SHAKUDO

How to Use Agentic AI in Health Care

A Strategic Guide to Deploying Autonomous AI Agents for Clinical and Operational Excellence

January 29, 2026

White Paper

Table of Contents

Executive Summary

Healthcare organizations are witnessing an unprecedented acceleration in AI adoption, with 86% of hospitals and health systems now utilizing artificial intelligence—a dramatic increase from just 19% two years ago. The next frontier is agentic AI: autonomous systems capable of reasoning, decision-making, and taking action with minimal human intervention.

Unlike traditional AI that analyzes data and provides recommendations, agentic AI independently executes multi-step workflows—from processing prior authorizations to coordinating care pathways to managing revenue cycle operations. Industry analysts project that by 2028, 33% of enterprise software will include agentic AI capabilities, enabling 15% of day-to-day work decisions to be made autonomously.

For healthcare executives, this represents both an opportunity and an imperative. Leading health systems like Kaiser Permanente and Mayo Clinic are investing over $1 billion in AI initiatives spanning documentation, imaging, and clinical decision support. The business case is compelling: improved margins, reduced operational costs, enhanced clinician satisfaction, and better patient outcomes.

However, success requires more than technology adoption. Healthcare leaders must address unique challenges including regulatory compliance, data sovereignty requirements, clinical validation, and integration complexity. Organizations that establish robust AI governance frameworks, prioritize high-ROI use cases, and deploy on platforms that maintain data privacy while accelerating time-to-production will lead this transformation. This whitepaper provides a strategic roadmap for deploying agentic AI across clinical and operational domains while navigating the distinctive requirements of healthcare environments.

Overview

Agentic AI represents a fundamental shift from reactive to proactive artificial intelligence. Traditional AI systems—whether predictive models or generative tools—require human oversight at every step. They analyze patient data and flag sepsis risk, but a clinician must review and act. They draft clinical notes, but a physician must approve and sign. Agentic AI, by contrast, operates with goal-directed autonomy.

These systems combine large language models with robust tool integration, enabling them to perceive their environment, reason about optimal actions, and execute multi-step processes independently. An agentic AI system might monitor real-time patient data streams, detect deterioration patterns, automatically page the appropriate specialist, retrieve relevant medical history, and prepare a preliminary assessment—all before a human intervenes.

The technology is emerging now due to the convergence of several factors:

- **Generative AI maturity**: Large language models like GPT-4 and Claude provide the reasoning engine that enables agents to interpret complex scenarios and plan action sequences
- **API proliferation**: Modern healthcare systems increasingly expose data and functionality through APIs, giving agents the "hands" to take action across disparate systems
- **Computational accessibility**: Cloud infrastructure and specialized AI accelerators have made the compute requirements for agentic systems economically viable
- **Workforce pressures**: Clinician burnout and administrative burden have reached crisis levels, creating urgent demand for automation that actually works

Adoption is accelerating rapidly. Healthcare organizations are deploying domain-specific AI tools at seven times the rate of 2024, with 22% already implementing specialized healthcare agents. This far outpaces the broader economy, where only 9% of companies have adopted AI. Major health systems are moving from pilots to production: Kaiser Permanente deployed ambient documentation across 40 hospitals and 600+ medical offices, while Advocate Health implemented 40 live AI use cases after reviewing 225 potential tools.

Yet healthcare faces unique constraints. Unlike retail or finance, healthcare AI must navigate HIPAA compliance, FDA oversight for clinical algorithms, and the life-or-death consequences of errors. Data sovereignty is non-negotiable—patient information cannot leave controlled environments for cloud-based SaaS processing. Integration is complex, with agents needing to orchestrate actions across electronic health records, imaging systems, lab interfaces, and administrative platforms.

For organizations seeking to deploy agentic AI while maintaining regulatory compliance and data control, platforms like Shakudo provide pre-integrated tool ecosystems that run entirely within customer environments—whether private cloud, VPC, or on-premises infrastructure. This enables rapid deployment of AI agents without the 6-18 month infrastructure buildout typically required, while ensuring data never leaves the organization's security perimeter.

The question for healthcare leaders is no longer whether to adopt agentic AI, but how to do so strategically, safely, and at scale.

