, and they ultimately settled on a two-year cycle.

   Governments can also use pilots to stress-test governance requirements before formalizing them. Running a small number of AI projects through draft governance processes and deliberately capturing friction points generates real-world evidence for policy refinement. The U.S. Department of Homeland Security created responsible use working groups for each of its initial cohort of three generative AI pilot projects, where cross-functional teams helped design responsible use approaches based on the unique needs of each pilot before formal governance requirements existed. These efforts then helped inform formal policy development.

4. **Leverage existing policy mechanisms as much as possible.** Embedding AI governance into established systems makes compliance easier and reduces duplication. Using formats agencies already know — such as IT directives, procurement checklists, or existing governance boards — helps normalize AI oversight as part of day-to-day government operations.

   As Matthew Graviss, former Chief Data and AI Officer at the U.S. Department of State (now the Chief Technology Officer (CTO) for the Public Sector at Atlassian), explained: "The major policy handbook [at the State Department] is the Foreign Affairs Manual … we established [our AI policy] into it. Our folks were already best friends with the people who manage the Foreign Affairs Manual, so we could start pumping AI policy through that system."

5. **Use standard definitions of AI, and be intentional about scope.** Definitions set the scope for governance. Relying on widely recognized standards helps avoid confusion across agencies, vendors, and legislatures. The most commonly used definition of AI, for example, is found in the National Artificial Intelligence Initiative Act of 2020 (and is also adopted by NIST): "The term 'artificial intelligence' means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments."[18]

   Equally important is deciding up front whether policies will apply to all types of AI, including machine learning and predictive analytics, or more narrowly, for example to generative AI only. Nikhil Deshpande, Chief Digital and AI Officer for the State of Georgia, emphasized this point: "We relied on the NIST definition, but across the street in our legislature, they came up with their own version. I had to make a few trips and work with senators to wordsmith it so we'd align. They were defining AI only from the lens of GenAI, and we had to explain it's a suite of technologies, including NLP, computer vision, and machine learning."

By the numbers

Table 1: Policy vehicles

	States	Cities
AI governance enshrined in law	11	
Issued formal, binding administrative policies, such as Executive Orders or Directives	24	6
Issued informal, nonbinding guidance or strategies	8	2
We could not identify any policy	7	

Table 2: Policy scope

	States	Cities
Policies apply equally to both traditional and generative AI	20	3
Policies apply to both, but with additional GenAI requirements	14	2
Policies apply only to generative AI	9	3

18 National Artificial Intelligence Initiative Act of 2020, Title 15 U.S.C Section 9401 (2020), accessed October 10, 2025.

Stage 2: Leadership and Resourcing

Policy documents need champions and institutional backing to have impact. Governments often designate senior officials to lead AI governance, supported by governance boards or councils that convene stakeholders from across legal, security, technology, and program areas. These leaders provide accountability, coordinate across fragmented bureaucracies, and keep attention on long-term priorities. Clear leadership signals from cabinet-level officials or senior executives can help embed AI governance as an organizational priority rather than a niche technical issue.

The resourcing side of leadership is just as critical. Many governance teams are small and stretched thin across intake, review, training, and policy development. Without adequate funding and personnel, even the best AI governance frameworks may be perceived as barriers rather than enablers. Building durable capacity means staffing beyond a single champion, investing in training across agencies, and creating forums for practitioners to share expertise. Over time, governments are finding that governance must be treated as a core function of technology management, not an unfunded add-on.

Best Practices

1. **Put someone in charge, regardless of job title.** AI governance requires a designated leader who can coordinate across fragmented bureaucracies and ensure accountability. While the Biden AI Executive Order required every federal agency to name a Chief AI Officer, most agencies dual-hatted an existing official, such as a Chief Information Officer, Chief Technology Officer, or Chief Data Officer.

   States and cities have followed similar approaches, often assigning AI leadership responsibilities to existing technology or policy officials rather than creating new roles. What matters most is not the title but clear ownership: someone who has the authority and accountability to drive governance across the organization.

   Jonathan Mayer, former Chief Technology and AI Officer at the U.S. Department of Justice, explained that at first, “the department wasn’t configured to do the kind of AI governance it should be doing”, due to its highly decentralized approach to technology management. His appointment directly in the Attorney General’s office empowered him to coordinate across the bureaucracy.

