

The Enterprise AI Transformation Guide for Healthcare & Life Sciences

Accelerate your enterprise AI transformation with
proven strategies from Anthropic's customers.



Contents

Foreword	3
Step 1: Lay the foundation	4
Step 2: Launch a pilot	9
Step 3: Scale impact	12
Building an AI-first enterprise: real-world examples from Pfizer and Novo Nordisk	14
Your AI transformation starts today	17

The AI imperative for HCLS: why the time to act is now

Healthcare and life sciences (HCLS) executives face a defining moment: integrate AI or risk falling behind in one of the world's most competitive and regulated industries.

AI is already transforming outcomes, from accelerating drug discovery timelines to improving clinical trial efficiency and enhancing patient care. **PwC reports** that 88% of executives plan to increase AI investment in their teams this year, with 53% of industries increasing AI usage year-over-year.

Yet many HCLS leaders find the path to AI transformation complex and difficult to navigate, particularly given regulatory requirements and patient safety considerations. According to **McKinsey**, while 92% of companies plan to invest in generative AI (GenAI) over the next 3 years, only 1% believe their investments have reached full maturity.

The good news? You're not alone.

Working alongside industry leaders, Anthropic has developed a proven three-step approach to drive enterprise AI transformation at scale:

- **Lay the foundation:** Organizations begin by building a solid AI foundation that includes strategy development, stakeholder alignment, and governance.

- **Launch a pilot:** Next, enterprises launch carefully selected pilots that demonstrate value quickly while building expertise.
- **Scale impact:** Finally, successful pilots scale across the business through structured rollout programs and centers of excellence.

This guide covers essential elements for for successful AI transformation each step of the way.

You'll learn how to identify high-impact use cases specific to HCLS organizations—from R&D and clinical operations to regulatory affairs and commercial functions. You'll discover how to build effective business cases that secure cross-functional support while addressing compliance requirements. We'll share proven approaches for managing change, AI upskilling, and driving adoption across diverse teams. You'll also see how teams at organizations like **Pfizer** and **Novo Nordisk** deploy AI at scale to drive better clinical and research outcomes.

Let's dive in.



Chapter 1

Step 1: Lay the foundation

Step 1: Lay the foundation

Driving enterprise-wide AI transformation in HCLS starts with proactive organizational groundwork. In this section, we discuss how to define your AI strategy for HCLS-specific use cases, drive alignment among clinical and administrative stakeholders, establish robust governance frameworks that address HIPAA and other regulations, and apply best practices from HCLS organizations that have successfully navigated this journey.

Driving leadership and stakeholder alignment

Leaders must understand both the opportunities AI presents—including 20-30% gains in productivity, significant reductions in clinical documentation time, and demonstrable improvements in patient outcomes—alongside challenges such as substantial upfront investments, regulatory compliance complexity, data integration obstacles, and the critical need to maintain clinical judgment at the center of care delivery.

Building a coalition across stakeholders

In HCLS, stakeholder alignment means securing support from, among others:

- **Executive leadership** (CEO, CFO, COO) who control resources and set strategic priorities
- **Scientific leadership** (CSO, VP of R&D, Head of Medical Affairs) who understand research realities and influence scientific teams
- **Clinical leadership** (CMO, CNO, CMIO) who understand care delivery realities, influence physician adoption, and ensure patient safety

- **IT leadership** (CIO, CISO) who manage technical infrastructure and security
- **Compliance and legal teams** who ensure regulatory adherence
- **Department heads** in key pilot areas who can champion adoption within their teams
- **Frontline clinicians** who will ultimately determine whether AI tools succeed or fail

Technical solutions alone cannot drive transformation in healthcare—people, processes, and clinical workflows must evolve alongside these implementations. Organizations that treat change management as an afterthought struggle with adoption, particularly when clinicians perceive AI as adding burden rather than reducing it.

Starting with deep listening

The most successful AI transformations begin with deep listening rather than technology evangelism. You should seek to answer the questions: What are your clinicians' biggest frustrations? Where do administrative tasks pull them away from patient care? Which workflows feel broken? What keeps them working late into the night?

