White Paper

The Human Edge of AI: Building Trust, Safety and Value in Regulated Healthcare Systems

A strategic framework for responsible AI governance

MARIE FLANAGAN, Director, Product, Regulatory & AI Governance SME, IT Design & Development, IQVIA
MIKE KING, Senior Director, Product & Strategy, IQVIA
JANE REED, Director, Life Sciences, IT Design & Development, IQVIA

Table of contents

Executive summary

As AI transforms healthcare, organizations face a critical challenge: harnessing AI's capabilities while preserving the human expertise, ethical judgment, and institutional knowledge essential to patient safety. This framework synthesizes IQVIA's insights for responsible AI governance in regulated healthcare.

The pharmaceutical, medical device and in-vitro diagnostics industries stand at an inflection point. AI promises revolutionary advances, yet without adequate governance, poses significant risks to patient safety and compliance. The challenge isn't whether to adopt AI but how to implement it responsibly.

This paper provides Quality Assurance and Regulatory Affairs (QARA) professionals and clinical leaders a roadmap to navigate AI transformation while maintaining the human judgment healthcare demands, balancing innovation with safety, efficiency with oversight and technological capability with ethical responsibility.

AI should amplify human judgment, not replace it. Success demands a hybrid intelligence model where machine precision enhances clinical expertise, ethical reasoning and regulatory interpretation.

Five interconnected pillars of responsible AI governance

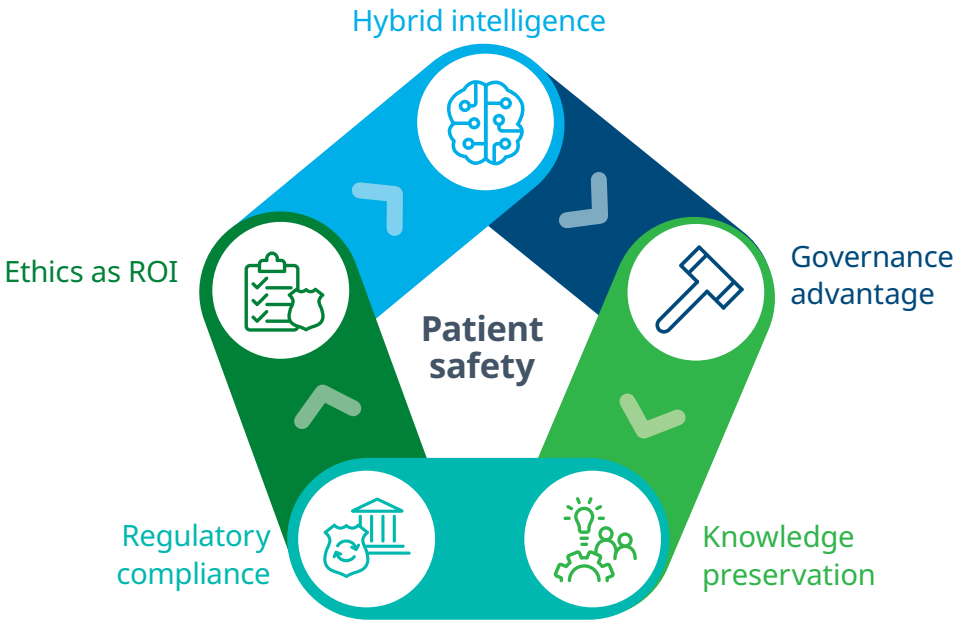

Key points:

- **Human-in-the-loop is non-negotiable:** Regulators worldwide mandate meaningful human oversight, making it the trust layer that legitimizes AI in clinical settings. This requirement reflects irreplaceable human capabilities in clinical judgment, ethical reasoning, and regulatory interpretation. AI should amplify these human strengths, never replace them.
- **Governance builds competitive advantage:** Transparency, validation and auditability create market differentiation and accelerate regulatory approval. Organizations with mature governance frameworks enjoy faster market access, stronger clinician trust and reduced regulatory friction. The pillars of institutional knowledge — clinical judgment, ethical judgment and regulatory interpretation — remain the gold standard that AI must enhance, not replace.
- **Preserving institutional knowledge is critical:** As AI automates routine tasks that traditionally provided learning opportunities, organizations risk losing tribal knowledge and contextual expertise that decades of experience have built. This competency cliff threatens patient safety when less experienced staff lack judgment to appropriately question AI recommendations or recognize when escalation is warranted.
- **AI literacy is now a regulatory mandate:** The EU AI Act and FDA guidance require organizations to demonstrate workforce competency, transforming training from operational choice to compliance requirement. Organizations must prove staff understand AI limitations, can recognize bias and exercise appropriate oversight. This mandate creates both compliance obligations and strategic opportunities for capability development.
- **Ethics and beneficence define ROI:** Healthcare AI success must be measured by patient outcomes and safety improvements, not automation speed or cost reduction alone. The principle of beneficence demands AI systems demonstrably improve patient care, not merely avoid harm. This patient-centric measure of success creates necessary discipline in resource allocation.

The regulatory landscape is evolving rapidly. While AI capabilities advance exponentially, regulatory frameworks struggle to keep pace. Organizations must work within current standards while remaining flexible enough to accommodate future guidance. Success requires starting with proven governance approaches and adapting systematically as requirements crystallize.

The hybrid future: AI as amplifier, not replacement

The most fundamental insight from healthcare AI deployment: AI exponentially expands human expertise rather than replacing it. Across medical device innovation, pharmaceutical manufacturing and Post-Market Surveillance (PMS), successful organizations embrace hybrid intelligence where machine precision and human judgment leverage capabilities that neither could achieve alone.

This partnership model represents a fundamental departure from traditional automation. Where automation has replaced human labor in repetitive tasks, AI identifies patterns, generates insights and makes predictions that inform human decision-making. The human role shifts from execution to oversight, from routine processing to strategic judgment.

The human-machine partnership model

Medical device innovation and development

AI enables new depth in product development and post-market intelligence. Organizations mine manufacturing records, service logs, clinical notes and user feedback across millions of data points to uncover patterns invisible to manual analysis. Regulatory documentation that previously required weeks — risk assessments, standards mapping, global requirement comparisons — can now be dramatically accelerated through AI assistance, though expert review remains essential.

PMS represents perhaps the most transformative application. Traditional systems operate retrospectively: events occur, reports are filed, analysis happens, responses are implemented. AI enables a shift toward real-time monitoring where statistical anomalies, emerging patterns and clusters of similar issues can be identified immediately rather than discovered months later during periodic reviews.

Pharmaceutical manufacturing

AI transforms quality control, process optimization and regulatory compliance. Predictive maintenance prevents equipment failures before they impact production. Real-time monitoring detects process deviations instantly. Pattern recognition identifies quality issues across batches that human reviewers might miss. Yet humans remain essential for interpreting results within regulatory contexts, making critical decisions about batch release and managing exceptions that AI systems cannot handle.

Clinical trials

AI accelerates patient identification by analyzing electronic health records for enrollment criteria matches. Protocol optimization uses historical data to predict recruitment challenges and optimize design. Safety monitoring flags potential adverse events across larger datasets faster than manual review. However, human investigators maintain ultimate accountability for patient safety, protocol adherence and ethical oversight. In short, the machine provides signals; humans provide judgment.

Why humans remain essential

Three capabilities distinguish human judgment from machine intelligence in healthcare: clinical expertise for contextual interpretation, ethical reasoning for navigating moral complexity and regulatory interpretation requiring nuanced understanding of evolving frameworks. These capabilities cannot be automated — they represent the irreplaceable human edge.

The regulatory imperative reinforces this reality. U.S. FDA guidance, EU Medical Device Regulation and emerging global AI-specific frameworks worldwide mandate meaningful human oversight. This requirement reflects recognition that healthcare decisions carry consequences too significant to delegate entirely to algorithmic systems, regardless of their accuracy.