2. **Use governance boards for alignment, not micromanagement.** Every federal agency is required to have an AI Governance Board, and most states have created similar bodies. These groups are valuable in bringing together diverse stakeholders — from IT, legal,

privacy, civil rights, and mission leaders — to share perspectives and align strategy. But these boards work best as deliberative forums, overseeing decision-making at a high level, with smaller teams empowered to handle day-to-day operational decisions.

Josh Williams, former Senior Enterprise Data Architect with the Colorado Governor's Office of Information Technology, described how state leaders repurposed and strengthened an existing forum, the Government Data Advisory Board, to include AI oversight. The Board includes representatives from across state agencies and has helped inform the implementation of their AI governance process.

3. **Resource governance as a core function, building on and strengthening IT governance.** Governance should be staffed and budgeted as a permanent function, not treated as an unfunded mandate. The most effective approaches embed AI into existing IT or data governance processes with which agencies are already familiar, while still dedicating permanent resources. If such processes do not already exist or are immature, AI governance offers an opportunity to build and strengthen them.

   Dave Larrimore, former Chief Technology Officer in the U.S. Department of Homeland Security, explained that he, "carved out about 30% of my team that already performed governance on IT investments and programs across the department to focus on AI." This gave them a head start as they could leverage existing relationships and processes, but also created a downstream challenge with capacity. "We never really had the people to operationally run the governance," Larrimore said. "The same folks doing oversight kept getting pulled onto the next urgent thing, and over time governance just lost priority."

4. **Signal prioritization from top leadership.** When senior leaders consistently reinforce that AI governance is a priority, it changes how seriously agencies engage with the process. Political and executive backing provides legitimacy and urgency that no policy memo alone can achieve.

   In the U.S. Department of Justice, Mayer credits the role of former Deputy Attorney General Lisa Monaco, who "came to our AI governance meetings and made it very clear: this is a priority, we need to take it seriously, we need to do it right."

5. **Build communities of practice to scale expertise.** Central teams cannot carry the full burden of governance. Practitioner networks, training cohorts, and regular convenings help spread expertise and create a culture of shared responsibility.

   California State Chief Technology Officer Jonathan Porat described their monthly "AI Community" meetings: "We had 200–400 people showing up regularly … it created a

space where employees could surface use cases and share lessons with each other, making governance a collaborative effort rather than a top-down directive."

By the numbers

Table 3: Governance Boards and Committees

	States	Cities
Formal governance bodies with decision-making power	24	2
Advisory-only committees	6	4
No identified governance or advisory board	20	2

Stage 3: Intake and Inventory

Centralized submission and review processes give governments visibility into proposed and ongoing AI uses. By requiring agencies to submit applications for review prior to using AI, they establish a single pathway to assess risk, encourage responsible planning, and prevent uncoordinated experimentation. Inventories built from these processes offer a comprehensive picture of AI adoption and help identify common use cases that may benefit from shared standards or pre-approved solutions. Though administratively demanding at first, inventories create the foundation for transparency and long-term governance.

Beyond adding visibility, intake and inventory systems also shape culture and practice. They signal to agencies that AI cannot be adopted informally and encourage early consultation with governance bodies. The design of intake processes matters: overly burdensome requirements risk discouraging disclosure, while overly permissive review processes may fail to surface risks. Governments are increasingly exploring phased intake approaches, lighter requirements for low-risk experimentation, and digital platforms that reduce paperwork and streamline compliance. Done well, intake and inventory processes not only track adoption but also serve as enablers of safe innovation.

Best Practices

1. **Understand the key difference between use cases and systems.** Traditional technology governance looks primarily at IT systems, with clear boundaries where one system ends and another begins, while AI governance frequently, but not always, looks at use cases. These are not synonymous: AI use cases may involve one or more IT systems, and one IT system could be a part of multiple AI use cases. Use cases are defined based on how AI is used in the context of broader operations, not system boundaries. Decide which one you are inventorying and build processes around that. AI governance processes organized around systems are generally more focused on technical capabilities and outputs, while processes organized around use cases are centered on how people use and are affected by AI.

   At the Department of Homeland Security, a single face recognition system was used in multiple contexts, from airports to border checkpoints. After the initial inventory covered a single, overly broad use case, the team divided the system into over two dozen different uses for every possible variation. This proved both unwieldy internally and hard to explain to stakeholders, and officials finally landed on six use cases based on common operating environments and user expectations.