Starting with these pain points rather than leading with technology builds trust and ensures your AI strategy addresses real needs rather than theoretical possibilities. When clinicians see that AI initiatives target the problems that frustrate them daily, they become advocates rather than resisters

Addressing clinical skepticism directly

HCLS professionals, from doctors and nurses to clinical researchers, have seen many technology initiatives that promised to make their lives easier but instead added burden. They've endured electronic health record implementations that increased documentation time. They've experienced clinical trial management platforms that increased data entry requirements without improving trial oversight. This history creates legitimate skepticism about new technology promises.

Address this skepticism head-on by:

- Acknowledging past technology disappointments rather than ignoring them
- Committing to measuring actual impact on clinical workflows, not just technical metrics
- Establishing clear mechanisms for clinicians to provide feedback and influence implementation
- Promising to sunset AI applications that don't deliver on their value proposition
- Emphasizing that AI will enhance clinical judgment, not replace it

This honest engagement builds credibility. Clinicians respect leaders who acknowledge challenges rather than making unrealistic promises.

Assembling your AI steering committee

The key to successful change management begins with assembling a steering committee that represents critical business functions and decision-making authority. This group should include the CEO or another C-suite sponsor who can remove organizational obstacles, functional leaders who understand operational realities, technology executives who grasp implementation requirements, finance representatives who track ROI and manage budgets, and legal or compliance leaders who can establish governance frameworks.

Building AI implementation champions

In addition to assembling a larger steering committee, it's important to identify and empower change champions at every organizational level. These individuals should include respected managers who influence their peers, technical experts who understand both legacy systems and AI capabilities, early adopters enthusiastic about innovation, and skeptics whose questions reveal legitimate implementation concerns. Provide champions with additional training, direct access to leadership, and recognition that elevates their status while making their advocacy visible across the organization.

Regulatory alignment

Healthcare AI governance must address regulatory frameworks that other industries don't encounter. Building the necessary compliance into your agent architecture from day one is essential—retrofitting compliance after deployment is costly and may be impossible without rebuilding systems entirely. Below are some examples of regulatory frameworks that may apply to your organization.

HIPAA compliance for AI systems

The Health Insurance Portability and Accountability Act establishes baseline requirements for protecting patient data. Recently added cybersecurity guidelines specifically address AI systems that process Protected Health Information (PHI), requiring:

- **Business Associate Agreements (BAAs)** with all AI vendors who will access PHI. Never deploy AI systems that process PHI without proper BAAs in place—doing so creates personal liability for executives and institutional liability for the organization.
- **Robust observability and accountability mechanisms** that track AI access to patient data. You should be able to detect inappropriate access within minutes, not discover it months later during an audit.

- **Risk assessments specific to AI systems** documenting potential privacy and security risks, including what patient data the AI system will access, how data flows through the system, where data is stored, who has access, what security controls protect data, and what could go wrong.
- **Technical safeguards**, including:
 - Encrypting PHI at rest and in transit using industry-standard encryption
 - Implementing multi-factor authentication for AI system access
 - Using secure APIs with authentication and authorization
 - Monitoring network traffic for anomalous patterns
 - Implementing intrusion detection and prevention systems
- **Comprehensive audit trails** showing when AI systems accessed patient records and what actions resulted. These must be tamper-evident, retained for at least 6 years, and readily searchable.
- **Breach notification procedures** specific to AI-related security incidents with clear definitions of what constitutes a breach, who investigates, notification timelines, and remediation steps.

EU AI Act compliance

For healthcare organizations operating in Europe or serving European patients, the EU AI Act classifies most healthcare AI systems as "**high-risk**," triggering extensive compliance requirements, including:

- **Risk management systems** with documented risk assessments and mitigation strategies throughout the system lifecycle. Organizations must identify foreseeable risks, document how risks will be mitigated, monitor systems after deployment, and update assessments as systems change.
- **Data governance** ensures training data is relevant, representative, and free from bias—particularly critical in healthcare where biased training data can lead to disparate health outcomes.

- **Technical documentation** maintained throughout the system's lifecycle, including system design, training data characteristics, testing results, change logs, and user instructions.
- **Human oversight mechanisms** built into system architecture so humans can understand AI outputs, override operations when needed, and receive adequate training for oversight responsibilities.
- **Robustness, accuracy, and cybersecurity measures** appropriate for high-risk systems, with continuous monitoring after deployment.
- **Conformity assessments** before deployment, typically involving internal testing, technical documentation review, and potentially third-party assessment.

