Our People Are Ready for AI. Are Our Pools?

October 28, 2026

Author: Matthew Burleson

Every pool in NLC-RISC is having some version of the same conversation right now. Someone on staff is already using AI. Someone in leadership is worried about it. A vendor just added AI features to software we already pay for. And a member city is about to call and ask what we think.

At this year's NLC-RISC Staff Conference, I led an interactive session on leveraging AI in pool work. Pool staff voted anonymously on their phones as we worked through realistic situations: member communications, a storm response, an HR complaint, sensitive data, and the building blocks of an AI policy. I went in expecting to spend most of the hour easing nerves. I came out with a very different read on where our industry stands.

Pool staff are more ready for AI than most of us assume, and they bring exactly the right instincts to it. What's lagging is the structure around them. That's good news, because structure is something pools know how to build.

Readiness isn't the problem

There's a common assumption that pool employees are wary of AI and need to be eased into it. The room didn't support that. Genuine worry was rare. Most people were curious, and many were already using AI in some form at work. Over the course of a single hour, the room grew noticeably more enthusiastic, and nobody left more nervous than they arrived. By the end, almost no one felt unprepared to make AI decisions, and many said they could help set guardrails at their own pool.

That shift tells me something important. What holds pool staff back isn't fear. It's the lack of a shared way to think about AI. Without a framework, every request is a judgment call, and people either freeze or quietly do what seems reasonable to them. Give them a simple way to separate routine work from high-stakes work, and confidence comes quickly.

This matters for how pools plan. If you believe your staff are afraid, you build a strategy around persuasion: awareness campaigns, cautious pilots, long timelines. If you recognize they're already curious and experimenting, you build around direction instead. That means approved tools, clear expectations and practical training. Only one of those plans matches what's actually happening in our offices.

Pool staff also bring practical judgment about cost. Most people understand that a routine announcement doesn't need the most powerful, most expensive AI model available. That instinct will matter as AI subscriptions and add-ons start multiplying across departments.

Our instincts are already right

Pools have a real advantage here. Our people think like risk managers because that's the job. We spend our days asking what could go wrong, who's responsible when it does and how to keep it from happening again. Those same instincts translate directly to AI, and they showed up consistently.

The clearest one is accountability. Pool staff are close to unanimous that using AI doesn't move responsibility away from the person who used it. In our world, "the AI wrote it" will never be an acceptable answer to a member, a board or a court, and people already know that.

The second is verification. Pool staff are comfortable letting AI help draft member communications, even urgent ones during a disaster, but they expect every fact to be checked and a person to review it before it goes out. They see AI as a fast first draft, not a finished product.

The third is that decisions about people stay human. Pool staff draw a firm line at letting AI decide the outcome of an HR complaint, a disciplinary action or anything similar. They're open to AI organizing information for a human decision-maker, but not to AI being the decision-maker.

The fourth may be the most valuable: healthy skepticism. One of the real dangers of AI is that it sounds sure of itself even when it's wrong. It will quote the wrong section of a policy or smooth over a conflict instead of flagging it. Pool people are trained to read the fine print, and they're quick to catch a polished answer that doesn't hold up.

That's exactly the culture you need to adopt AI safely. The catch is that at most pools, it lives in people's heads. It isn't written into policy or built into training, and instincts don't scale. They vary from person to person, they leave when someone retires, and they're hard to point to when an auditor or board asks how you're managing AI risk.

The gray zones are where pools need guidance

Good instincts carry you most of the way, but they break down in the gray areas. Those are the places where well-meaning staff at the same pool will make different calls on the same request, and that's exactly where written guidance does the most good.

We tend to judge risk by department instead of by data. Anything labeled "HR" feels risky, even when the information is completely public. Using AI to help write interview questions from a posted job description is low-risk work, yet it makes a lot of people hesitate. Meanwhile, requests that sound routine, like summarizing incident reports or cleaning up a spreadsheet, can carry sensitive personal information without anyone noticing. Without a framework, people go with their gut, and the gut often focuses on the wrong thing. Easy wins get blocked while real risks slip through.

"Anonymized" doesn't mean the same thing to everyone. Most pool staff want member and organizational data kept inside approved, organization-controlled tools. But a meaningful share believe anonymized or aggregated data is safe to use anywhere. That's a bigger deal for pools than for most organizations. Many of our members are small cities and towns. Remove the name from a workers' comp claim in a small community, and the job title, injury date and department may still point to exactly one person. The same is true for loss runs and litigation details. Each pool needs to define what de-identified actually means for its data.

