Coates' Canons
NORTH CAROLINA LOCAL GOVERNMENT LAW BLOG

September 9, 2026

Developing and Updating AI Use Policies: 14 Recommendations for Local Governments

Kristi Nickodem

Back in 2024, I blogged about developing artificial intelligence (AI) guidelines in the government context. Since then, widespread use of AI has accelerated rapidly and the AI tools government employees are using have evolved. Nonetheless, many local governments across the country still lack policies or guidelines governing how employees and officials may use AI in their work. This post updates and revises my list of recommendations for government AI policies and highlights why units of local government need to have these guardrails in place.

Why Do Local Governments Need a Robust AI Use Policy?

Local governments may assume that if they have not formally invested in or endorsed the use of AI tools, employees must not be using them. That assumption is likely incorrect, considering the growing phenomenon of **"shadow" AI usage**. In a **2025 survey** conducted by KPMG and the University of Melbourne, 58% of employees reported intentionally using AI tools in their work on a regular basis, and 70% of those employees reported using free, publicly available AI tools. A startling 48% of employees reported that they had "uploaded sensitive company information, such as financial, sales, or customer information, or copyrighted material, into public AI tools." Moreover, 59% of employees reported they had used AI tools at work without knowing whether it was allowed by their employer. Over half (57%) of employees reported that they had used AI in "non-transparent" ways in the workplace, including presenting AI-generated content as their own or hiding when they had used AI tools to complete their work.

More specific to the government context, a recent **survey of public sector workers** in the United States revealed that a whopping 72% use AI at work, while only a third (32%) of those workers report having access to enterprise-grade AI tools. More than one in three public servants surveyed were unclear as to whether

their organization had a formal AI policy, and fewer than half reported that they received clear direction from leadership.

Meanwhile, in addition to data privacy concerns, cybersecurity risks from semi-autonomous or fully-autonomous **agentic AI** tools continue to grow as well. This summer, the UK-based AI Security Institute (AISI) evaluated frontier large language model (LLM) agents to test their cybersecurity capabilities. The **researchers found** that the LLM agents—specifically, Anthropic's Mythos 5 and OpenAI's GPT-5.6 Sol models—took unsanctioned actions 10 times out of a total of 122 instances in which they had to solve a cybersecurity challenge. The agents "attempted to deceive and target real people and to plant and prompt-inject malicious code."

In recent months, OpenAI and Anthropic have both acknowledged separate cybersecurity incidents during internal testing involving their models. In July, **OpenAI announced** that two of its AI models had escaped their testing sandbox environments and hacked into an online platform called Hugging Face to steal answers to a test the agents were being graded on. Just days later, Anthropic **announced** that it had identified incidents in which three different Claude models had escaped test environments, accessed the internet, and gained unauthorized access to the systems of three different organizations. Claude "compromised the impacted organizations' infrastructure using basic techniques, such as exploiting weak passwords and unauthenticated endpoints."

Now more than ever, units of local government face legal, reputational, and security risks based on how employees are using AI to carry out their work. Robust policies for internal AI use can help to mitigate these risks.

Recommendations for Drafting and Updating Government AI Use Policies

1. Don't just copy and paste. I've seen many state and local government AI policies from around the United States circulating online over the past few years, and I've noticed that a number of them contain substantially similar terms—including terms that sometimes seem impractical, illogical, or unnecessary. Templates from national organizations or models from other government entities can be a useful starting point, but a unit of government should be thoughtful about its own unique legal, technical, and practical issues around AI when drafting its own policy. Those issues may differ widely based on applicable state law, technical capabilities, staff capacity, fiscal resources, jurisdiction size, and risk tolerance.

Likewise, if local governments want to use AI in their policy drafting process, it may be tempting to simply copy and paste policy language generated by a tool like ChatGPT, Claude, Copilot, or Gemini. While these tools may offer a helpful starting point for a policy, users should expect their outputs will require a *significant* amount of editing and revision. For example, a prompt to ChatGPT to write an AI use policy for a local government in North Carolina produced a bloated policy with 28 different sections, including a significant amount of redundant and unnecessary language. Moreover, ChatGPT made multiple judgments in that draft policy about whether particular actions should be prohibited or authorized by the local government, along with judgments about risk levels for various AI use cases. These are judgments that should be made thoughtfully by humans who have evaluated the potential issues and risks in the local context, not by an AI system that lacks the institutional knowledge, legal accountability, and contextual understanding necessary to make those policy choices.