FDA oversight considerations

AI systems that meet the definition of medical devices require FDA clearance or approval before deployment. This includes:

- Software that diagnoses, treats, prevents, or mitigates disease
- Software that analyzes medical images or patient data to support clinical decisions
- Clinical decision support software that goes beyond simple reference information

Even if your initial AI applications don't qualify as medical devices, plan your governance framework with the possibility of future FDA oversight in mind. Key considerations include pre-market submission requirements, clinical validation demonstrating safety and effectiveness, labeling requirements, post-market surveillance, adverse event reporting, and managing modifications.

State and specialty-specific regulations

Don't overlook additional regulatory layers, including:

- **State medical board regulations** around clinical decision support and telemedicine—requirements vary by state

- **Specialty board guidelines** for AI use in specific clinical contexts (radiology, pathology, emergency medicine, etc.)
- **Institutional review board requirements** for AI systems used in research contexts
- **Accreditation standards** from Joint Commission, NCQA, and other bodies that may address AI use

Establish an AI governance framework

With your compliance frameworks in place, you're ready to establish governance that operationalizes these requirements. AI governance frameworks are a critical component of any enterprise AI transformation by enabling innovation while managing risk through carefully designed policies that balance protection with productivity. Here are some critical components:

- **Access controls** determine who can use AI systems and what data they can access, implementing role-based permissions that align with job responsibilities and information sensitivity levels.
- **Usage guidelines** clarify acceptable applications while explicitly prohibiting problematic uses such as processing personally identifiable information in public AI models, making employment decisions without human oversight, or generating content that could expose the organization to legal liability.
- **Quality standards** specify review requirements and accuracy thresholds, defining when AI outputs require human verification versus when they can proceed autonomously.
- **Compliance protocols** ensure AI use meets regulatory requirements across all relevant jurisdictions, from GDPR in Europe to HIPAA, FDA regulations, and GxP standards in life sciences.

The earlier you prioritize governance and compliance, the more successful your AI transformation will be in the long-term.

The importance of AI governance

AI governance is core to our DNA; in fact, we were the first AI company to achieve the **ISO 42001** for responsible AI. Many of our customers have found the following resources useful when it comes to better understanding the foundations of AI safety and governance:

- **Claude's Constitution** - Details the principles that guide Claude's behavior through our Constitutional AI approach
- **Constitutional AI: Harmlessness from AI Feedback** - Technical foundation for training AI systems with embedded values and principles



Chapter 2

Step 2: Launch a pilot

Step 2: Launch a pilot

Successful pilots deliver quick wins while building organizational capability. Choose projects carefully, showcase cross-functional potential, and learn rigorously from every experiment.

Choosing your pilot projects

AI transformation in healthcare begins with carefully selected pilot projects aligned to both business objectives and clinical priorities. The key is identifying where AI delivers the greatest impact while minimizing risk to patient safety and regulatory compliance.

Unlike other industries where "move fast and break things" is acceptable, HCLS demands a more measured approach. Focusing on lower-risk, high-value applications to start builds organizational capability and trust before tackling more complex clinical decision support systems.

Here are three low-risk, high-value pilot projects for your organization's first foray into AI:

1. AI-powered scientific documentation and report generation

AI tools that transform research data, study protocols, and clinical findings into comprehensive scientific documentation—from regulatory submissions to research reports—represent an ideal initial investment. Scientific documentation is considered "low-hanging fruit" in life sciences AI because it accelerates existing workflows without altering scientific judgment and can be piloted with small teams.

Automating documentation and report generation can also help reduce researcher burnout, improve time spent on actual scientific work, and deliver quick ROI through increased productivity.

2. Revenue cycle management: claim scrubbing

Another ideal first application for AI in healthcare orgs is an AI system that automatically reviews claims before submission to identify and correct errors, ultimately reducing insurance denials. According to the American Hospital Association, 46% of hospitals already use AI in RCM operations.

Value: AI-powered revenue cycle management could eliminate \$200-360 billion in organizational healthcare spending through effective deployment, according to McKinsey analysts. Delivers measurable financial impact through reduced denials and faster reimbursements—making ROI easy to demonstrate.