High-Impact Use Cases Across Clinical and Operational Domains

Healthcare organizations are deploying agentic AI across a spectrum of use cases, from front-office automation to direct clinical care. Understanding where autonomous agents deliver the highest return on investment helps leaders prioritize implementations and build organizational confidence.

Administrative and Revenue Cycle Automation represents the most common entry point. Agentic systems are processing prior authorizations by automatically retrieving patient records, cross-referencing insurance policies, generating required documentation, and submitting requests—reducing what once took staff hours to minutes. Claims processing agents monitor submissions in real time, detect anomalies that could trigger denials, and proactively correct issues before claims leave the building. CMS has deployed machine learning algorithms to flag fraudulent billing patterns among Medicare agents, demonstrating how autonomous monitoring can protect program integrity.

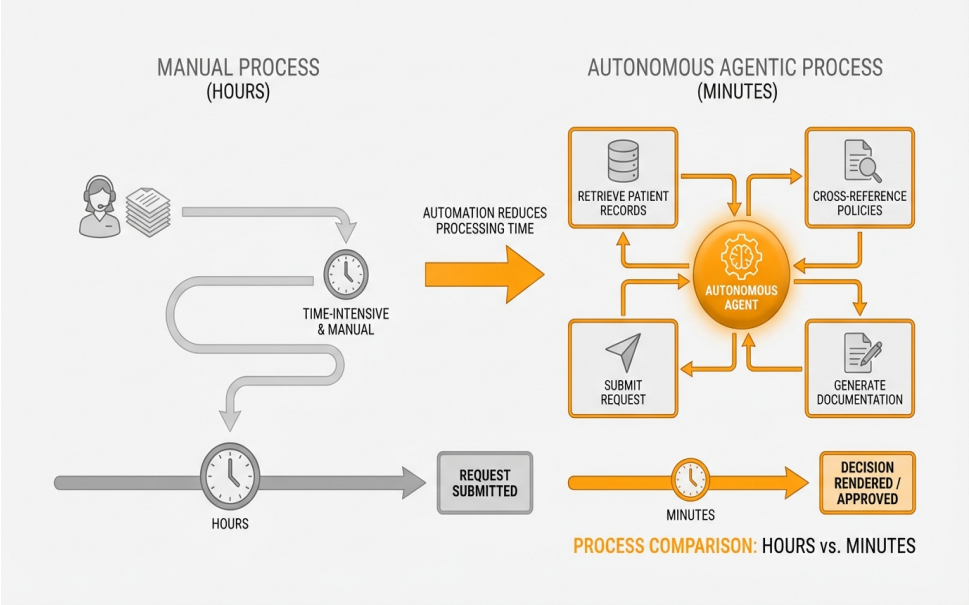

How agentic AI transforms prior authorization from hours of manual work into minutes of automated processing.

Ambient clinical documentation has become the largest category of generative AI deployment in healthcare. However, agentic versions go beyond transcription. These systems listen to patient encounters, generate structured notes, automatically populate discrete EHR fields, order appropriate follow-up tests based on clinical guidelines, and schedule next appointments—all without physician intervention beyond approval. UC San Diego Health reports that predictive AI deployed in their emergency departments saves approximately 50 lives annually by autonomously identifying deterioration patterns and alerting rapid response teams.

The most impactful administrative use cases include:

- **Prior authorization processing**: Autonomous agents retrieve records, verify coverage, generate documentation, and submit requests without human handoffs

- **Claims denial prevention**: Real-time monitoring agents analyze claims pre-submission, detect likely denial triggers, and automatically remediate issues
- **Patient intake automation**: Agents collect insurance information, verify eligibility, send intake forms, and populate registration systems before patients arrive
- **Call center augmentation**: Autonomous systems handle appointment scheduling, prescription refills, and basic triage, escalating only complex cases to human staff

Clinical decision support and care coordination represent higher-complexity use cases with significant clinical impact. Agentic AI systems are analyzing imaging studies not just to detect abnormalities, but to automatically trigger appropriate care pathways. When an agent identifies a suspicious lung nodule, it can order guideline-concordant follow-up imaging, schedule pulmonology consultation, retrieve smoking history, and generate patient education materials—orchestrating an entire care sequence.