Figure 1: The three pillars of non-replaceable human judgment

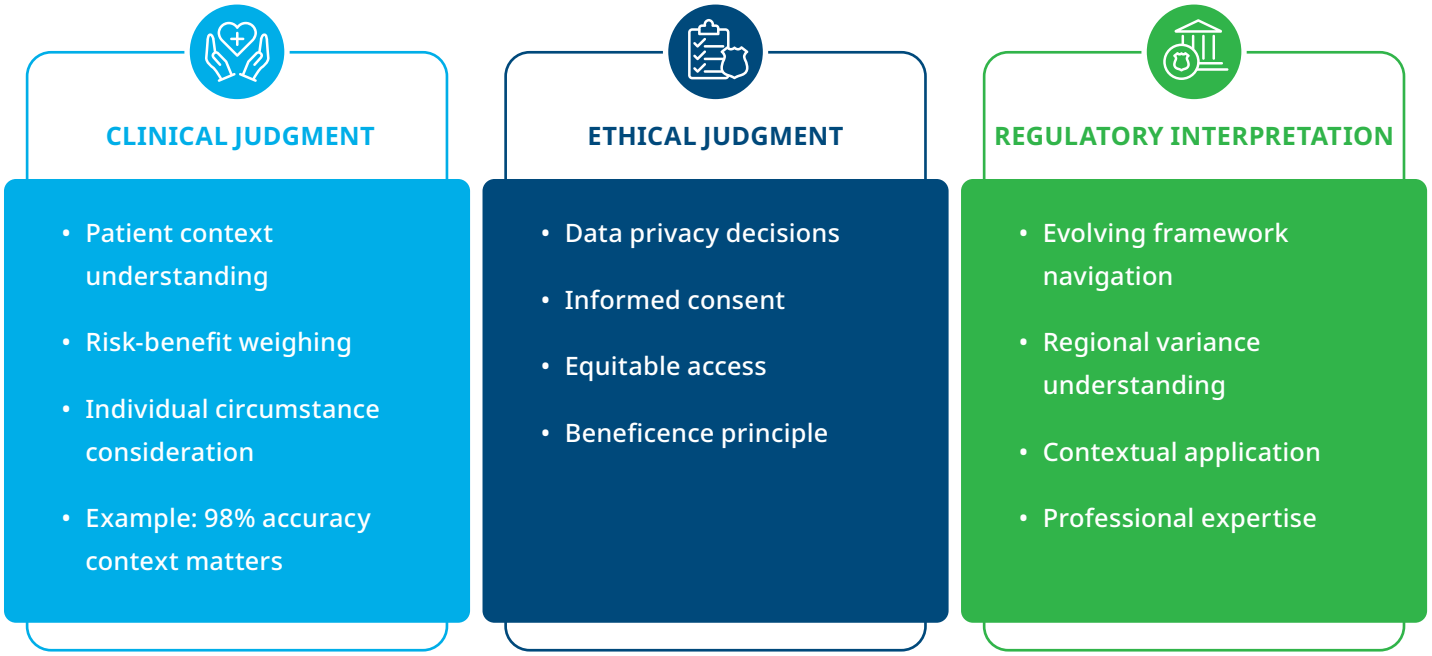

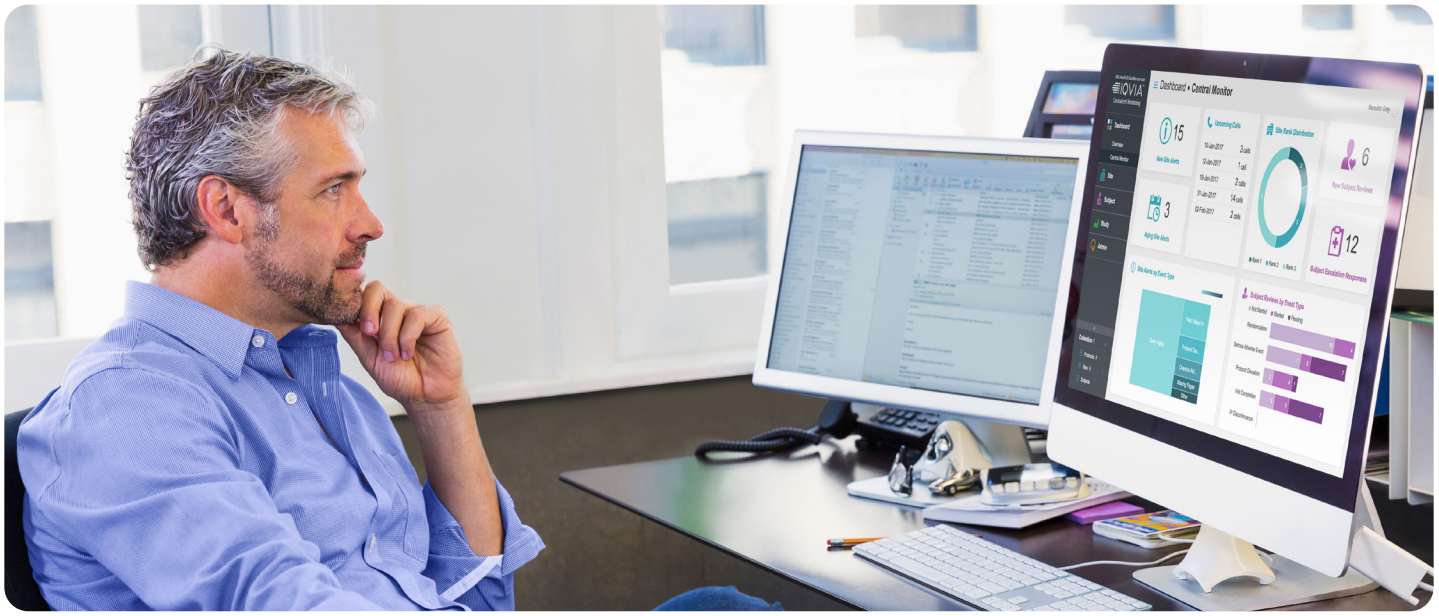

Governance as trust infrastructure

Effective AI governance isn't regulatory overhead — it's strategic infrastructure enabling innovation while protecting patients and ensuring compliance. Organizations with mature governance frameworks enjoy faster regulatory approvals, stronger clinician trust and reduced risk. Three pillars define robust governance:

TRANSPARENCY: Making AI decisions understandable
Stakeholders must understand how AI systems reach decisions. This requires clear documentation of data sources, algorithmic logic, validation methods and limitations. Transparency builds trust with regulators who must approve systems, clinicians who must rely on recommendations, and patients whose care depends on AI-informed decisions.

The business case for transparency extends beyond compliance. Organizations that can clearly explain AI decision-making gain faster regulatory approval, easier clinical adoption and stronger competitive positioning. When competitors struggle to demonstrate how their "black-box" systems work, transparent "white-box" approaches win market access.

VALIDATION: Proving performance and safety
Rigorous validation proves AI systems perform as intended across diverse scenarios and populations. This includes testing for bias, validating against ground truth, establishing performance benchmarks and continuous monitoring post-deployment. Validation rigor must match risk levels - higher-risk applications demand more extensive testing.

Traditional software validation approaches require adaptation for AI. Static systems can be fully tested against specifications. AI systems that learn from data require ongoing validation as they encounter new scenarios. Organizations must establish protocols for initial validation, continuous monitoring, periodic revalidation and managing model drift.

AUDITABILITY: Creating complete documentation
Complete audit trails document decisions, changes and oversight activities. This enables regulatory inspection, internal quality reviews, root cause analysis and continuous improvement. Version control tracks model changes, data lineage documents training inputs and decision logs capture AI recommendations alongside human actions.