3. Automated patient scheduling and communication

As non-clinical tasks with minimal risk and easy human oversight, handling routine communications, appointment request triage, and scheduling optimization are ideal workloads for AI agents.

Value: Automated patient communications could reduce readmission rates by 30% and reduce time spent reviewing patients by up to 40% while improving patient satisfaction and freeing staff for higher-value work, according to research from the World Economic Forum.

Showcase cross-functional potential

These successful pilots create momentum - but momentum alone won't scale transformation. You need systematic approaches to knowledge transfer. Design pilots that reveal possibilities beyond their immediate scope. When clinical teams see research's success with literature review automation, they begin imagining clinical documentation applications. When drug safety demonstrates adverse event pattern detection, regulatory affairs envisions

accelerated submission preparation. When medical affairs showcase evidence synthesis, pharmacovigilance teams recognize opportunities for signal detection. This cross-pollination of ideas drives organic expansion beyond initial implementation plans.

Create structured opportunities for cross-functional learning that accelerate insight transfer across the organization, such as a monthly "AI Showcase" where pilot teams highlight their projects. These sessions should include live demonstrations of AI capabilities in clinical or research workflows, before-and-after comparisons of documentation time or analysis speed, and open discussion about regulatory considerations, patient safety protocols, and validation approaches discovered during implementation.

Aligning on clear success metrics

Before launching any pilot, establish concrete success metrics that stakeholders understand and accept. These metrics typically span four dimensions:

- **Adoption metrics** track how many users engage with AI tools and how frequently they use them—for example, measuring daily active users, feature utilization rates, and session frequency across different departments.
- **Efficiency measures** document time saved and productivity improvements, such as reducing chart preparation time from hours to minutes or enabling care coordination teams to handle more patient inquiries per hour.
- **Quality metrics** ensure AI outputs meet organizational standards through measures like 95% accuracy in data extraction, error rates below 5%, or first-pass approval rates for AI-generated content.
- **Clinical impact metrics (where applicable)** track improvements in scientific outcomes, such as faster drug discovery timelines, improved clinical trial recruitment and retention rates, enhanced protocol compliance, earlier identification of safety signals, or higher-quality regulatory submissions.

- **Satisfaction scores** capture user experience and identify areas needing improvement via Net Promoter Scores (NPS), task difficulty ratings, and willingness to recommend the tool to colleagues.

Track these metrics weekly to catch issues early, review monthly to identify trends, and adjust your approach based on data rather than assumptions. This creates accountability while building confidence in your AI initiatives.

Conduct pilot post-mortems

Once each pilot concludes, the real work of learning begins, moving beyond surface metrics to understand the deeper story of what unfolded and why.

Start by examining the numbers but also evaluate quantitative metrics. In fact, these “anecdotal” proof points are often more informative, for instance, moments when users discovered unexpected benefits, the friction points that emerged in practice, the workarounds teams invented when the technology didn't quite fit their workflow.

On the technical side, probe system reliability, integration snags, data quality surprises, and infrastructure needs that only became apparent under real-world conditions. Understanding user adoption requires detective work—why did some teams embrace the technology while others quietly resisted? What practical barriers emerged?



Chapter 3

Step 3: Scale impact

Step 3: Scale impact

Moving from successful pilots to enterprise-wide transformation requires structured training programs that build genuine AI capability, centers of excellence that include clinical domain experts, and governance that scales with adoption while maintaining patient safety and regulatory compliance.

Transform pilots into a launchpad for AI upskilling

Scaling success demands more than replicating what worked in your pilot. It requires developing deep AI capability across every role in your HCLS organization, from the C-suite to the bedside. Think of training not as mandatory compliance modules but as a strategic investment in your workforce's ability to deliver better patient care and research outcomes with less administrative burden.

Don't overlook the power of experiential learning. Hackathons inject energy and experimentation into what could otherwise feel like mandatory corporate training. When teams compete to solve real business problems using AI, learning happens organically while building genuine enthusiasm that top-down mandates can never achieve.

Consider creating peer learning and mentorship programs that pair experienced AI users with colleagues just beginning their journey. These mentorship relationships build skills while creating support networks that increase adoption. These might include regular check-ins, shadowing opportunities to observe AI use in real workflows, or problem-solving sessions.

