USF | Muma Case Review
A publication of the Muma College of Business | University of South Florida

Volume 10 Number 5 17 DECEMBER 2025

ABBIGAIL SANCHEZ, CHRISTINA ZELEZNIK, ELISA WILSON, KRISTIE MARTINEZ, MARISSA MORTON

TAMPA GENERAL HOSPITAL: ARTIFICIAL INTELLIGENCE - DREAM OR DILEMMA?

How far should a hospital go to embrace AI while protecting patients and maintaining public trust?

It was late on a Thursday afternoon at Tampa General Hospital (TGH). The executive floor had started to quiet down, but in a corner office overlooking the bay, Anthony Zalet, Director of Digital Innovation, remained at his desk, facing one of the most consequential decisions of his career. TGH was not simply choosing among three Artificial Intelligence (AI) vendors. Instead, Anthony was leading the hospital through a far more complex question: How should a major academic medical center responsibly establish an enterprise-wide process for evaluating, deploying, and governing AI tools?

A white paper from one vendor lay open before him, filled with promises: predictive diagnostics, automated workflows, and optimized operations. But Anthony understood that TGH's challenge was bigger than selecting a product. The hospital needed a structured, scalable process for integrating AI, one that could guide adoption while navigating risks such as privacy, bias, compliance, and clinician trust.

Earlier that day, in a high-stakes leadership meeting, executives debated how to balance the pressures of modernization against the responsibilities of care. Peer institutions were already piloting AI in radiology and sepsis detection. Yet voices like the Chief Medical Officer urged caution: "We can't just plug in a black box and hope it improves outcomes. What happens if it gets something wrong, or introduces bias into our care?"

Anthony realized TGH lacked a comprehensive AI governance structure. While the hospital had a history of innovation, from robotic surgery to precision medicine, AI brought new and unfamiliar risks. Cybersecurity warned of vulnerabilities. Compliance officers flagged regulatory complexity under Health Insurance Portability and Accountability Act (HIPAA) and Food and Drug Administration (FDA) guidelines. Ethics advisors cautioned about transparency and accountability. And yet, the opportunity was undeniable: AI promised better decision-making, faster diagnoses, and more personalized care.

As the sun dipped below the Tampa skyline, Anthony stared out the window and returned to the central dilemma: If TGH embraces AI, what should the process look like? Who needs to be at the table? And how can the hospital ensure that innovation advances without compromising safety, equity, or public trust? This case examines the strategic crossroads facing TGH: whether to proceed with vendor-led AI adoption, pause to build internal governance capacity, or pursue in-house development of AI tools. It explores how the hospital can create a responsible, repeatable framework for AI evaluation that protects patients, supports clinicians, and upholds public confidence in a rapidly evolving digital healthcare landscape.

AI in Healthcare

AI was increasingly integrated into healthcare, supporting a wide range of functions from administrative operations and diagnostics to patient engagement and medication discovery. Driven by advancements in machine learning, AI solutions were expected to significantly enhance both clinical outcomes and operational efficiency (Zhang & Zhang, 2023). Common applications in the healthcare sector included clinical decision support (CDS) tools, predictive risk scores, natural language processing for medical documentation, imaging diagnostics and pathology screening, and patient engagement bots (Mount Sinai Health System, 2025; Ghassemi et al., 2024). These tools were already improving patient throughput, increasing diagnostic accuracy, and streamlining administrative workflows in leading hospitals across the country.

The overarching goal of AI adoption was to reduce provider burden, accelerate diagnosis, and create more personalized and data-informed care, addressing both workforce shortages and patient satisfaction challenges. Institutions like Mayo Clinic and Cleveland Clinic had successfully deployed AI to improve scheduling, triage, and support early intervention (Cleveland Clinic, 2024).

Despite the benefits, AI presented considerable risks including data privacy, algorithmic bias, lack of transparency and blurred lines of accountability. A successful governance framework was complex as it relied on sensitive patient data and had to prioritize fairness, interpretability, and compliance (Coalition for Health AI, 2023; Hassan et al., 2025). In response to these challenges, regulators and policymakers continued to articulate clearer expectations.

About TGH

Anthony Zalet understood that more was at stake than simply implementing a new AI tool. For over a century, TGH had stood as a pillar of healthcare excellence in Florida and beyond. As one of the largest academic medical centers in the state, TGH was nationally recognized for its cutting-edge treatments, innovative care models, and unwavering commitment to community service (see Exhibit 1).

TGH had long embraced innovation, from robotic surgery to precision medicine. But in today's rapidly evolving healthcare landscape, the hospital faced mounting pressure to adopt AI and other digital transformation technologies. While AI held the potential to improve diagnostic accuracy, operational efficiency, and patient outcomes, it also raised pressing questions about privacy, equity, and safety, questions that TGH, as a trusted healthcare leader, had to proactively address.

Under Anthony's leadership, the Digital Innovation Office at TGH focused on advancing care through technologies that align with three guiding principles (Arnold, 2024):

- Clinical Excellence
- Patient-Centered Innovation
- Operational Efficiency

Each innovation, whether in predictive analytics, patient flow optimization, or clinical decision support, was assessed not only for its technical capabilities but also for its alignment with TGH's mission to provide world-class care to all patients. Rather than outsourcing innovation, TGH believed in co-creating solutions through strategic partnerships with technology vendors, academic researchers, and internal clinical teams (TGH, 2025). This collaborative model ensured that every solution was rigorously tested, ethically sound, and tailored to the diverse needs of the hospital's patient population.

TGH's emphasis on governance, safety, and ethical deployment set it apart from other health systems racing to implement AI. At TGH, innovation wasn't just about being first, it was about getting it right. Through interdisciplinary task forces, continuous clinical validation, and strict adherence to healthcare regulations such as HIPAA and FDA guidelines, TGH was building a model for AI integration that patients, clinicians, and regulators could trust (Parker, Economou, & Silcox, 2024).

In an era where many hospitals were struggling to balance innovation with integrity, Tampa General continued to lead by example, anchoring every decision in its core values of compassion, accountability, excellence, and safety (TGH, 2024). The hospital's long-standing reputation and its commitment to putting patients first provided a powerful foundation for shaping the future of healthcare, ethically, responsibly, and equitably.