Figure 2: Governance trust stack

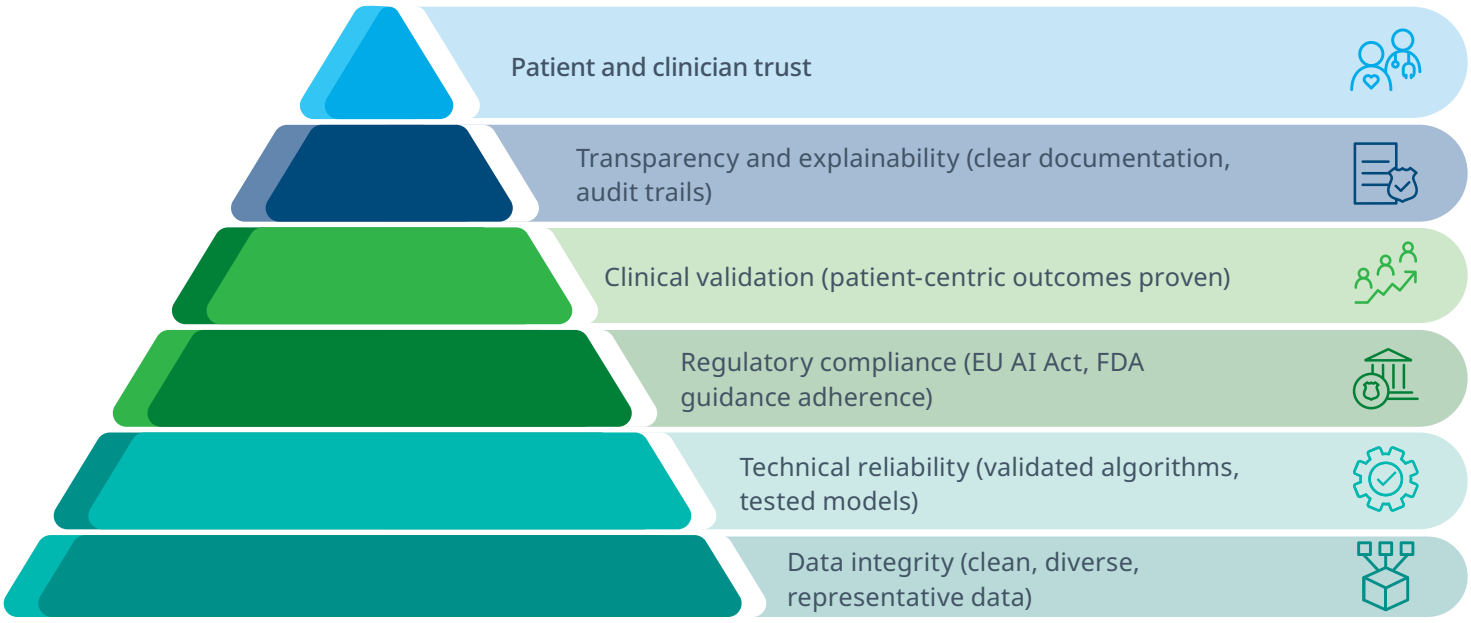

The challenge: balancing comprehensive documentation with practical usability. Audit trails must be detailed enough for regulatory scrutiny yet accessible enough for daily operations. Automated logging, structured data capture and clear escalation procedures make auditability sustainable rather than burdensome.

You can't have AI governance without business, quality and tech at the same table together.

Governance structures that work

Effective governance requires multidisciplinary teams with genuine authority: quality and regulatory professionals who understand compliance requirements, clinical experts who interpret medical implications, data scientists who explain algorithmic behavior and ethicists who navigate moral complexity. This team reviews AI applications before deployment, monitors performance continuously and makes binding decisions about implementation.

Organizations that treat governance as formality rather than substance face predictable consequences: regulatory delays, clinical resistance, safety incidents and erosion of stakeholder trust. Those that invest in governance infrastructure realize competitive advantages that compound over time.

Figure 3: The three-legged stool — multidisciplinary governance

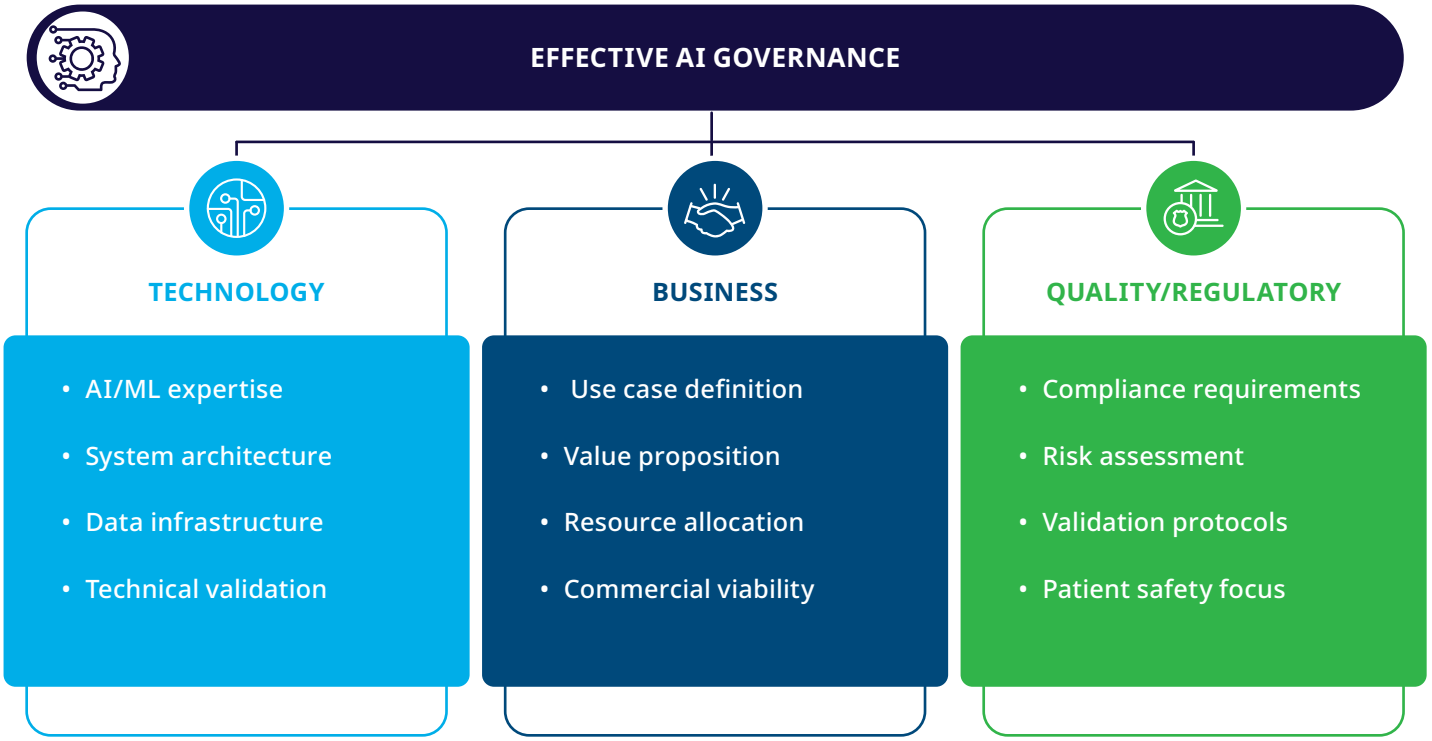

The evolving regulatory landscape

The regulatory environment for healthcare AI remains dynamic. The EU AI Act establishes comprehensive requirements for high-risk AI systems, including mandatory conformity assessments, technical documentation and ongoing monitoring. FDA guidance continues evolving, with frameworks for Software as a Medical Device (SaMD) and predetermined change control plans enabling limited algorithm modifications without full resubmission.

Organizations must navigate this uncertainty strategically:

- Begin with governance approaches exceeding current requirements to build margin for future regulation
- Establish processes flexible enough to accommodate new guidance without wholesale restructuring
- Participate in industry working groups shaping regulatory frameworks
- Document decision rationale thoroughly so retrospective compliance becomes manageable when requirements crystallize

The organizations succeeding in this environment treat regulatory uncertainty as a competitive opportunity rather than an implementation barrier. While competitors wait for perfect clarity that may never arrive, leaders deploy AI within robust governance frameworks that position them for success regardless of how regulations evolve.

Table 1: Global regulatory landscape comparison

Region	Definitional approach	Submission pathways	Self-learning algorithm policy	Timeline
United States (FDA)				
European Union (EMA)				
United Kingdom (MHRA)				
Asia-Pacific (PMDA, TGA)				
Rest of world				

Established | Evolving | Unclear/Fragmented

The competency challenge: Preserving institutional knowledge

Organizations face a strategic paradox: AI automates routine tasks that traditionally provided learning opportunities for junior staff, while experienced professionals who possess critical institutional knowledge approach retirement. This creates a “competency cliff” — a sudden loss of expertise as automation eliminates developmental pathways and veterans exit simultaneously.

The risk is concrete. Junior quality professionals who have never manually reviewed manufacturing deviations may lack judgment to recognize when AI flags require escalation versus dismissal. Regulatory specialists who never drafted submission documents from scratch may miss nuances in AI-generated content. Clinical reviewers who never personally evaluated adverse event reports may trust AI assessments without appropriate skepticism.

While the full potential of AI is still emerging, the industry must ensure a strong pipeline of junior professionals who can grow into future experts — providing the clinical context and professional judgment required for AI-enabled, regulated activities.

Strategic responses to the competency cliff

Structured knowledge transfer programs

Organizations must deliberately capture and document institutional knowledge before it walks out the door. This includes creating decision trees that codify expert judgment, documenting the rationale behind "tribal knowledge" practices, recording case studies of complex scenarios AI cannot handle and establishing formal mentorship programs pairing veterans with junior staff.

Redesigned development paths

Junior staff need exposure to complex scenarios AI cannot handle. This requires intentional rotation programs ensuring contact with challenging cases, assignment of edge-case investigations requiring deep thinking, participation in cross-functional problem-solving teams and graduated responsibility with appropriate oversight rather than immediate AI delegation.