Mayo Clinic is investing over $1 billion across 200+ AI projects that extend from administrative automation to direct clinical care, with agents helping clinicians explain AI model outputs and accelerating precision medicine research. In maternal care, AI agents are making ultrasonography more accessible by providing real-time guidance to less-experienced technicians, effectively augmenting human capability.

For value-based care organizations, agentic AI offers particular advantages. These systems continuously monitor population health metrics, identify patients falling out of compliance with care protocols, automatically reach out via preferred communication channels, and escalate cases that require clinical intervention. The automation of data aggregation and analysis frees healthcare leaders to focus on outcome improvement rather than report generation.

Fraud detection and program integrity applications leverage agentic AI's ability to monitor vast data streams and take automated action. Beyond CMS's work, health systems are deploying agents that analyze billing patterns across departments, flag statistical outliers, cross-reference against clinical documentation, and generate investigation reports for compliance teams. SimonMed, a major radiology group, is piloting over 50 AI systems across intake, documentation, and revenue cycle management.

Healthcare organizations deploying these use cases face a common challenge: tool fragmentation and integration complexity. A typical health system might evaluate 225 different AI vendors, as Advocate Health did, only to face the daunting task of integrating disparate systems with different data models, security requirements, and deployment architectures. Organizations using Shakudo's pre-integrated ecosystem of 200+ data and AI tools can deploy agentic workflows in days rather than months, with built-in governance and the flexibility to combine best-of-breed components without vendor lock-in. This is particularly valuable in healthcare, where regulatory requirements demand data sovereignty and audit trails that many cloud SaaS solutions cannot provide.

The highest-performing organizations take a portfolio approach, implementing quick-win administrative use cases to fund more complex clinical applications, while maintaining rigorous governance over autonomous decision-making in patient care contexts.

Architectural Requirements for Healthcare-Grade Agentic Systems

Deploying agentic AI in healthcare demands architectural patterns that differ significantly from general enterprise applications. The stakes of autonomous decision-making in clinical contexts require robust technical foundations that balance autonomy with safety, integration with security, and performance with auditability.

The core architecture of healthcare agentic AI consists of four layers: perception, reasoning, action, and memory. The perception layer ingests data from multiple sources—EHR systems, medical devices, patient portals, and external databases. Unlike batch-oriented analytics, agentic systems require streaming data integration to enable real-time responsiveness. Sixty percent of companies emphasize real-time data integration needs when implementing AI agents, reflecting the shift from periodic reporting to continuous monitoring.

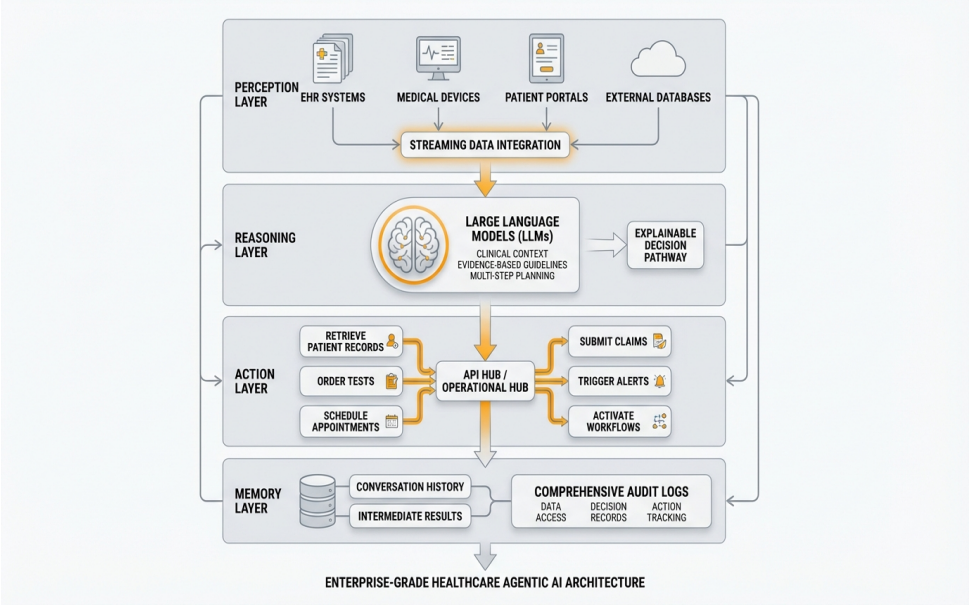

The four-layer architecture of healthcare-grade agentic AI systems, from data perception through autonomous action.

