



Governing AI where it meets people.

Turning responsible AI principles into practical governance for cities and public institutions.

PRESENTED BY Jaime Wascalus • Chief Information Officer • GovAI Coalition Board Member

LEANDER, TEXAS

A fast-growing city.
No AI standards.

Then its CIO found 3,000 peers already
building the playbook.

0

TO

90

LEANDER DID NOT HAVE TO START ALONE.

A real city need

A peer network

A stronger starting point

0

TO

90

MEET THE PRACTITIONER

Jaime Wascalus

Chief Information Officer

City of Saint Paul, Minnesota

**GovAI Coalition Board Member &
Advisor, Use Cases Committee**

**One
practitioner
with the
experience of
3,000 more
behind her.**

Built inside government.
Tested in the real world.
Shared across
jurisdictions.

The work begins with public purpose.

AI governance is a part of building a city that works for everyone.

**SAINT PAUL,
MINNESOTA**




“Join me in building a city that works—for all of us.”

Former Mayor Melvin Carter, City of Saint Paul

From a pilot to shared public infrastructure.


Three years of practitioner demand created the foundation for an independent institution.



The next chapter: Global connection.

A network becomes valuable when experience travels.

3,000+ government practitioners and partners.



Rooted in local government.

Relevant far beyond it.

The network has begun connecting operating experience across borders.



THE IMPLEMENTATION GAP


**AI is evolving faster
than government's
ability to govern it.**

From tools that draft → agents that act

77%

**of coalition
members want
governance
and risk
frameworks
for agentic AI
before they
scale.**

GovAI is a practitioner-led operating system for collective learning.



HOW IT BEGAN: A REQUEST FOR INFORMATION.

In the heart of Silicon Valley, a small team asked a giant for some facts.



Three years later, San Jose, CA, was overwhelmed by response to an AI RFP.

Practitioner-led by design.

The people writing the standards are the people who will use and defend them.



Working committees • Monthly practice • Taskforces for emerging needs

Proof that standards can be usable.

Three shared tools replace three recurring reinventions.

01	AI FACTSHEET	One plain-language governance questionnaire	Instead of every agency asking vendors different questions.
02	AI REGISTRY	Peer reviews and reusable vendor submissions	Instead of every government vetting alone.
03	AI SBOM	A machine-readable ingredients list	Instead of technical risk staying invisible.

The standard is successful when a smaller agency can actually use it.



AI Governance Starter Guide

Step 1

Review the 0-to-90 AI Governance framework

The GovAI Coalition's plug-and-play resources can help you take your AI governance program from 0 to 90. The framework provides a high-level overview of the Coalition's governance resources and how they can be used in conjunction with one another.

Governance Framework

1

How do I prepare my agency to use and govern AI?

- AI Policy
- AI Governance Handbook
- AI Upskilling Course

2

What can I do with AI that adds value?

- Public Procurement Government AI Applications
- Quick AI Assistants for Government Agencies
- Use Cases Repository

3

Who can I work with for good, responsible AI solutions?

- AI Registry
- AI Contract Hub

4

How do I ensure vendors follow my expectations?

- Vendor Agreement
- Guide to Measuring AI Performance

5

How do I prepare for when things go wrong?

- AI Incident Response Plan Template

Step 2

Identify the AI-related roles in your organization.

Identify the people that fit the following roles: someone who handles cybersecurity, someone who handles legal issues, someone who handles purchasing, someone who oversees technology, and someone willing to be the organization's AI expert. The AI expert doesn't need to be an expert on AI today, but can gain the necessary skills through the Coalition.

Note on AI expert:

The AI expert will typically be someone in the agency who is close to data and/or privacy initiatives. For example, a data manager, privacy officer, tech policy lead, or other person willing to learn would suit the AI expert role.

Step 3

Align on the AI Policy within your organization

Discuss the AI Policy with relevant stakeholders. Align with your executive sponsor and broader leadership on your agency's responsible AI principles and overall approach to AI governance.

Note on AI Policy:

Consider what your organization's AI "red-lines" should be. Are there any applications of AI that do not align with your organization's values? If so, confirm with your departments that your organizations is not currently implementing those use cases, and add them to the Prohibited Uses section of the AI Policy. You can reach out to relevant departments, particularly Public Safety, using our [template](#).

SEPTEMBER 2024

Guide to Measuring AI Performance

This guide aims to equip those assessing AI systems in the public sector with an understanding of the possible metrics at their disposal to evaluate AI performance.

GovAI
COALITION

How do vendors measure performance?

Vendors often use traditional AI metrics when assessing their tools. These metrics are based on established AI model types like regression (continuous outputs), classification (categorical outputs), and natural language processing (e.g., LLMs). Figure 1 outlines for each model type: a simple definition, common metrics, and use case examples.

As you view the chart, keep in mind that many AI systems draw from multiple model types, so multiple metrics across model types may be needed in their assessment.

