

The Ultimate Guide to Implementing AI in Healthcare Operations



Table of Contents

Key Takeaways	3
Introduction: How Has Healthcare Evolved?	4
How Do Health Systems Respond to Complexity Today?	5
How AIs Augment EHRs and RTLS	6
What Operational and Workflow Problems Does AI Solve for Hospitals?	8
The Different Types of AI	12
Assess Your Hospital's AI Maturity Level	16
Do Hospitals Need a Chief AI Officer?	19
Why Do Many AI Implementations Fail?	20
Partnering With Your Vendor	21
Getting Your Staff on Board	22
Integrate AI Into Your Workflows and Systems	23
Deployment	25
What a Sample Successful Deployment Could Look Like	28
Start Small, Scale Big	29
Closing: Where to Begin?	30
References and Further Reading	31
Contributors	32




Key Takeaways

1. Trust, between leaders and employees, and between employees and their AI solution, is vital to any implementation effort.
2. Systems must partner with vendors to tailor AI for their needs. The more context, customization, and data that an AI receives, the more effective its outputs will be.
3. Systems must have clearly defined goals and metrics with which to judge success. Agentic implementations must move past anecdotes and survey sentiment, and towards measurable operational impact.

Introduction: How Has Healthcare Evolved?

The healthcare landscape has become increasingly complex. Hospital populations have become more difficult to care for, with multiple comorbidities, more acute medical issues (*up to five by one count*), and higher likelihoods of unplanned readmission after initial discharge (*Naik et al., 2024*).

In addition to more patients and more complicated cases, hospitals must also support more treatments and devices. The National Comprehensive Cancer Network (*NCCN*) guidelines mean page count rose from 26 to 198 pages between 1996 to 2019, a 762% jump, an increase of nearly 7.5 pages annually (*Kann et al., 2020*). Demonstrating how disease and care management complexities have increased following advances in personalized medicine, novel therapeutics, and technology. As care shifts, so too must the operations, policy, workflows and education that support it.



How Do Health Systems Respond to Complexity Today?

To ensure the delivery systems that support care that are compounding in complexity, health systems turn to new, innovative tools like real-time location systems (RTLS) and electronic health records (EHRs). These tools not only bring solutions to address immediate challenges and drive impact but also bring the promise of continued innovation. Tools that predict, grow, and respond as the future unfolds.

RTLSes have earned their place in hospital operations, providing staff immediate, accurate visibility into the real-time locations of devices, patients, and providers on live maps. Nurses can pull up their interfaces and locate the nearest IV pump or wound vac in seconds. Properly implemented, a RTLS solution can reduce search times by up to 50% (*Pendrill et al., 2021*), freeing nurses for other duties, primarily patient care.

EHRs were created to provide clinical context around patients, treatments, and conditions. They record which medicines a patient has taken; previous treatments, diagnoses, and lab results; and any notes or visit summaries from past providers and procedures.

As hospitals grow more complex, adding more specialists, more handoffs, and more data generated at every visit, that context has become harder to keep track of by memory or paper chart alone. EHRs are one of the ways health systems have kept pace: by centralizing a patient's history in one place, they let clinicians make faster, better-informed decisions even as the number of variables involved in a single patient's care keeps increasing.

How AIs Augment EHRs and RTLS


So why do hospitals need AI when they already have RTLS and EHR?

The answer is simple: operational orchestration. While there are some exceptions, most EHR and RTLS solutions were not designed to help hospitals coordinate and manage their operations in real time. EHRs and RTLS represent significant technological solutions, yet in their original designs, the insight, visibility, and benefit they provide still depend on frequent, and often intensive, manual engagement. As a result, their core functions fall short of autonomous, predictive action.

Within a health system, the complexity present at the level of care units, and even individual patients, can readily overwhelm staff capacity. When a function is instead made automatic and autonomous, and is designed thoughtfully into existing workflows, it addresses a need before that need is consciously felt. Similarly, it is unreasonable to expect any individual to continuously monitor an entire system, each unit, and each patient for operational insight; orchestrated intelligence removes this requirement altogether.

After all, RTLS sensors simply display dots on a map, while EHRs organize clinical knowledge and patient care journeys. If a nurse still has to sprint across a campus to retrieve a wound vac, or a patient flow coordinator has to manually call environmental services and notify them that discharges are complete and to start cleaning rooms, hospital operations are still bottlenecked.

This is not to say that RTLS and EHR data are useless; fusing these data types together greatly improves the efficacy of AI outputs, and opens up new possibilities:





After all, RTLS sensors simply display dots on a map, while EHRs organize clinical knowledge and patient care journeys. If a nurse still has to sprint across a campus to retrieve a wound vac, or a patient flow coordinator has to manually call environmental services and notify them that discharges are complete and to start cleaning rooms, hospital operations are still bottlenecked.


