Coates' Canons
NORTH CAROLINA LOCAL GOVERNMENT LAW BLOG

March 11, 2026

The Intersection of Artificial Intelligence and the Public Records Act

Kristi Nickodem

North Carolina's Public Records Act, codified as **Chapter 132 of the North Carolina General Statutes**, establishes that the public records created or received by state agencies and local governments are "the property of the people." G.S. 132-1(b). Generative artificial intelligence (AI) creates new ways for the government to make and receive records. How does the public's right of access—and the government's responsibility for retaining and producing public records—intersect with new types of records created while using generative AI?

The Broad Applicability of the Public Records Act

Members of the general public have a right to inspect or copy any record that is "made or received pursuant to law or ordinance in connection with the transaction of public business by any agency of North Carolina government or its subdivisions," unless an exception to that right of access applies. **G.S. 132-1**; **132-6**. North Carolina law defines a public record to include "documents, papers, letters, maps, books, photographs, films, sound recordings, magnetic or other tapes, electronic data-processing records, artifacts, or other documentary material, *regardless of physical form or characteristics*." G.S. 132-1(a). This definition includes a wide variety of different physical and digital records.

If a state or local government official or employee is using a generative AI tool (e.g., ChatGPT, Gemini, Claude, Copilot) to carry out the duties of their role, then both the prompts the individual gives to the AI tool (records "made" in connection with the transaction of public business) and the outputs the individual receives from the AI tool (records "received" in connection with the transaction of public business) 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, which would all likewise be subject to the public records law if they were "made or received…in connection with the transaction of public business." This has several distinct implications for government officials and employees in North Carolina.

First, it means that these AI prompts and outputs must be made available for inspection or copying in response to a public records request unless an exception applies. **G.S. 132-6(a), (a1)**.There are *many* exceptions to the public's right of access to records established in state law, including statutes that exclude certain categories of records from the definition of public records (e.g., criminal investigation records, per G.S. 132-1.4(a)) and other statutes that make certain information confidential (e.g., social services records under G.S. 108A-80). As with any other type of records, governments are not required to disclose generative AI records to the public if (1) federal or state confidentiality laws prohibit such disclosure, or (2) the records fall within a category specifically excluded by state law from the definition of "public records."

Second, these AI-generated records (and records of prompts/inputs into an AI platform or tool) will be subject to records retention requirements. Public officials and employees are generally prohibited from destroying, selling, loaning, or disposing of any public record except as allowed by a records retention and disposition schedule established by the North Carolina Department of Natural and Cultural Resources (NC DNCR). **G.S. 132-3(a)**; **G.S. 121-5**. These schedules determine how long certain categories of records must be retained and when they may be disposed of or destroyed. The current retention and disposition schedules are available on **NC DNCR's website**.

Is Actual Possession of the Record Required?

Many generative AI tools used by government employees are maintained by third parties (OpenAI, Anthropic, Google, Microsoft, etc.), meaning the records created on those platforms are often stored on third-party servers. Does this mean that the records created on those tools do not fall within the scope of North Carolina's public records law, since they are not in the physical custody of the government agency? Arguably, the answer is no, so long as a government official or employee has made or received the record and has the legal right to retrieve the record at any time. Consider the implications if, instead, the government's lack of *physical* custody of a record meant that the record was categorically excluded from coverage under the public records law: it would mean that any government

emails maintained on a third-party server would also fall outside the scope of the public records law due to the government's lack of actual custody. Such a result would seem antithetical to the purpose of the Public Records Act, which is "intended to be liberally construed to ensure that governmental records be open and made available to the public." *DTH Media Corp. v. Folt*, 374 N.C. 292, 300 (2020).