Evolved veteran roles

As AI handles routine work, experienced professionals should shift toward exception handling and quality oversight, training and mentorship of junior staff, validation and governance activities, complex problem-solving requiring deep expertise and cross-functional leadership bridging technical and clinical domains. This evolution preserves institutional knowledge while leveraging automation's efficiency.

Deliberate capability development

Organizations must resist short-term cost pressures to eliminate developmental opportunities. Assign some complex cases manually to junior staff with oversight rather than routing everything through AI. Maintain capacity for manual processes as backup and training ground. Create simulation exercises that develop judgment without risking patient safety. Measure staff capability development as key performance indicator alongside efficiency metrics.

It's easier to teach AI to a domain expert than to teach decades of healthcare expertise to an AI scientist.

The competency challenge isn't theoretical future concern — it's present reality requiring immediate action. Organizations that fail to preserve institutional knowledge while implementing AI face long-term capability erosion that undermines the very judgment AI systems depend on for effective oversight.

Figure 5: Skills hierarchy for AI-enabled healthcare

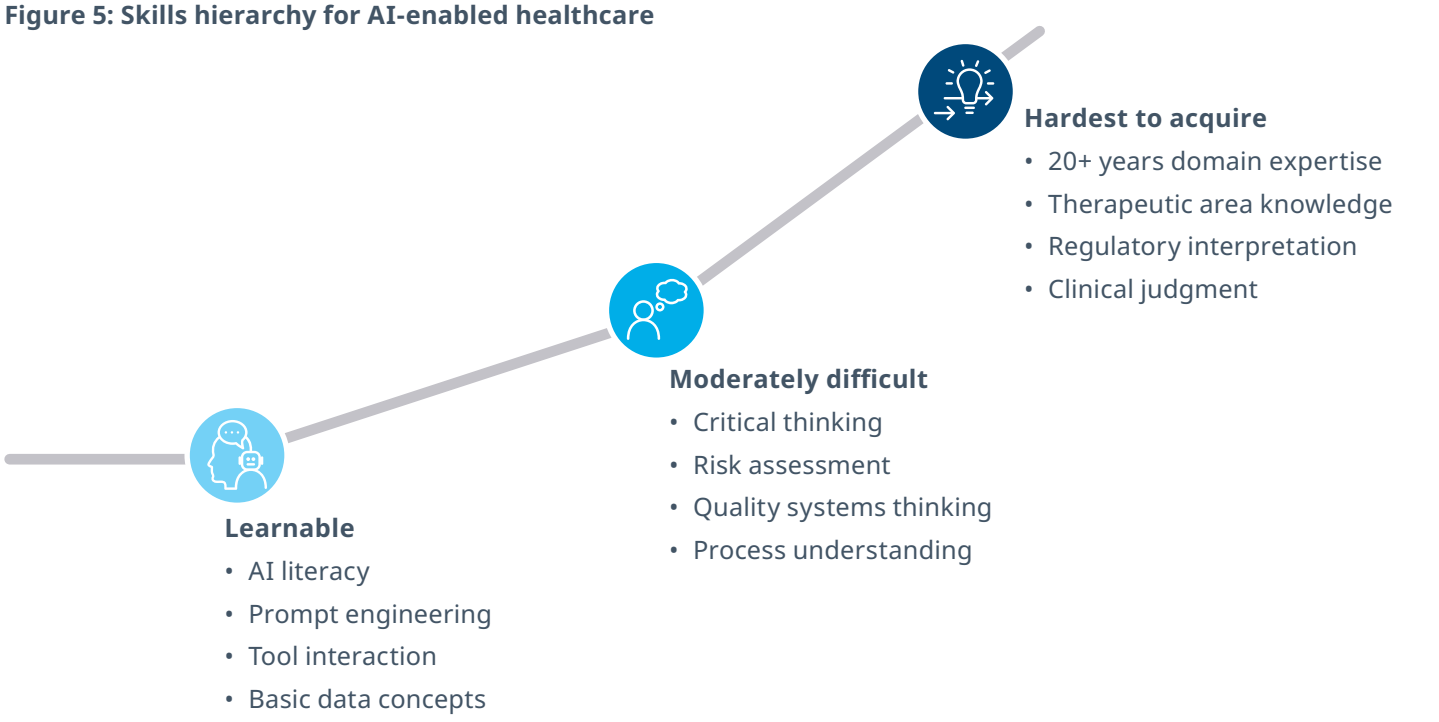

AI literacy as regulatory compliance

The EU AI Act transformed AI literacy from competitive advantage to legal requirement. Organizations deploying high-risk AI systems must demonstrate workforce competency in understanding, using and overseeing these technologies. This regulatory mandate creates both compliance obligations and strategic opportunities.

Figure 6: The three training pillars

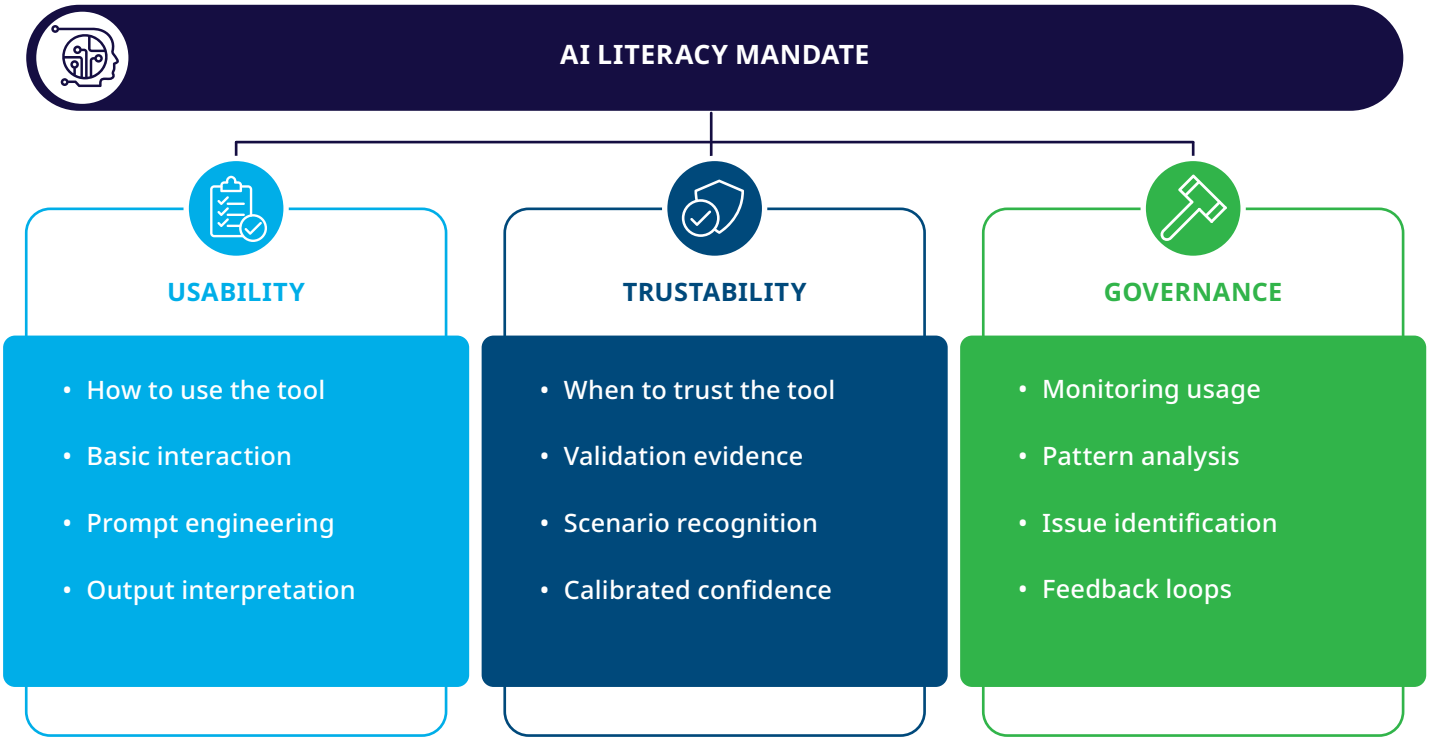

Three dimensions of AI literacy

USABILITY: Effective system interaction
Staff must understand how to interpret system outputs, recognize when results warrant skepticism, execute appropriate oversight procedures and escalate concerns through proper channels. This goes beyond basic technical proficiency to include judgment about when and how to apply AI recommendations in practice.

TRUSTABILITY: Critical assessment capability
Users must recognize potential bias in training data and algorithmic outputs, understand validation methods and their limitations, question anomalous or unexpected results and identify when AI confidence levels should influence decision-making. This dimension separates appropriate reliance from dangerous over-trust.