This is not to say that RTLS and EHR data are useless; fusing these data types together greatly improves the efficacy of AI outputs, and opens up new possibilities:

- Charge nurses and leaders can understand resourcing at a unit level, why it matters clinically, and what needs to happen next.
- ED boarding and wait times for inpatient beds fall once room turnover, staff location, and care progression are orchestrated automatically.
- Infection control becomes traceable, letting teams reconstruct who occupied which room and when, all within minutes instead of days.
- Equipment shortages are predicted and addressed before the impact is felt. Staff spend less time searching for equipment and more time doing the clinical work they trained for.
- An outpatient clinic can see that their patient is still in their previous appointment and adjust their schedule to account for the delay before it happens.
- Discharge dates are supported with real-time care progressions, barriers to discharge are flagged, and ultimately, patients are able to return home sooner. This removes clinically unnecessary days, reducing overall length of stay, freeing up beds, and generating more revenue for hospitals and clinics.



What Operational and Workflow Problems Does AI Solve for Hospitals?

Supply Chain Management

Whether it's consumables (*like vaccines and medicines*) or devices (*such as infusion pumps or wound vacs*), hospital supply chains are multifaceted and complicated. Consumables must be stored in compliant fridges with regularly calibrated temperature probes (*usually every 1-3 years*), while other mobile assets such as beds and pumps carry their own maintenance cycle, utilization pattern, and cost of downtime.



To complicate matters, hospitals also need to track it all across departments, units, floors, and campuses. Previously this tracking was either done manually (*a next-to-impossible task for any team, especially at scale*), or simply left undone. With AI, teams can easily pull asset location, usage, and maintenance data together to automatically flag shortages before they surface; immediately identify idle equipment when it is sitting unused elsewhere in the system; and dynamically adjust PAR levels to account for hospital demand and census



Patient Flow

On paper, admitting patients, coordinating treatment across multiple service lines, and eventually discharging a patient seems like a simple journey. However, a patient's care progression is heavily dependent on coordinating beds, staff, consults, transport, scheduling, and cleaning crews simultaneously and across departments. Currently, much of this work is done by hand: coordinators reach out to different teams by phone, manage and hand off multiple patients at once, and track everything using paper or digital forms that they must fill out.

AI can close the gap between operational response and clinical documentation. For one, it can identify operational delays that hinder care progressions for an individual patient and recommend interventions to remove these barriers.

In addition, AI can help automate the room turnover process, which at many hospitals remains highly dependent on inputs from nurses and other teams. AI can pinpoint which patients have actually been discharged with RTLS data without relying on staff manually entering the discharge note, and alert environmental services to quickly clean a room after patients have physically entered the discharge lounge.

Ambulatory Access

Template optimization, scheduling, and provider capacity all determine whether patients get seen in a timely manner, and whether clinics can fill open slots and increase room utilization. Rather than reviewing templates manually, AI can surface which blocks are underbooked, which visits run over time, and which providers have additional capacity, all in order to make recommendations to optimize the template settings.

This closes the gap between what is on a provider or resource template, and what is actually happening in the clinic. This drives insights and optimization that no amount of constant interviews and observations can replace and ensures your templates are optimized as the workforce turns over (*providers retiring, going on maternity or paternity leave, or new hires joining the practice*).

Predicting Demand

Emergency department surges, seasonal patterns, and capacity crunches are easier to plan for when a system can forecast them before they occur, rather than react to them after the fact. By analyzing reams of data, AI can surface insights on past performance, and recommend solutions to improve future operations, including optimizing float staffing decisions or asset allocation.

General Paperwork

AI can also automate away manual but vital tasks, such as rounding documentation or temperature monitoring logs, all of which consume clinician time that could instead be used towards direct patient care. Importantly, clinicians can offload these tasks onto AI while ensuring the accuracy and integrity of records for future audits and compliance.





Patient Satisfaction and Experience

As industry focus and investment in patient experience and satisfaction become a growing priority for health systems, operational decisions and tools that support positive patient experiences are correspondingly gaining value. While simple RTLS integrations support functions such as patient wayfinding, automated arrival and check-in processes, and the automatic display of care team member information at a patient's bedside throughout an inpatient visit, AI capabilities extend patient experience benefits further still.

Real-time, staff-facing visibility and status boards, for instance, can flag when a patient has gone without staff contact for a prolonged period; a metric with significant bearing on patient satisfaction, particularly when a patient is roomed and awaiting their provider. This capability supports proactive service recovery, in which even a brief check-in can meaningfully improve the patient's experience. Such systems can also display estimated wait times in waiting areas, helping to manage patient expectations and reduce perceived delays.

The Clinical/Operational Divide for AI

Clinical AI use cases include diagnostic imaging analysis, flagging potential abnormalities, artifacts, or fractures on a scan; clinical decision support that alerts providers to drug interactions or sepsis risk based on vitals trends; and predictive risk scoring that identifies patients likely to deteriorate or be readmitted.

Operational AI covers asset and supply chain management, clinic access, patient flow, and predictive analytics. While clinical AI tools augment clinical judgment at the point of care, operational AIs, the focus of this guide, augment the systems and workflows that determine whether care can be delivered efficiently in the first place.