Governance Criteria

As part of developing a comprehensive and repeatable process for evaluating AI solutions at TGH, Anthony Zalet began by meeting with key internal and external stakeholders. These conversations focused on defining the broader criteria and governance structures necessary to guide AI adoption across the organization. Without clear criteria, decisions risked being driven by vendor marketing rather than institutional needs or patient safety.

Anthony engaged with leaders such as Jason Swoboda (Director of Innovation for TGH Ventures), Justin Norden (CEO & Co-Founder of Qualified Health), and Tom Kirby (CFO, Coalition for Health AI, Inc.) to surface insights on how healthcare systems should approach AI governance. These discussions revealed several foundational themes, chief among them was the recognition that advanced AI applications, such as generative models, often relied on large, diverse datasets to function optimally. However, healthcare regulations imposed strict constraints on data usage, making it difficult to realize the full potential of such tools without jeopardizing compliance.

Stakeholders strongly advocated for the establishment of formal governance processes to assess, monitor, and validate any AI tool considered for use at TGH. This included calls for a standing AI governance committee responsible for overseeing all phases of the AI lifecycle: evaluation, deployment, ongoing performance monitoring, and model updates. Their consensus was clear: a strong governance infrastructure is essential to ensuring AI tools remain safe, effective, and aligned with clinical and ethical standards.

They also emphasized that successful AI integration required collaboration across the healthcare ecosystem. Cross-sector partnerships with vendors, academic researchers, investors, and regulatory bodies would be critical to establishing industry standards, sharing best practices, and fostering widespread, responsible adoption of AI in clinical environments.

The competitive AI landscape remained fragmented. Some companies prioritized regulatory compliance, others focused on real-time model integration or post-deployment monitoring. While this diversity created a challenge for decision-makers, it also highlighted the need for TGH to implement a structured, criteria-driven approach that could evaluate vendors consistently and transparently, regardless of market shifts or technological advancements (see Exhibit 2).

As Anthony and his team reflected on these findings, it became clear that TGH needed to build a long-term governance strategy that could support safe, scalable AI adoption across the organization. With

foundational governance needs defined, Anthony turned his attention to the practical challenges TGH now had to address.

Implications for TGH

Building on the governance foundations discussed, TGH began considering practical implications for vendor engagement and long-term oversight. Rather than making a one-time vendor decision, TGH was setting the foundation for a broader institutional framework—one that would allow the hospital to evaluate, select, and monitor AI tools in a consistent, repeatable, and compliant manner. In building this process, several key considerations had emerged:

1. Use Case Viability: Given healthcare's regulatory constraints around data sharing, not all AI use cases were feasible. TGH had to prioritize applications that offered clear clinical value while remaining fully compliant with laws such as HIPAA and FDA guidelines.
2. Data Pooling Requirements: Some AI vendors relied on pooled data across multiple institutions to refine their models. TGH had to carefully assess whether participating in such arrangements would compromise its data protection obligations.
3. Contractual Safeguards: Any engagement with a third-party vendor had to include robust contractual provisions to protect regulated data by ensuring proper data usage rights, auditability, and adherence to regulatory standards.

Outstanding Questions for TGH's Governance Path Forward:

Despite this early vendor analysis, several important questions remained:

- Given data privacy constraints, could TGH fully leverage the advanced capabilities of AI models that required vast, pooled datasets?
- Should TGH continue engaging with vendors like Specta Health, even amid uncertainty about client onboarding timelines?
- How should TGH assess the total cost of ownership—including downstream integration, clinician training, and ongoing compliance—not just upfront vendor pricing?
- What structures had to be put in place to ensure that any AI solution integrated seamlessly into existing workflows and would be supported long-term by both internal and external teams?

These questions went beyond a single procurement cycle—they underscored the deeper challenge TGH faced: how to build a sustainable, trustworthy, and future-ready approach to evaluating and adopting AI in healthcare.

Risk Assessment & Compliance

To effectively integrate AI healthcare solutions at TGH, it was essential to establish a robust risk assessment and compliance framework (see Exhibit 5). This framework would be pivotal in ensuring that AI technologies were implemented in a way that minimized risks associated with patient data security, clinical outcomes, and adherence to healthcare regulations. The lack of a formal AI governance structure at TGH presented significant risks, including potential breaches in patient data security, biases in AI-driven decisions, and violations of regulatory standards. Addressing these risks through a comprehensive AI risk assessment and compliance strategy was therefore imperative for maintaining trust, patient safety, and operational integrity.

One of the core concerns with the adoption of AI technologies in healthcare was the security of sensitive patient data. AI systems required vast amounts of personal health information to train algorithms and provide actionable insights. Without proper oversight, there was an increased risk of data breaches or unauthorized access, potentially violating HIPAA and other data protection laws. An AI governance framework for TGH had to prioritize data encryption, ensuring that all patient data, both in transit and at rest, was encrypted using industry-leading protocols (Parker, Economou, & Silcox, 2024). Strict access controls should limit patient data access based on user roles. Furthermore, continuous monitoring of AI systems was essential to detect any unusual activity that may indicate a security breach. To mitigate these risks, TGH had to partner with external cybersecurity firms and leverage AI-driven security solutions designed to safeguard patient data in real-time. Additionally, periodic audits and vulnerability assessments would help identify and address potential weaknesses in the system.

Another critical concern was the potential for bias in AI healthcare solutions, which could lead to inequities in care. For instance, algorithms trained on skewed data sets could produce inaccurate or biased outcomes that could adversely impact patient health, particularly in underrepresented populations (Stetler, 2024). The AI governance framework should include processes for conducting regular bias audits to identify and mitigate biases in AI algorithms, especially those affecting decision-making in diagnosis, treatment recommendations, and patient outcomes. Transparency in algorithm development was also crucial to ensure that models were explainable and that clinicians could understand the decision-making processes behind AI-generated recommendations. Additionally, AI systems should be trained on diverse and representative data sets to minimize the risk of systemic biases in treatment outcomes. By incorporating these strategies, TGH could reduce the risk of AI perpetuating healthcare disparities and ensure that its AI solutions provided equitable care for all patients.