GOVERNANCE: Regulatory and quality framework
Staff need understanding of regulatory requirements governing AI use, documentation and audit trail expectations, change management and version control procedures, incident reporting and response protocols and their individual accountability within the governance framework. This ensures compliance becomes organizational capability rather than individual burden.

Building literacy programs that work

Effective programs require role-based training tailored to specific responsibilities, hands-on experience with actual systems rather than theoretical instruction, regular refresher training as systems evolve, competency assessment demonstrating understanding rather than mere attendance and integration with existing quality and regulatory training rather than treating AI as separate domain.

The strategic opportunity: upskilling experienced professionals creates hybrid expertise more valuable than hiring AI specialists lacking domain knowledge.

A quality professional who understands both AI capabilities and pharmaceutical manufacturing brings perspective no pure technologist can match. A regulatory specialist who combines submission expertise with AI literacy can navigate novel approval pathways competitors cannot access.

Organizations that treat AI literacy as compliance checkbox rather than capability development miss this opportunity. Those that invest in comprehensive workforce development position themselves for sustained competitive advantage as AI becomes ubiquitous across healthcare operations.

Ethics and beneficence: The patient-centric imperative

Healthcare AI ethics center on beneficence — systems must demonstrably improve patient care, not merely avoid harm. This principle elevates the bar beyond traditional software validation. While conventional systems succeed by functioning as specified without causing injury, healthcare AI must actively enhance outcomes, safety or experience to justify deployment.

Figure 7: Redefining ROI — from efficiency to patient outcomes

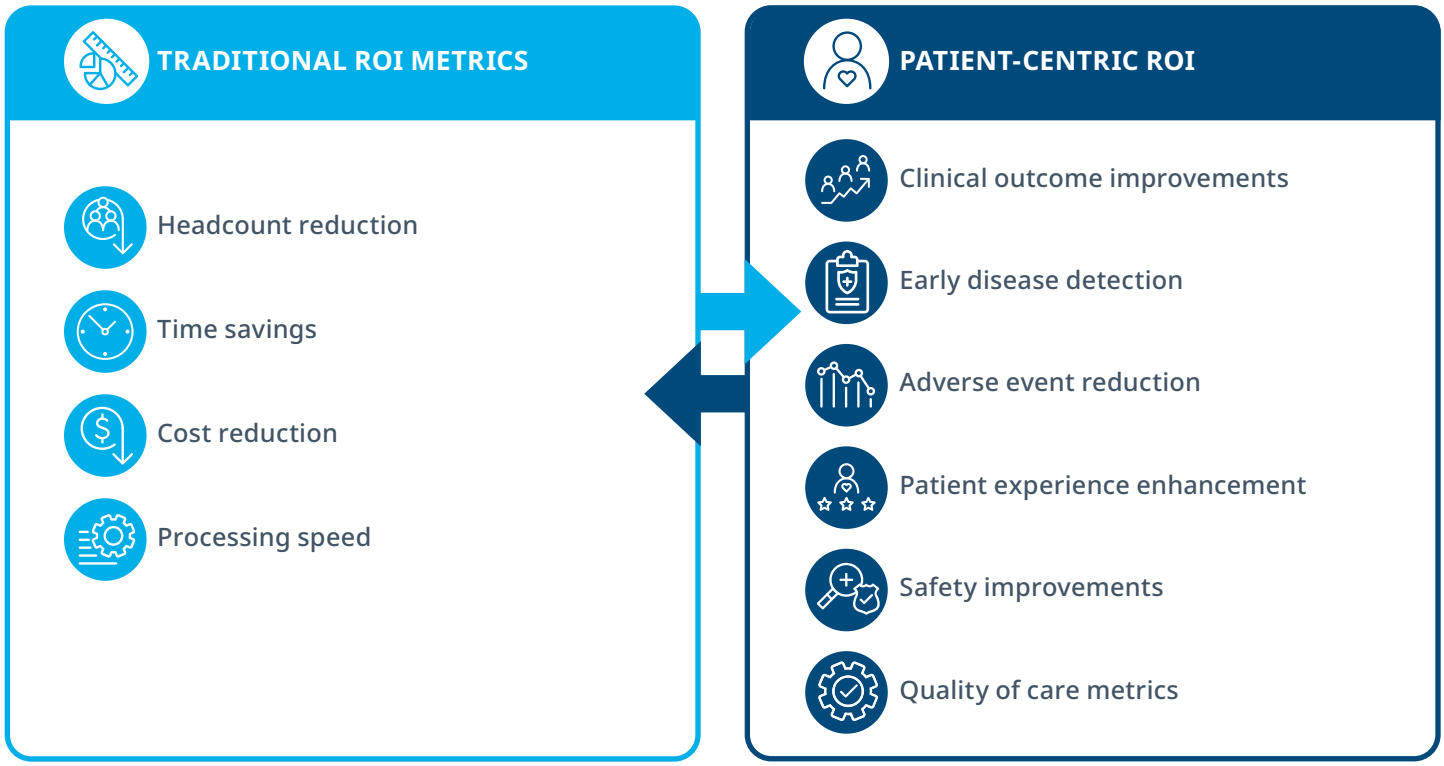

Four ethical principles for healthcare AI

Patient benefit as primary measure

Every AI application must answer: Does this improve patient outcomes? Not does it reduce costs, accelerate processes or demonstrate innovation - does it benefit patients? This criterion forces discipline in resource allocation, ensuring investment flows toward applications with genuine clinical impact rather than impressive technology lacking medical value.

Bias detection and mitigation

AI systems trained on historical data inherit historical biases. If training data underrepresents certain demographic groups, system performance may degrade for those populations. Organizations must actively test across demographic categories, monitor real-world performance by subgroup, implement correction strategies when bias emerges and maintain transparency about known limitations affecting specific populations.

Transparency with patients

Patients deserve to know when AI influences their care. This includes clear communication about AI's role in diagnosis, treatment planning or monitoring decisions,

explanation of how AI recommendations factor into clinical judgment, disclosure of limitations or uncertainties in AI assessments and opportunity to question or appeal AI-informed decisions. Transparency respects patient autonomy and builds trust in AI-enabled healthcare.

Continuous monitoring and accountability
Deployment doesn't end ethical responsibility — it begins ongoing obligation to monitor real-world impact on patient populations, investigate unexpected outcomes or performance degradation, maintain mechanisms for reporting concerns and demonstrate responsiveness to emerging safety signals. This continuous accountability ensures systems that validated successfully in controlled settings perform ethically in clinical practice.

This is more than just cost-saving...this is about making sure patients experience a high quality of healthcare.

Measuring patient-centric ROI

Traditional ROI metrics — cost savings, efficiency gains, productivity increases — inadequately capture healthcare AI value. Patient-centric ROI measures clinical outcome improvements such as reduced complications or faster diagnosis, safety enhancements including earlier adverse event detection, patient experience benefits like reduced wait times or improved communication and health equity advances addressing disparities in care quality or access.

Healthcare AI must not only avoid harm but actively improve patient outcome.

Organizations that measure AI success primarily through operational efficiency miss the ethical imperative and strategic opportunity. Those that focus on patient benefits in decision-making build sustainable programs aligned with healthcare's fundamental purpose: improving human health and wellbeing.

