

# AI, the Audit Trail, and the Extended Enterprise

A Governance Framework for Healthcare Directors

By M. Eric Johnson and Ed Magee

---

*“I never knew you had green eyes.”*

Fifteen years of visits. Fifteen years of a physician typing into a screen while his patient talked. The first time he used ambient AI clinical documentation, he hit record, completed the visit, and asked her how the experience felt. Her answer was those seven words. He said it hit him like a ton of bricks.

That very human story was told at Vanderbilt’s Owen Graduate School of Management on April 16 by Greg Aaron, Global Vice President, Market Leader, US Provider Enterprise Group at Oracle Health. He offered it as an illustration of something he believes directly: “Healthcare is still very, very personal. AI, oddly enough, can make it personal again.”

That experience was counterbalanced by a contrasting pair of AI implementation stories. Two AI deployments from the same healthcare organization governed by the same board, side by side. In one application, ambient physician documentation was successfully scaled past 75% ER adoption. The other case involved a technically sophisticated tool that ran into operational headwinds that neither the technology nor the capital solved. Same board. Same technology class. Different outcomes. The variable was not the technology. Rather, the challenges arose from people and governance.

We find that the variable that separates AI wins from AI failures in many organizations is the governance architecture around the capital allocation decision. For public company directors, that architecture now *has* to stretch across something far larger than the enterprise as traditionally defined.

## The Extended Enterprise Has a New Problem

Emeritus Dean Johnson has spent thirty years studying what happens when information technology deploys across an organization’s network of partner relationships. Security failures create friction throughout the extended enterprise. Value creation flows through the same network. The extended enterprise is not a metaphor. It is how technology actually operates in a modern public company.

Healthcare AI is an extended-enterprise challenge by construction. Clinical AI lives in the vendor’s cloud. The training data comes from somewhere upstream. The agent that submits a prior authorization routes through a payor’s API. The ambient documentation tool captures audio through a device, transcribes it through a model, and stores in a vendor cloud. Each step sits outside the hospital’s four walls.

Public company directors have been trained to govern what sits inside the enterprise. The AI pressure test is different. Directors are now accountable for technology decisions that execute across a network they do not directly control. Extended enterprise governance is not a refinement of internal governance. It is a different discipline.

A year ago, at the same program at Vanderbilt Owen, directors were asking whether to pilot AI. The question has moved. AI pilots have scaled, with some becoming embedded into day-to-day operations. Agents that did not exist twelve months ago are making decisions that directors are personally responsible for. That shift is the year-over-year governance marker. Mark Williams of the Vanderbilt AI Law Lab put it plainly on the regulatory panel: “Increasingly, it’s less about what happens when the AI goes wrong because it didn’t work as intended, and more about what happens when the AI works exactly as it was intended.” In other words, the AI did what management told it to do. The failure was upstream. *“It wasn’t the AI’s fault,”* Williams said. *“It was all of the data governance and permissions.”*

Those permissions are extended enterprise artifacts. Who granted the agent its scope? What contractual terms govern the vendor’s use of protected health information? Which version of which model was active when the transaction posted to the general ledger? Most boards don’t have access to those execution details and cannot answer those questions today. The ones that are putting governance in place to tackle that issue are defining the standard.

## The Audit Trail Problem

The question that landed hardest on April 16<sup>th</sup> was about the audit trail. AI makes a decision that touches a financial statement. What documentation does the audit committee need to see? What should the walkthrough expose? Who is personally accountable?

Josh Trusley, EY Nashville Managing Partner and US Assurance Growth Markets Leader, answered directly. *“If the client is using AI in the reporting of financial transactions, it creates a trigger for us. As part of our audit process, we have procedures to make sure we identify, risk assess and then respond to the use of AI. We want to make sure we understand how companies are responsibly using AI and how they are making sure there is appropriate human-in-the-loop involvement and oversight. The regulators are also understandably focused on that too and it will continue to get more complex as the pace and scale of AI increases.”*

The operational implication is sharp. Audit firms are running walkthroughs. The walkthrough only works if management has been transparent about where AI is actually operating in the financial reporting workflow. *“You have got to be able to identify that,”* Trusley said. *“That creates risk for everybody if you’re not able to identify the use and understand how it is being used and controlled.”* Directors carry personal responsibility for financial statements regardless of whether AI participated in producing them.

Mike Schlosser, Senior Vice President and Chief Transformation Officer at HCA Healthcare, described how his organization governs the investment side. HCA is spending roughly two hundred to two hundred fifty million dollars all-in on AI infrastructure against seventy-five billion in annual revenue. The board supported it because HCA came to them with a five-year plan. *“They reacted positively specifically because we weren’t thinking about quick wins,”* Schlosser said. On the ambient documentation return, he was candid: *“We’ve been able to recognize an ROI on ambient listening, enough to pay for the technology. It has not been a financial home run, but we’ve given doctors 1.5 to 2 hours back.”*

That combination sharpens the audit expectation. A disciplined AI investment is not self-auditing. Even the most disciplined programs require the board to ask the harder questions about what the technology is doing, what it costs, and where the value actually shows up. Transparency with the audit committee is not a courtesy. It is the control.

