The AI Catalyst Pulse

March 5th, 2025

Upcoming AI Catalyst Events

Mark your calendars

- **Mar 26: Executive Deep Dive: How AI is Turbocharging Patient Flow and Throughput** | *Register here*
    - Join our Executive Deep Dive to learn how health systems are utilizing AI to drive significant ROI within patient flow and throughput. We will share how AI Catalyst members are approaching ROI calculations for AI solutions in the patient flow space, as well as implementation tips to ensure your AI investments are improving ClinOps efficiency and care coordination.
- **Apr 10: Candid Vendor Conversations - Perspectives on Palantir** | *Register here*
    - Join a candid discussion with executive peers who have implemented or evaluated Palantir's healthcare AI solutions. Get unfiltered insights about implementation challenges, cost considerations, and realized benefits. You'll hear direct feedback about integration complexity, staff adoption experiences, and how Palantir's solutions compare to alternative vendors in the market.
- **Apr 15: Epic Systems and the Future of AI-Assisted Nursing Documentation** | *Register here*
    - Learn about Epic's evolving AI nursing documentation solutions, their impact on clinical workflows, and real-world implementation guidance. We'll be joined by Epic's development experts for a candid discussion of Epic's development roadmap, as well as how to address key issues relating to workforce buy-in and adoption of these solutions.

Central Oversight vs. Frontline Freedom: Your Second Critical AI Decision in 2025

This is the second in a four-part series examining the critical strategic dilemmas healthcare executives must navigate in 2025. Read part one here.

Every AI strategy must balance control against innovation. This creates healthcare's second fundamental AI dilemma in 2025: Do you maintain tight central oversight of AI deployment, or do you empower your frontline teams with the freedom to experiment and innovate?

As with all dilemmas in this series, “central oversight vs. frontline freedom” is to some extent a false binary; every health system needs some combination of both approaches. However, our conversations with AI Catalyst members consistently reveal that budgets, attention spans, and implementation capacity are limited — forcing tough choices about strategic emphasis.

The case for central governance: One unified AI strategy

The basic argument for centralizing AI governance is simple: The most reliable way to create a consistent, safe, and compliant AI strategy is to put *one* body in charge of AI decisions.

Most health systems seem to be moving in this direction. When we surveyed 2024 AI Bootcamp attendees, 52% told us their systems have "a formalized, system-wide AI steering committee." Among the benefits of this approach:

- **An empowered AI steering committee can make coordinated, strategic bets.** According to McKinsey, 75% of organizations list digital transformation as a high priority but admit they haven't sufficiently planned or resourced these initiatives. Without central coordination, departments risk pursuing duplicative AI projects even as strategic priorities go unfunded.

- **A central "veto point" can save you from dangerous AI blunders.** To give one scary example: Texas's attorney general recently alleged that an AI vendor "made deceptive claims about the accuracy of its healthcare AI products, putting the public interest at risk." If you have a well-informed governance committee reviewing every new AI tool, you're more likely to catch these risks before they impact patients.

- **A well-structured oversight body can ensure that important voices are represented in *every* AI decision.** For instance, HCA's recent contract with National Nurses United provides nurses formal representation on the system's AI Governance Committee — guaranteeing that frontline perspectives shape all AI adoption decisions.

- **Central governance structures make it easier to comply with fast-changing legal requirements.** Lawmakers are considering an ever-expanding list of requirements for AI in healthcare, such as standardized disclosures of how AI tools are trained. A central governance committee makes it easier to meet these demands.

- **An organization-wide oversight body can identify cross-departmental benefits that individual teams might miss.** When departments pursue AI in isolation, they naturally focus on their own pain points — and may miss opportunities to make an enterprise-wide impact. For example, while clinical documentation AI tools are often championed by physicians for burnout reduction, they can also deliver big revenue cycle benefits through improved coding.

So what might a centralized AI governance committee look like in practice? Consider UW Health's Clinical AI & Predictive Analytics Committee, which includes representatives from bioethics, law, biostatistics, data science, clinical operations, and DEI. For implementation, the organization created specialized "algorithm sub-committees" that work directly with requesters to oversee specific use cases — maintaining central control while engaging domain expertise.