2. **Intake early and often.** The most effective governance processes encourage agencies to submit AI projects as early as possible, not just as part of a pre-launch review. Keeping intake low-burden enables visibility across the lifecycle and helps identify risks and opportunities before systems are too far along.

   Nikhil Deshpande described the approach used by the State of Georgia: "If you have an idea, if you have a problem, if you have a vendor knocking down your door, trying to sell you something, just fill out this form and talk to us, and then collectively, we'll come together to see if this truly has potential."

3. **Iterate intake to collect more information for higher-risk or later-stage uses.** Not every AI project warrants the same level of scrutiny. Many governments use tiered intake, where early pilots provide basic details and higher-risk or production systems must submit full documentation and risk assessments.

   Federal agencies are required to collect 12 basic pieces of information for all AI use cases, such as the intended purpose and AI outputs. Use cases at later stages of development must include up to 22 additional data points "regarding development, data and code, and AI enablement and infrastructure," and the highest-risk use cases require 13 final fields "regarding the AI's compliance with the risk management practices."[19]

4. **Make intake collaborative, not just compliance.** When intake is framed as a pathway to advice and support, agencies are more willing to disclose projects. Treating intake as an enabler, not a barrier, helps governance teams build trust and surface more use cases.

   According to Nishant Shah, Maryland deliberately worked to ensure that intake was viewed "not as a blocker but as a tool for risk categorization and responsible AI enablement." State AI leaders embarked on an outreach tour to meet agencies where they were and help them navigate requirements. This helped build trust, encourage disclosure, and prevent AI uses from slipping past the oversight process.

5. **Automate intake by embedding checks in procurement, budget, and security processes.** Voluntary submissions alone will miss "shadow AI" that agencies adopt informally, outside official channels. Stronger systems catch use cases by embedding AI checks into existing processes like procurement, budget, and cybersecurity reviews, ensuring agencies cannot bypass governance.

19 "Guidance for 2024 Agency Artificial Intelligence Reporting per EO 14110," U.S. Office of Management and Budget, August 14, 2024.

At the State Department, Matthew Graviss explained, "we injected AI reviews into the cybersecurity authorization process, into the budget request process, and into procurement. Those are three of the key ways to capture AI pursuits across the entire department. We even used GenAI to scan budget requests from all 35–50 bureaus to tease out any mentions of AI."

By the numbers

Table 4: Use Case Inventory Requirements

	States	Cities
Inventories all AI use cases	30	6
Inventories only select AI uses	5	1
No inventory requirement identified	15	1

Stage 4: Risk Assessment and Management

Risk frameworks help governments allocate attention and resources where they are most needed. By categorizing AI uses into risk tiers, governance systems can identify where to apply proportionate oversight. Low-risk use cases may proceed with minimal review, while high-risk uses trigger deeper assessments, privacy and security checks, or senior-level approval. This approach allows innovation to move quickly where stakes are low, while ensuring caution in areas with potential impacts to individuals' rights or safety.

Managing risk should be an ongoing process, not just an upfront approval step. Governments are developing tools such as algorithmic impact assessments, mitigation plans, and requirements for human oversight to manage risk throughout a system's lifecycle. The challenge is ensuring that risk processes are rigorous enough to build trust without becoming so complex that they stall adoption. When risk assessments are poorly structured — with overly broad risk categories, inflexible approval processes, or excessive documentation requirements — they can stifle innovation and experimentation in precisely the areas where governments most need to learn by doing. Iterative reviews, ongoing monitoring, and continuous refinement of risk categories are emerging as key practices.

Best Practices

1. **Define risk tiering using clear, practical signals.** Effective AI governance avoids abstract criteria and instead uses practical signals — such as data sensitivity, external exposure, or decision criticality — to classify risk. This approach helps agencies focus their limited resources on the highest-impact cases without overwhelming the system. It also reduces the need for non-experts to make complex judgments about rights or safety impacts, instead surfacing the information domain experts need to make those determinations effectively.

   As an example, leaders from Ohio developed a "qualifier heat map," which maps AI use cases based on "data classification, scope (internal versus external), potential impact of the decisions made (legal, financial, human resources, legislative, reputational, social, or regulatory), and ties to a critical state function." In the federal government, the Biden-era requirements in OMB Memo M-24-10 created classifications of "rights-impacting" and "safety-impacting" AI, but even with lengthy rules, agencies found these difficult to consistently identify in practice. The updated requirements in Memo M-25-21 issued in 2025 opted instead for a simpler "high-risk" tier.