Figure 8: Beneficence framework — AI ethical hierarchy

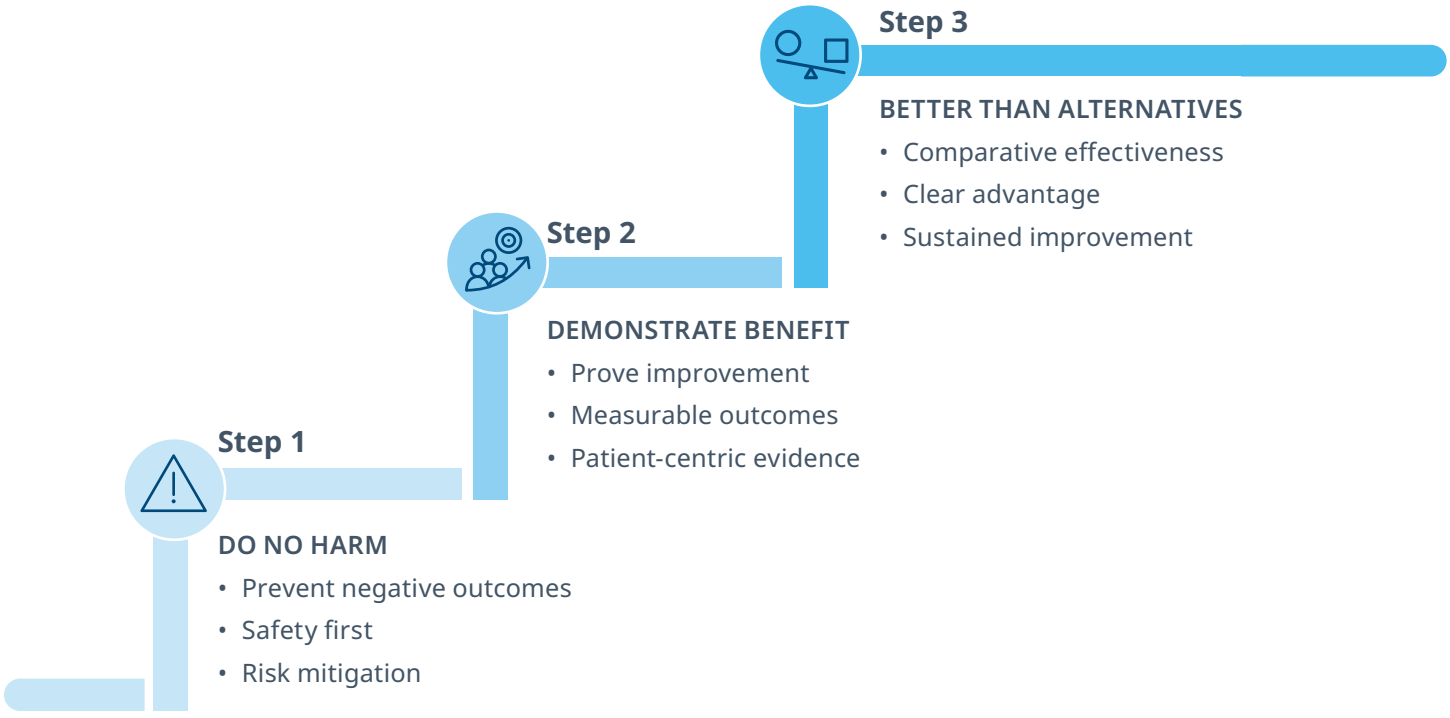

The implementation reality: Common challenges

Organizations implementing AI governance frameworks encounter predictable challenges. Understanding these obstacles enables proactive mitigation rather than reactive crisis management.

Organizational resistance and change management
AI disrupts established workflows and professional identities. Quality professionals who built careers on manual review expertise may view AI as threat rather than tool. Clinical staff accustomed to autonomous decision-making may resist oversight requirements AI systems impose. IT teams managing validated systems may fear audit exposure from AI's inherent uncertainty.

Effective change management requires transparent communication about AI's role and limitations, involvement of frontline staff in system design and validation, training that builds confidence rather than compliance, demonstrated leadership commitment beyond technology pilots, and celebration of early wins building momentum for broader adoption. Organizations that treat AI implementation as purely technical project fail. Those recognizing it as organizational transformation succeed.

Data quality and availability
AI systems require high-quality data in substantial volumes. Healthcare organizations often discover their data exists in incompatible formats across disconnected systems, lacks standardization enabling aggregation, contains gaps or errors undermining training, or faces privacy restrictions limiting use. Addressing these challenges requires data governance infrastructure many organizations lack.

The pragmatic approach: start with data-rich domains where quality issues are manageable, invest in data infrastructure as strategic asset rather than operational cost, establish data governance alongside AI governance, accept that initial applications may address narrow use cases while building foundation for broader

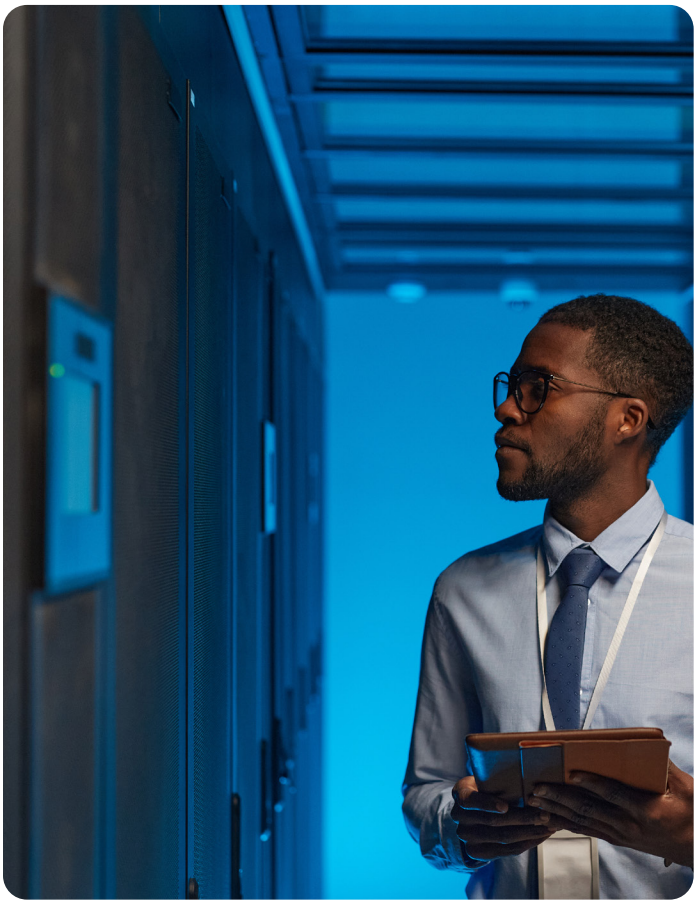

deployment. Perfect data is unattainable. Sufficient data is achievable.

Resource constraints and competing priorities
AI governance requires sustained investment in specialized expertise, training programs, validation infrastructure and ongoing monitoring systems. Organizations face competing demands for resources and pressure to demonstrate rapid ROI. This tension creates temptation to cut corners, for example, skip comprehensive validation, minimize training and defer monitoring investments.

The false economy of inadequate governance becomes apparent when regulatory delays exceed cost savings, clinical adoption fails due to trust concerns or safety incidents trigger costly remediation. Organizations that treat governance as discretionary expense rather than essential infrastructure pay higher costs through slower deployment, greater risk and missed competitive opportunities.