2. Define the covered technology. "Artificial intelligence" is an incredibly broad term that covers a wide array of different technologies. Among other subcategories, this term could include tools that fall into the following categories.

- **Generative AI:** AI systems that create text, images, audio, video, code, or other new content based on a user's prompt. This could include using a tool like ChatGPT, Claude, Copilot, or Gemini to draft a document, edit a policy, summarize a meeting, or create a new website.

- **Predictive AI:** AI systems that analyze existing data to forecast, estimate, rank, or score likely future or unknown outcomes. Predictive AI generally does not create new content; instead, it uses patterns in existing data to make predictions about potential outcomes. Predictive AI might be used, for example, to predict service demands, estimate the likelihood of equipment or system failures, flag applications for government programs based on potential fraud or error risks, or predict rates of recidivism for individuals involved with the criminal justice system.

- **Automated decision-making or decision-support systems:** AI systems that evaluate information and recommend, classify, prioritize, or make decisions about people, applications, eligibility determinations, or other matters. For example, this could include systems that prioritize or screen applications for review, identify cases or individuals for intervention or investigation, classify or rank candidates, flag potentially fraudulent

transactions, or recommend or determine eligibility for a government program.

- **Agentic AI:** AI systems that can pursue a goal and take actions to achieve that goal within digital systems, rather than simply producing information or recommendations. An AI agent may interpret information, determine what steps are needed, use software or other tools, and take actions independently with limited human intervention. For example, some agentic AI systems could monitor an inbox and route messages, schedule appointments, update a database, initiate a workflow, or use multiple software tools to complete a task autonomously.

In reality, AI tools and systems are not always neatly confined to one of these categories. A single platform or tool could include generative, predictive, decision-support, and agentic capabilities or features.

For those wanting to craft an AI policy that applies to a wide-ranging spectrum of different technologies, I like the definition of artificial intelligence used by the **Stanford Institute for Human-Centered AI**: "a broad term for computer systems that can perform tasks with human-like intelligence, such as understanding language, recognizing images, learning from data, reasoning, and making decisions." Other units of government may want to specifically focus their policies specifically on **generative AI**, as opposed to AI writ large.

3. Require appropriate IT review and approval for new software or subscriptions. Due to the cybersecurity and data privacy risks around publicly available AI platforms, local governments may want to require that employees use standard, pre-existing procedures for acquiring new software before downloading or subscribing to any AI tools or services for work purposes. For example, the Town of Chapel Hill's **Generative AI Procedures document** requires all AI software services (including those that are free) to be acquired by following the standard procedures established by the town's Technology Governance Policy. It also requires employees to get approval from the town's Technology Solutions Department before creating a new account for a generative AI service or otherwise using a generative AI system to perform work-related tasks. Creating some sort of centralized review and clearance process for AI tools—whether across an entire unit of government or within individual departments—may help to avoid "shadow" AI usage by employees and ensure consistent mitigation of cybersecurity and data privacy risks.

4. Address boundaries around agentic AI. As described above, generative AI has transitioned beyond mere conversational chatbots, as tools that can autonomously or semi-autonomously execute a series of tasks through agentic workflows are increasingly available to the public. In some cases, AI agents may be able to access datasets and computer systems as necessary to independently carry out their objectives, with little to no human intervention.

Rogue activity from AI agents has the potential to do a significant amount of damage when those agents are integrated into an organization's data systems. In February, a **Meta employee reported** that the AI agent OpenClaw (which can be integrated into email systems and messaging apps) autonomously bulk deleted hundreds of her emails without her permission. And in April, the **Guardian reported** that Cursor, an AI agent powered by the Claude Opus 4.6 model, deleted a company's entire production database and its backups in less than nine seconds.

An AI policy can establish clear boundaries on whether agentic AI systems may be deployed on government devices and the extent to which they are allowed to take actions or access systems autonomously. Local governments wanting to guard against potential data destruction and system security risks should consider requiring the centralized review and approval process for AI tools and systems described above, to ensure that any tools downloaded or accessed on a government-owned device have been vetted and approved by IT staff. A policy alone won't necessarily prevent damage from a rogue AI agent, but IT professionals can't mitigate risks that they don't know about in advance.

5. Prohibit improper disclosures of confidential and sensitive information. Some government employees handle information that is confidential under state or federal law, or sensitive information that is shielded from disclosure under state public records law. An AI use policy should prohibit employees from entering confidential or non-public information into any publicly accessible AI tools. Ideally, a policy should provide specific examples of confidential information that might be handled by local government employees, such as personnel information (**G.S. 153A-98**; **G.S. 160A-168**); protected health information held by **HIPAA-covered entities**; substance use disorder information (**42 C.F.R. Part 2**); social services information (**G.S. 108A-80**); and communicable disease information (**G.S. 130A-143**), among others.