The case for frontline freedom: Innovation needs room to breathe

Despite these benefits, there are important reasons *not* to centralize your AI governance. If you want to ensure your AI applications solve real-world clinical and operational problems, you need to give your teams the freedom to experiment and iterate.

A more frontline-driven approach addresses several key challenges:

- **When central approval processes become too cumbersome, frontline workers may create (potentially dangerous) workarounds.** When your official AI processes require months of committee reviews, clinicians and administrators will find unauthorized shortcuts — creating far greater risk than controlled departmental experimentation.

- **The people closest to the work can best understand which AI tools will make a real difference.** A nurse manager or service line director knows *exactly* where documentation burdens are heaviest or where communication breakdowns occur.

- **Many central governance frameworks demand resources health systems simply don't have.** One AI executive we spoke with put it bluntly: "In the world I work in, I do not have the resources to implement the level of governance coming out of [current frameworks]."

Providence illustrates what a more decentralized approach to AI can look like in action. Their governance model establishes basic checkpoints while empowering frontline "idea champions" to drive AI projects forward. These champions identify use cases, complete basic templates, and partner with technical guides throughout development — treating governance as guidance rather than a permission-granting bureaucracy.

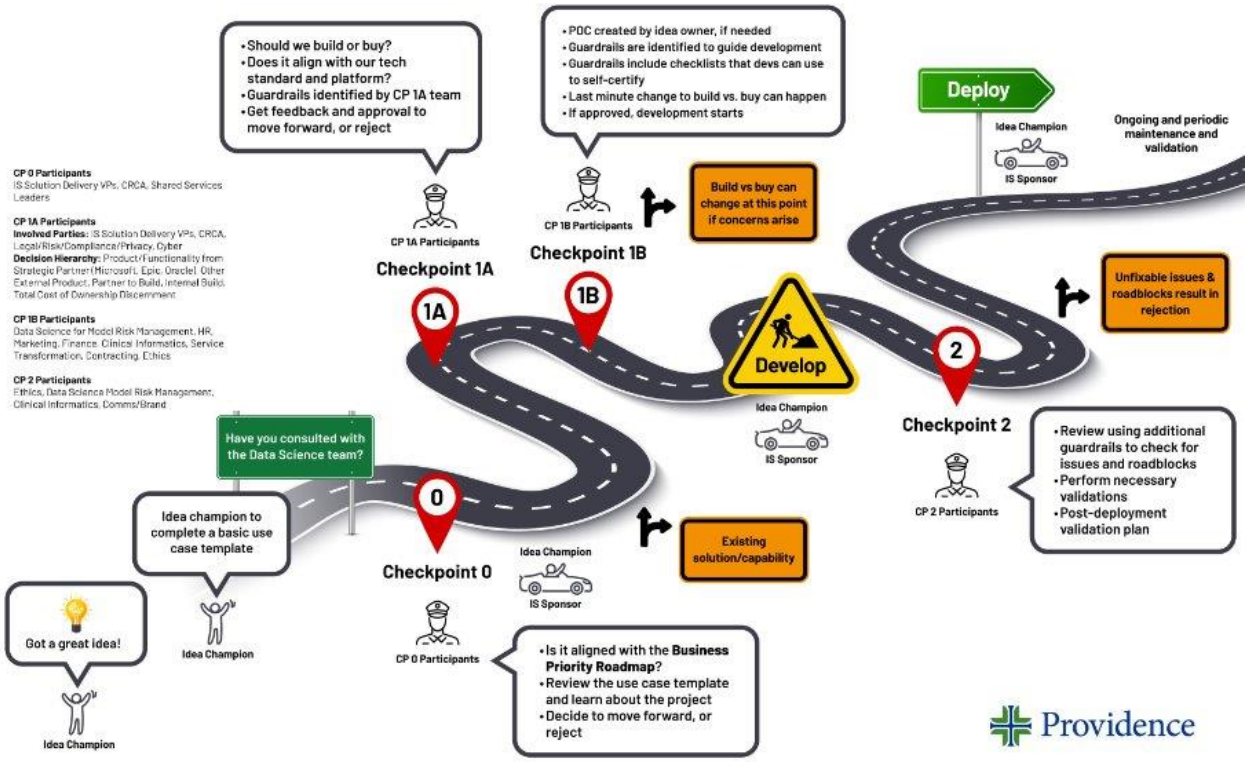

Which should you choose: Central oversight or frontline freedom?