For public company directors, the technical components of an AI audit trail are becoming specific.

- A model registry that inventories every AI system touching financial reporting.
- Model cards documenting purpose, data sources, training methodology, known limitations, and risk ratings.
- Decision logs that capture input, output, model version, and the identity of any human in the loop.
- Provenance chains that can trace any AI-touched transaction back through the extended enterprise to the source.
- A tested kill switch with a documented last-test date.
- Vendor contracts that preserve audit rights over training data, especially where HIPAA-subject information is involved.
- And directors' and officers' coverage reviewed in light of agentic AI risk, because most D&O policies were written before agentic systems existed.

Casey Santos, who joined the regulatory panel as a sitting director across DTE Energy, Omnicom, and Horizon Blue Cross Blue Shield of New Jersey, framed the director test in one line: *“Would we let a junior employee do this on their own?”* That is the governance standard for every AI agent operating inside the extended enterprise. If the answer is no, the guardrails are not yet in place. She added a related warning about benchmarking. Josh Trusley described the trap: management presenting to the board that the organization is ninety percent ahead of peers on AI maturity. “Trusley continues, “remaining competitive may also include ensuring that your sector comp set is not too narrow. If your benchmark doesn’t include some of the more advanced organizations in your industry or broader ecosystem, then you may be years behind the market leaders and not know it.” The comp set is the wrong comp set. Directors should force the benchmarking question.

Bobby Frist, chairman and CEO of HealthStream, summarized the resulting pressure on the board calendar. “Audit committee meetings have doubled in their length of time because they took on security, and the security quickly turns into AI and then AI governance and data access. It all starts to blend together. And all of a sudden, the task of the audit committee is so great that you need a five-hour audit committee meeting several times a year.” He added a humility line that the whole room felt. “Management doesn’t really have answers. And I don’t think the board really has answers. The technology is in front of government policy. It’s in front of governance.” And the framing that positions the stakes: “Over the next five years, this is as big a change as the internet.”

## Four Governance Implications

From the extended enterprise and the audit trail, four directives close the day for healthcare directors.

**1. Boards must have a multi-year AI thesis.** Funding must be dedicated and patient - not a pilot budget. A thesis longer than the fiscal quarter tolerates, covering compounding value, planned capital, and the accountability structure. If management cannot articulate a long-term vision, the organization is left with experiments, not a strategy.

**2. Change management must be modeled and owned before deployment.** Surveying the troops after a rollout is too late. India Mullady of EY framed the stakes on the workforce panel: *“People do need to be managed, but no one really wants to be managed. They want to feel*

*invested. They want to feel ownership. They want to feel purpose.” AI transformation is not incremental. “It’s no longer enough to incrementally nibble around the edges of jobs, five or ten minutes at a time. You have to fundamentally transform how work is done.”*

**3. Patient-centric outcome metrics must be part of every AI update.** Financial metrics alone are insufficient for a healthcare board. The Green Eyes story is the measurable end. Ambient documentation returned time to physicians, improved note quality, and restored the doctor-patient interaction. All of these must be measured. Directors who do not require patient-outcome metrics alongside financial metrics are accepting an incomplete thesis.

**4. Boards need a standing AI governance cadence distinct from the audit committee.** Many audit committees have already absorbed cybersecurity. They cannot also absorb AI governance, data governance, and financial statement traceability indefinitely. The governance structure has to catch up with the extended enterprise reality.

## De-aging

A year of experience and research has sharpened a different question for directors. Is AI an aging agent for companies, or a de-aging agent? The aging-agent framing is familiar. The race to keep up with efficiency gains, cost reductions, and maintain a business model beyond its natural expiration. The de-aging framing is harder and more consequential. AI can restore a business model to a prior state of capability. It can return a doctor to his patient.

Green eyes is what de-aging looks like at the bedside. The doctor-patient relationship is restored by the technology. Public company directors who protect that moment while they govern the technology are doing the work. Directors who optimize the technology and leave the moment to chance are running experiments.

The technology is moving faster than government policy. Faster than most committee charters. Public companies in healthcare are the pressure test for every sector that will face the same dynamics within five years. The boards that close the governance gap in the extended enterprise will define the standard for everyone else.

That is the work.

---

*M. Eric Johnson is Ralph Owen Dean Emeritus and Bruce D. Henderson Professor of Strategy at Vanderbilt University’s Owen Graduate School of Management. He serves as President of the Production and Operations Management Society (POMS) for the 2025–2027 term and serves as a director for the Swiss EoM Foundation. His research and teaching focus on digital strategies and the impact of information technology on the extended enterprise, with particular attention to healthcare quality, cost, and information risk. He is the author of *The Economics of Financial and Medical Identity Theft* (Springer) and editor of *Managing Information Risk and the Economics of Security* (Springer). His work has appeared in *Harvard Business Review*, *Wall Street Journal*, *Health Services Research*, *Management Science*, and the *Journal of the American Medical Informatics Association*.*

*Ed Magee is Chair of the NACD Nashville Chapter of the National Association of Corporate Directors, an independent board director at WD-40 Company (NASDAQ: WDFC), and COO of the Thomas F. Frist, Jr. College of Medicine at Belmont University.*