In some cases, the types of information that may be entered into a specific tool—even an enterprise-level tool—may depend on the legal requirements associated

with that information. For example, some tools may not meet **HIPAA's security requirements** for protected health information. San Franscisco's **Generative AI Guidelines** provide an example of how different data restrictions could apply to different enterprise AI tools (see the "Data Protection Requirements" section).

6. Require fact-checking and accountability for outcomes. Consider requiring employees to fact-check outputs from LLM-powered text generation tools (e.g., ChatGPT, Gemini, Copilot, Claude) before relying on them for decision-making or using them in any internal or external communication. LLM-based text generation tools are notoriously prone to "**hallucinations**," meaning they will sometimes confidently assert facts that are not true or even make up fake citations or sources to support those false statements. An AI policy could require that any text outputs created by generative AI tools should be reviewed by a human for accuracy and tone before being used in internal or external communications. Moreover, an AI policy provides an opportunity for a unit of government to remind officials and employees that *they are ultimately individually responsible for the results of the work they produce*, regardless of whether or not AI tools were involved in preparing that work product.

7. Decide whether department-specific policies or provisions are needed. Some local government agencies have unique AI use cases that may warrant their own separate departmental policies, or at a minimum, their own section in a unit of government's broader AI use policy. For example, some local law enforcement agencies use an array of AI tools to carry out their duties, including surveillance cameras with AI capabilities, AI facial identification tools, and AI report drafting technology integrated with body-worn cameras. These tools may raise issues outside the scope of a standard local government AI use policy, including concerns regarding the potential for **wrongful arrests**, **internal misuse**, and **mistakes**, which may need to be addressed with separate policy provisions.

8. Remind employees of public records requirements. Generative AI creates new ways for a local government to make and receive records. Under North Carolina's public records law, any record made or received by a local government official or employee "in connection with the transaction of public business" is a public record subject to disclosure upon request, unless an exception applies (**G.S. 132-1**). If a local government official or employee is using a generative AI tool to carry out the duties of their role, then both the prompts and other data the individual gives to the AI tool and the outputs the individual receives from the AI tool would be subject to North Carolina's public records law. In addition to creating text records like chatbot logs, some officials and employees may also

make or receive other types of records using generative AI tools, including AI-generated images, videos, audio files, and transcripts or summaries of online meetings.

Ideally, an AI policy should require employees to use work accounts (e.g., enterprise accounts provided by the government employer or accounts created on free platforms using the employee's government email address) for work-related purposes, to ensure that public records are easily retrievable and not comingled with personal records. If an employee uses a free, publicly available version of a chatbot (e.g., ChatGPT), those chatbot logs may exist only in a transient browser cache or a personal account history that the local government cannot easily access. For more information about how North Carolina's public records law applies to AI-related records, including retention requirements for those records, see **this blog post**.

9. Address AI transcription, notetaking, and recording. Units of government may want to address, in policy or guidelines, when and how AI-generated meeting transcription tools or features may be used. Popular online meeting platforms like Zoom and Microsoft Teams have AI transcription features available to users. Other services, such as Read AI, Fireflies AI, or Otter.ai's Otter Notetaker, may allow an individual to send an AI bot to "attend" an online meeting, including recording and transcribing that meeting. An employee using a chatbot or other AI tool to record a meeting that the employee is not attending could potentially run afoul of North Carolina's wiretapping law (**G.S. 15A-287**) by "intercepting" the communications in the meeting, unless a human participant in the meeting agrees to the recording. A **putative class action lawsuit** is currently pending against Otter.ai in California, alleging that its Otter Notetaker violates state and federal privacy and wiretap laws. My colleagues Kristina Wilson and Phil Dixon have written more about the legal issues involving recordings by government officials in **this blog post**.

Even assuming that consent to record a meeting is not an issue, an AI policy could be an opportunity to remind employees that when an AI tool records or transcribes a meeting, it may have created a record that will be subject to public disclosure under North Carolina's public records law, unless an exception applies. And given that generative AI transcription features are also prone to errors (see **this article** and **this story** for examples), the AI-generated record of a virtual meeting may inaccurately reflect the contents of the meeting.