2. **Develop structured, formal impact assessments and allow self-review for low-risk uses.** Formal assessments like algorithmic impact evaluations provide a consistent way to surface bias, transparency, and recourse issues. But they should be scaled, with full reviews for high-risk cases, and lighter-touch or self-assessments for low-risk uses.

   Canada's Algorithmic Impact Assessment tool is likely the first example of this, and their website also includes a library of completed assessments.[20] California CTO Jonathan Porat described a process that empowers agency technology leaders to complete and certify a structured GenAI Risk Assessment, and only requires consultation from the California Department of Technology for moderate or high-risk use cases.[21]

3. **Tie risk acceptance to non-technical business owners.** Unlike traditional IT risks, many AI risks are reputational, ethical, or mission-specific, including potential impact to individuals' rights and safety. Decisions about acceptable risk in these areas can involve a more significant role for subject-matter experts, not just technologists.

   At the U.S. State Department, for example, Matthew Graviss recalled that he "didn't feel comfortable authorizing the use of AI for particular use cases in a mission space that belonged to other executives. For example, the Under Secretary for Management had the say-so on letting employees use generative AI for performance appraisals, and the Office of Language Services" determined that AI use was acceptable for translating a local news article but not for a treaty.

4. **Make risk management continuous, not a one-time checkpoint.** AI systems can drift, degrade, or introduce new risks after deployment. Leading practices include ongoing monitoring, regular status updates, and repeat reviews.

   Vanitha Zacharias described how, in Ohio, quarterly status updates are used to update the AI inventory, and the information is then shared with agencies. When new information emerges that affects the categorization of a use case, it opens the door for further assessment and review, ensuring that evaluations remain accurate and effective.

5. **Avoid framing high risk as inherently bad.** High-risk use cases are often among the most valuable applications of AI if managed responsibly. But without careful communication, agencies may interpret "high risk" as "too risky," leading them to avoid precisely the projects they should be taking on. In bureaucracies, where incentives reward compliance over innovation, extra oversight can feel like punishment. Labeling a project as high risk may trigger longer reviews or political attention, prompting

20 "Algorithmic Impact Assessment tool," Government of Canada, accessed October 10, 2025.
21 "GenAI risk assessment workflow at a glance," State of California, accessed October 10, 2025.

officials to seek the path of least resistance — redefining, reclassifying, or abandoning complex uses altogether.

When asked about low numbers of identified high-risk use cases, some of the leaders interviewed cited agencies' early stage of technical maturity, or appropriate self-restraint, which requires teams to ask whether AI is truly necessary for a given problem. But others questioned whether the numbers were low because risk classification systems were being gamed, or whether the policies themselves discouraged consideration of AI in sensitive but high-impact domains. Effective governance reframes risk as a neutral descriptor, not a judgment. It treats high-risk use cases as priorities for deeper support and mitigation.

By the numbers

Table 5: Risk Management Practices

	States	Cities
Requires the same risk management practices for all AI use cases	11	3
Requires risk management practices for all AI use cases, with stricter practices for high-risk uses	25	4
Requires risk management practices only for high-risk AI use cases	6	1
No risk management requirement identified	8	

Stage 5: Publication and Engagement

Transparency builds trust in government AI use. Many governance frameworks include requirements to publish inventories of AI systems, summaries of assessments, or explanations of how risks are managed. Public disclosure allows stakeholders to understand where AI is being applied and what safeguards are in place, and it provides a check against misuse or overreach. Plain-language tools and internal communication materials also help employees and constituents understand rules and expectations.

Engagement with external stakeholders adds another dimension to transparency. While some governments have developed policies largely internally, others have actively involved civil society, academia, and affected communities in shaping their frameworks. External engagement can surface concerns that might otherwise be overlooked and strengthen legitimacy. Governments face tradeoffs between moving quickly and spending time consulting broadly with external stakeholders, but over time, openness to feedback is proving critical. Publishing not only lists of systems but also evaluation results, governance standards, and lessons learned helps create a virtuous cycle of trust and accountability.

1. **Invite stakeholders into the policy and implementation process.** Transparency extends beyond publishing inventories of AI uses. Opening AI policy drafts to employees, civil society, and academia increases legitimacy and surfaces issues and edge cases that a government alone might miss. This approach also helps develop internal champions who understand and can explain the jurisdiction's approach to AI to their communities.