Compliance with healthcare regulations such as HIPAA, FDA guidelines, and local healthcare laws was non-negotiable for any AI solution adopted by TGH. Following best practices from the National Institute of Standards and Technology (NIST) AI Risk Management Framework and Coalition for Health AI, TGH's governance framework should address transparency, fairness, and lifecycle risk mitigation (Coalition for Health AI, 2023; NIST, 2023). Failure to comply with these regulations could lead to significant legal and financial consequences, as well as damage to TGH's reputation.

The governance framework had to include a clear mapping of relevant regulatory requirements for AI in healthcare, both at the federal and state levels, to ensure that any AI tool adopted complied with the necessary laws and standards. To support internal governance discussions, a summary of key legal and ethical frameworks was prepared to guide AI evaluation and oversight (see Exhibit 6). Additionally, ongoing compliance checks were essential to ensure that AI systems remained aligned with regulations. If TGH partnered with third-party AI solution providers, it also had to ensure that these partners adhered to the same regulatory standards and that any software or system deployed had undergone rigorous testing and approval. Engaging legal and regulatory experts specializing in healthcare AI would be essential for regularly reviewing AI deployments and ensuring that all systems remained in full compliance.

The goal of implementing AI in healthcare was to improve patient outcomes and enhance operational efficiency. However, without proper evaluation, AI solutions could fail to meet clinical needs or improve workflows, leading to inefficiencies, misdiagnoses, or adverse effects on patient care. The governance framework should incorporate a rigorous pre-implementation validation process for evaluating AI technologies before they were integrated into clinical workflows, ensuring they met the standards of clinical effectiveness. Post-implementation monitoring was also crucial, with continuous oversight of AI systems after deployment to ensure they operated as intended, along with regular feedback loops to refine algorithms and improve outcomes. Including clinical staff, data scientists, and operational teams in the

evaluation process would ensure that the AI tools aligned with hospital priorities and improved patient care delivery.

TGH would also need to collaborate with specialized AI and healthcare governance partners to ensure that its AI risk assessment and compliance efforts were comprehensive and effective. Identifying reputable partner organizations that specialized in healthcare AI governance, risk management, and regulatory compliance would be crucial to the success of this initiative. These partners could offer expertise in areas such as developing and implementing customized risk assessment frameworks, providing guidance on navigating the complex healthcare regulatory landscape, and offering training programs to build internal expertise on AI governance, compliance, and security best practices. Through these partnerships, TGH could develop a strong, sustainable AI governance framework that not only mitigated risks but also fostered innovation in healthcare technology.

Establishing a structured risk assessment and compliance framework for AI at TGH was essential for the safe, ethical, and compliant integration of AI technologies. By addressing patient data security, bias, regulatory adherence, and clinical effectiveness, TGH could mitigate potential risks while maximizing the benefits of AI in healthcare. Collaborating with external partners would further strengthen these efforts, ensuring that TGH remained at the forefront of AI innovation while safeguarding patient safety and operational integrity. Ultimately, the vendor TGH selected not only had to meet today's compliance and integration needs but also align with long-term governance structures, including post-deployment monitoring, clinician training, and ethical revalidation. With risk mitigation strategies outlined, the focus turned to validating AI tools before real-world deployment.

Deployment & Validation

With vendor comparisons in progress, Anthony was now focused on developing a scalable validation process that TGH could apply to any AI tool, both current and future. Although vendor selection remained important, the greater challenge lay in creating a reliable, repeatable method for evaluating AI performance in real-world clinical settings.

Rather than rushing into deployment, Anthony was proposing a phased, institutional approach grounded in the core principles of safety, transparency, and clinician trust. Drawing lessons from Mayo Clinic and Cleveland Clinic, he envisioned a validation model that could serve as a blueprint for all future AI integrations ensuring that each solution, regardless of origin, underwent rigorous clinical and technical scrutiny (Cleveland Clinic, 2024; Ghassemi et al., 2024).

The proposed validation plan began with simulation testing. Using retrospective data, the team would assess how well a given AI model performed on key predictive tasks such as identifying high-risk patients. This phase would also stress-test the model against edge-case scenarios and ambiguous inputs to evaluate its robustness in variable clinical environments. Following FDA's Total Product Lifecycle (TPLC) guidance for AI/ML-based Software as a Medical Device (SaMD), the team was also considering implementing a Predetermined Change Control Plan (PCCP) outlining how AI models would be monitored and revalidated over time as updates were introduced (FDA, 2021).

If simulation benchmarks were met, the next phase would explore a controlled pilot implementation. Anthony envisioned launching pilots within select inpatient units, with close collaboration from clinical leadership, innovation staff, and frontline providers. The pilot would be structured to evaluate not just model performance, but practical concerns such as usability, documentation burden, and impact on

workflow. Clinician feedback would be collected through structured debriefs and interactive dashboards (see Exhibit 7), helping to avoid the pitfalls of opaque, 'black-box' AI systems (Ghassemi et al., 2024). This validation process would be informed by Hassan et al.'s governance framework, which stressed explainability, interpretability, and model robustness. The team would explore integration of Explainable AI (XAI) techniques to make model recommendations transparent and understandable to clinicians—aligning with the World Health Organization's ethics principles for AI (WHO, 2021). Additional scenario-based testing would assess how the AI system performed with rare diagnoses or atypical clinical cases to determine its reliability across a diverse patient population.

All models undergoing evaluation would be subject to the principles of a Model Risk Management Framework (MRMF), with independent validation, comprehensive documentation, and full transparency in algorithmic design. TGH would track inputs, model architecture, and decision boundaries through audit logs, ensuring regulatory readiness and institutional accountability (Zhang & Zhang, 2023; NIST, 2023). To meet both ethical and legal standards, Anthony's team would explore documentation tools such as audit trails and explainability reports. These would provide clear insights into how models make decisions—essential for building clinician trust. Quantitative benchmarks such as sensitivity, specificity, and confidence intervals would be included in validation reports to demonstrate clinical efficacy and operational improvements.