Every health system needs both central guidance and frontline innovation, but you need to choose an emphasis that fits your organization.

Our guidance: Consider prioritizing central oversight if you have established governance structures you can build upon (such as a Digital Health Taskforce), or if your organization is taking its first steps into AI

and needs to carefully manage early risks. This approach makes particular sense for organizations focused on just a few big, focused AI bets. If you're pursuing a centralized approach:

- Identify strong functional representatives from clinical, IT, revenue cycle, and other key areas — and expect meaningful time commitment from these individuals.
- Create a risk-tiered review process where high-risk clinical applications receive full committee scrutiny, while lower-risk tools follow a more streamlined path.
- Consider establishing a controlled "sandbox" environment for low-risk frontline experiments to prevent shadow AI use while still enabling innovation.

On the other hand, consider prioritizing frontline freedom if your staff demonstrates strong AI literacy and enthusiasm, or if your organization has a culture that traditionally rewards responsible experimentation. This approach works best when you trust your divisional leaders to understand and enforce appropriate guidelines. In that case:

- Create standardized, practical tools for frontline leaders — such as vendor assessment questionnaires and risk scoring templates — that help them make good decisions independently.
- Invest heavily in AI literacy education to ensure department-level decision-makers understand both AI's potential and its limitations.
- Implement regular auditing processes to identify both successful use cases worth scaling and potential risks requiring intervention.

So where do top health system leaders come down on the "centralized governance" vs. "frontline freedom" dilemma? When we posed the question at The Health Management Academy's recent Trustee Summit, 72% of health system CEOs and board members preferred a centralized strategy.

Yet a few health systems are finding ways to split the difference. Mass General Brigham, for instance, combines a system-level AI Governance Committee with subcommittees integrating frontline technical and clinical considerations. Adventist Health demonstrates another hybrid approach: Clinical AI applications must undergo rigorous, centralized scrutiny, while operational AI tools follow a more streamlined path.

Ultimately, AI oversight requires bidirectional communication. Department leaders need to understand system-level risks, but just as importantly, system leaders must pursue strategies that capture benefits within and beyond individual departments.

Questions to consider:

1. If you conducted a candid audit of your governance approach, what gaps would you find between your formal structures and how AI is actually being implemented?
2. Where have you seen frontline teams using AI successfully without central oversight? What lessons could you draw from these examples?
3. How effectively does your current governance approach identify AI use cases with benefits that span multiple departments?

4. Does your AI governance approach reflect your organization's overall culture and decision-making style, or are you imposing a clashing governance style?

Overcoming Resistance to AI-Powered Sepsis Tools at Adventist Health

On February 19, AI Catalyst convened a cross-system group of health system executives for a "challenge huddle" where we brainstormed solutions to thorny AI implementation challenges. The first challenge we tackled together: Adventist Health's workforce resistance to a new AI-driven tool for early sepsis detection.

California-based Adventist has been working to implement KATE, a nurse-facing AI tool that links sepsis detection with triage scores using the Emergency Severity Index (ESI) scale. So far, Adventist leaders have encountered three distinct forms of clinical resistance:

- **Skepticism of AI accuracy:** Clinicians have disputed the accuracy of KATE's predictions, questioning whether the AI could truly identify sepsis risk factors that experienced nurses would miss.
- **Perceived threat to clinical judgment:** Staff have expressed concerns that the AI undermines their autonomy and critical thinking, with some worrying that early-career nurses in particular might over-rely on the technology.
- **Concern about "gaming the system":** Some clinicians worried that the system could be manipulated, potentially altering triage outcomes and compromising patient safety.

These concerns aren't unique to Adventist clinicians. When AI Catalyst recently surveyed 403 bedside nurses about their AI concerns, we found that 73% cite AI's potential impact on patient care quality as their top worry, followed by concerns about AI accuracy and data integrity (59%).

Three proven strategies for navigating workforce pushback to AI

Participants in the challenge huddle shared three strategies that, in their experience, have worked to overcome these kinds of clinician pushback.

1. Position AI as revealing "hidden risks," rather than replacing clinician judgment

Irene Louh, VP of Data Science at Baptist Health Jacksonville, shared that she has found clinicians are more open to AI tools when they're framed as providing support to harried professionals as they try to identify difficult-to-detect risk factors.

