Infosys topaz

WHITE PAPER

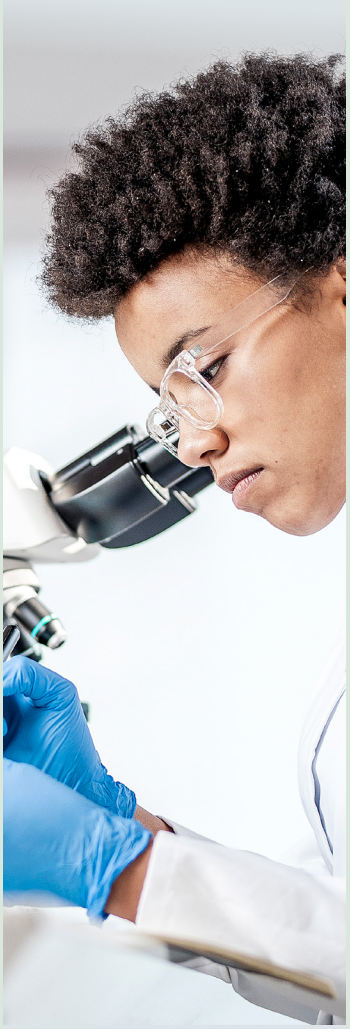

AI IN HEALTHCARE+: GOVERNING THE NEXT FRONTIER

ADAPTING GOVERNANCE FOR AI

Infosys®
Navigate your next

Introduction:
Building Trust in AI-enabled HealthCare

Artificial Intelligence (AI) is reshaping health and care across diagnostics, treatment, and operations. But alongside opportunity come challenges: misdiagnosis risks cause harm and erode trust, while privacy breaches threaten physical and mental safety. Traditional clinical governance, long established in healthcare to ensure safety, quality, and accountability, provides a necessary foundation for navigating these challenges.

To responsibly integrate AI into health and care, we propose adapting clinical governance into a broader trusted "HealthCare+ Framework" that is agile enough to support any emerging technology, such as AI. Such a framework must be built on trust, transparency, and shared responsibility. At its core are five considerations that reflect the values of safe, ethical, and person-centred health and care. Together, these offer a unified approach to governance that includes AI and supports innovation while safeguarding the well-being of those we serve. They are not rigid rules but adaptable foundations that can be tailored to the unique context of each organisation.

Key Stakeholders in Trusted HealthCare+ Governance for AI

Successfully integrating AI into health and care requires a shared commitment to safety, ethics, and continuous improvement across three key stakeholder groups: health technology solution providers, healthcare providers, and policymakers. Each plays a distinct yet interconnected role in shaping a trustworthy and effective governance framework.

Health Technology Solution Providers

Health technology solution providers must go beyond technical innovation to embed safety and integrity across the product lifecycle. Clear use cases, rigorous assessments, and close collaboration with healthcare implementers ensure solutions are ethical, usable, and context aware. Embedding governance in development builds transparency, trust, and a culture of safety-first.

HealthCare Providers

Healthcare providers and their organisations are the frontline users of AI-enabled health and care technologies, bringing extensive experience in clinical governance and healthcare systems designed to ensure safety, quality, and continuous improvement in the delivery of care. As AI becomes embedded in clinical workflows, these governance systems must evolve to include the technologies that assist, enable or influence the delivery or access to health or care. Rather than creating separate frameworks, healthcare providers can refresh existing governance structures to incorporate emerging technologies such as AI, ensuring that innovation aligns with patient safety and care standards. Their role is critical in designing and evaluating AI tools, monitoring performance, and maintaining human oversight to prevent automation bias and compromised decision-making.

Policymakers

Policymakers set the guardrails for responsible AI in health and care. By embedding healthcare governance into law, policy, and culture, they safeguard trust, ensure accountability, and provide system-wide leadership as AI becomes part of national health infrastructure.

Towards a Unified Approach

Understanding the distinct roles of health technology solution providers, healthcare providers, and policymakers is essential, but it's only the starting point. To build a truly trustworthy, agile and effective HealthCare+ Governance Framework, all three must align around an approach that includes the foundations for safe, ethical, and person-centred innovation, guiding how AI technologies (and future emerging technologies) are designed, implemented, and monitored across the entire health and care ecosystem.

Here, we explore five foundational considerations. These will naturally overlap and reinforce one another; they will look and feel different in each organisation, though together they create a good foundation for the procurement and implementation of safe, ethical, and effective AI and technology integration in health and care. We will explore how they can be adapted and applied to AI development, offering a practical roadmap for responsible governance and highlighting the risks of falling short.

1. Culture, Leadership, and People

A strong culture begins with leadership that visibly commits to safety, quality and the continuous improvement of technology. This leadership extends beyond executives, and it includes health and care leaders, managers, and frontline staff who collectively shape the organisation's governance culture.