Before any full deployment, the validation process would include a cross-functional review involving legal, compliance, cybersecurity, and clinical leadership. These stakeholders would verify that the solution met HIPAA standards, FDA SaMD regulations, and internal expectations around safety and transparency (U.S. Department of Health & Human Services, 2022; FDA, 2021).

In parallel, bias audits would be conducted to ensure that the model treated all patients equitably. Tools and strategies for detecting algorithmic bias would be embedded into the validation phase, ensuring that AI deployment did not reinforce disparities in care (Hassan et al., 2025). Models would only move forward if they met fairness criteria and showed minimal risk of harm across diverse patient populations. If a tool successfully cleared all phases, the PCCP would guide post-deployment changes. Clear criteria would be defined to trigger re-validation when updates were made ensuring that the model adapted safely to evolving patient populations or clinical protocols without compromising performance or transparency. Through this carefully structured process, TGH was not just evaluating a tool; it was creating a durable validation framework: One that would ensure that every AI solution used in the hospital was clinically sound, ethically grounded, and aligned with the institution's mission to deliver safe, equitable, and high-quality care.

Post-Monitoring & Oversight

As part of its broader AI governance framework, TGH had to establish a robust, institution-led system for post-deployment monitoring to ensure that any AI tool used in clinical or operational workflows continued to perform reliably, ethically, and safely over time. Anthony understood from global case studies, such as those from the NHS and Mount Sinai Health System, that even well-validated models could degrade after deployment due to factors such as data drift, changing patient populations, or unanticipated clinical use cases (Mount Sinai Health System, 2025; NHSX, 2020).

Rather than relying solely on vendor-provided dashboards, Anthony envisioned a layered oversight strategy that empowered TGH to track ongoing model performance through internal review protocols, independent audits, and real-time alert systems. This approach would ensure that TGH maintained control and transparency, regardless of the vendor selected.

To ensure consistency, all future AI deployments would be required to conform to a set of vendor-agnostic performance monitoring standards. These would include continuous tracking of key performance indicators (KPIs), such as diagnostic accuracy, false positive and negative rates, model drift, and fairness across demographic subgroups. Any decline in performance or evidence of biased outcomes would trigger a formal review and, if necessary, temporary suspension of the tool.

Vendors partnering with TGH would be contractually obligated to support this monitoring infrastructure, including provisions for:

- Real-time performance reporting,
- Transparent model versioning and changelogs,
- Ongoing vulnerability assessments, and
- Demonstrated adherence to HIPAA, FDA SaMD, and other regulatory requirements.

Importantly, Anthony's oversight framework was designed to minimize institutional blind spots. Independent monitoring conducted by interdisciplinary teams including clinical leaders, data scientists, and compliance officers would provide objective analysis beyond vendor assurances. This layered approach ensured accountability not only to the developers but also to the patients and clinicians who relied on the tools.

By institutionalizing a continuous monitoring and oversight process, TGH aimed to close the loop on AI governance, transforming what might otherwise have been a one-time technology rollout into a dynamic system of stewardship, risk mitigation, and iterative improvement. In doing so, TGH could set a precedent for responsible AI use in healthcare: one that prioritized not only innovation, but the long-term trust of its patients and providers.

Oversight Protocols and Response Framework

To maintain accountability and agility, Anthony would implement a multi-layered oversight structure. Monthly AI oversight audits, led by the Innovation Department and overseen by Rachel Feinman, VP of Innovation, would review bias reports, performance logs, and clinician feedback. Drawing inspiration from the Mount Sinai case, he would also deploy a Bias & Equity Dashboard that flagged potential disparities in AI recommendations based on race, gender, age, or socioeconomic status (Mount Sinai Health System, 2025). An Incident Escalation Process would be established to ensure that any model behavior suggesting potential harm or systemic bias triggered an automatic halt in usage, followed by a rapid-response review by clinical and data ethics teams. Additionally, a Feedback Loop Mechanism would allow clinicians to submit real-time feedback through a dedicated portal, with reports triaged weekly and integrated into governance huddles to support continuous learning and responsiveness.

Continuous Documentation

All post-deployment activity would be logged for compliance and auditing purposes. Performance Monitoring Logs would track key metrics such as accuracy, model confidence intervals, override rates, and adverse events. Bias & Fairness Reports would compare model performance across patient subgroups to ensure demographic parity. Audit Trails would document and timestamp every decision, review, or retraining event for regulatory transparency. The Version Control Register would catalog each model update, including performance benchmarks, retraining methods, and stakeholder approvals.

Anthony would ensure that each model iteration underwent a mini-validation process before redeployment. This would include clinician sign-off, fairness auditing, and updated documentation to ensure continuous alignment with regulatory and ethical standards.

To further support accountability, Anthony planned to create an AI "Model Registry" at TGH. The registry would serve as an institutional memory, facilitating transparency and reducing knowledge silos during audits or leadership changes (see Exhibit 8). This registry would include metadata such as intended use, training data summaries, validation history, monitoring logs, and stakeholder contacts, ensuring both institutional memory and regulatory readiness for inspections or audits. In addition, provisions for real-time monitoring of patient data integrity and user access would be implemented through AI-enabled cybersecurity solutions. These systems would help detect unauthorized access or anomalous behavior, enhancing compliance with HIPAA security rule requirements (U.S. Department of Health & Human Services, 2021).

Through these strategies, Anthony would not only deploy innovative AI tools but also establish a precedent for responsible AI governance at TGH, one that balanced innovation with transparency, clinical safety, and equity. His structured validation and post-monitoring approach would ensure that AI evolved responsibly and worked in harmony with clinical care.

Vendor Evaluation: Applying the Framework to Vendors

Vendors would remain an essential part of TGH's AI strategy. These third-party providers served as critical partners in delivering, deploying, and maintaining AI tools tailored to healthcare environments (Bannister, 2024). But selecting vendors was only one element of the equation. TGH's long-term success depended on embedding evaluation and oversight practices into the organization's core innovation strategy.

As a pilot exercise in applying this framework, Anthony initiated a structured comparison of three leading vendors—Specta Health, Newton's Tree, and Qualified Health—each offering unique strengths aligned with different governance goals (see Exhibit 3). Though no final selection had been made, the comparison process had provided valuable insights into how TGH might evaluate future vendors based on standardized criteria, including: data security, customization & scalability, vendor support & training, cost efficiency, and timeline to launch.