The reasoning layer employs large language models fine-tuned on medical knowledge to interpret clinical context, evaluate options against evidence-based guidelines, and plan multi-step action sequences. This is where the agent's "brain" resides, determining what actions to take based on current state and goals. However, healthcare reasoning cannot operate as a black box. Every decision pathway must be explainable, with the agent capable of articulating why it chose a particular action over alternatives.

The action layer is where APIs become critical infrastructure. By combining LLMs with robust tool integration, APIs enable agents to act as operational hubs—retrieving patient records, ordering tests, scheduling appointments, submitting claims, and triggering alerts. Yet only 10% of organizations fully document their APIs, and 75% of production APIs diverge from their formal definitions. This documentation gap creates integration risks that can cause agents to malfunction in unpredictable ways.

The memory layer maintains conversation history, tracks intermediate results, and stores learned patterns. For healthcare agents, this must include comprehensive audit logs capturing every data access, every decision made, and every action taken. HIPAA compliance demands this level of traceability, but it also enables quality improvement—analyzing agent decisions that led to poor outcomes and refining models accordingly.

Key architectural requirements include:

1. **Data sovereignty and secure deployment**: Patient data cannot be transmitted to external cloud services for processing, eliminating most general-purpose SaaS AI platforms from consideration. Agents must run entirely within the organization's controlled environment—private cloud, VPC, or on-premises infrastructure. This is non-negotiable for HIPAA compliance and essential for earning clinician trust.

2. **EHR integration without vendor lock-in**: The agent architecture must connect to Epic, Cerner, Meditech, and other core systems through HL7 FHIR APIs and proprietary interfaces. However, tightly coupling to a single EHR vendor's AI tools creates dependency and limits flexibility. A better approach uses abstraction layers that allow agents to work across multiple EHR platforms.

3. **Human-in-the-loop controls with variable autonomy**: Not all agent actions carry equal risk. Scheduling an appointment can be fully autonomous; ordering a CT scan should require clinical approval; adjusting medication dosages demands explicit physician authorization. The architecture must support granular control over what agents can do independently versus what requires human confirmation.

4. **Clinical validation and bias monitoring**: Healthcare agents must undergo validation processes analogous to medical devices, with performance monitoring that detects model drift and algorithmic bias. The architecture needs built-in mechanisms for A/B testing, holdout validation, and fairness auditing across demographic groups. Cedars-Sinai's CIO emphasizes calibrating AI adoption pace to ensure generative AI is deployed fairly, appropriately, validly, effectively, and safely.

5. **Multi-model orchestration**: No single AI model excels at all tasks. Healthcare agents typically orchestrate specialized models—one for medical image analysis, another for clinical documentation, a third for predictive risk scoring. The architecture must manage model selection, chain inference calls, and aggregate results coherently.

A significant implementation challenge is standing up this architectural stack without consuming 6-18 months in infrastructure development. Building from scratch requires assembling streaming data pipelines, container orchestration platforms, model serving infrastructure, API gateways, security controls, and monitoring systems—then integrating them into a coherent whole.

For organizations seeking to accelerate deployment while maintaining data sovereignty, Shakudo provides a unified operating system for AI that includes these architectural components pre-integrated and enterprise-hardened. Teams can deploy streaming data ingestion, model orchestration, and API management in days rather than quarters, with everything running in their own VPC or on-premises

environment. This eliminates the build-versus-buy dilemma: organizations get SaaS-like ease of deployment with on-premises data control.

The architectural foundation determines what's possible with agentic AI. Organizations that invest in flexible, secure, well-integrated infrastructure can iterate rapidly on use cases. Those that compromise on architecture find themselves locked into vendor ecosystems or unable to meet regulatory requirements, ultimately limiting AI's strategic value.