The Different Types of AI

One major concern shared by systems and individuals alike, regardless of industry, is establishing trust in the insights and recommendations generated by AI. While AI models are improving rapidly, they remain prone to occasional hallucination or misinterpretation. This raises a critical question: how are these risks weighed against the technology's potential benefits?

However, AI is a broad umbrella that covers tools with diverse architectures, limitations, and failure modes. As a result, health systems may be conflating very different tools, and using the wrong tool for the wrong task. In order to effectively implement an AI product, health system leaders must have a clear understanding of their goals and workflows, as well as the various kinds of AI and where they excel.

How do large language models and predictive analytics differ?

Large language models (LLMs) generate language (*Google for Developers, n.d.*). When prompted, they use mathematical reasoning to produce the most statistically likely sequence of words. This quirk of their design is precisely why they can sound so confident, even while stating something blatantly false.

A 2026 study tested nine large language models on two hospital administration tasks: counting patient records and filtering them by criteria, using 50,000 real emergency department visits (*Klang et al., 2026*). Prompting a model to answer directly produced poor results across all nine models. Chain-of-thought reasoning, where the model shows its work before answering, helped only a little, and accuracy collapsed as the record set grew larger. Even GPT-4o, the strongest model tested, fell from roughly 95% accuracy on small tables to under 60% on larger ones.

Findings like this help explain why healthcare leaders remain cautious about applying language models to problems where a single fabricated detail could steer an operational decision in the wrong direction (*Ahmad et al., 2023*).

However, predictive analytics models function differently. They take structured, quantifiable data, admission timestamps, room sensor pings, lab values as discrete fields, and learn statistical relationships between defined variables.

Because these models draw on defined variables instead of generating open-ended language, their outputs stay anchored to the data they were trained on, which is one reason why hospitals lean on this class of tool for day-to-day operational forecasting (*Tello et al., 2022*).

Importantly, predictive analytics are already common across various industries and use cases, such as optimizing maintenance or forecasting drill failure events for oil and gas companies. As a result, leaders can adapt this proven technology to healthcare settings with more confidence.


Where predictive analytics excels

Predictive analytics has an extensive, research-backed record of success in forecasting and coordinating hospital operations, whether it's a head-to-head test against an LLM, predicting ICU flow and demand, or booking inpatient beds in advance.

Researchers at Mount Sinai tested GPT-4 against a traditional machine learning solution on a purely operational question: which emergency department patients would end up admitted as inpatients? Using records from more than 864,000 ED visits across seven New York City hospitals, the ML solution, built from XGBoost and Bio-Clinical-BERT, reached an accuracy of 82.9% and an AUC of 0.88 (*Glicksberg et al., 2024*).

GPT-4, reasoning on its own from the same records, performed considerably worse. It only improved its accuracy once researchers stopped asking it to reason unassisted, and instead fed it retrieved real-world examples and the ML model's own probability scores to lean on. Left to reason from scratch, on a task that decides how many beds a hospital needs to hold open, the language model was the weaker tool.

Other researchers have found similar success using machine learning. Researchers built an interpretable predictive model from 63,432 hospital admissions and used it to forecast short-term discharges, flag long-stay patients, and anticipate ICU flow with accuracy above 80% (*Bertsimas et al., 2022*).



In 2026, scientists assessed the performance of ML models for predicting elective surgical ICU admission at a UK NHS Trust. In a prospective evaluation, only 71.6% of patients who required an ICU bed had one electronically requested in advance (*Hunter et al., 2026*). The researchers then tested the same model against the hospital's existing electronic system, with both forecasting how many ICU beds the whole hospital would need the next day. The new model's outputs came closer to what actually happened, meaning it made smaller errors than the system already in place.

Where LLMs succeed

While LLMs are poorly suited for predicting, orchestrating, and optimizing hospital operations on their own, they can be effective when incorporated into a larger system.

In one study, researchers built a patient scheduling system to coordinate appointments across 12 anesthesiologists and surgeons, 12 operating rooms, 15 patients, and three surgery types. While their solution included a LLM, its tasks were limited to translating user queries and algorithmic outputs for a separate model (*not an LLM*) that actually ingested data and carried out the core analysis work. This integrated system landed within 1% of the optimal solution found by branch-and-bound search, and outperformed two of the hospital's existing scheduling approaches (*Chen et al., 2026*).

Another study attempted something similar, using seven LLMs, including ChatGPT-4o, LaMDA 2, PaLM 2, and Qwen, on allocating operating rooms, postoperative beds, and surgeons. Each LLM was part of a larger framework that combined prompt engineering, hospital data retrieval, and API calls to dedicated tools, and handled orchestration, rather than computation (*Wan et al., 2025*).

In this experiment, researchers leaned heavily on the strengths of LLMs: summarizing and interpreting complex data, and natural language processing. Rather than crunching data and creating predictions, LLMs simply served as a coordination and communication layer, helping administrators, clinicians, and patients understand, improve, and navigate the care process.