You can also create certification programs that validate competency while fueling ambition for continued growth. Design multi-level credentials from foundational literacy through advanced practitioner status, but insist on demonstrated proficiency through real-world assessments rather than multiple-choice tests that prove nothing.

Pro-tip: Recognition matters: celebrate certified users in internal communications, provide them priority support access, and most importantly, consider certification status in promotion decisions. Nothing signals organizational commitment like tying advancement to AI proficiency. This transforms training from "nice to have" into "essential for career growth"—exactly the mindset shift you need to build true institutional capability that outlasts any individual pilot's success.

Establish centers of excellence

Create specialized teams – or centers of excellence – dedicated to sustaining and expanding AI capabilities. These centers of excellence develop best practices for implementation across functions, ensuring consistent approaches and knowledge sharing, provide technical support and troubleshooting when users encounter challenges, and systematically experiment with new use cases.

Structure centers of excellence with clear accountability and cross-functional representation that prevents siloed thinking. Include technical architects who understand system integration and data flows, domain experts from each major function who translate business needs into AI opportunities, and data scientists who optimize model performance and identify emerging capabilities.

Pro-tip: To facilitate additional upskilling, establish regular rotation programs that bring functional experts into the center for 3-6 month increments, building their AI expertise while ensuring the center stays connected to evolving business needs.



Chapter 4

Building an AI-first enterprise: real-world examples from Pfizer and Novo Nordisk

Building an AI-first enterprise: real-world examples from Pfizer and Novo Nordisk

Moving beyond isolated pilots into building AI-native HCLS operations requires reimagining entire workflows around AI capabilities rather than implementing point solutions. The most successful transformations demonstrate a common pattern: organizations that achieve breakthrough results in one function use those wins to catalyze adoption across the enterprise.

The examples below showcase how organizations use Claude to transform operations across multiple functions simultaneously—and how early wins in one area create momentum for enterprise-wide change.

Research and development: From discovery to documentation

HCLS institutions achieve their most dramatic AI gains when they connect related workflows across the research and development lifecycle. Two pharmaceutical leaders and one healthcare system demonstrate how AI transforms both ends of the R&D process.

Drug discovery acceleration: Pfizer

Pfizer, one of the world's largest pharmaceutical companies, integrated Claude in Amazon Bedrock into their internal VOX platform to accelerate drug development while maintaining rigorous standards of excellence.

The challenge: manual search processes slowing drug discovery

Pfizer's research teams faced a critical bottleneck in their drug development process. Scientists were spending countless hours navigating through more than 20,000 documents for each drug discovery project, manually searching

through vast databases to find relevant information. Against the backdrop of an industry where bringing new treatments to market faster can save lives, Pfizer needed to accelerate their research workflows while maintaining their high standards for safety and scientific rigor.

The solution: VOX platform powered by Claude in Amazon Bedrock

Pfizer integrated Claude through Amazon Bedrock into their internal VOX platform, carefully designed to enhance research capabilities through natural language search and voice command functionality. The VOX platform leverages Claude's natural language understanding to transform how scientists access and interact with research information. Rather than constructing complex search queries or navigating through layered database structures, researchers can use conversational language and voice commands to quickly locate relevant documents, studies, and data across their extensive knowledge repositories.

Results and impact

- Prototype-to-MVP timeline reduced by 50% (from over three months to six weeks)
- 55% reduction in infrastructure costs while improving capabilities
- 16,000 hours reclaimed annually—scientists redirected time toward critical research activities
- Accelerated drug development through faster access to information

Regulatory documentation: Novo Nordisk

While Pfizer accelerated research workflows, Novo Nordisk tackled the opposite end of the development pipeline: the documentation bottleneck preventing treatments from reaching patients.

The challenge: documentation delays blocking patient access

Each new treatment Novo Nordisk develops requires mountains of documentation: clinical study reports running hundreds of pages, highly technical device verification protocols, and patient guides explaining complex treatments in consumer-friendly language. Producing a clinical study report was a grueling multi-month cycle, with staff writers averaging only 2.3 CSRs per year. The stakes couldn't be higher: each day of delay in bringing a new medicine to market can cost pharmaceutical companies up to \$15 million in potential revenue, and more importantly, patients with diabetes, obesity, and other chronic conditions continue waiting for treatments that could transform their lives.