Each vendor represented a distinct approach to AI integration. Newton's Tree offered a vendor-neutral platform focused on seamless AI integration and ongoing performance management. Specta Health emphasized regulatory compliance and adoption processes, providing robust support for validation and transparency. Qualified Health specialized in governance and ethical oversight of generative AI, with a particular focus on post-deployment monitoring and responsible model management (see Exhibit 4).

By testing its evaluation criteria through this comparative process, TGH had begun to identify critical gaps and strengths in the market insights that would inform future procurement decisions and internal governance strategies.

TGH at a Turning Point

As the sun lowered over TGH, Anthony knew that his decision would shape not just TGH's adoption of AI, but its very identity as a leader in responsible innovation. The potential benefits of AI were promising: enhanced clinical decision-making, operational efficiency, and personalized patient care. Yet

the road ahead was layered with uncertainty. Without a formal governance structure, TGH faced critical questions about how to ensure transparency, safety, and regulatory compliance amid rapid technological advancement.

Over the past few weeks, Anthony and his team had shifted their focus from products to processes, recognizing that TGH's long-term success with AI depended not on a single solution but on a thoughtful framework that could guide implementations for years to come. They had initiated the development of evaluation criteria, engaged key stakeholders, and begun to outline oversight strategies; however, significant decisions still lay ahead as the process continued.

Anthony reflected on TGH's mission to be the safest, most innovative system in America, a model for responsible use of AI in healthcare. There was uncertainty around how the team could identify vendors that truly aligned with TGH's mission and values, and how they might structure internal oversight beyond implementation. At the same time, TGH had to find ways to mitigate risks like data privacy violations, regulatory non-compliance, and the threat of algorithmic bias while following HIPAA, the FDA, and NIST AI Risk Management Frameworks.

Even if a tool met these standards, the question of deployment remained: How could TGH validate AI tools in a way that built, rather than eroded, clinician trust? What did safe implementation look like in a clinical setting where patient outcomes were on the line? And once deployed, how would the hospital track model performance, catch early signs of drift or inequity, and ensure that updates were properly managed over time?

The urgency was real. Competitors were moving fast and vendors were eager to collaborate. Anthony knew that his decision would impact the organization for years to follow. TGH wasn't merely deciding whether to adopt AI; it was determining how to govern innovation in a way that balanced technological advancement with accountability, equity, and public trust.

The Dilemma

Anthony was at a crossroads, and TGH had to decide when and how to act. Each potential path offered opportunity, but carried its own risks.

TGH could engage a vendor at this stage to support the development of the governance framework, achieving immediate improvements in operational efficiency and clinical care while signaling organizational momentum and strategically positioning the hospital alongside peer institutions leading in AI adoption. However, without a fully formed governance framework, this move could expose TGH to serious risks, including vendor dependance, inefficiencies in model performance, diminished trust, regulatory violations, or unintended bias.

The second path was to pause by temporarily suspending AI vendor selection and deployment, instead focusing on building an internal framework. TGH had the opportunity to build a more resilient foundation, incorporating a multidisciplinary governance team, standardized protocols, and value-aligned policies. This could slow hospital innovation in the short term, but when AI was implemented, it would be done with intention, safety, and transparency.

Lastly, TGH could also walk away from vendor-driven AI solutions entirely. Instead, TGH could invest in proprietary internal AI capabilities more closely aligned with its unique institutional priorities. While

this would give TGH the highest level of control, it also came with longer timelines and the risk of falling behind industry peers. It also required investment and technical expertise. This could delay AI benefits and create operational gaps if resources were stretched.

The decision now rested in Anthony's hands: How should TGH proceed with AI adoption? How could it ensure that implementation was guided by the hospital's core values of safety, equity, and operational integrity? The outcome of Anthony's decision would define not only when and how AI was adopted, but also what kind of organization TGH would become in a rapidly evolving healthcare environment.

References

American Medical Association. (2021). *AMA adopts new policies to guide ethical use of AI*. https://www.ama-assn.org/delivering-care/public-health/ama-adopts-new-policies-guide-ethical-use-ai

Arnold, S. (2024, September 23). Building a culture of innovation at Tampa General Hospital [Podcast episode]. Healthcare IT Leaders. https://www.healthcareitleaders.com/podcast/building-a-culture-of-innovation-at-tampa-general-hospital/

Bannister, T. (2024, September 4). *Efficient collaboration between AI vendors and hospitals with Specta Health*. Specta Health. https://www.spectahealth.com/post/efficient-collaboration-between-ai-vendors-and-hospitals-with-specta-health

Cleveland Clinic. (2024, August 27). *Consider this comprehensive approach for evaluating AI*. Consult QD. https://consultqd.clevelandclinic.org/consider-this-comprehensive-approach-for-evaluating-ai

Coalition for Health AI. (2023). *Blueprint for Trustworthy AI Implementation Guidance and Assurance for Healthcare*. Coalition for Health AI. Retrieved from https://www.chai.org/news/chai-tm-unveils-blueprint-for-trustworthy-ai-in-healthcare

Food and Drug Administration. (2021). *Artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD) action plan*. https://www.fda.gov/media/145022/download

Ghassemi, M., McCoy, L. G., Majumder, M. S., Thadaney-Israni, S., & Patel, V. (2024). Embedding internal accountability into health care institutions for artificial intelligence. *NEJM AI, 1*(1), 100083. https://doi.org/10.1016/j.nejai.2024.100083

Hassan, M., Borycki, E. M., & Kushniruk, A. W. (2025). Artificial intelligence governance framework for healthcare. *Healthcare Management Forum, 38*(2), 125–130. https://doi.org/10.1177/08404704241291226

Mount Sinai Health System. (2025, April 7). *Is AI in medicine playing fair?* MountSinai.org. https://www.mountsinai.org/about/newsroom/2025/is-ai-in-medicine-playing-fair

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)*. https://www.nist.gov/itl/ai-risk-management-framework