An AI for every task

From these studies, one conclusion emerges: language models, left to reason on their own about admissions, schedules, or bed allocation, underperform predictive algorithms that were built for purpose. Given a narrower niche, such as translating outputs into natural language for end users or calling the right tool, LLMs perform much better.

After all, a model trained on structured data can be checked against a benchmark and corrected, but LLMs designed to generate language cannot be optimized the same way. If leaders treat these two AI types as interchangeable, then their deployments will encounter friction. Only by keeping predictions and language processing separate, can leaders play to the strengths of AI.



Assess Your Hospital's AI Maturity Level

Ramping Up

Before evaluating outputs, hospitals need to evaluate how this new solution will fit into their workflows and what the ramp up period will look like. AI needs time to learn a hospital's specific environment and establish operational baselines before it can perform reliably. This adjustment period is normal, and should be an expected part of the roadmap, rather than seen as a sign that the deployment has failed.

Machine-Generated Data vs. Human-Generated Data

Today, 65% of hospitals now use predictive AI models, the vast majority of which were built by the same company that developed the hospital's electronic health record. Clearly, AI has arrived in healthcare (*Nong et al., 2025*). Its actual effectiveness is a separate question, and the evidence suggests the answer is complicated.

A survey of 67 health systems found that 77% of respondents identified immature AI tools as their single greatest obstacle to widespread adoption (*Poon et al., 2025*). A problem common across many industries: algorithms are only as capable as their data inputs, and right now, there isn't enough high-quality data to feed the machines.

When health systems talk about leveraging AI, the conversation almost always turns to the EHR, the largest repository of patient and operational data within most hospitals, making it the logical starting point for any data-driven initiative. While they hold massive volumes of data, training AIs solely on EHR data is insufficient. This is simply because EHRs were never designed as AI inputs. Instead, EHRs were created to document care, support billing, and satisfy regulatory requirements, purposes that produce data with very different characteristics than what AI models need to perform well.

The gap between data that best supports AI models, and the data that already exists in our EHRs, stems from how the data is produced in the first place. A comprehensive review published in ACM Computing Surveys (*Sarwar et al., 2022*) examined EHR data quality across more than a decade of research and found that raw EHR records are routinely prone to inaccuracies, inconsistencies, sparsity, and bias, due to several factors:

- Clinicians, pressed for time, type quickly and make errors
- Notes get copied forward from previous visits without review
- The same patient receives different diagnostic codes across different encounters
- Data recorded across departments and facilities rarely conforms to a common standard.

As a result, feeding this data into an AI model only produces unreliable outputs, and in cases where the training data carries embedded bias, it produces results that actively mislead the end users (*who tend to be clinicians*).

The reflexive response to unreliable data is to ask for better, more complete documentation. But registered nurses already spend nearly 25% of every shift on EHR documentation and review (*Yen et al., 2018*), while physicians spend close to 49% of their working hours on administrative tasks (*Sinsky et al., 2016*); in both cases, this time comes directly at the expense of patient care. These figures reflect a workforce that has reached the limit of what human abilities can sustain.

Asking clinicians to generate more data, of higher quality and at greater volume, is not the path to better AI, especially given that the documentation burden already contributes to burnout, which then leads to high turnover, costing the average hospital up to \$6.2 million annually in recruitment and replacement costs for nurses alone (*NSI Nursing Solutions, 2026*).

The ceiling on EHR data quality is a human capacity problem, and it cannot be solved by placing additional demands on people already stretched thin.

Hospitals can instead determine how to offload this burden onto technology, identifying operational data that can be collected continuously, accurately, and at scale without placing more demands on clinical staff. Sensors, location solutions, and automated data pipelines already exist in many hospital environments, and the next step is using this infrastructure to generate the ambient, structured, continuous data that AI models need in order to function well.

The contrast with human-generated documentation is significant across three dimensions.

Accuracy increases when collection is automated, because sensors do not mistype, copy forward, or record data hours after the fact under time pressure, so the data produced reflects what actually happened, as well as when and where it happened.

Consistency also improves when collection is programmatic, because sensors can be configured to capture data at defined intervals and in specific formats, producing the structured, regular inputs that predictive modeling depends on.

Lastly, volume increases without adding any burden on clinical staff, because a single location sensor passively generates continuous data across an entire shift, without requiring a clinician to touch a clipboard or enter extra information into the EHR.



Do Hospitals Need a Chief AI Officer?

As with every emerging technology, best practices around AI are still being established. One controversial question has been making the rounds: do health systems need Chief AI Officers, or are they better served by putting AI under the purview of the CIO?

Part of the controversy stems from the already full portfolios that CIOs manage. They are tasked not only with modernization, but also with maintaining legacy systems, upon which hospitals depend, a job that grows more difficult each year as vendors go out of business or deprecate support for older products.

As a result, some health systems have taken the burden off CIOs and instead appointed Chief AI Officers to take on this task. Mayo Clinic recently named a Chief AI Implementation Officer, while Kaiser Permanente, City of Hope, and Children's National Hospital have created equivalent roles under different titles. All of them share similar responsibilities: overseeing AI governance, implementation, and oversight.