FIGURE 9: THE RESPONSIBLE AI IMPLEMENTATION ROADMAP

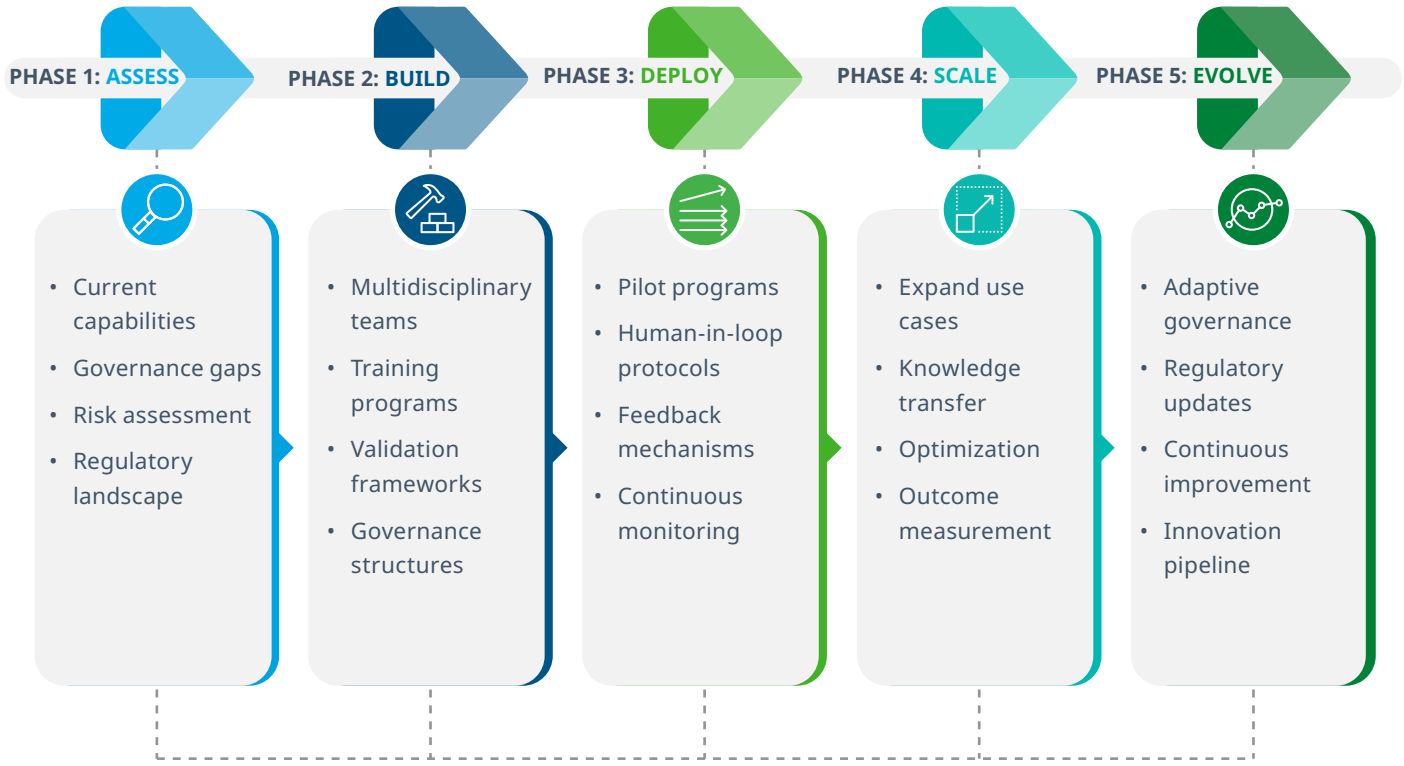

Practical deployment strategies

Moving from framework to reality requires practical strategies addressing how organizations actually deploy AI in regulated environments. These approaches emerge from collective experience of organizations navigating the messy realities of implementation.

STARTING SMALL: The pilot approach

Organizations achieve greatest success beginning with limited-scope pilots in domains where AI provides clear value without excessive risk. Document management systems that organize regulatory submissions, surveillance tools identifying potential signals in post-market data, quality control applications detecting obvious defects — these applications deliver tangible benefits while building organizational capability for more complex deployments.

The pilot approach enables learning governance practices in manageable contexts, demonstrating value to skeptical stakeholders, identifying technical and organizational challenges before high-stakes deployment, and building internal expertise through hands-on experience. Organizations rushing to deploy AI across critical processes without this foundation face predictably higher failure rates.

Building cross-functional teams

Successful AI deployment requires perspectives traditional organizational silos cannot provide. Quality professionals understand compliance requirements but may lack AI technical expertise. Data scientists grasp algorithmic capabilities but often miss regulatory nuances. Clinical experts provide medical context but may not appreciate validation requirements. Bringing these perspectives together creates teams capable of navigating AI's unique challenges.

Effective teams include quality and regulatory affairs professionals ensuring compliance, clinical or domain experts providing medical/scientific context, data scientists explaining algorithmic behavior and limitations, IT professionals managing infrastructure and

integration, and ethics representatives addressing moral implications. This diversity enables holistic evaluation of AI applications from technical feasibility through ethical acceptability to regulatory viability.

Managing stakeholder expectations

AI generates unrealistic expectations — technology leaders promise revolutionary transformation while skeptics predict disaster. Managing these extremes requires transparent communication about what AI can and cannot do, realistic timelines acknowledging validation and governance requirements, clear articulation of roles and responsibilities in human-AI partnership, and honest discussion of risks alongside benefits.

Organizations that overpromise AI capabilities face inevitable disappointment, damaging credibility and hindering future initiatives. Those setting appropriate expectations while delivering incremental value build stakeholder trust enabling sustained investment in AI capabilities.

Creating feedback mechanisms

AI systems require continuous improvement based on real-world performance. This demands structured mechanisms for capturing user feedback about system usability and accuracy, monitoring performance metrics detecting degradation, analyzing error patterns revealing systematic issues and implementing improvements through controlled change management. Organizations treating AI deployment as one-time implementation fail. Those building continuous improvement into system design succeed.

The most effective feedback loops make reporting issues easy for frontline users, ensure reports receive timely response demonstrating organizational commitment, track patterns across reports revealing systematic rather than isolated problems and close the loop by communicating improvements back to reporting users. This virtuous cycle builds user confidence while improving system performance.

Measuring success and driving continuous improvement

What gets measured gets managed. Organizations need clear metrics demonstrating AI value while identifying areas requiring improvement. The challenge lies in balancing quantitative performance measures with qualitative assessments of clinical impact and safety.

Performance metrics that matter

Put simply, traditional IT metrics — system uptime, processing speed, user adoption rates - provide necessary but insufficient measures of AI success in healthcare. Organizations must track clinical outcome improvements including earlier disease detection, reduced adverse events, improved treatment efficacy or enhanced diagnostic accuracy. Operational efficiency gains matter but only alongside quality improvements: reduced processing time for regulatory submissions, accelerated post-market signal detection or decreased manufacturing defect rates.

Safety metrics require particular attention. Track false positive and false negative rates understanding their clinical implications, monitor near-miss incidents where AI recommendations required human override, analyze error patterns revealing systematic issues and measure time to identify and address safety signals. These metrics ensure AI systems maintain safety profiles that justify their deployment.

User experience metrics reveal whether AI truly augments human capability or creates friction. Measure user confidence in AI recommendations, frequency of overrides indicating trust issues, time required for human review of AI outputs and user-reported usability concerns. Systems technically accurate but practically unusable fail to deliver promised value.

Building learning loops

Metrics become meaningful only when organizations act on insights they reveal. Effective learning loops establish regular review cadences examining performance trends, conduct root cause analysis when metrics indicate problems, implement corrective actions through controlled change management and measure whether changes achieve intended improvements. This disciplined approach transforms monitoring from passive observation to active optimization.

The frequency of review cycles should match system risk and stability. High-risk applications warrant more frequent assessment. Newly deployed systems require closer monitoring than mature, stable ones. Organizations must balance thoroughness with resource constraints, focusing detailed analysis where it generates greatest value.

Sharing learnings across the organization

AI deployment generates valuable lessons applicable beyond individual applications. Organizations capturing and sharing these insights accelerate subsequent implementations while avoiding repeated mistakes. This requires documentation of what worked and what didn't in deployment processes, technical approaches enabling or hindering success, governance practices proving effective or burdensome and training methods building versus merely checking competency boxes.

Creating forums for cross-functional teams to share experiences — regular communities of practice, lessons-learned sessions after major deployments, accessible repositories of implementation guidance and mentorship programs pairing experienced implementers with new initiatives — builds organizational capability transcending individual projects. Organizations treating each AI deployment as isolated initiative waste resources relearning known lessons. Those systematically capturing and sharing knowledge compound capabilities over time.

Practical implementation: Case vignettes

Abstract governance principles become concrete through real-world applications. These vignettes illustrate how organizations navigate regulatory constraints, balance competing priorities and build trust through demonstrated governance.

Table 2: Case study comparison matrix

CASE STUDY	CHALLENGE ADDRESSED	AI APPROACH	HUMAN ROLE	OUTCOME
IQVIA self-learning validation	• Continuously evolving algorithms	• Disable self-learning, validate fixed models	• Human oversight of version control	• Regulatory compliance achieved
Manufacturing quality control	• Inconsistent visual inspection	• AI defect detection	• Human decision-making on flagged items	• Improved consistency + maintained safety

Case vignette 1: Turning off self-learning

Every AI demonstration IQVIA conducts for QARA professionals raises the same question within the first ten minutes: "How do you validate a self-learning, continuously evolving algorithm within a quality management system governed by 21 CFR Part 11, ISO 13485, and EU Medical Device Regulation?"

The question reflects genuine dilemma. Traditional validation assumes stable systems with deterministic behavior - define requirements, test against specifications, document results, maintain validated state through change control. Self-learning AI systems violate this assumption fundamentally. Every patient interaction potentially modifies the algorithm. How do you maintain validation when the system changes continuously?

The pragmatic solution

IQVIA's approach is to turn off self-learning. This sacrifices some AI adaptability but gains critical advantages including regulatory compliance through traditional validation pathways, audit trail integrity with documented system versions, performance predictability enabling meaningful testing, manageable risk through controlled updates and commercial viability with acceptable regulatory timelines.

When this approach is explained, QARA professionals visibly relax. The lesson: effective AI governance sometimes means constraining technological capabilities to fit regulatory realities. This isn't compromise but strategic decision-making prioritizing deployment viability over theoretical optimization.

Case vignette 2: Manufacturing quality control partnership

Pharmaceutical manufacturing has employed AI-powered visual inspection for decades, creating extensive learning about where AI excels, where humans remain essential, and how to structure optimal partnership of the two.