10. State explicitly if any AI use cases or tools are completely prohibited. A unit of government that wants to allow its employees to use AI may nonetheless have some potential uses of AI that it deems unacceptable in all circumstances. For example, a local government might decide to prohibit employees from using AI to draft sensitive or high-profile external-facing communications; make employment-related decisions; make final decisions about a permit or application without human review; or impersonate a real person through the production of "deepfake" videos, audio recordings, or photos. Local governments in North Carolina may want to cross-reference their existing policies prohibiting employees from accessing pornography on government devices and networks under **G.S. 143-805** (see **this blog post**), since using an AI tool to generate pornographic content would violate that prohibition.

Conversely, a local government may want to consider providing examples of *allowable* use cases for AI in its policy as well (for example, **New Jersey's generative AI policy** for state employees describes six broad potential use cases for generative AI, while also providing "dos and don'ts" for each use case).

11. Be thoughtful about the extent to which transparency requirements are needed. Many government AI policies and guidance documents I've seen from around the country require some sort of disclosure or labeling from employees when they use generative AI to perform their work. Some policies require employees to affirmatively state or label when work product was created using generative AI. These types of requirements are intended to promote government transparency, but may create implementation problems in practice. Will employees have to put such statements on internal communications, external communications, or both? What if an employee uses generative AI for initial idea generation or a first draft, but does a substantial amount of work to subsequently build on that idea or refine that draft? If an employee goes through dozens of iterations of various prompts into a generative AI tool to create a final product, should every iteration be reported to a supervisor, or labeled as AI-generated for the public? Vague requirements around transparency, disclosure, or citation will inevitably lead to many questions around implementation.

Local governments generally do not require employees to report, cite, or label every time they use *other* types of technology to do their daily work, so they may want to consider whether their approach to AI tools (or certain subcategories of AI tools) should be any different, and if so, how. This conversation may also include discussing how local community members will perceive the extent and nature of the local government's AI usage, regardless of whether employees are required to

affirmatively disclose it. In a 2026 **NBC News poll**, 57% of registered voters said they believe the risks of AI outweigh its benefits, and only 26% of voters say they have positive feelings about AI. This public perception issue may impact how community members react to communications from a local government that are labeled as "AI-generated."

12. Make sure the right people are involved. The process of drafting an AI policy for a government entity should ideally involve both legal counsel and IT leaders. Understanding internal AI usage issues from a legal, technical, and practical perspective is essential to drafting any robust AI policy.

Local governments may also want to seek input from department heads and other supervisors about how AI tools and systems are already being used by their employees (if at all). Desired use cases and associated risks will vary widely from department to department. Some local government departments may even be using state technology platforms or systems that have AI features built in.

When drafting or updating an AI policy, units of government could consider starting with an AI use case inventory to determine how employees are already using AI in their work. This process could involve asking questions such as:

- What AI tools are employees currently using?
- What information is being entered into AI tools? Are any records being uploaded to AI tools or used to train AI tools? Is any of this information confidential or sensitive?
- What work products are affected by AI usage?
- Are employees using personal accounts for work-related purposes?
- Is AI being used to classify, prioritize, make recommendations, or make decisions about people, applications, program/service eligibility, or other matters?
- Are employees using any AI tools that are capable of taking actions on government computer systems autonomously?

13. Establish a periodic AI policy review process. As we've seen over the past few years, AI policies can quickly become obsolete. A local government should make

sure it is staying ahead of changes by reviewing its AI policy on a routine basis—for example, at least annually—and whenever there is a significant change in technology, law, or the government's use of AI. That periodic review could consider questions such as:

- What AI tools are officials and employees using now?
- What new AI capabilities and features have emerged since our last policy update?
- Have any concerning incidents related to AI use occurred since the last policy update?
- Has state or federal law around AI use changed?
- How has community perception around AI changed, and does that impact our internal policy?
- Do the existing restrictions still make sense?

14. Ask employees if they can understand and apply the policy. Policies are only effective if employees understand exactly what they allow and prohibit. Ambiguous language may create confusion and lead to more covert AI usage. For example, some government AI use policies classify potential AI use cases into “low risk,” “medium risk,” and “high risk.” This type of classification could help employees understand possible risk levels, but without further guidance and clarification, it does not provide employees a clear understanding of what they are allowed to do. Are “high risk” uses of AI always prohibited? What different guardrails should an employee have in place when conducting a “medium risk” activity as opposed to a “low risk” one? Are “low risk” use cases always permissible, regardless of the type of AI tool used? Seeking feedback from select employees on a proposed (or existing) AI use policy can help identify areas of confusion or ambiguity, while also potentially identifying practical policy implementation issues before they arise.