Many of these networks also share another key quality: they are large health systems, often with multiple facilities. Given their size and complexity, creating a dedicated executive position to handle AI is a logical step, especially since AI is already in widespread use across their campuses and workforce.

Regional and community health systems may not have adopted AI at the same scale, which means a formal C-suite role would be premature for many of them. Even so, these health systems still need to answer key governance questions, such as who is responsible for each tool or workflow, who approves new changes, and who is tasked with identifying model drift.

The key is ensuring that someone in leadership can answer these questions and be accountable for AI. Whether that means appointing a dedicated AI officer depends on the needs of the specific hospital or health system.

Why Do Many AI Implementations Fail?

A 2024 systematic review in the *Journal of Medical Internet Research* examined the enablers and barriers behind real-world AI deployments in hospitals and found that most failures trace back to process, not the underlying technology itself (Kamel Rahimi et al., 2024). Four patterns recur across these failed implementations.

Success requirements are often left unclear from the outset. Teams roll out a tool without agreeing in advance on what „working“ actually looks like, so six months later no one can say whether the deployment succeeded or simply became part of the routine. A defined metric, tied to an operational outcome such as reduced search time or shorter length of stay, needs to exist before go-live, not after.

Staff distrust the technology, particularly when they were not consulted before it arrived. This distrust is often a reasonable response rather than pure resistance to change. Recall that 77% of health systems surveyed cited immature AI tools as their biggest obstacle to adoption. When staff have already encountered inconsistent outputs, skepticism toward the next tool is earned, not irrational.

Poor data quality undermines results before a single output ever reaches a clinician or coordinator. Kamel Rahimi et al. found that insufficient data quality, incomplete data sets, and a lack of rigorous performance assessment were among the most common barriers identified across the studies reviewed. This connects directly to the gap between EHR-generated data and the ambient, sensor-generated data described earlier in this guide. A model trained on inconsistent or incomplete inputs will produce inconsistent or incomplete recommendations, regardless of how sound its underlying algorithm is.

Finally, tools that are not woven into existing workflows tend to fail even when they perform well. If using the tool feels like one more system to check rather than something embedded in the job, staff route around it, and the intended outcomes never materialize. Training that mirrors real screens and real workflows, discussed later in this guide, is one direct way to prevent this pattern from taking hold.

Partnering With Your Vendor

The vendor relationship determines how effective an AI deployment becomes. Hospitals should confirm that the AI receives the data it needs to produce accurate outputs, that it is trained on the hospital's own context and terminology (for instance, that „Tower A“ and „North Tower“ refer to the same place), and that the underlying data is clean enough to support reliable outputs. Where gaps exist, hospitals should work with vendors to build out new features or customize inputs and outputs to fit their specific environment, rather than adapting their workflows to fit a generic product.



Getting Your Staff on Board

Understanding Their Point of View

Staff buy-in starts with understanding clinicians' point of view and addressing common fears directly rather than dismissing them. While predictive algorithms that forecast equipment usage or optimal maintenance periods have long been used in industrial sectors, they aren't yet widespread in hospital operations. As a result, many clinicians may associate AI with commercially available large language models like ChatGPT or Claude, neither of which were designed for healthcare operations.

Overcoming this mistrust through education is critical. Once end users, whether clinicians, team leads, or technicians, understand that AI is there to support their judgment rather than replace it, it becomes important to train them in best practices around AI usage. Because these best practices will continue to change, training has to be updated, either through refresher courses or short, one-day workshops that will not set staff further behind.

Mythbusting Common Fears and AI as Partner

AI should be positioned as a partner that removes tedious work and removes barriers, not a replacement for clinical judgment. Onboarding should make the actual role of AI concrete: taking away administrative burden so clinicians, techs, and coordinators can spend more time on the work they trained for. Implemented well, AI streamlines daily work. Implemented poorly, it becomes another system competing for staff attention.

Rather than framing training sessions as another mandatory duty, leaders can position AI as an intelligent assistant that surfaces recommendations and automates away key tasks. Clinicians no longer need to spend long stretches on the phone with other departments or dig through records to find the information they need; instead, this time can go back to bedside care.

Integrate AI Into Your Workflows and Systems

From Reactive to Real-Time

Most hospitals take a reactive approach today, analyzing historical data for insights, reports, and decision making. A charge nurse might review staffing patterns only at the end of the quarter, once attendance data has been collated and analyzed from EHRs or other devices. This means the same charge nurse, or even the CNO, cannot surge extra nurses or technicians in response to unexpected occurrences.

To be clear, this isn't the fault of any leader or hospital. It has only recently become possible to build the devices and data tools needed for real-time visibility and insight. As healthcare evolves, this after-the-fact methodology falls short. Compared to even 20 years ago, hospitals now manage a more challenging patient population, with more comorbidities and acute issues, and longer, more complex clinical treatments.