NHSX. (2020). *A buyer's guide to AI in health and care*. https://transform.england.nhs.uk/documents/38/NHSX_A_Buyers_Guide_to_AI_in_Health_and_Care.pdf

Parker, V. J., Economou, N. J., & Silcox, C. (2024, October 28). *White paper: AI governance in health systems—Aligning innovation, accountability, and trust*. Duke-Margolis Center for Health Policy. https://healthpolicy.duke.edu/publications/ai-governance-health-systems-aligning-innovation-accountability-and-trust

Stetler, C. (2024, November 4). *AI algorithms used in healthcare can perpetuate bias*. Rutgers University. https://www.newark.rutgers.edu/news/ai-algorithms-used-healthcare-can-perpetuate-bias

Tampa General Hospital. (2024, June 11). *Tampa General Hospital deploys artificial intelligence tools to elevate patient experience*. https://www.tgh.org/news/tgh-press-releases/2024/june/tgh-deploys-nuance-artificial--intelligence-tools-elevate-patient-experience

Tampa General Hospital. (2025). *About Tampa General Hospital*. https://www.tgh.org/about-tgh

U.S. Department of Health & Human Services. (2022). Fiscal Year 2021 Annual Performance Report. U.S. Department of Health & Human Services. Retrieved from https://www.hhs.gov/sites/default/files/fy2021-performance-report.pdf

World Health Organization. (2021). *Ethics and governance of artificial intelligence for health: WHO guidance*. https://www.who.int/publications/i/item/9789240029200

Zhang, J., & Zhang, Z. (2023). Ethics and governance of trustworthy medical artificial intelligence. *BMC Medical Informatics and Decision Making, 23*(1), 7. https://doi.org/10.1186/s12911-023-02103-9

Biographies

Abbigail Sanchez earned her Bachelor's degree in Business Administration from St. Petersburg College and is currently pursuing her MBA at the University of South Florida. She brings a strong background in financial services, having worked in personal banking at two major U.S. institutions. Abbigail currently serves in the fraud department at USAA Federal Savings Bank, her professional experience in banking operations and fraud analytics informs her interest in responsible AI governance.

Christina Zeleznik is a venture capital intern for Tampa General Hospital and a current MBA & MHA student at the University of South Florida. Set to graduate in May 2025, Christina provides her knowledge of experience in healthcare investments, analytics, and operations. She has worked for several large healthcare institutions including Tampa General Hospital, Johns Hopkins All Childrens' Hospital, and Atlantic Health System. Christina is passionate about improving the US healthcare systems and patient care experiences.

Kristie Martinez is a cardiac registered nurse and advanced cardiac sonographer with a strong foundation in cardiovascular imaging and patient care. She holds a Bachelor of Science in Nursing from Keiser University, an undergraduate degree in Liberal Studies from Mercer University School of Medicine, and a specialized degree in Cardiac Sonography. Currently pursuing her MBA with a concentration in Healthcare Leadership, Kristie is passionate about combining clinical excellence with strategic insight to improve healthcare systems, elevate team performance, and enhance patient outcomes.

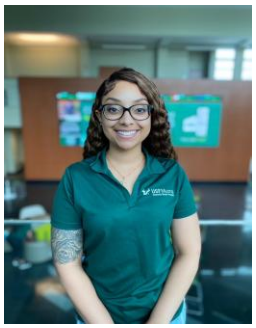

Elisa Wilson is an MBA Associate at Moffitt Cancer Center, working in the Entrepreneurship and Innovation department under the Chief Business Officer of Finance. She earned her bachelor's degree in Entrepreneurship and Innovation and is currently pursuing her MBA at the University of South Florida. Elisa brings a strong interest in business strategy and innovation, with a growing focus on impactful, mission-driven ventures.

Marissa Morton graduated with honors from Midwestern State University with a Bachelor's degree in Business Administration. She is currently completing a Master's degree in Business Administration at the University of South Florida. Since 2019, Marissa has served as a Human Resources Officer in the United States Army, where she has gained valuable experience in leadership and personnel management.

Exhibit 1: TGH - Nationally Recognized

Source: Tampa Bay Business & Wealth. (2024, July 16). *Tampa General Hospital: The Region's Only Academic-Based Health System and No. 1 Hospital in the Region Helps Lead the Nation in World-Class Care*. Retrieved from TBBWmag.com: https://tbbwmag.com/2024/07/16/tampa-general-hospital-the-regions-only-academic-based-health-system-and-no-1-hospital-in-the-region-helps-lead-the-nation-in-world-class-care/

