Rockin' HIT Sales

Episode Transcript

AI Governance in the Real World: What Vendors Need Before Provider Review

Guest: Yasir Tarabichi, MD, MSCR, Head of Digital & CMIO, Ovatient; Chief Health AI Officer, MetroHealth
Release Date: June 17, 2026
Transcript edited lightly for clarity and readability. Intro and outro omitted.

David Hacker (00:00)
Dr. Tarabichi, welcome to Rockin' HIT Sales. I really appreciate you joining me today to help our listeners better understand how health systems think about AI governance, responsible adoption, and what vendors need to know before bringing an AI-enabled tool into a provider environment.

Yasir Tarabichi (00:22)
Thanks for having me.

David Hacker (00:24)
Your roles span digital health, clinical informatics, Ovatient, and MetroHealth AI strategy. How do you describe your role in the AI governance conversation?

Yasir Tarabichi (00:36)
The common theme across that portfolio is digital transformation — and I would put a finer point on it as responsible digital transformation. I lead the organization's health AI strategy and governance structures. On the startup side, I am also part of Ovatient, a virtual care startup where we are bringing a solution built on Epic to enhance virtual care options across customer sites. The goal is a single platform, across several states, integrated into other platforms so patients receive contiguous care regardless of where they are. Information, orders, and expectations are shuttled across organizations in a meaningful way. As far as patients are concerned, it is all one system. I do a lot of work in interoperability and governance, bringing different groups together toward a common goal.

David Hacker (01:40)

When the term AI governance is used, what does that mean in practical terms inside the four walls of a health system?

Yasir Tarabichi (01:49)

It depends on who you are talking to. My definition is that AI governance is a way of working — almost an operating system. It is deciding how we are going to approach this technology from a digital transformation perspective. A lot of it comes down to clear expectations, accountability, and transparency for the organization. When I think about a strong AI governance structure, I expect everybody in the organization to know what we are trying to do, where we are headed, how we are going to get there, and who is responsible for what. We are doing this at scale and in a resource-constrained setting, so people need to understand where their contributions are and what they should or should not do.

David Hacker (02:39)

Why has AI governance become so urgent now compared with traditional digital health or clinical decision support oversight?

Yasir Tarabichi (02:51)

It has been a fascinating acceleration. In many ways, AI is the unlock we have been looking for in several spaces, including clinical decision support. In the past, we ran a lot of this through a serial process. We did waterfall development in healthcare systems, built solutions at the point of care, decided how they would work, defined the population and expectations, and then layered more and more solutions on top of the electronic health record. That process was slow, and over time it became onerous to manage and maintain. AI has come along and scaled the problem faster. Generative AI creates the ability to understand and use data locked away in the EHR in a much more fluid way. It pushes us to be more agile in our development process. The need is to be quick and resourceful, even though many of the problems we are trying to solve are the same ones healthcare has been working on for over a decade. We have been doing these things; now we need to do them better and faster.

David Hacker (04:33)

What does the front door for an AI product look like in a health system? How does an idea, internal use case, or vendor solution enter that process?

Yasir Tarabichi (04:34)

It varies a lot, and that is something we have been trying to clamp down on. Our evolution started by leading with the problems we were trying to solve. Early AI initiatives often came from squeaky wheels or leaders who saw a compelling product and wanted to understand the technology. For example, we needed to understand generative AI, so we looked at tools for summarizing patient charts. We needed to understand conversational agents, so we looked at voice solutions for specific patient situations. We were dipping our toes in the solution space.

Now, like many organizations, we have taken several steps back to think about overall strategy from a higher vantage point. We need to structure responsibility and compliance, bring the right people to the larger table, and make sure people from across the organization bring both subject-matter expertise and an understanding of what AI may unlock. We have separated steering and strategy into one group. That group asks what the organization should be thinking about over the next six, twelve, or eighteen months: what to invest in, how to upskill the workforce, and what solutions we want. That is a proactive approach rather than reacting every time someone sees a compelling product.