Real-time visibility closes the gap between what leadership assumes is happening on the floor and what is actually happening. A charge nurse can see immediately whether a shift change is working and where a new shortfall is appearing, and can drill into the specific cause rather than waiting weeks for a retrospective. That shift also changes how she spends her time. Instead of piecing together the floor's status through phone calls, in-person check-ins, and EHR lookups, she can see the picture directly and spend that time deciding what to do next.

When compounded across multiple hospitals, real-time visibility helps a health system become more nimble, delivering in a fraction of the time what used to take months or years, simply because a steady stream of fresh insight is now flowing to leaders at every level of the organization. Nurses, technicians, team leads, and executives will not be replaced by this shift. They will simply spend less time wading through reports to find inefficiencies, shuttling between units and campuses, or reconciling perception with reality, and more time doing the work itself.



What This Means for Staffing Models, Roles, and Governance

Unless a health system is set up properly, it will not be able to use the full potential that real-time insight brings. This requires more flexibility in workflows. As Matt Cannell, Kontakt.io's Senior Director of Professional Services, wrote in his implementation article: leaders either have to deprecate current procedures and build new ones, or work a new solution into what a hospital already has.

A charge nurse, for example, needs enough flexibility to shift a technician, move a float nurse, or call in additional staff from outside sources such as staffing agencies. Without an altered operating model, she is powerless to act on real-time recommendations, however accurate they are.

Governance has to be built in from day one. This means establishing, before rollout, who holds decision-making authority when an AI-generated recommendation conflicts with existing protocol, who reviews outcomes after a change takes effect, and how escalations move through the organization when the system flags something unusual. If every reallocation of staff or equipment has to wait on approval from a manager or director, then the real-time insight cannot be utilized in time, and goes to waste.

Frontline staff should be involved early and kept engaged in decisions that will affect their daily routines, since many of these decisions land directly on their shift. Compatibility with existing systems, particularly EHRs, needs to be confirmed before rollout rather than assumed.

Alerting systems deserve particular attention: how should interventions reach staff and leaders?

Deployment

Training and Education

Any real onboarding initiative begins with mapping out and addressing the impact of AI on each end user's role and responsibilities. That means sitting down with each department before rollout, determining what normal shifts and workflows look like, and identifying where new tools can improve these routines. Without this step, training remains generic and staff stay unclear on how the new solution fits into their existing workflow.

Training also has to be tailored to each role and should mirror actual workflows, complete with screenshots, dummy interfaces, and other tools that reflect real life as closely as possible. Technicians, nurses, charge nurses, and leadership will each need different materials. Take the example of a supply chain AI that helps orchestrate equipment across a health system: a supply technician would train on a user interface, usually an EHR app, that prioritizes work orders and would learn to use a live map in case requested equipment is moved. A VP of Supply at the same health system would likely use different tools, including a real-time dashboard that displays equipment utilization, distribution, quantity, and device flows between hospitals in the same network.

It's also important to be realistic about end user impacts and avoid over-promising. If a supply chain AI saves a technician several minutes on a task, or helps a VP of Supply coordinate a system-wide equipment fleet more efficiently, the AI should be framed as exactly that: an orchestration tool that helps everyone work more effectively and reduces spend.

A super-user or champion model, where select frontline staff support their peers after go-live, helps sustain adoption.



Change Management and Communication

Setting expectations before go-live, covering what is changing, what is not, and why, reduces resistance. Workflow disruption concerns should be addressed proactively, especially for staff whose routines shift the most, and skeptics should be identified and engaged early rather than after problems surface.

Technical Readiness and Integration Validation

EHR and RTLS integration points need testing before go-live. Hardware, including phones, tags, and displays, should be provisioned and checked ahead of time. A fallback or manual process should exist in case of technical issues during the transition.


Assess Output Quality and Success

Measuring effectiveness requires hard data, rather than survey sentiment or anecdote. Across a hospital's operational use cases, devices already record data that can sharpen AI performance.

Asset location sensors track the movement of infusion pumps, wheelchairs, and monitoring equipment continuously, generating utilization, distribution, and availability data that no nurse charting system captures, and reducing search times; right now, nurses spend up to 60 minutes per shift searching for equipment, time that could go toward patient care instead (*Yen et al., 2018*).

Staff badges record location data, documenting where care interactions occur and how long they take, filling gaps in the clinical record and correcting skewed EHR timestamps. They also serve as duress buttons: in dangerous situations, nurses can subtly press the button several times to summon help to their real-time location. This safety data is also logged for later analysis and auditing.

Room occupancy sensors capture patient flow data in real time, creating the granular, continuous journey model that discharge planning algorithms require to function accurately.



It looks like you entered a patient room without sanitizing!
Please use the hand sanitization unit.

Thank you!

Hand hygiene monitoring systems record compliance events at a scale no human observer can match; one study captured over 630,000 events in a single study period, compared with 480 recorded manually (*McCalla et al., 2017*). Hospital-acquired infections account for 72,000 deaths and 687,000 infections annually, which raises the stakes of that gap considerably (*Odoom et al., 2025*).