   Public participation can extend into the design of individual AI use cases. Gathering community input during the risk assessment process allows the surfacing of concerns from affected populations before deployment, rather than after. This is particularly important for uses that touch sensitive domains like housing, benefits, or enforcement, where communities may identify risks that internal reviewers miss.

   In the City of Boston, CIO Santiago Garcés described his approach to developing guidelines for the use of generative AI: "I put an initial draft out, and then started adding people as commenters and reviewers: community members, academics, employees. We ended up with hundreds of comments and suggestions. It really became a community-shaped document."

2. **Ensure that disclosures accurately reflect the reality of AI plans.** While inventorying ideas and early-stage projects can be valuable, publishing about AI uses without clarifying their status can mislead stakeholders into thinking that exploratory, unrefined concepts are real, funded programs. This challenge can be managed by including clear

status information or establishing specific criteria so that ideas become public only once they are concrete enough to be understood.

At the Department of Homeland Security, the first inventories of AI uses collected dozens of ideas that were never resourced or, in some cases, never even approved for early consideration. Leaders speculated this occurred because program staff incorrectly believed that submitting these ideas to the inventory could help them gain resourcing support. This led to unnecessary confusion, so the team updated the inventory to "clearly indicate which use cases are resourced and approved, and which were only ideas or in research and development."

3. **Differentiate transparency for experts versus end-users.** AI governance outputs support a range of activities throughout the AI lifecycle, from evaluation and oversight to day-to-day use. Not all audiences and uses need the same depth or style of disclosure. Employees and the general public benefit from plain-language guidance that is easy to act on, while auditors and civil society need technical documentation and detailed evaluations. Layering and tailoring disclosure for different audiences makes governance usable in practice without losing rigor for expert review.

   Solomon Abiola, Former Director of AI Policy and Governance for the State of Maryland, experimented with a novel approach of creating "governance cards" for state employees that include abbreviated rules and guidance for usage of a specific AI system, and ensuring employees can access these quickly and easily before they begin AI use.

4. **Extend transparency even to sensitive domains like law enforcement and national security.** While domains such as law enforcement and security may not always allow for revealing all operational details, governments can still disclose purpose, safeguards, and oversight mechanisms related to the use of AI. Leading jurisdictions have moved from a posture of assuming law enforcement use cases of AI cannot be shared to instead defaulting to at least partial disclosure, and have found that this strengthens their safety and security missions by building public confidence in their work.

   Jonathan Mayer explained that, at the Department of Justice, some law enforcement leaders started with the view that they could not disclose anything about certain AI use cases, but after review (including identifying other areas where those same uses were already public), they were able to include them in their inventory.

   The Department of Homeland Security's 2024 inventory made maximizing transparency a priority: "The AI Governance Board directed that [the Department] disclose as many AI use cases as possible, even if some details about certain use cases cannot be publicly shared. We re-reviewed previously undisclosed sensitive use cases and deter-

mined that we can now disclose at least some information about every AI use case subject to M-24-10."[22]

5. **Sustain an ongoing dialogue after publication.** Publishing an AI use case inventory should be viewed as the start of a conversation, not the end of a transparency process that is not revisited until the next update. Continual engagement about AI uses can help clear up questions and build trust.

   The White House Office of Science and Technology Policy set a clear tone by inviting civil society leaders to share their views at the very first meeting of the Federal Chief AI Officers Council. And California CTO Jonathan Porat described deliberately building a continual process of engagement with external stakeholders throughout their AI governance journey.

By the numbers

Table 6: Transparency Requirements

	States	Cities
Requires publication of all AI use cases	1	
Requires publication of only select AI use cases	32	5
No transparency requirement identified	17	5

22 "AI at DHS: A Deep Dive into our Use Case Inventory," U.S. DHS, December 16, 2024.

Conclusion

AI governance in the public sector has moved from concept to practice with remarkable speed. Yet the adoption of policies around AI use is still in its early days: oversight structures still vary across governments, definitions diverge, and many programs remain thinly resourced. The experience of leading jurisdictions points to a balanced path forward: governments should be ambitious about AI's benefits, disciplined about risk, and humble about the need to iterate.