Governance Frameworks and Regulatory Compliance

The autonomous nature of agentic AI makes governance not just important but existential. Unlike supervised AI where humans review every output before action, agents make decisions and execute tasks independently. In healthcare—where errors can harm patients and violations can trigger million-dollar fines—robust governance separates successful implementations from catastrophic failures.

AI governance in healthcare requires a formal committee structure with representation from clinical, technical, legal, compliance, and operational domains. Leading organizations establish AI governance councils that evaluate every proposed use case against a consistent framework before deployment approval. This isn't bureaucracy for its own sake; it's essential risk management in an environment where autonomous systems interact with patient care.

The governance framework should address several core dimensions. First, use case evaluation: Which autonomous actions are appropriate for agents versus requiring human intervention? The committee must define autonomy boundaries clearly. Second, model validation: What evidence is required to demonstrate that an agent's decisions are clinically sound and statistically robust? Third, bias and fairness: How will the organization detect and mitigate algorithmic bias across patient demographics? Fourth, monitoring and accountability: Who "owns" each deployed agent, and what metrics trigger escalation or shutdown?

Healthcare organizations mapping out AI governance typically establish these foundational elements:

- **Use case intake and prioritization process**: Standardized mechanism for proposing agent applications, with evaluation criteria including clinical impact, regulatory risk, integration complexity, and expected ROI
- **Clinical validation requirements**: Evidence standards for different autonomy levels, from documentation assistance (lower bar) to clinical decision-making (higher bar requiring clinical trial-grade validation)
- **Bias auditing protocols**: Regular testing of agent outputs across demographic groups, with statistical methods to detect disparate impact and processes to retrain models when bias is identified
- **Incident response procedures**: Clear escalation paths when agents make errors, including immediate remediation, root cause analysis, and determination of whether to pause deployment
- **Regulatory compliance mapping**: Explicit documentation of how each agent use case complies with HIPAA, FDA guidance on clinical decision support, state medical board regulations, and payer requirements

The regulatory landscape for healthcare AI is evolving rapidly. The FDA has issued guidance distinguishing between clinical decision support tools that require regulatory oversight and those that don't, based primarily on whether the tool is intended to replace clinician judgment. Agentic systems that operate fully autonomously in clinical contexts may trigger FDA review, while agents that augment clinicians with automatic recommendations subject to approval typically fall outside device regulation. However, this distinction is nuanced and fact-specific.

HIPAA compliance presents unique challenges for agentic AI. Every agent action that accesses, uses, or discloses protected health information must have appropriate authorization and be logged in audit trails. When agents integrate with external systems or third-party AI models, Business Associate Agreements must be in place. The shift toward agentic systems intensifies these requirements because autonomous actions generate far more data access events than human-mediated workflows.

Data sovereignty emerges as a critical compliance enabler. When patient data must leave an organization's environment to be processed by cloud-based AI services, HIPAA risk increases dramatically. Even with BAAs in place, many healthcare legal teams are uncomfortable with patient data flowing to external AI vendors, particularly given uncertainties about how those vendors secure data and whether they use it for model training.

This is where deployment architecture and governance intersect. Organizations that can run agentic AI entirely within their own infrastructure boundaries—using platforms like Shakudo that operate in customer VPCs or on-premises environments—dramatically simplify compliance. Data never leaves the organization's control perimeter, eliminating a major category of HIPAA risk and satisfying even the most conservative legal interpretations. For regulated health systems, this isn't just a technical preference; it's often a governance requirement.

Stakeholder engagement is the final critical governance element. Clinical staff must be involved in designing agent workflows, not just presented with completed systems. Soliciting feedback from physicians, nurses, and administrative staff who will work alongside agents improves both adoption and safety. These frontline users often identify edge cases and failure modes that governance committees might miss.

Establishing robust governance doesn't mean moving slowly. In fact, well-designed governance accelerates safe deployment by creating clear pathways for approval, standardized validation approaches, and confidence among leadership that risks are being managed appropriately. The organizations struggling most with AI adoption are often those with either no governance (leading to paralysis from legal concerns) or governance so cumbersome that innovation stalls. The goal is structured flexibility: clear principles and processes that enable rapid iteration within appropriate guardrails.