These models can also shorten length of stay and improve hospital revenue. In 2023, a three-month study across 50 New York hospitals found 992 patients experiencing discharge delays of over two weeks, resulting in an estimated \$167 million in costs (*HANYS, 2023*). Even small interventions carry outsized impact: reducing unnecessary stays by a single day could recover up to \$2,373 per patient (*Kozma et al., 2010*).

Useful effectiveness signals, then, include reduced equipment search time, less time spent on paperwork tied to rounding or documentation, and improvement in length of stay, bed turnaround, and same-day discharge rates. An effective deployment tracks how specific interventions move these metrics over time, alongside patient access, patient satisfaction, and continued staff adoption, rather than relying on survey sentiment alone.

What a Sample Successful Deployment Could Look Like

After a health system deploys an agent, the work does not end at go-live. Feedback loops matter: frontline staff need a structured way to report friction points, and thresholds or alerts should be iterated on quickly based on real usage rather than waiting for a formal review cycle. Teams should distinguish training gaps from genuine system or workflow issues, since the fix for each looks different.

Sustainment planning follows. Someone needs to own ongoing training for new hires after the initial rollout, a refresher training cadence should be set, and governance should determine who approves future changes to thresholds or workflows.



Start Small, Scale Big

Change is intimidating, and the pressure to roll out AI quickly can be overwhelming. Just as any journey starts with a single step, an AI implementation can begin with a single workflow. Health systems can select one procedure, ideally one that already uses real-time data, and task one person with overseeing the success of the altered, AI-powered workflow, another with identifying which threshold triggers escalation, and a third with reviewing the decision and its efficacy afterward.

Beginning with one workflow is a more manageable starting point than a system-wide rollout or a new executive hire. It also gives leadership a concrete result to justify further changes to stakeholders, including clinicians, insurers, and patient advocacy groups alike.





Closing: Where to Begin?

The requirements outlined above, deep integration with existing RTLS and EHR data, clear success metrics, and a genuine partnership model with staff and vendors alike, describe what a mature AI deployment in healthcare operations looks like.

However, few systems arrive with everything already in place, so that is not a reason to wait. To make progress, health systems should treat this ebook as a starting place, rather than a finish line, working through each requirement one at a time rather than searching for a single vendor or product that satisfies all of them at once.

In addition, not every health system will start from the same place; variables include its size, existing infrastructure, and the degree to which it already captures and uses real-time data. At times, it might be easier for an outsider to accurately assess where a health system is starting from; this isn't only due to experience, but also arises from the fact that it might be easier to gauge the landscape from fresh eyes.

If you want help thinking through where your organization stands, or which workflow makes sense as a starting point, book a conversation with our team.