Exhibit 2: AI Applications in Healthcare

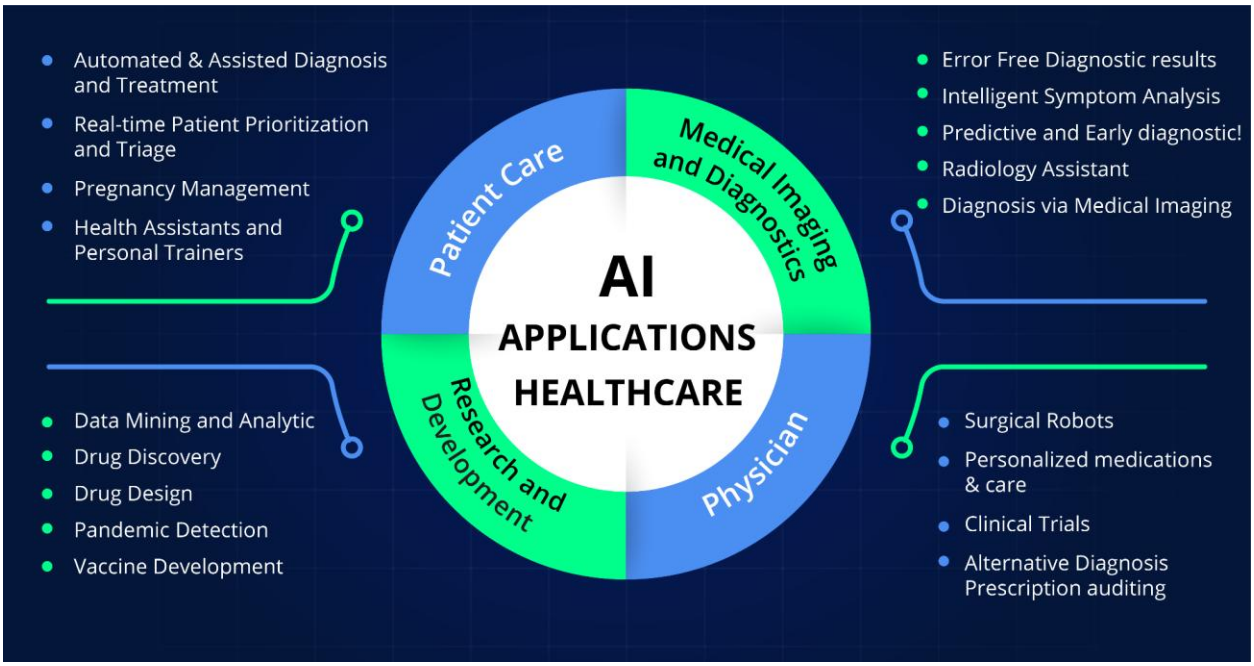

Source: Delveinsight. (2022, February 2022). *Most Promising Applications of Artificial Intelligence (AI) in Healthcare Segment* . Retrieved from Delveinsight.com: https://www.delveinsight.com/blog/top-applications-of-artificial-intelligence-in-healthcare

Exhibit 3: Patient Flow With/Without AI

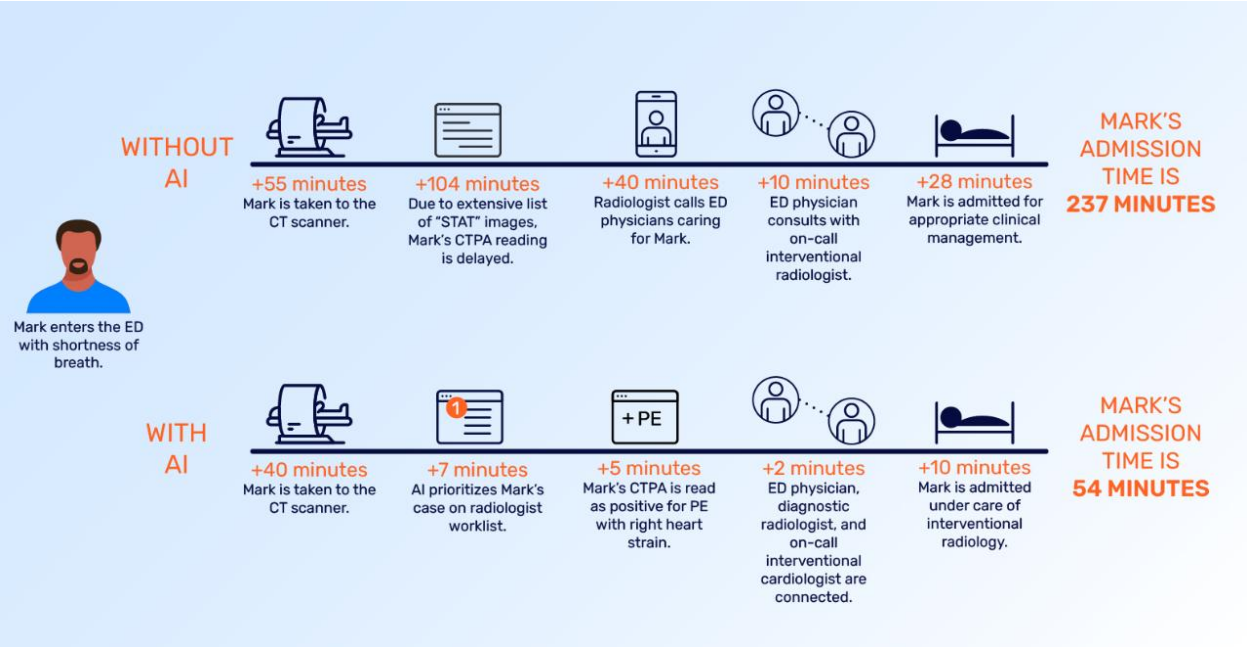

Source: Kah, L. (2024, May 16). *AI in Healthcare: The Ultimate Guide*. Retrieved from Aidoc.com: https://www.aidoc.com/learn/blog/ai-in-healthcare/

Exhibit 4: AI Vendor Comparison Summary

Criteria	Specta Health	Newton's Tree	Qualified Health
Data Security	Uses advanced encryption (AES-256), secure cloud storage, and MFA to protect patient data. Provides real-time monitoring to detect and prevent cyber threats.	Uses SSL/TLS encryption, secure storage solutions, and an incident response plan that complies with healthcare standards.	Applies end-to-end encryption (AES-256), ongoing cybersecurity audits, and access control. Prioritizes data anonymization for privacy.
Customization & Scalability	Offers customizable AI solutions tailored to different specialties. Scales from small clinics to large hospitals.	Flexible and scalable for various healthcare environments. Customizable for monitoring and data analysis.	Scalable AI tools with customization based on organizational needs. Designed for long-term utility.
Vendor Support & Training	Provides 24/7 support and training programs. Ensures users are equipped to effectively use AI tools.	Offers continuous training resources and full customer support, including system integration assistance.	Provides onboarding and training to ensure staff proficiency. Dedicated support teams assist with troubleshooting and system issues.
Cost Efficiency	Competitive pricing, with flexible models for small and large institutions. Maintains quality while reducing cost.	Pay-as-you-go model allows scalable use with strong clinical value.	Subscription pricing ensures predictable costs and affordability for smaller providers.
Target Audience	Healthcare organizations prioritizing regulatory compliance and AI adoption.	Hospitals seeking to evaluate, integrate, and monitor AI tools.	Organizations needing a governance platform for managing AI ethically.
Timeline to Launch	Waitlist – no confirmed timeframe.	Available immediately.	Available immediately.

