




THE AI COWBOYS FOUNDATION

RESEARCH BRIEF 01 · GOVERNMENT

Responsible AI Adoption in the Public Sector



A plain English brief for government teams, grounded in NIST guidance and what Texas already requires.



Responsible AI Adoption in the Public Sector



Public trust is not a soft concern. It is the thing between a good tool and its use.

A plain English brief for government teams moving from interest to action. It draws on the NIST AI Risk Management Framework, federal AI actions, and what Texas already requires of its own agencies.

The opportunity, and the duty

Public sector teams face two pressures at once. Modernize service and protect the people they serve. These are not in conflict. Responsible adoption is what makes fast progress safe enough to keep.

Government also carries a duty that private teams do not. Decisions made with public tools affect rights, benefits, and safety. That raises the bar for transparency and oversight, and it is exactly why a clear approach matters before the first pilot.

Adoption is no longer the question

Stanford University's [2025 AI Index](#) reports that 78 percent of organizations used AI in 2024, up from 55 percent a year earlier. Federal AI policy is active and evolving, tracked at [ai.gov](#). The question for a public agency is no longer whether AI arrives. It is whether it arrives with rules, training, and accountability already in place.



78%

of organizations used AI in 2024, up from 55%


3

Texas laws now govern public sector AI use

90

days is enough to run a safe first pilot

Organizations using AI



Source: Stanford HAI, 2025 AI Index

Adoption is already here. For a public agency the task is not to decide whether to use AI, but to make sure it arrives with rules and oversight in place. Source: Stanford HAI, 2025 AI Index.



REAL WORLD: WHAT TEXAS REQUIRES OF ITS AGENCIES

Texas has moved further than most states, and it shows what mature public sector AI governance looks like in practice:

Training. Under Government Code 2054.5191, state and local government employees who work with AI systems complete a certified AI awareness training program every year, with programs certified by the [Texas Department of Information Resources](#).

Standards. Under 2054.702, DIR maintains a statewide AI code of ethics and minimum risk management and governance standards for AI systems under heightened scrutiny.

Transparency. Under 2054.711, agencies must post a standardized public notice when AI is public facing or a controlling factor in a consequential decision.

Templates. DIR publishes an [AI acceptable use policy and templates](#) any Texas government entity can adopt.

A city department or school district does not need to invent responsible AI from scratch. The state has already published the pattern.

Start with low risk, high value work

The safest first uses are the ones where a person still owns the decision and the tool simply saves time. Drafting routine correspondence, summarizing long documents, translating materials for residents, and organizing information are strong starting points. They build staff confidence without putting sensitive decisions in the hands of a machine.

Set a use policy before you scale

A short policy that names approved tools, sets rules for handling data, and requires human review for decisions that affect people keeps experimentation inside safe boundaries. Adopt it, train your team once, and revisit it as your work grows. A policy nobody can find protects nobody, so keep it plain and short. Our free AI Use Policy Template gives you a working draft.



Align to a recognized framework

The [NIST AI Risk Management Framework](#) gives government, legal, and technical teams a shared language for risk. It organizes the work into four functions: Govern, Map, Measure, and Manage. Using it signals to leadership and the public that your approach is deliberate, not improvised. The companion [NIST AI Resource Center](#) provides the playbook and supporting guidance free.

Keep a human in the loop


For any decision that affects a person, a qualified reviewer approves the outcome. AI may draft or suggest. It does not decide. Write this into the policy and make it real in the workflow, not just on paper.

Measure and be honest

Pick one number per use case and track it. Check outputs for accuracy on a schedule, and review results across the communities you serve, not only on average. Log problems and near misses. Honest measurement is what earns the trust that lets you do more.

A simple first 90 days

A responsible first 90 days



A responsible first 90 days. Set up, approve one tool and brief staff, run a single low risk pilot, then review against the baseline and decide.



Days	Action
1 to 15	Name an owner. Adopt a short use policy modeled on the Texas DIR template.
16 to 30	Approve one tool with a signed data agreement. Deliver a one hour staff briefing.
31 to 75	Run one low risk pilot with a single success measure and human review.
76 to 90	Review results against the baseline. Fix what broke. Decide: expand, adjust, or stop.

Sources and further reading

- [NIST AI Risk Management Framework](#), National Institute of Standards and Technology
- [NIST AI Resource Center and Playbook](#), National Institute of Standards and Technology
- [AI and Innovation at Texas DIR](#), Texas Department of Information Resources
- [Statewide AI awareness training](#), Texas Department of Information Resources
- [America's AI strategy and federal actions](#), The White House
- [2025 AI Index Report, Economy chapter](#), Stanford Institute for Human Centered AI



PUT THIS TO WORK

Want help applying this in your organization?

The AI Cowboys Foundation delivers free briefings, workshops, and readiness assessments for business, government, classrooms, and veteran programs. Tell us what you are working on and we will point you to the right next step.

[Contact us](#)

roger@theaicowboys.com · theaicowboysfoundation.org/contact