References and Further Reading

- Ahmad, M. A., Yaramis, I., & Roy, T. D. (2023). *Creating trustworthy LLMs: Dealing with hallucinations in healthcare AI* (arXiv:2311.01463). arXiv.
- Bertsimas, D., Pauphilet, J., Stevens, J., & Tandon, M. (2022). Predicting inpatient flow at a major hospital using interpretable analytics. *Manufacturing & Service Operations Management*, 24(6), 2809–2824.
- Chen, T. T., Chiu, M., & Wu, H. (2026). *Solving patient scheduling problems in hospitals using a dedicated GenAI system*. *Health Informatics Journal*, 32.
- Davenport, T. H. (2018). *The AI advantage: How to put the artificial intelligence revolution to work*. MIT Press.
- Glicksberg, B. S., Timsina, P., Patel, D., Sawant, A., Vaid, A., Raut, G., Charney, A. W., Apakama, D., Carr, B. G., Freeman, R., Nadkarni, G. N., & Klang, E. (2024). *Evaluating the accuracy of a state-of-the-art large language model for prediction of admissions from the emergency room*. *Journal of the American Medical Informatics Association*, 31(9), 1921–1928.
- Google for Developers. (n.d.). *Large language models: Transformers*. *Machine Learning Crash Course*.
- Healthcare Association of New York State (HANSY) *"The Scope of Complex Case Discharge Delays in New York State."* February 2023.
- Hunter, J., Vaidya, H., Crowe, S., Utley, M., King, Z., Li, K., & Harris, S. (2026). *Development of a machine learning model predicting elective surgical ICU bed demand using electronic health record data at a large UK National Health Service Trust*. *BJA Open*, 17, 100513.
- Kamel Rahimi, A., Pienaar, O., Ghadimi, M., Canfell, O. J., Pole, J. D., Shrapnel, S., van der Vegt, A. H., & Sullivan, C. (2024). *Implementing AI in hospitals to achieve a learning health system: Systematic review of current enablers and barriers*. *Journal of Medical Internet Research*, 26, e49655
- Kann, Benjamin H., et al. *Changes in Length and Complexity of Clinical Practice Guidelines in Oncology, 1996–2019.* JAMA Network Open, vol. 3, no. 3, 13 Mar. 2020, e200841,
- Klang, E., Sorin, V., Korfiatis, P., Sawant, A. S., Freeman, R., Charney, A. W., Nadkarni, G. N., & Glicksberg, B. S. (2026). *Large language models are poor clinical administrators: An evaluation of structured queries in real-world electronic health records*. *PLOS Digital Health*, 5(5), e0001326.
- Kozma, C. M., Dickson, M., Raut, M. K., et al. (2010). *Economic benefit of a 1-day reduction in hospital stay for community-acquired pneumonia (CAP)*. *Journal of Medical Economics*, 13(4), 719–727.
- Maleki Varnosfaderani, S., & Forouzanfar, M. (2024). *The role of AI in hospitals and clinics: Transforming healthcare in the 21st century*. *Bioengineering*, 11(4), 337.
- McCalla, S., Reilly, M., Thomas, R., Young, A., *An automated hand hygiene compliance system is associated with improved monitoring of hand hygiene.* American Journal of Infection Control, 45(5):492–497, 2017. DOI: 10.1016/j.ajic.2016.12.015.
- Naik, Hiten, et al. *Population-Based Trends in Complexity of Hospital Inpatients.* JAMA Internal Medicine, vol. 184, no. 2, 1 Feb. 2024, pp. 183–192,
- Nong, P., Apathy, N. C., Adler-Milstein, J., et al. (2025, January). *Current use and evaluation of artificial intelligence and predictive models in US hospitals*. *Health Affairs*, 44(1).
- NSI Nursing Solutions, Inc. (2026). *2026 NSI National Health Care Retention & RN Staffing Report*.
- Odom, A., & Donkor, E. S. (2025, July 9). *Prevalence of healthcare-acquired infections among adults in intensive care units: A systematic review and meta-analysis*. *Health Science Reports*, 8(7), e70939.
- Omar, M., Sorin, V., Collins, J. D., Reich, D., Freeman, R., Gavin, N., Charney, A., Stump, L., Bragazzi, N. L., Nadkarni, G. N., & Klang, E. (2025). Multi-model assurance analysis showing large language models are highly vulnerable to adversarial hallucination attacks during clinical decision support. *Communications Medicine*, 5, Article 330.
- Pendrill, L., Espinoza, A., Wadman, J., Nilsask, F., Wretborn, J., Ekelund, U., & Pahlm, U. (2021). *Reducing search times and entropy in hospital emergency departments with real-time location systems*. *IJSE Transactions on Healthcare Systems Engineering*, 11(4), 305–315.
- Poon, E. G., Lemak, C. H., Rojas, J. C., Guptill, J., & Classen, D. (2025). *Adoption of artificial intelligence in healthcare: Survey of health system priorities, successes, and challenges*. *Journal of the American Medical Informatics Association*, 32(7), 1093–1100.
- Sarwar, T., Seifollahi, S., Chan, J., Zhang, X., Aksakalli, V., Hudson, I., Verspoor, K., & Cavedon, L. (2022). The secondary use of electronic health records for data mining: Data characteristics and challenges. *ACM Computing Surveys*, 55(2), Article 33.
- Sinsky, C., Colligan, L., Li, L., Prgomet, M., Reynolds, S., Goeders, L., Westbrook, J., Tutty, M., & Blike, G. (2016). Allocation of physician time in ambulatory practice: A time and motion study in 4 specialties. *Annals of Internal Medicine*, 165(11), 753–760.
- Tello, M., Reich, E. S., Puckey, J., Maff, R., Garcia-Arce, A., Bhattacharya, B. S., & Feijoo, F. (2022). Machine learning based forecast for the prediction of inpatient bed demand. *BMC Medical Informatics and Decision Making*, 22, Article 55.
- Wan, F., Wang, K., Wang, T., Qin, H., Fondrevelle, J., & Duclos, A. (2025). *Enhancing healthcare resource allocation through large language models*. *Swarm and Evolutionary Computation*, 94, 101895.
- Yen, P.Y., Kellye, M., Lopetegui, M., et al. *Nurses' Time Allocation and Multitasking of Nursing Activities: A Time Motion Study.* AMIA Annual Symposium Proceedings, 2018.

Contributors



Olivia Osborn Zavalick

Olivia Zavalick (formerly Osborn) is a Clinical AI Success Consultant at Kontakt.io and an implementation and operations expert specializing in patient access and experience, ambulatory care, and Epic EHR integration. Olivia's career spans the clinical and the operational; she has worked as an EMT, in the emergency room, as an administrator at Boston Children's Hospital, and as a Huron Consulting Group consultant. She has helped lead large implementations at leading academic medical centers around the country, delivering care operations that are easier to access and better to experience.

Olivia's unique perspective shapes how she approaches implementation, going beyond simply rolling out a new technology and towards helping organizations rethink how their care teams work. In her current role at Kontakt.io, she continues to put providers and patients at the heart of everything she does, all while keeping a holistic view of the system.

Get in Touch

Kontakt.io

214 Sullivan Street, Floor 5

New York, NY 10012

+1 415-295-4558

northam@kontakt.io