The solution: NovoScribe AI-powered documentation platform

Novo Nordisk developed NovoScribe, their generative AI platform built on Amazon Bedrock and MongoDB Atlas, with Claude models as the frontier intelligence driving the entire system. The company used Claude Code to develop the platform, combining retrieval-augmented generation with domain expert-approved text and case-specific variables to produce accurate, compliant documentation. The platform initially focused on perfecting CSR production, then expanded to device protocol documentation and patient materials.

Results and impact

- 10+ weeks to 10 minutes for clinical study documentation
- 95% reduction in resources for device verification protocols (from entire departments to single users)
- 50% fewer review cycles through improved clinical accuracy
- Instant patient materials (complete study booklets produced in under one minute)
- Regulatory validation achieved with positive feedback from regulators
- Strategic team efficiency: 11-person development team maintains agility while dramatically expanding capabilities



Chapter 5

Your AI transformation starts today

Your AI transformation starts today

Successful transformation begins with an accurate assessment of your current capabilities and a clear understanding of the path forward.

Organizational readiness assessment

Before embarking on your first few pilots, conduct an honest evaluation of where your organization stands.

Use this matrix to assess readiness across critical dimensions and determine your optimal transformation path:

Dimension	Building foundation (1-2)	Growing capability (3-4)	Transformation ready (5-6)	Your Score
Executive commitment	AI viewed as IT project	CEO interested; competing priorities	CEO championing; multi-year commitment	--
Data infrastructure	Siloed systems; manual processes	Centralized warehouse; basic governance	Modern platform; automated pipelines	--
Technical capabilities	Legacy systems; limited cloud	Hybrid cloud; basic DevOps	Cloud-native; strong engineering team	--
Change management	Initiatives frequently stall	Mixed track record; moderate adoption	Proven success; high trust culture	--
Cross-functional collaboration	Departmental silos	Regular meetings; shared goals emerging	Integrated teams; aligned incentives	--
AI/ML maturity	No AI experience; exploratory	Initial ML models in production	Multiple AI applications deployed	--
Risk & compliance	Reactive; manual controls	Established program; regular audits	Proactive; automated controls	--
Budget & resources	Project-based funding	Annual AI budget established	Multi-year investment secured	--

Scoring guide:

30-48 points (High Readiness): Launch comprehensive transformation with multiple pilots across functions.

16-29 points (Moderate Readiness): Begin with 3-5 strategic pilots while addressing foundational gaps.

8-15 points (Building Readiness): Secure executive sponsorship and establish governance before launching 1-2 narrow pilots.

The vision: AI-first enterprises that reshape healthcare & life sciences

Six months from now, your organization could look radically different: Research teams accelerating drug discovery timelines by months as AI agents synthesize decades of scientific literature and identify novel therapeutic targets. Regulatory affairs teams preparing submissions in weeks instead of quarters, with AI assistants that understand CFR, ICH guidelines, and global requirements. Clinical operations resolving protocol deviations and safety signals in real-time, with AI that learns from every trial interaction.

But the real transformation goes deeper than operational efficiency. AI-first HCLS organizations bring therapies to patients faster, identify safety signals earlier, and unlock insights buried in decades of clinical and real-world data. They empower every scientist to focus on discovery, every medical affairs professional to synthesize evidence at scale, and every pharmacovigilance team to detect patterns invisible to manual review.

The question isn't whether AI will transform healthcare and life sciences—it's whether your organization will lead that transformation or follow competitors who moved first.

Ready to begin your AI transformation?

Join enterprises like Eli Lilly, Abbvie, and Sanofi who are already transforming their operations with Claude.

Start building today:

Contact Anthropic's Sales team to discuss your specific use cases and implementation strategy

Learn about our **Claude for Life Sciences** solution, trusted by organizations like Eli Lilly, AbbVie, and Sanofi to accelerate scientific discovery with enterprise AI

Access Claude through our **API**, **web interface**, or **Claude for Work** to begin piloting with your teams

Review our enterprise documentation and safety guidelines at **[docs.anthropic.com](#)** and our **[Trust Center](#)**

Explore **Claude Code** to automate clinical engineering workflows with agentic coding



<https://claude.com/product/overview>