- **AI strengths:** Consistency across millions of inspections without fatigue, throughput enabling 100% inspection rather than sampling, detection of subtle variations human eyes might miss, objective application of criteria without subjective drift
- **AI limitations:** Cannot distinguish truly problematic defects from cosmetic imperfections requiring contextual judgment, struggles with novel situations outside training data, requires retraining to adapt to process changes, lacks accountability for decisions

The optimal partnership:

AI handles detection of potential defects across 100% of production. Humans review all flagged items and make final acceptance/rejection decisions based on contextual knowledge of manufacturing conditions, product specifications and regulatory requirements. Configurable threshold settings enable quick verification of obvious cases while flagging marginal situations for deeper review. Sample testing of AI-accepted items provides independent validation. Feedback loops where human decisions inform system refinement enable continuous improvement.

This partnership achieves what neither participant could accomplish alone: the consistency and throughput of automated inspection combined with the contextual judgment and accountability of human oversight. The result is higher quality assurance at greater efficiency than either approach independently.

Conclusion: The path forward

The AI transformation of healthcare is neither optional nor reversible. Quality, regulatory and clinical leaders face not whether to adopt AI but how to implement it responsibly, preserving human expertise, ethical judgment and institutional knowledge that keep patients safe.

Five pillars for responsible implementation:

Pillar 1
Hybrid future

AI should amplify human judgment through partnership, not replacement. Structure roles, training and systems around this complementary relationship. Clinical judgment, ethical reasoning and regulatory interpretation cannot be automated — they represent the irreplaceable human edge.

Pillar 2
Governance as trust infrastructure

Transparency, validation and auditability create strategic advantages beyond compliance. Organizations with mature governance stand to gain faster regulatory approvals, stronger clinician trust, enhanced patient confidence and competitive differentiation in markets demanding governance evidence.

Pillar 3
Knowledge preservation

As AI automates routine tasks, seek to deliberately preserve institutional knowledge through structured transfer programs, the documentation of tribal knowledge, evolved veteran roles, development opportunities for junior staff and resistance to short-term cost-cutting pressures that erode capability.

Pillar 4
AI literacy as regulatory compliance

The EU AI Act transformed workforce competency from competitive advantage to legal requirement. Effective programs address usability, trustability and governance. The strategic opportunity: upskilling experienced professionals creates hybrid expertise more valuable than hiring AI specialists lacking domain knowledge.

Pillar 5
Beneficence and ethics as ROI

Healthcare AI success must be measured by patient outcomes - clinical improvements, experience enhancement and safety gains. This patient-centric ROI creates discipline ensuring resources flow toward applications with genuine clinical impact rather than impressive technology lacking medical value.

Seven strategic imperatives for success:

Establish multidisciplinary governance early: Create teams with genuine authority combining quality/regulatory professionals, clinical experts, data scientists and ethicists. Review AI applications before deployment, monitor performance continuously and make binding implementation decisions.

Invest in workforce development systematically: Implement role-based training, hands-on experience with actual systems, regular refresher training, competency assessment and integration with existing quality programs. Treat AI literacy as strategic capability not a compliance checkbox.

Start with high-value, low-risk applications: Build momentum through early wins demonstrating tangible benefit. Gain organizational experience before tackling complex, high-stakes deployments. Learn governance practices in manageable contexts.

Build feedback loops from inception: Design systems capturing user feedback, monitoring real-world performance which enables continuous improvement. Make adjustment capacity a core feature rather than an afterthought.

Prioritize validation rigor matching risk levels: Higher-risk applications demand more extensive testing. Establish protocols for initial validation, continuous monitoring, periodic revalidation and managing model drift. Adapt traditional validation approaches for AI's unique characteristics.

Communicate transparently to build trust: Be clear with regulators, clinicians and patients about AI's role, limitations and governance. Transparency creates foundation for stakeholder confidence enabling clinical adoption and regulatory approval.

Preserve institutional knowledge deliberately: Document decision rationale, create mentorship programs, redesign development paths, evolve veteran roles and maintain capability for manual processes. Resist short-term pressures eliminating learning opportunities.

The human edge of AI isn't a limitation to overcome. It's the foundation upon which responsible healthcare innovation must be built.

The ultimate measure: Patient benefit

Every AI implementation decision should answer one question: Does this benefit patients? Not does it reduce costs, improve efficiency or demonstrate innovation... does it benefit patients?

When organizations maintain this focus — measuring success by clinical outcomes, preserving human judgment for patient safety, implementing governance building trust and investing in workforce capability enabling appropriate oversight — they position AI to fulfill its transformative potential responsibly.

Organizations implementing this framework don't just deploy AI safely - they establish sustainable competitive advantage built on trust, expertise and demonstrated patient benefit. The path forward combines machine precision with human judgment, creating healthcare AI that serves its ultimate purpose: improving patient care while honoring the irreplaceable value of human expertise, ethical reasoning and clinical wisdom.

Figure 10: IQVIA differentiators

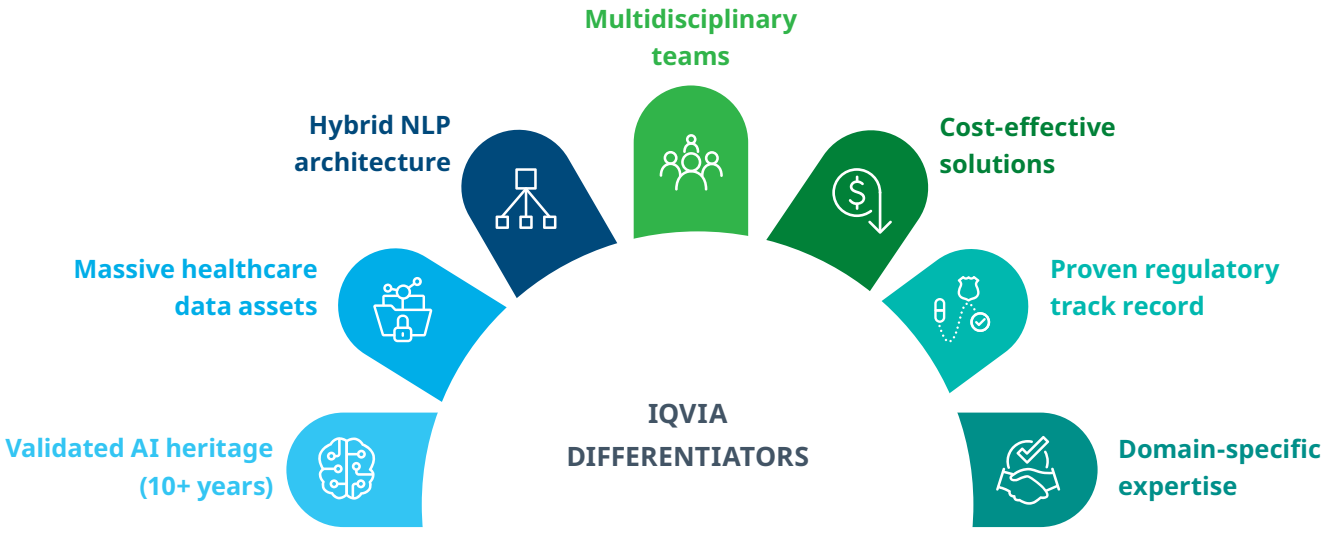

Next steps: Advancing your AI governance readiness

Take a strategic approach that guarantees the compliance of your quality management and regulatory activities with the applicable standards to your target markets.

100+ installs partnering with QA/RA professionals in design & development, quality management, supplier control, global market access and post-market activities:

IQVIA SmartSolve® eQMS: Automate and optimize enterprise-wide quality management processes, from design through to manufacturing and distribution through to post-market, all within a single, AI enabled, globally compliant platform.

IQVIA SmartSolve® RIM: Manage global regulatory submissions, product registrations, health authority interactions, and country submission documentation with confidence — fully integrated with your quality workflows.

IQVIA Regulatory Intelligence: Access real-time global regulatory requirements and updates in 89+ countries, receive insights and alerts on new or updated regulations, and streamline compliance.

IQVIA Detect: Identify potential product quality issues and adverse events using NLP AI to analyze sources such as social media, audio files, and engineering service records.

IQVIA Collect: Collect and manage safety data efficiently to support pharmacovigilance and regulatory compliance.

IQVIA NLP: Extract insights from real-world data using NLP to analyze unstructured sources and accelerate evidence generation.

IQVIA DDAI: Ensure your data and AI solutions are trustworthy, compliant, and aligned with evolving healthcare and regulatory standards.

CONTACT US

eQMS@iqvia.com

iqvia.com/smartsolve

IQVIA TECHNOLOGIES

02.2026.CS. BCS2025-4024-12DEC