"It's really that threat to clinical judgment or the lack of trust," Louh explained. "It's where does clinician autonomy fit in? This is another support tool in our armamentarium to help us make a better decision versus deferring clinical judgment over to AI. That's a very different model and mindset."

To implement this guidance at Adventist, leaders might consider framing KATE as giving nurses a second chance to reconsider assessment while maintaining their position as the primary decision-maker.

"You get better adoption when it's AI augmenting my clinical decision versus [replacing it]," Louh noted.

2. Invest in finding your AI "change champions" – especially by targeting mid-career nurses

Meagan Rose, Principal Product Manager for AI and Innovation at Ochsner Health, shared her team's success in identifying the right clinical champions for AI adoption. "We have a group of eager beavers who are always there and excited to try things," noted Meagan Rose from Ochsner. "We know that they're going to give great feedback ... we'll sometimes start with those folks."

So where should you look for your own organization's ideal "change champions"? You might imagine that younger, newer nurses will be the most tech-receptive – but that's not what our survey found. Nurses with 4-6 years of experience actually showed the highest level of support for AI implementation (36%), higher than new graduates (27%).

These mid-career nurses have enough experience to feel confident in their clinical skills but remain adaptable to new workflows, making them ideal early adopters and peer advocates.

3. Design escalation pathways for AI-driven decisions that mirror your current processes

Adventist also shared their strategy for managing staff oversight when clinicians disagree with one of KATE's recommendations. If a nurse dismisses a sepsis alert for a high-risk case, the system waits a customizable period (typically 10 minutes) before notifying the charge nurse.

This creates a clinical safety net that maintains individual nurse autonomy while ensuring critical information isn't lost. Because it mimics existing clinical escalation protocols, adoption feels more natural to staff.

"It's a teaching opportunity," explained an Adventist representative, adding that "the charge nurse, who has all eyes on the floor," can then provide support or validation.

Questions to consider:

1. How might you reframe your AI tools to reduce any perceived threat to clinician autonomy?
2. How can you align AI-enhanced escalation pathways with your existing clinical safety nets and communication flows?
3. Beyond technology enthusiasts and formal leaders, how might you identify and empower mid-career clinicians as change champions?

AI Outcomes Snapshot: Emory's EmChat

Clinicians spend countless hours searching for internal policies and protocols – and even when they find what they're looking for, the information is too often wrong or outdated, resulting in non-evidenced based (or at least non-standardized) care delivery.

To address this challenge, Emory partnered with Impact Advisors to create an AI-powered chatbot, and they ultimately succeeded in reducing clinician administrative burden and reinforcing system-wide standards of care.

About Emory's EmChat pilot

Emory's nursing and Emory Digital teams partnered with Impact Advisors in 2024 to create EmChat, an AI-augmented LLM chatbot. EmChat is designed for two purposes:

1. To support bedside staff to quickly find organizational protocols and policies

2. To allow infection preventionists to rapidly verify CLABSI cases

Leaders conducted proof-of-value pilots with several service lines cross the organization. Nine months later, Emory not only established proof of value, but fully integrated EmChat within the system. CLABSI cases have rolled out the clinical-reference development to one of their hospitals.

Internal policy look-up takes nurses up to 45 minutes

Emory discovered that it could take nurses up to 45 minutes to locate internal policies and protocols. To not delay care for their patients, nurses typically relied on their peers or memory for answers rather than system guidelines. To provide an easier and more accurate alternative, nurses worked directly with Impact Advisors to create, test, and deploy EmChat, a clinical reference AI-augmented LLM chatbot. When nurses ask EmChat a question regarding their organization's internal policies, EmChat leverages AI to locate the correct information from Emory's internal policy manager system and Epic learning home dashboard.

EmChat is designed to streamline workflows by saving nurses time and ensuring greater consistency in care delivery. Emory leaders also plan to use EmChat's software policy look-up to proactively identify differences in policies between sites. With this information, Emory plans to create greater policy consistency across their system.