A similar issue was addressed by the Court of Appeals in *Gray Media Grp., Inc. v. City of Charlotte*, 290 N.C. App. 384 (2023). In *Gray Media*, the City of Charlotte had entered into a contract with Ernst and Young ("EY") to perform some work for the City, including having EY develop and send out a survey to City Council members. The City argued that the Council members' survey responses were not public records because they existed solely on EY's servers and were only sent to Council members as a hyperlink to EY's website. The Court of Appeals rejected this line of reasoning:

> To accept the argument that a hyperlinked survey instead of an attached survey removes the document from the universe of public records requires us to read the statutory language much too narrowly. Such a reading would defeat the purpose of the statute, creating a clear path to hide huge swaths of governmental work from public scrutiny. Instead, we note that the statute includes broad language including "all documents … electronic data-processing records … *regardless of physical form or characteristics*." N.C. Gen. Stat. § 132-1(a) (emphasis added). Further, the Public Records Act has been repeatedly interpreted to provide liberal access to public records.... Therefore, we decline to adopt the City's narrow interpretation that a hyperlink to EY webspace does not constitute a "document or electronic data processing record."

Gray Media Grp., Inc, 290 N.C. App. at 395.

The City also argued that it did not have to produce the survey results in response to public records requests because it did not have "actual possession" of the records or "substantial control" over EY to demand the records. Once again, the Court of Appeals rejected this interpretation, highlighting that "actual possession" does not appear in **G.S. 132-6** and finding that "actual *or constructive* possession is sufficient to meet the requirement for custody." *Id.* at 639-40 (emphasis added) (citing *Fordham v. Eason*, 351 N.C. 151, 155 (1999) ("Constructive possession is a legal fiction existing when there is no actual possession, but there is title granting an

immediate right to actual possession."). In *Gray Media*, the City's constructive possession was readily apparent, since the contract between the City and EY required EY to "promptly provide the Contract data to the City in machine-readable format upon the City's request at any time while the contract is in effect or within three years from when the contract terminates." *Id.* at 398.

In some cases involving records and generative AI, the government's constructive possession may not be as clearly established as it was under the City's contract in *Gray Media*. However, if the AI tool or platform at issue is being accessed via an account, license, or contract established by a government official or employee, the terms and conditions of use might explicitly establish actual or constructive possession of the records created using the tool or platform. The contours of constructive possession over records in the digital space are likely to be explored and defined by North Carolina's appellate courts for years to come.

Records Retention Issues

As with many other forms of digital records, retention gets tricky with AI-related records. How long must a government entity retain chatbot logs and other records created using AI? NC DNCR's records retention schedules generally set different retention requirements based on the *content* of the record. It is the substance of the record—not its format—that determines how long it must be retained.

Both the **Functional Schedule for North Carolina State Agencies** and the **General Records Schedule for Local Government Agencies** include sections discussing "transitory records," defined as records that "may have little or no long-term documentary or evidential value to the creating agency." Among those transitory records are "drafts and working papers," defined as materials gathered or created to assist in the creation of another record. Though drafts and working papers may be public records, the retention schedules state that "many of them have minimal value after the final version of the record has been approved and may be destroyed after final approval if they are no longer necessary to support the analysis or conclusions of the official record." According to the schedules, this includes drafts and working papers for internal and external policies, internal administrative reports, internal non-policy-level documents, and presentations and workshops.

Some records created using generative AI—including prompts and outputs—may be transitory records. The nature of generative AI tools is that they generally

require multiple, iterative prompts to arrive at a final, useful output. And after multiple prompts and outputs, someone may take an output from one of these tools and edit it further manually to achieve a final draft. In some cases, it's possible that the government official or employee is only required to retain the final, approved version of the document that was created with AI assistance—if the prior prompts and outputs from the AI tool "are no longer necessary to support the analysis or conclusions of the official record." Keep in mind, however, that guidance in the General Records Schedule for Local Government Agencies explains that a *seemingly* transient record might actually be a record with a required retention period in the schedule. In such a case, the records custodian should comply with the applicable retention period rather than deleting the records. When in doubt, the schedules advise retaining the record in question and seeking guidance from a **NC DNCR records analyst**.

In some cases, a conversation with a chatbot or other record created using AI will clearly *not* be a draft, working paper, or otherwise "transitory." Consider, for example, a government employee who uses a tool like ChatGPT or Claude to help make an important decision like choosing between otherwise qualified candidates for an open position. The record of the chatbot's analysis and dialogue underlying that decision may need to be maintained under the "Personnel Records" category on the retention schedules. As another example, consider a municipality that has installed a public-facing **LLM-based chatbot** on its website and receives a citizen complaint about municipal services via the chatbot. The record of that complaint would need to be retained according to the "Constituent Comments, Complaints, Petitions, and Service Requests" category on the retention schedules.