The higher-level steering committee looks at direction, which solutions we may want to pursue through an RFP or potentially build, and how to get our people to a place where they can build or co-develop solutions. That is separate from the responsible AI oversight group, which acts more like an IRB. It sets expectations at the organizational level for reasonable risk, how we define high, moderate, and low risk, what fits our portfolio, and what we need to monitor and mitigate risk when implementing solutions.

Then our operating groups and subject-matter experts use that guidance and rubric to think about solutions they want to bring forward. If it touches their area of work and they know the workflow, they do not necessarily need the higher levels of the organization involved as long as they follow the rules of the road for our AI

portfolio. There is a top-down strategy for expensive or high-impact things, infrastructure, education, and organizational priorities. There is also a bottom-up path. If you are a radiologist, you know your workflow better than I do. If the use case fits within the guardrails, then move forward rather than waiting for a broad approval.

The triage question becomes: Is this reasonable for the organization? Is it consistent with our goals and portfolio? Do we already have a solution that does this? If the answer is positive, then the question becomes prioritization. The steering committee may say, this is a niche solution that will take half of our resources for two years, and we cannot afford that because other priorities are larger. Or maybe we need outside help or more staffing. The path depends on where the solution sits. An enterprise-wide solution with broad impact needs to enter at a higher level. A lower-risk solution that solves a specific pain point, is affordable, fits our technology stack, and has a good monitoring and mitigation plan will not meet the same level of resistance.

David Hacker (10:30)

There are a lot of moving parts that listeners need to pay attention to, depending on where they fit in the process. How do you balance the need for innovation and speed with discipline, safety, and operational control?

Yasir Tarabichi (10:58)

It is all about risk. As soon as an application or solution makes its way into the system, an assessment is made about the risks associated with that intervention. If it is a back-office solution that automates something fairly deterministic and carries low potential risk to the organization, legal, compliance, and patients, then assuming the budget is there, it is not going to be a major issue — especially if it fits our portfolio and roadmap.

Anything that touches patient care, interfaces with patients, assists clinical decision-making, or carries legal risk is going to be flagged for a fuller review. That includes HR systems as well. Those solutions go through the responsible AI advisory committee and are also reviewed at the higher executive level through the steering committee. We try not to say no. Usually the answer is potentially, or yes, but. That is where the burden falls back on the developer or vendor to explain how they have mitigated the risk, why the risk may be lower than we think, and how they would monitor, enhance, or change the solution going forward. The level of attention something receives is a combination of cost, implementation cost, and risk to the organization.

David Hacker (13:04)

Are there things vendors tend to underestimate or steps they miss — whether related to security, privacy, integration, clinical risk, or something else?

Yasir Tarabichi (13:27)

Most vendors have figured out security and privacy. If a vendor does not know that placing patient data into a large language model can risk training that model on patient data, that is a problem. If a vendor oversimplifies privacy and says they can scrub a few PHI identifiers and therefore HIPAA is no longer relevant, they do not understand how combinations of data can re-identify patients. That usually happens when someone who is not healthcare-native steps into healthcare.

I would flip the question and ask: what are the signs of maturity? Vendors who impress us come with clarity about the problem they are solving and how it affects us as an organization. They also understand the lift involved in implementing and responsibly operationalizing the solution. They show not only that the solution has historically made a difference, but how they will show us, during implementation, whether it is making a difference. How are we monitoring? How are we identifying risk? Do we have a plan and pathway for mitigating that risk?

For example, if we notice a drop-off in accuracy for certain demographics or subgroups in our patient population, how will the vendor respond? That level of maturity always impresses me. It shows insight beyond saying the model is fine or the solution is nifty. The niche becomes change management and operationalizing the solution: we can get this into your system, make a difference, and work hand in hand with you. That is a phenomenal relationship when it happens. People often get stuck at, 'I built a really impressive model.' That is only the beginning of the journey. In healthcare, it is not just the size of the AUC that matters; it is how you

use it. When a vendor brings a plan and can explain how they have done this and how they would do it with us, that builds trust early.