EmChat also aids in CLABSI verification

In addition to its policy look-up function, EmChat also helps the system's infection preventionists quickly find the data they need to verify infection. This process begins when a pre-existing alert mechanism within the EHR signals to IPs that a patient is at high risk of CLABSI. IPs ask EmChat "Does this patient have a CLABSI?", which triggers the chatbot to pull the relevant clinical indicators IPs need to detect and verify the presence of a CLABSI (informed by the NHSN CLABSI criteria). Rather than having to search through the patient chart to identify CLABSI criteria, EmChat makes the information readily available.

As Nancye Feistritzer, VP of the Center for Care Delivery and Innovation at Emory Healthcare, said, "We're not relieving the clinician of the obligation to review and confirm accuracy, but this has dramatically reduced the amount of time associated [with searching for policies]."

Outcomes of the EmChat pilot to date

Emory has completed the pilot proof of value and has now moved into the industrialization phase, with plans to scale both models widely across the health system.

Their outcomes to date include:

- Reduction in RN and clinician time spent searching for internal clinical reference
- Reduction in time spent to verify the presence of a CLABSI
- 15+ min estimation of clinician time saved per EmChat-augmented task

Policy Updates

Understanding the evolving AI regulatory and legislative landscape

Over the past few weeks, the Trump administration and the Department of Government Efficiency initiated layoffs of staff across federal health agencies, including about 700 probationary employees at the FDA, as part of a broader effort to reduce the size of the federal workforce.

But industry stakeholders pushed back, citing healthcare AI as a major concern. Among their worries was that, if FDA lacked the staff needed to approve AI-powered devices and tackle emerging AI questions, the result could be slower product approvals and broader delays in AI adoption. (Another fear: If FDA couldn't offer regulatory guidance, health systems would ultimately bear more risk and burden themselves.)

This week, the FDA appears to have reversed course, rehiring at least some employees – but it's not yet clear how many were offered their jobs back, or how many would accept the offer. We'll continue to monitor implications for AI policy across future editions of the *Pulse*.

- **FDA** moves to rehire medical device staffers days after mass firings
- **FDA** layoffs may shift regulatory burden of AI to health systems
- **FDA** staff cuts impact AI oversight as federal investment surges

AI Strategy Quick Hits

Noteworthy moves from peers to implement AI technologies

Catch up on health systems who have joined the transition to AI-powered clinical documentation:

- **Cleveland Clinic** adopts **Ambience Healthcare's** AI-powered notes to streamline outpatient care
- **Inova Health** partners with **Abridge** to automate clinical documentation
- **Akron Children's Hospital** adopts **Abridge AI** to cut clinician cognitive load

Other strategy quick hits:

- **Philips** & **Mass General Brigham** join forces to integrate live health data with AI
- **Memorial Sloan Kettering Cancer Center** partners with **AWS** to accelerate AI-powered cancer research and drug discovery
- HIMSS25 forum: New tactics for integrating AI in healthcare

Emerging Use Cases

New capabilities that indicate AI's potential

- **University of Michigan Health System** harnesses digital twins for enhanced heart failure prognosis
- **UCHealth** monitors 22K hospital beds with AI to reduce mortality rates
- **Mass General Brigham** introduces HTN-AI for hypertension detection and risk stratification
- **Mass General Brigham** streamlines heart trial enrollment with AI-powered RECTIFIER
- How LLMs can enhance antibiotic prescribing

Market Moves

A round-up of AI company announcements and stories

- **Teladoc Health** reports a $1B loss in 2024, following **BetterHelp's** market woes

- **Vitalchat** secures $6M to expand its AI-powered virtual care platform
- **Memorial Hermann** and **Northwell Health** invest $17.75M in imaging data startup **Avandra**
- FDA veteran Troy Tazbaz returns to **Oracle** as a Senior Vice President
- **Abridge** raises $250M in Series D funding to scale AI-powered clinical documentation

The News in Numbers

An interesting data point that caught our eye

66% of physicians in a recent AMA study report that they are using some form of healthcare AI in 2024.

61% of physicians in a new AMA study are concerned that the increased use of AI by health systems will increase the number of prior authorization denials.

Expert Insights

For further reading, articles, videos, and podcasts that we found insightful

- *NPJ*, "Clinical trials informed framework for real world clinical implementation and deployment of artificial intelligence applications"
- *NEJM*, "From Clinical Notes to GPT-4: Dr. Emily Alsentzer on Natural Language Processing in Medicine"