Member disclosure is an open question. Most pool staff don't see a need to tell members when AI helped create something they receive, as long as a person reviewed it first. That may be the right answer. We don't disclose templates or spell-check either. But coverage letters, claims communications and legal guidance carry real weight for our members, and I suspect very few pools have made this call deliberately. In a business built on trust, it's a question worth answering on purpose, and in writing, before a member or a reporter asks it for us.

"Human review" needs a definition. Pool staff agree AI output should be reviewed, but they differ on whether that means every output, every time, or only higher-risk work. Both are reasonable positions. What isn't reasonable is leaving it undefined and letting each employee decide.

Shadow AI is already here

When pool staff are asked what AI questions they're facing at work, the answers are revealing. They want to know whether they can use a personal AI account for work tasks. Whether they can upload loss data that includes employee names into a tool that's "supposedly closed." How to anonymize a long, complicated claim before asking AI to summarize it. Whether AI should help write a disciplinary memo, or play a role in setting experience modifiers.

A few of these are easy. Most aren't. They touch personal information, employment decisions, claims data and rating, which are the areas where pools carry the most risk. And several of them make one thing clear: staff aren't waiting for permission. They're experimenting now, often before their organization has given them an approved tool or a clear answer.

That isn't a discipline problem. It's a signal. When the official path doesn't exist, people build their own, and their own path usually runs through a free consumer account with few data protections. The fastest way to reduce that risk isn't a ban. It's giving people a better, approved option and clear rules for using it.

A pool AI policy doesn't have to be complicated

When pool staff work through the basic questions of an AI policy together, a fairly clear consensus emerges. It's simple enough to fit on one page, and it makes a solid starting point for any NLC-RISC pool:

1. **Approved tools for pool work.** Staff generally favor limiting work involving member or organizational data to tools the organization has approved. Some room for open brainstorming is reasonable, but sensitive information belongs in approved systems.
2. **A person reviews AI output before it's used.** The open question is scope: everything, or only work that touches data, HR, legal matters or high-impact decisions. Pick one and write it down.
3. **Sensitive information stays in-house.** Member rosters, claims details, coverage decisions and financial data shouldn't leave approved, organization-controlled environments. Define clearly whether anonymized data is an exception.
4. **The employee owns the result.** AI is a tool, not a shield. Whoever uses it is accountable for what it produces.
5. **A deliberate position on disclosure.** Decide when members should know AI helped, and apply it consistently.

None of these are radical. Most pool staff already believe them. The value of writing them down is that everyone works from the same rules, new employees learn them on day one, and the pool can show its board and members exactly how it's managing the risk.

What NLC-RISC pools can do now

The gap between staff readiness and organizational readiness is very closable. Most of the work is policy and communication, not technology, and none of it requires a big budget.

1. **Give staff an approved tool.** Shadow AI grows where there's no sanctioned option. Start by checking what you already own, since many tools pools pay for now include AI features. Then make sure staff know what they can use and how.
2. **Write the one-page policy.** Use the five points above as a starting draft. Staff already agree on most of it, so this is mostly about putting common sense in writing.
3. **Classify by data, not by department.** A simple green, yellow, red framework helps people sort requests quickly. Green is routine work with public or non-sensitive information. Yellow touches member data, goes out to members or informs a decision, so it needs safeguards and review. Red involves protected personal information or a decision about a specific person.
4. **Define "anonymized" for your data.** Be specific about claims, loss and litigation data, and remember how easily information from small members can be traced back.
5. **Make the disclosure call deliberately.** Whatever your pool decides, document it and apply it consistently.

6. **Check your vendors.** AI is arriving inside claims systems, document tools and email, often through quiet updates. Ask vendors what AI features they've added, what happens to your data and whether features can be turned off.
7. **Train with scenarios, not slides.** Working through realistic pool situations together builds shared judgment far faster than a policy memo. Confidence grows quickly once people practice.
8. **Get ready to help members.** Member cities are facing the same questions, usually with fewer resources. Pools are well positioned to offer model policies, training and guidance on AI vendor contracts, the same way we've helped members with cyber security and law enforcement liability. AI-related claims will follow AI adoption, and pools that understand how members are using it will be better prepared to anticipate those losses.

NLC-RISC has always been strongest when pools share what they learn. AI is a great place to keep doing that. Our people are ready. If we give them the structure to match, and share our policies and lessons with each other along the way, our pools and our members will be in a much better position than if we wait.

These observations come from an interactive session Mr. Burleson led at the 2026 NLC-RISC Staff Conference, where pool staff voted anonymously on realistic AI scenarios and policy questions. Participants chose to attend, so this is one read on engaged pool professionals, not a statistical survey of every NLC-RISC pool.