Implementation Roadmap: From Pilot to Production at Scale

Moving agentic AI from promising pilot to production-scale deployment requires a structured approach that builds organizational capability while delivering measurable value. Healthcare organizations that successfully scale AI follow a consistent pattern: start with high-ROI use cases, establish technical foundations, prove value, then expand systematically.

Phase 1: Foundation and Quick Wins (Months 1-3)

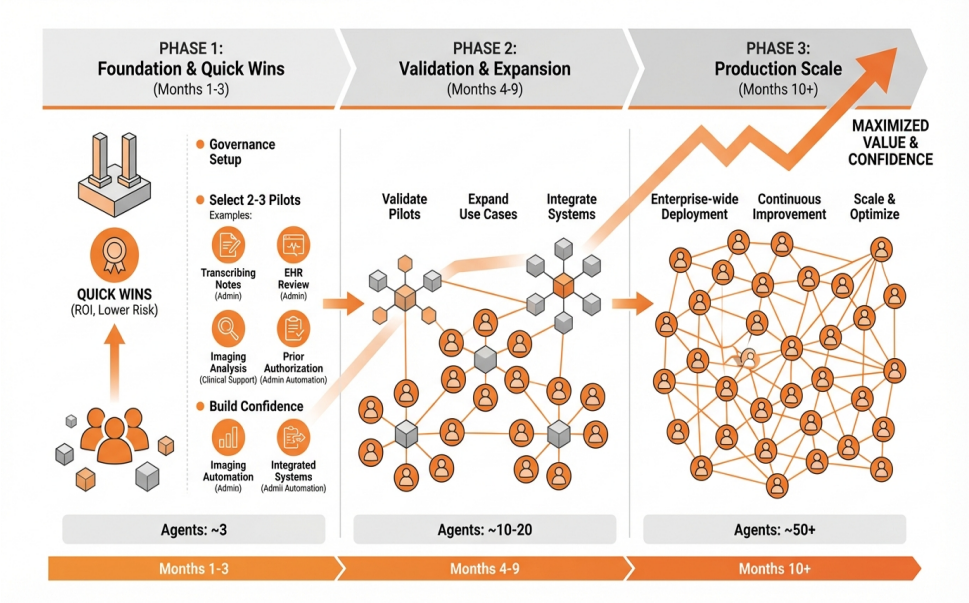

Strategic roadmap for scaling agentic AI from pilot projects to enterprise-wide production deployment.

Begin by establishing governance structures and selecting initial use cases with high probability of success. The HIMSS survey found that the most common healthcare AI applications are administrative tasks like transcribing notes, reviewing EHRs, and analyzing imaging studies. These are ideal starting points because they deliver immediate value, have lower clinical risk, and build organizational confidence in AI capabilities.

Select 2-3 pilot use cases that meet three criteria: clear ROI, manageable integration scope, and enthusiastic clinical champions. Administrative automation often provides the easiest wins. Prior authorization agents, for instance, can demonstrate value within weeks—measurable reduction in processing time, decreased staff burden, and improved authorization approval rates.

During this phase, make critical architectural decisions. Will you build infrastructure in-house, buy point solutions from multiple vendors, or adopt a unified platform approach? Organizations that choose to build face 6-18 months of infrastructure development before deploying their first agent. Those that adopt multiple point solutions encounter integration nightmares as different vendors' agents can't share data or coordinate actions.

A platform approach offers a middle path. Shakudo enables organizations to deploy production-grade AI infrastructure in days by providing pre-integrated tools, container orchestration, and governance controls

that run entirely in the customer's environment. This accelerates time-to-value for pilot use cases while establishing architectural foundations that scale to dozens of agents.

Critical activities in Phase 1:

1. **Assemble cross-functional AI steering committee** with clinical, IT, legal, and operational representation to own governance and prioritization
2. **Conduct use case inventory and prioritization** using a scoring framework that weighs clinical impact, ROI, integration complexity, and regulatory risk
3. **Stand up technical infrastructure** including data pipelines, model serving platforms, API integration layers, and monitoring systems
4. **Deploy 2-3 pilot agents** in controlled environments with extensive logging and human oversight to validate functionality
5. **Establish baseline metrics** for efficiency, accuracy, user satisfaction, and clinical outcomes that pilots will be measured against

Phase 2: Validation and Expansion (Months 4-9)

Once initial pilots demonstrate value, focus shifts to rigorous validation and expanding to additional use cases. This is where many organizations stumble—they declare victory after a successful pilot but struggle to scale because they haven't validated the agent's performance rigorously or built the infrastructure to support multiple concurrent deployments.

