AI is changing healthcare technology. But who's deciding whether it actually helps?

An executive brief for rural, community and critical-access hospital leaders

Which AI tools are running in your building today, and who decided each one should be?

AI does not come into a small hospital through a strategy meeting. It comes in through three doors. A vendor switches on a feature in an upgrade. A department buys a tool for documentation, coding or prior authorization. Someone on staff pastes a draft into a free chatbot.

None of those doors has a sign-in sheet.

Picture upgrade night. The IT director is working through the release notes, and somewhere around page 40 a line says a predictive model is now available. Nobody on the clinical side will read page 40. That is how a lot of EHR capability has reached small hospitals over the years, and AI could follow the same path. The greatest risk isn't falling behind on AI. It's letting vendor availability drive implementation instead of clinical priorities, measurable value and patient safety.

The rules just moved, and not the way most people expect

Two kinds of rules matter here, and they are easy to mix up. Some govern the companies that build health IT. Others fall on the hospital.

RULES ON VENDORS

On the vendor side, the federal program that certifies EHR software finalized a rule called HTI-5 in October 2026. It removed the requirement that certified EHR software support standard transparency information for predictive AI tools: what a model was built for, what it uses, and how it was tested. That requirement sat on the software developer, never on the hospital. If you want that information now, put it in the contract.

DUTIES ON YOUR HOSPITAL

On the hospital side, the duties did not go anywhere. Since May 1, 2025, any hospital taking federal funds has had to make reasonable efforts to find the patient care decision tools that use race, color, national origin, sex, age or disability as inputs (Section 1557, 45 CFR 92.210). Then it has to reduce the risk they discriminate. That covers AI and non-AI tools alike. HHS's June 2026 notice on the partial vacatur of certain Section 1557 provisions did not remove that duty or the broader nondiscrimination framework. Texas now requires practitioners to tell patients when AI is used in their care and to personally review what it produces. California requires a disclaimer on AI-written patient messages.

The certification program stepped back from checking these tools. The job of deciding landed on the hospital.

The evidence for checking locally

In 2021, Michigan Medicine tested a widely used sepsis prediction model. It ran the model against 38,455 of its own hospital stays. The model missed 67 percent of the patients who had sepsis. It also alerted on 18 percent of all patients. The vendor disputed the study's method and has since revised the model (Wong et al., *JAMA Internal Medicine*, 2021).

67%

of sepsis cases missed · 18% of patients alerted · 38,455 stays · Michigan Medicine, 2021

The lesson is not about one vendor. A model can be live at hundreds of hospitals and still never have been tested against the patients who walk into yours. Rural hospitals are exactly where that check matters, because the people in your beds may look nothing like the population the model learned from.

What good looks like in a small hospital

1. **One front door.** Every AI tool, feature or pilot gets logged before anyone signs a contract, switches it on or tries it, including free trials and features that arrive in an upgrade.
2. **Review sized to risk.** A tool that drafts job postings gets signed off in a week. A tool that flags deteriorating patients gets tested on your own patient data, then a trial run where it works but nobody acts on it, then a sign-off from the medical staff.
3. **Test it here before you trust it.** Vendor results are where the review starts, not where it ends.
4. **A named owner after go-live.** Every live tool has an owner and a check-in date.
5. **A list the board sees.** Once a year, the board gets every AI tool in use, its risk level, its owner and what went wrong.

This is built for the staff you actually have. The committee is five or six people who already sit in your leadership meetings, often wearing two hats, and they meet monthly for an hour. Low-risk requests never reach the meeting.

What I am not saying

Let me be clear about what I am not saying. AI is not the problem. A good tool, run well, gives a thin team real time back. Most of the people building these tools are serious. And a 25-bed hospital does not need a health system's governance office. Copy it down and it stops good ideas at the door.

The failure pattern is different. It is not too little process. It is no front door at all: no list of what is running, no one who owns each tool after go-live, and no moment where someone asks whether it works here.

Funding depends on your state's plan

The Rural Health Transformation Program is $50 billion. It goes to states over five years, fiscal 2026 through 2030. Its approved uses name artificial intelligence and the training to adopt it. The money flows through each state's approved plan, so whether it pays for this work depends on what your state chose to fund.

AI readiness belongs inside work you're already doing

AI capabilities are increasingly becoming available through EHR vendors, managed-service providers and technology partners. For a smaller organization, that raises one question worth asking every time: is this capability being introduced because it addresses an identified clinical or operational need, or simply because it's available?

I look at technology decisions through the lens of patient safety, workflow effectiveness, operational readiness and measurable value. AI readiness and governance are not another standalone initiative. They belong inside the technology transformation, integration and implementation work you are already doing, with a structured 30-, 60- or 90-day approach when the work calls for one.

The objective isn't to adopt AI faster. It's to make better decisions about where AI belongs, what problems it should solve, and how success will be measured.

So here is the question worth asking this week: which AI tools are running in your building today, and who decided each one should be?

If your organization is navigating these decisions, I'm always happy to pull up a chair and talk through what you're seeing.

ABOUT THE AUTHOR

I have worked in healthcare IT and operations since August 2012. That has put me on most sides of the table: implementing EHRs for health systems at Cerner, leading technology deployment inside a revenue cycle managed-services business at R1 RCM, consulting at S&P, and leading IT integration for a value-based primary care platform at Mosaic Health. My deepest rural work was at Cerner: more than 90 community, rural, district, and critical-access hospitals in 34 states, including 62 critical-access hospitals.

CONTACT

Jeff Paul, Founder & Principal
JP Healthcare Advisors
jeff@jphealthcareadvisors.com
816-820-8849
jphealthcareadvisors.com

SOURCES

- ONC, HTI-5 final rule (released October 26, 2026; Federal Register October 28, 2026)
- 45 CFR 92.210, Section 1557 nondiscrimination rule (patient care decision support tools), and the HHS June 2026 notice on the partial vacatur of certain provisions
- Texas SB 1188 (effective September 1, 2025) and HB 149 (effective January 1, 2026)
- California AB 3030 (effective January 1, 2025)
- Wong A, et al. External validation of a widely implemented proprietary sepsis prediction model. *JAMA Internal Medicine.* 2021;181(8):1065-1070
- CMS Rural Health Transformation Program, medicaid.gov/resources-for-states/rural-health

Regulatory status as of November 12, 2026. This brief is general information, not legal advice; confirm with counsel before acting.