David Hacker (16:45)

Who is actually monitoring the AI to make sure there is not drift? Have there been situations where an AI product needed to be paused or taken a step backward until issues were resolved?

Yasir Tarabichi (17:02)

Yes, that has happened. We want to see a little bit of the sausage making. For example, we had a vendor develop a model to predict a patient's risk of readmission or unplanned admission, and some variables looked directionally incoherent. That gave us pause. We stopped, went back into the model, figured out how it was trained, tried to understand why we were seeing the signal we were seeing, and even made recommendations about how to develop a better solution.

That level of give and take is important, as long as there is transparency and a true eagerness not to oversell but to co-develop. That improves trust and credibility. Black boxes are not going to work. There are situations where we have a wealth of data and need to work with the solution developer to marry the data together in a way that lets us understand impact. Some things are not immediately reportable from the vendor perspective, so we establish pipelines of data exchange. Some of our most successful implementations had data and dashboarding first, where we could see good, bad, or deviant signals early. It is often a mashup of vendor data and data surfaced within our system, matched to patient or workflow, so we can look at the impact together.

David Hacker (18:53)

At what point in the sales process should an AI-enabled Health IT company start addressing AI readiness and governance expectations? Early conversations are often with operational or clinical leaders who may not be deeply technical, so how should founders explain readiness, risk, monitoring, and governance in a way that is credible without overwhelming the team they are speaking with?

Yasir Tarabichi (19:27)

It depends entirely on the audience. I will give a clinical example. If a solution can scan an ultrasound of the heart and predict the likelihood of a certain cardiac condition, what I want the head of cardiology, echocardiography, or imaging to think through is: how does this work for me, my patients, our population, and our workflow? And how are we different? As a safety-net healthcare system, MetroHealth is different. We often sit on the wrong side of validation assumptions about how a solution will perform for our patient population.

That is part of the upskilling. When people come to us and say they want to talk about AI solutions, we explain the things we care about as a safety-net system. I have seen that work and not work. When it works, the clinical leader asks the vendor, how does this work for my patients? The vendor may answer that they have run it on highly diverse demographics and have a process to monitor it in our system. When it does not work, the vendor says, don't worry about it.

We like to place our responsible AI intake process as early in the conversation as possible. One thing that may come up is that we already have a solution, or we have one that is better than what the vendor is presenting. That brings us back to governance as an operating system: everyone needs to understand what we are doing and what we already have. Vendors, even if the end users do not ask the AI-readiness question, can still position themselves in a healthy way so that when the AI-informed part of the organization hears about the conversation, it is clear that these issues were considered from the start.

If my end user comes back and says the vendor has the data, says it will work for all segments of our patient population, and will help us monitor to make sure it does, those few sentences are highly valuable. They show that we are aligned on expectations and that the conversation is likely to work. What does not work is when vendors check boxes on intake material and then, in the next conversation, admit they do not actually do what they claimed. If you say you check for bias, what does that mean? How does it carry over to us? How can we trust it? We need to be able to look over your shoulder and see the work. These solutions require more

conversation and collaboration than static solutions. The level of involvement required for AI solutions is much heavier than what we have talked about in the past.

David Hacker (23:03)

We have reached what I call the lightning wrap. What should companies just stop saying or doing when they are talking to a health system about an AI solution?

Yasir Tarabichi (23:21)

Leading with the fact that it is an AI solution.

David Hacker (23:27)

Complete this sentence: a healthcare AI company is not ready for a health system review until...

Yasir Tarabichi (23:40)

...it can explain how the solution will be implemented, monitored, and governed, not just how it performs.

David Hacker (23:47)

Dr. Tarabichi, thank you so much for joining Rockin' HIT Sales. This was incredibly timely and practical. I know our listeners are walking away with a much clearer understanding of what AI-enabled Health IT companies need to prepare before they sit in front of providers: the evidence, the workflow context, the risk thinking, the monitoring plan, and the ability to explain AI readiness in a way that both technical and operational leaders can trust. Thank you for sharing your time and perspective.

Yasir Tarabichi (24:24)

Thank you for having me. It has been a pleasure.