There is no single "right" organizational chart or canonical process for AI governance. The jurisdictions progressing fastest are those that (1) articulate a mission-first vision, (2) designate accountable leadership and provide resourcing, (3) keep the processes lightweight but reliable, (4) use risk assessments to focus and prioritize their efforts, and (5) publish and engage continuously. Future work can expand the development of cross-jurisdictional standards (e.g., common intake fields, shared risk indicators, and reusable assessment templates), strengthen integration with existing IT governance and cybersecurity frameworks, and improve mechanisms for continuous monitoring and evaluation of deployed AI systems.

Practitioners can take immediate action by benchmarking their current programs against the best practices outlined in this paper, identifying the two or three gaps with the highest leverage, and prioritizing those for their next update. For those just starting, the most reliable foundation remains: designate a leader to oversee the AI governance process, stand up an easy-to-use intake process for new AI uses, and publish what you learn.

Done well, AI governance becomes an engine for responsible innovation, supporting governments in earning public trust and delivering measurable improvements in people's lives.

Acknowledgments

I am grateful to the UC Berkeley Executive Fellowship in Applied Technology Policy program and its sponsors — the Goldman School of Public Policy and the School of Information — for providing the necessary funding and administrative support for this project. Thank you to Deirdre Mulligan for her leadership, the opportunity to participate in the fellowship, and being an invaluable collaborator on responsible AI adoption, both while we served together in government and at UC Berkeley. Thank you also to Dan Zhukov for his support throughout the fellowship.

Special thanks go out to my research assistants — Prakash Krishnan, Omar Morales, and Isabelle Qian — for their contributions to this project's planning, research, drafting, editing, and final publication.

Thank you to the current and former government AI leaders who were interviewed for this project and contributed invaluable insights:

- U.S. Department of Homeland Security: David Larrimore, Former Chief Technology Officer; Anil Dewan, Former AI Task Force Director
- U.S. Department of State: Matthew Graviss, Former Chief Data and AI Officer
- U.S. Department of Justice: Jonathan Mayer, Former Chief Technology Officer
- State of California: Jonathan Porat, Chief Technology Officer
- State of Colorado: Josh Williams, Former Senior Enterprise Data Architect
- State of Georgia: Nikhil Deshpande, Chief Digital and AI Officer; Donna Sumner, Operations and Delivery Manager
- State of Ohio: Vanitha Zacharias, Chief AI Strategist and Administrator, Investment andTechnology Governance; Kaylani Thota, IT Enterprise Architecture Manager; Brooke Speert, Senior IT Policy Analyst
- State of Maryland: Nishant Shah, Former Senior Advisor for Responsible AI; Solomon Abiola, Former Director, AI Policy and Governance
- Government of Canada: Stephen Burt, Chief Data Officer
- City of Boston, MA: Santiago Garcés, Chief Information Officer

Partway through my fellowship, I began a new full-time position at Salesforce as Senior Vice President, Chief AI & Transformation Officer for Legal & Corporate Affairs. I am grateful to

Salesforce for supporting my completion of this independent research. This white paper and other research products do not reflect the views of Salesforce.

We made use of generative AI tools for interview transcription, document analysis, drafting, and editing. All AI outputs were reviewed and edited by a human before inclusion in this white paper or other research products.

About the Author

ERIC HYSEN is an Executive Fellow in Applied Technology Policy at UC Berkeley's Goldman School of Public Policy and School of Information. This white paper is the product of his fellowship research on public-sector AI governance.

He currently serves as Senior Vice President and Chief AI & Transformation Officer for Legal & Corporate Affairs at Salesforce, where he leads efforts to responsibly integrate AI into the company's global legal and corporate functions.

Previously, Eric served as the Chief Information Officer and the first Chief AI Officer at the U.S. Department of Homeland Security from 2021 to 2025, overseeing $10 billion in annual IT spending and a workforce of more than 10,000 technology employees and contractors. In that role, he developed DHS's AI Roadmap, launched its first generative AI pilots, hired 50 AI specialists through the DHS AI Corps, and established governance and risk management practices across more than 160 AI use cases. He shared lessons from that work through DHS's public-sector Generative AI Playbook, on which this white paper builds.

Earlier in his career, Eric was a founding member of the U.S. Digital Service at the White House, helped launch the California Office of Digital Innovation, and worked as a software engineer and product manager at Google. He is a graduate of Harvard University.

UC BERKELEY EXECUTIVE FELLOWSHIP

IN APPLIED TECHNOLOGY POLICY

UC Berkeley School of Information

UC Berkeley Goldman School of Public Policy