One challenge for North Carolina agencies is figuring out how to capture and retain AI-generated records. What if a government employee uses his own personal AI account—for example, an account he has created on ChatGPT using his personal email address—to perform government work? It is the content and purpose of a record, not its location, that determines whether it is subject to North Carolina's public records law. For example, text messages or emails sent or received by government officials on their personal cell phones or through their personal email accounts may constitute public records if they involve "the transaction of public business." The same is true for records of inputs to and outputs from AI tools and platforms. Ideally, government agencies should adopt policies or guidelines that 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 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 agency cannot easily export.

Records Created Using Online Meeting Tools

Generative AI transcription and summarization features are now standard options across popular online meeting platforms (e.g., Zoom, Teams). Some generative AI tools, such as Fireflies or Otter.ai's OtterPilot, allow an individual to send an AI bot to "attend" an online meeting. In some cases, the bot can record audio from the meeting and transcribe the meeting without the individual who actually "sent" the chatbot being present. I've written about the wiretapping, consent, and privacy concerns with these autonomous notetaking tools in **this blog post** and **this blog post**. Even assuming that participant consent to record is not an issue, government agencies should be mindful that when an AI tool records or transcribes a meeting, it may have created a record—an audio recording, transcript, or summary—that is subject to public disclosure under North Carolina's public records law unless an exception applies.

Consider, for example, a county manager who sets up a Microsoft Teams meeting to talk with a county department head about a sensitive, politically charged issue involving the department's funding and services. The information discussed in the Teams meeting, while sensitive, is not confidential. The meeting would not be subject to North Carolina's **open meetings law** and therefore would not have to be open to the public. But if the manager or the department head enables the transcription feature in Teams, they have now created a record of the meeting that may have to be disclosed in response to a public records request.

The public records disclosure requirements raise another concern: what if these transcriptions are inaccurate? Like other generative AI tools, these AI transcription features and notetaking bots are also prone to errors (see **this article** for examples). According to this **story from The Guardian**, a recent study found that social workers using AI transcription tools in England and Scotland made errors in children's social work records—errors that could lead to potential harm if a concerning pattern of behavior was missed. As another example, per **this story from the AP**, a University of Michigan researcher conducting a study of public meetings found hallucination errors in eight out of every ten AI-generated audio transcriptions he inspected. Government agencies facing public records requests for AI-generated transcriptions and summaries may be in a double predicament:

worried about disclosing the contents of an internal conversation while simultaneously concerned that the record of the conversation does not accurately reflect what occurred.

Government agencies may want to address, in policy or guidelines, when and how AI-generated meeting transcription tools or features may be used. Maryland's Department of Information Technology provides one robust example of a **guidance document addressing AI-powered call recording & transcription tools**. My colleagues Kristina Wilson and Phil Dixon have written more about the legal issues (including public records implications) of government recordings in **this blog post**.

A Cautionary Tale

As government officials are considering the public records implications of using AI, they should also be mindful of what these records may reveal about how government employees are using these tools. In 2025, a **reporter filed records requests** with cities across the state of Washington, asking for records of government officials using ChatGPT. One chat log disclosed in response to the requests caught the reporter's attention. A city employee preparing a request for proposals (RFP) to solicit applications for a new city utility billing software asked ChatGPT to write contract requirements that would "exclude" a particular vendor, Origin SmartCity, "without explicitly saying so." The same employee then asked ChatGPT to write requirements that would "favor" a different vendor, VertexOne. Some of the language generated by ChatGPT appeared in the final RFP, and the city ended up awarding the five-year, $2.7 million contract to VertexOne. You can read more about the story **at this link**. Much like the advent of email and social media messages, generative AI creates a new way by which government employees might leave a virtual "paper trail" of their activities.

Looking Ahead

As generative AI becomes more integrated into the daily operations of North Carolina state and local government—from drafting policies to summarizing meetings—the "property of the people" will increasingly consist of AI-generated content. State and local leaders need to be mindful of the unique records challenges posed by generative AI to ensure that technology does not outpace the transparency required by the Public Records Act.