Note. All vendors adhere to HIPAA and FDA regulations. This table summarizes differences in core offerings, scalability, support, and deployment timelines

Sources: Spectra Health. (2025). *Product Overview.* Retrieved from Spectrahealth.com: https://www.spectahealth.com/s-projects-side-by-side, Newton's Tree. (2025, February 8). N*ewton's Tree Joins National Centre of Excellence in AI & Digital Health Regulation.* Retrieved from newtonstree.ai: https://www.newtonstree.ai/post/newtons-tree-joins-national-centre-of-excellence-in-ai-digital-health-regulation, Qualified Health AI. (2025). *Advancing human health.* Retrieved from qualifiedhealthai.com: https://www.qualifiedhealthai.com/

Exhibit 5: Healthcare AI Framework

Source: Abujaber, A. A., & Nashwan, A. J. (2024). *Ethical framework for artificial intelligence in healthcare research: A path to integrity.* Doha: Baishideng Publishing Group Inc. Retrieved from Wjgnet.com.

Exhibit 6: Legal and Ethical Frameworks for AI in Healthcare

Framework	Focus Area	Implications for AI in Healthcare
Health Insurance Portability and Accountability Act	Patient Data Privacy & Security	AI systems must comply with ePHI protection standards
Food and Drug Administration	Risk & lifecycle management of AI SaMD	AI tools require risk-based evaluation, regulatory approval and ongoing monitoring
Office of National Coordinator for Health IT	Health IT Standards, Interoperability & transparency	Ensures AI integration supports patient data access and fairness
American Medical Association	Ethical Use of AI in Healthcare	Supports fairness, accountability and physician oversight
World Health Organization	Global Ethical Guidelines for Health AI	Promotes transparency, safety, human oversight in AI systems

Note. This table summarizes key governance frameworks relevant to AI adoption in healthcare. Sources: HHS (2021); FDA (2021); ONC (2020); AMA (2021); WHO (2021).

Exhibit 7: AI Healthcare Dashboard

Source: Hugle, T. (2021, November 30). *What could an AI-supported treatment dashboard look like?* . Retrieved from linkedin.com: https://www.linkedin.com/pulse/what-could-ai-supported-treatment-dashboard-look-like-thomas-h%C3%BCgle

Exhibit 8: Model Registry

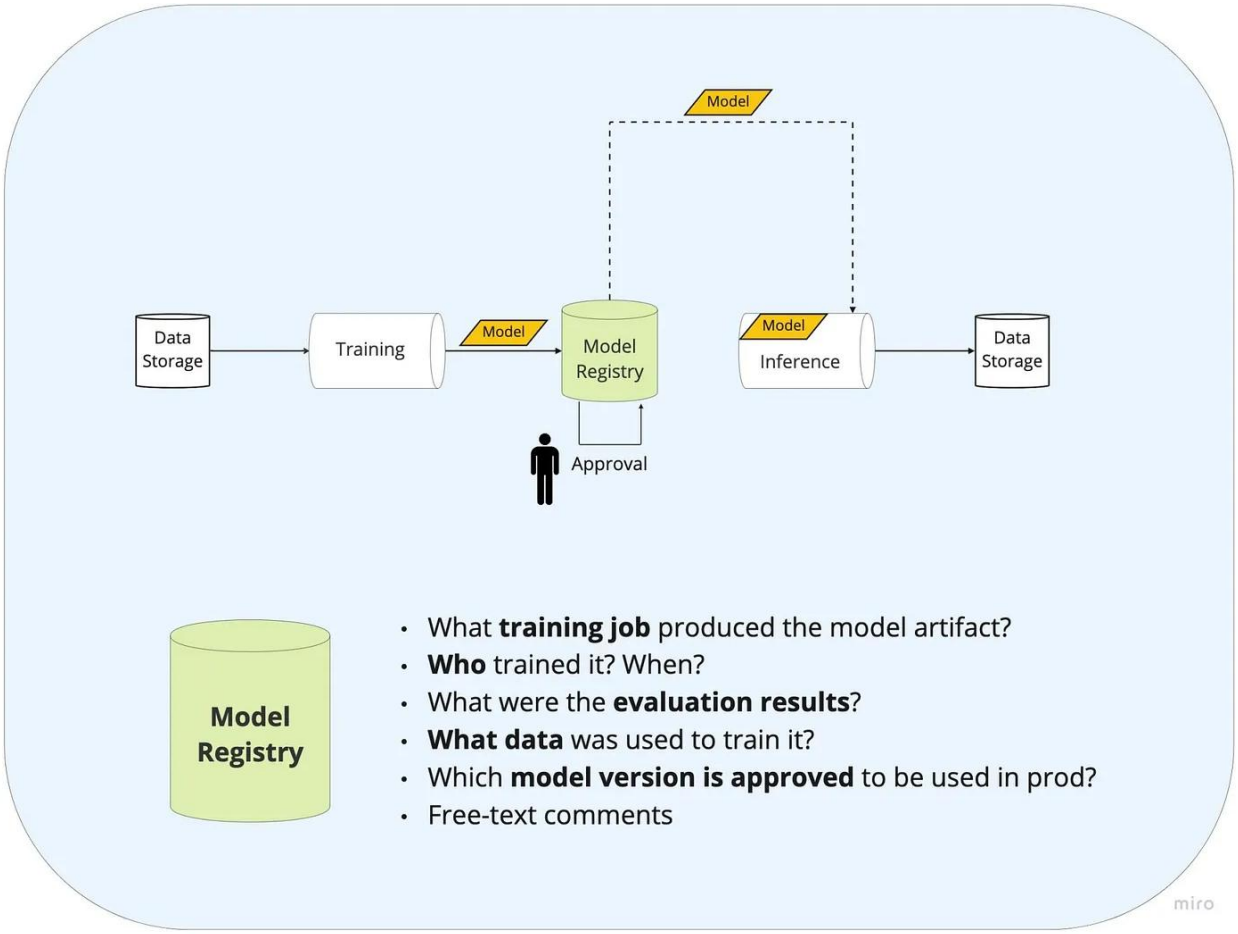

Source: Stiebellehner, S. (2023, July 31). *MLOps in a Nutshell: Model Registry, ML Metadata Store and Model Pipeline*. Retrieved from Medium.com: https://python.plainenglish.io/mlops-in-a-nutshell-model-registry-metadata-store-and-model-pipeline-3dc377babce1