Application to AI:

In the context of AI, this approach must be broadened to ensure that governance is embedded across the entire AI lifecycle from design to decommissioning.

Leadership commitment: Leaders can lead by example by actively promoting and modelling behaviours that prioritise AI safety and ethical use. This means setting clear expectations for responsible AI development and deployment, enabling resources for governance, and embedding oversight into strategic and operational decision-making involving AI components.

Workforce engagement: All people involved in AI development, implementation and use should understand their roles in AI governance. They must be empowered to raise concerns, report issues, and contribute to continuous improvement. Training in AI safety, ethics, and governance should be an integral part of both onboarding and ongoing professional development.

Accountability and transparency: A culture of curiosity, open communication and learning from mistakes is essential. This supports proactive identification and management of risks associated with AI technologies, including algorithmic bias, data misuse, and unintended clinical consequences.

Organisational alignment: The governance of AI must be integrated across all levels and functions, ensuring that oversight is not siloed but embedded in daily practice from clinical workflows to enterprise operations and vendor partnerships.

2. Person-centred or Person-Led Care

Our governance must provide a focus on person-centred care, designing products and services around individuals' needs, rights, and preferences. It integrates lived experience and active input into evidence-based technology development. Beyond health and care outcomes, it emphasises self-determination and data sovereignty, especially in AI interactions that capture individuals' health data. Person-led means engaging communities with dignity and respect to build systems that serve people, not the other way around.

Application to AI:

AI technologies must be developed and deployed in ways that uphold dignity, autonomy, and equity for the person receiving care.

Respect for autonomy and dignity: AI systems should empower individuals to make informed decisions about their health, not override or obscure them. This means ensuring transparency in how AI recommendations are generated and used in care delivery.

Privacy and data sovereignty: Personal health data used in AI must be handled with strict safeguards, clear consent processes, and respect for individuals' control over their information. AI systems should not collect or use data beyond what is necessary and ethical.

Accessibility and equity: AI tools must be inclusive, designed to serve diverse populations and reduce health disparities. This includes addressing algorithmic bias and ensuring that AI does not reinforce systemic inequities, especially for historically underserved communities such as Aboriginal and Torres Strait islander peoples.

Engagement and feedback: Patients, carers, and healthcare consumers should be actively involved in the design, evaluation, and governance of all technologies that enable, assist or influence the delivery of health and care, including AI. Their real-life experiences provide critical insights that ensure AI tools are relevant, usable, and aligned with real-world needs.

3. Partnerships and Collaboration

Effective governance in healthcare, especially when technologies are involved, requires strong and transparent collaboration across the entire product lifecycle. This includes healthcare providers, health technology solution providers, policymakers, and the people these systems are designed to serve. No single organisation can manage the complexity and risks of technologies in healthcare alone, especially when it comes to AI. This is a shared responsibility that goes beyond implementation.

Application to AI:

The governance of AI must be built on interdisciplinary and cross-sector partnerships that ensure shared responsibility for safety, quality, and ethical innovation.

Shared responsibility: All stakeholders, including AI developers, system implementers, regulators, and users, must collectively manage the risks and benefits of AI technologies. This means clearly defining roles, responsibilities, and accountability across the AI products and solutions lifecycle.

Interdisciplinary engagement: Governance approaches should include diverse expertise, including clinical, technical, ethical, data science, and consumer/patient perspectives, to ensure comprehensive oversight. AI governance is not just the domain of Chief Information Officers, and it must reach across disciplines and departments.

Transparent relationships: Open communication between partners builds trust and enables early identification and resolution of risks. This is especially critical in the relationship between health technology solution providers and healthcare organisations, where shared understanding of clinical and care context and safety expectations is essential when it comes to the application of AI.

Alignment with broader systems: The governance of AI should integrate with existing frameworks and policies, not operate in isolation. This requires an enterprise view that considers how AI tools interact with clinical workflows, data infrastructure, and regulatory environments.

4. Evidence and Lived Experiences

Evidence, enriched by real-world experiences, is critical. This ensures that the technologies that enable, assist, or influence health and care are not only safe and effective but also relevant and responsive to the needs of those they serve. Lived experience is evidence that can provide insights just as powerful as research.

Application to AI:

In AI development and deployment, evidence must extend beyond technical performance metrics to include clinical relevance, healthcare provider and patient feedback, and contextual understanding.

Evidence-based practice: AI systems should be built and validated using quality clinical data, peer-reviewed research and real-world outcomes. This includes external validation across diverse populations and healthcare settings.

Continuous learning: The governance of AI must support ongoing monitoring, recalibration, and improvement. As AI systems interact with dynamic health and care environments, they must evolve based on new data, emerging risks and lessons learned from both successes and failures.