Clinical validation becomes paramount as you move beyond purely administrative use cases. Agents that influence clinical decisions require evidence that their recommendations align with accepted standards of care and improve outcomes. Partner with clinical informatics teams to design validation studies, often using retrospective analysis where agent recommendations are compared against actual clinician decisions and outcomes.

Integration complexity multiplies as you deploy more agents. Early pilots might integrate with one or two systems; production deployments require agents to orchestrate actions across EHRs, imaging systems, lab interfaces, billing platforms, and scheduling systems. This is where API governance and documentation become critical. Organizations must catalog available APIs, document their behaviors, version them properly, and monitor for drift.

Expand thoughtfully to 5-10 active agent deployments, deliberately selecting use cases across different domains—clinical, administrative, and operational. This builds diverse expertise within your team and demonstrates AI's broad applicability. SimonMed's approach of piloting 50+ AI systems across intake, documentation, and revenue cycle illustrates aggressive but systematic expansion.

The financial model begins to clarify during this phase. Two-thirds of healthcare CIOs report their AI strategies align well with overall business strategy, with ROI measured primarily through improved margins (26%), cost reductions (24%), and staff productivity gains (16%). Track these metrics rigorously. The most successful organizations find an AI use case with strong ROI and continually reinvest the savings into expanding AI capabilities, creating a virtuous cycle of value generation and reinvestment.

Phase 3: Production Scale and Optimization (Months 10+)

At scale, the challenge shifts from proving value to operating efficiently. Organizations managing dozens of agents need robust MLOps practices: automated model retraining pipelines, continuous performance monitoring, A/B testing frameworks, and incident management processes. The infrastructure that supported three pilot agents often buckles under the demands of thirty production agents serving thousands of users.

This is where platform decisions made in Phase 1 pay dividends or exact tolls. Organizations that cobbled together point solutions find themselves drowning in tool fragmentation—different monitoring systems for different agents, incompatible data formats, security inconsistencies. Those that built in-house struggle with the operational burden of maintaining infrastructure while also expanding use cases.

Healthcare systems deploying agentic AI at scale using Shakudo benefit from unified operations across their entire portfolio of agents. A single platform provides centralized monitoring, consistent security controls, and the ability to compose new agents by combining pre-integrated tools—without teams needing to negotiate new vendor contracts or integration projects for each use case. This reduces total cost of ownership by 40-60% compared to managing disparate tools, while accelerating deployment of new agents from months to days.

By this phase, successful organizations have embedded AI capabilities throughout their operations. Kaiser Permanente's deployment of ambient documentation across 40 hospitals and 600+ medical offices exemplifies production scale. Mayo Clinic's $1 billion investment across 200+ projects shows the strategic commitment required. Advocate Health's implementation of 40 live use cases after systematically evaluating 225 tools demonstrates disciplined expansion.

The roadmap doesn't end. Healthcare AI is evolving rapidly, with new models, techniques, and use cases emerging constantly. Organizations that establish strong foundations—governance, infrastructure, skills, and culture—can iterate continuously, staying at the frontier of what's possible while managing risk appropriately. Those that rush deployment without building these foundations inevitably hit scaling ceilings or suffer governance failures that damage trust and stall progress.

Building Technical Capabilities and Managing Change

Technology alone doesn't determine success with agentic AI. Healthcare organizations must simultaneously build technical capabilities within their teams and manage the profound organizational change that autonomous AI represents. The human elements—skills, culture, and change management—are often the binding constraints.

Technical skills requirements are evolving rapidly. IT teams accustomed to managing traditional infrastructure must now understand machine learning operations, prompt engineering, and AI model behavior. Sixty percent of companies emphasize real-time data integration needs, requiring talent skilled in streaming data architectures and AI operations. Data engineers need fluency in vector databases, embedding models, and retrieval-augmented generation architectures that enable agents to access organizational knowledge.