Model type	Definition	Common metrics	Examples
Classification	Used to categorize data into predefined labels or classes.	• Accuracy • Precision • Recall • F1 Score • F2 Score • F-beta Score • Area Under the Curve-Receiver Operating Characteristic Curve (AUC-ROC)	• Medical diagnosis • Sentiment analysis • Spam email detection • Object detection
Regression	Used to find a relationship between variables and continuous numerical outcome.	• Mean Squared Error (MSE) • Root Mean Squared Error (RMSE) • Mean Absolute Error (MAE) • R-squared	• House prices • Projected sales • Cyber risk • Weather patterns • Future stock prices
Clustering	Used to categorize data without predefined labels.	• Silhouette score • Dunn index • Rand index	• Market segmentation • Search results • Anomaly detection • Identifying themes

Community Engagement Handbook

FOR ARTIFICIAL INTELLIGENCE



2–4 Weeks Before Engagement: Outreach and Promotion

Announce the engagement activities through local government city websites, social media, neighborhood associations, community partners, and flyers. Partner with trusted intermediaries—such as libraries, schools, or community-based organizations—to reach groups that might not otherwise participate. Be clear about the timeline, purpose, and how feedback will be used.

Engagement Period

Carry out the planned activities. This may include town halls, working groups, pop-up events, surveys, or participatory design sessions. Meet communities in their own spaces and document feedback thoroughly.

0–1 Month After Engagement: Synthesis and Reporting Back

Analyze the input received, identify key themes, and summarize what was heard. Share findings publicly and directly with participants, making the connection between their contributions and decision-making explicit.

1–2 Months After Engagement: Decision Integration and Follow-Up

Incorporate community input into program design or policy recommendations. Communicate next steps to the public and evaluate the engagement process to refine future efforts.

Longer-Term Engagement

Major initiatives, such as infrastructure or agency-wide technology deployments, may involve multiple rounds of engagement over a year or more, with iterative feedback loops and ongoing participation.



GovAI Coalition Publication

Upskilling Training Playbooks Series: Getting Started with AI Training in Your Agency

Version 1.0

Upskilling Working Group
GovAI Coalition Readiness & Adoption Committee
GovAI Coalition

This publication is available free
of charge at xxxxxxx



Key Considerations Before Implementing AI Training

Before implementing AI training, agencies should assess **organizational readiness for AI based on existing governance, policies, and values**. While training is a low-risk entry point for engaging with AI, it is most effective when grounded in clear goals, basic governance, and an understanding of organizational context.

Key considerations include clarifying what types of AI tools will be used for training, what guardrails already exist (or are needed), and how training aligns with broader workforce and technology strategies. Agencies should also consider how AI training fits alongside existing policies related to data privacy, security, records management, and acceptable technology use.

It is important to recognize that AI training does not require full AI implementation or procurement decisions to be finalized. However, even foundational training benefits from clarity around intent, whether the goal is risk awareness, productivity support, workforce readiness, or longer-term innovation.

Agencies may also want to consider:

- Whether leadership and key stakeholders share a common understanding of AI goals and risks.
- How training will be communicated to staff to set appropriate expectations.
- What internal teams (e.g., IT, legal, HR) should be involved or consulted.
- How training will be evaluated and updated over time.

Addressing these considerations early helps ensure that AI training efforts are purposeful, credible, and aligned with public-sector values.

Readiness and Governance Resources:

- **Government AI Coalition:** [AI Governance Handbook](#) Supports early thinking about guardrails, roles, and oversight.
- **National Institute of Standards and Technology (NIST):** [AI Risk Management Framework](#) Helps agencies consider risk even before tools are

WHAT'S NEXT

Govern agentic AI before it scales.

The Agentic AI Governance Taskforce

- Policy
- Governance handbook
- Incident response
- Procurement language
- Training

THE TEST

**Can a city attorney
put it into an RFP?**

**Can a CISO use it
before an incident?**

**Can a smaller
agency adopt it?**

Government-led does not mean government-only.

The Coalition protects practitioner authority while opening the work to partners who can help it scale.

GOVERNMENT

Defines the need and owns the decision

CIVIL SOCIETY

Builds trust and community connection

INDUSTRY

Responds to shared standards and evidence

PUBLIC

Collaborates on responsible AI adoption

PHILANTHROPY

Partners to build common goods and shared

capacity

Three ideas to carry forward.

01

FROM PRINCIPLES TO PRACTICE

Responsible AI moves through usable policy and contract language.

02

FOR GOVERNMENT, BY GOVERNMENT

The work is written by government practitioners who must implement and defend it.

03

GLOBAL CONNECTION

Local operating experience should plug into global standards.

Three steps to continue the work.

Move from conference connection to practitioner exchange.

01 CONNECT

Link the ITU-T Focus Group with GovAI Coalition practitioners.

02 ADAPT

Use the open templates, then improve and return them.

03 JOIN

Join us at the September 30 GovAI Coalition all-member meeting.

AN INVITATION

**Connect global standards
to real-world practice.**



Proposed next step

**Create a direct working connection between
the GovAI Coalition's practitioner network and
the ITU-T Focus Group on AI for Smart
Sustainable Cities.**



3,000+

**government
practitioners
ready to
learn, share
and build
with you.**

Thank you.



Jaime Wascalus

City of Saint Paul, Minnesota, USA

**GovAI Coalition Board Member &
Advisor, Use Cases Committee**

Jaime.Wascalus@ci.stpaul.mn.us