Incorporation of lived experience: Insights from people, their carers, healthcare providers, and other end users are essential to understanding how AI tools function in practice. These perspectives reveal usability challenges, unintended consequences and opportunities for refinement that data alone cannot capture.

Transparency of evidence: Clear communication about the evidence base supporting AI tools is critical - including their limitations, risks, and performance variability - to build trust and enable informed decision-making by healthcare providers and those they care for.

5. Safety, Quality, and a Learning Health System

Systems safety is a structured, prospective, and proactive approach to the lifecycle of health technologies. Bookended with continuous quality improvement and vigorous incident management, we create a learning health system. In health and care, this approach reduces the potential for technologies to cause avoidable harm in the delivery of health or care and enables the consistent evolution of technologies to meet changing healthcare needs. Health technologies can be as complex as healthcare itself and feel very opaque, with the potential to have large scale impact across health and care systems. If we design it wrong in technology, it will be flawed every time.

Application to AI:

AI technologies must be governed with a lifecycle approach that prioritizes safety by design through deployment and beyond.

Systems safety approach: AI safety depends on the interaction between technology, people, processes and organisational context. Governance, specifically systems safety assurance, must anticipate and prevent unintended consequences by embedding safety into every phase of the AI lifecycle.

Clinical safety management: Risks associated with AI, such as incorrect predictions, automation bias, or data drift, must be continuously identified, assessed, and mitigated. This requires transparent collaboration between health technology solution providers and implementers, with shared responsibility for clinical safety that continues after implementation.

Incident management: Organisations must establish clear procedures and controls for reporting, investigating, and responding to AI-related incidents and near misses. This includes recognising that technical failures and data security breaches can have direct clinical safety impacts.

Quality assurance and monitoring: AI systems must be continuously monitored for performance, fairness, and compliance with clinical and technical standards. Governance of AI should include mechanisms for recalibration, auditing, and updating models as new data and insights emerge.

Continuous improvement: Feedback from users, incident reports, and performance metrics should drive iterative refinement of both AI tools and governance practices. This ensures resilience and adaptability in the face of evolving healthcare needs and technological advancements.

Consequences of ignoring HealthCare+ Governance

Neglecting safety, quality, and continuous improvement in AI systems for health and care can lead to harm, regulatory non-compliance, outdated and biased models, poor adoption, and erosion of trust that ultimately compromising health and care outcomes and experiences, workforce wellbeing, and the reliability of technology-enabled care.

When safety, quality, and continuous improvement are embedded throughout the lifecycle of AI systems, they become trusted, adaptive, and clinically effective tools that enhance care delivery, keep people safe, and protect healthcare workforce well-being. Transparent governance, inclusive design, and proactive risk management build resilience, equity, and meaningful adoption - ensuring technology serves people, not the other way around.

Final Reflections: From Ideas to Practice

As AI continues to evolve and embed itself into the fabric of health and care, the urgency to govern it responsibly grows stronger. We should not be passive observers, but instead actively participate in a transformative experiment that touches every aspect of care delivery, patient experience and system design. The five considerations outlined in this paper are not theoretical ideas, but practical tools for shaping a future where technology enhances, rather than compromises, the values at the heart of health and care.

Embedding these approaches into everyday practice requires more than policy, and it demands leadership, collaboration, and a shared commitment to doing what is right, not just what is possible. It means listening to lived experiences, learning from failures and successes, and continuously improving how we design, deploy, and oversee AI technologies. It means recognising that trust is earned through transparency, accountability, and inclusion.

Ultimately, the success of AI in health and care will not be measured solely by its technical sophistication but by its ability to deliver safe, equitable, and meaningful health and care outcomes and experiences for all. By aligning with a unified governance approach, like a HealthCare+ Governance Framework, that embraces emerging technologies, we can move forward with confidence, knowing that innovation is guided by care and that every stakeholder has a role in shaping a future in health and care, with technologies that are safe and that we can trust.

Authors:

Dr. Catherine Lopes, AI and Data, Strategy and Advisory, APAC

Liz Keen (RN), Head of HealthCare & Governance – Innovation and Transformation

Infosys Topaz™ is an AI-first set of services, solutions and platforms using generative AI technologies. It amplifies the potential of humans, enterprises, and communities to create value. With 12,000+ AI assets, 150+ pre-trained AI models, 10+ AI platforms steered by AI-first specialists and data strategists, and a 'responsible by design' approach, Infosys Topaz™ helps enterprises accelerate growth, unlock efficiencies at scale and build connected ecosystems. Connect with us at infosystopaz@infosys.com.

For more information, contact askus@infosys.com

Infosys®
Navigate your next