Clinical informaticists face new demands as well. They must bridge between clinical workflows and AI capabilities, translating physician requirements into technical specifications for agents. They validate that agent behaviors align with clinical guidelines and detect when models are hallucinating clinically implausible recommendations. This role requires deep clinical knowledge combined with understanding of AI limitations—a rare combination.

Building these capabilities requires deliberate investment. CIOs should prioritize training developers and integration specialists on AI/ML concepts, prompt engineering for configuring agents, and data science fundamentals. Don't assume that strong software engineers automatically understand AI systems. The skills are related but distinct, and treating them as interchangeable leads to poor outcomes.

Cross-functional collaboration becomes crucial. IT must work closely with business units to identify suitable use cases for agentic AI rather than imposing technical solutions in search of problems. Finance teams need to understand AI well enough to build realistic ROI models. Legal and compliance teams must become conversant in AI risks to provide practical guidance rather than blanket prohibitions.

Some organizations are creating new roles specifically for AI. Chief AI Officers and Chief Health AI Officers—like Dr. Karandeep Singh at UC San Diego Health—provide executive leadership for AI strategy and governance. These leaders combine clinical or operational expertise with technical fluency, enabling them to bridge domains and drive adoption.

However, not every organization needs to build deep AI expertise in-house from scratch. Strategic partnerships and platform choices can augment internal capabilities. Organizations using comprehensive AI platforms inherit best practices embedded in the technology, reducing the expertise required for successful deployment. This is particularly valuable for mid-sized health systems that lack the resources to hire specialized AI teams but still need production-grade capabilities.

Change management challenges are substantial. Introducing autonomous agents into clinical workflows triggers anxiety among staff who worry about job displacement, loss of professional autonomy, or being held accountable for AI errors they didn't cause. These concerns are legitimate and must be addressed directly.

Transparent communication about AI's role is essential. Frame agents as augmentation, not replacement—technology that handles routine tasks so humans can focus on complex judgment and patient interaction. Share early results showing how AI reduces burden rather than replacing jobs. At UC San Diego Health, predictive AI saving 50 lives annually in emergency departments doesn't eliminate positions; it gives clinicians earlier warning to intervene effectively.

Engage frontline staff in agent design. Physicians and nurses who help define how agents should behave become advocates for adoption rather than resistors. Their input improves agent usefulness while building investment in success. Soliciting stakeholder feedback, as best practices recommend, turns potential opponents into partners.

Address the very real question of accountability. When an agentic system makes an error that harms a patient, who is responsible—the physician, the IT department, the vendor, the health system? Legal frameworks are still evolving, but organizations must establish clear policies. Generally, clinicians remain accountable for care delivery, which is why human-in-the-loop controls for high-risk decisions are essential. Agents should be positioned as tools that clinicians use, not autonomous actors making medical decisions independently.

Measure and celebrate wins. Share metrics showing how agents reduced prior authorization processing time from 4 hours to 15 minutes. Highlight clinician testimonials about regaining time for patient care. Quantify cost savings and reinvestment into additional staff or capabilities. These success stories build momentum and organizational confidence.

Provide ongoing support as agents are deployed. Training shouldn't be a one-time event but continuous learning as agents evolve and new capabilities are added. Establish clear channels for reporting issues and requesting enhancements, and be responsive to that feedback. Nothing kills adoption faster than users feeling their concerns are ignored.

The pace of AI evolution means capabilities will continue expanding rapidly. Organizations that build learning cultures—where teams expect continuous change and view AI as an evolving toolkit rather than a fixed solution—will adapt most successfully. Those that treat AI as a one-time implementation project will struggle as the technology advances beyond their initial deployment.

Ultimately, success with agentic AI in healthcare depends on maintaining the right balance: ambitious enough to capture AI's transformative potential, but disciplined enough to manage risks appropriately. Technical excellence, strong governance, and thoughtful change management must work in concert. Organizations that integrate all three dimensions position themselves to lead in the emerging era of autonomous intelligence in healthcare.

Ready to Get Started?

Shakudo enables enterprise teams to deploy AI infrastructure with complete data sovereignty and privacy.

shakudo.io

info@shakudo.io

Book a demo: shakudo.io/sign-up