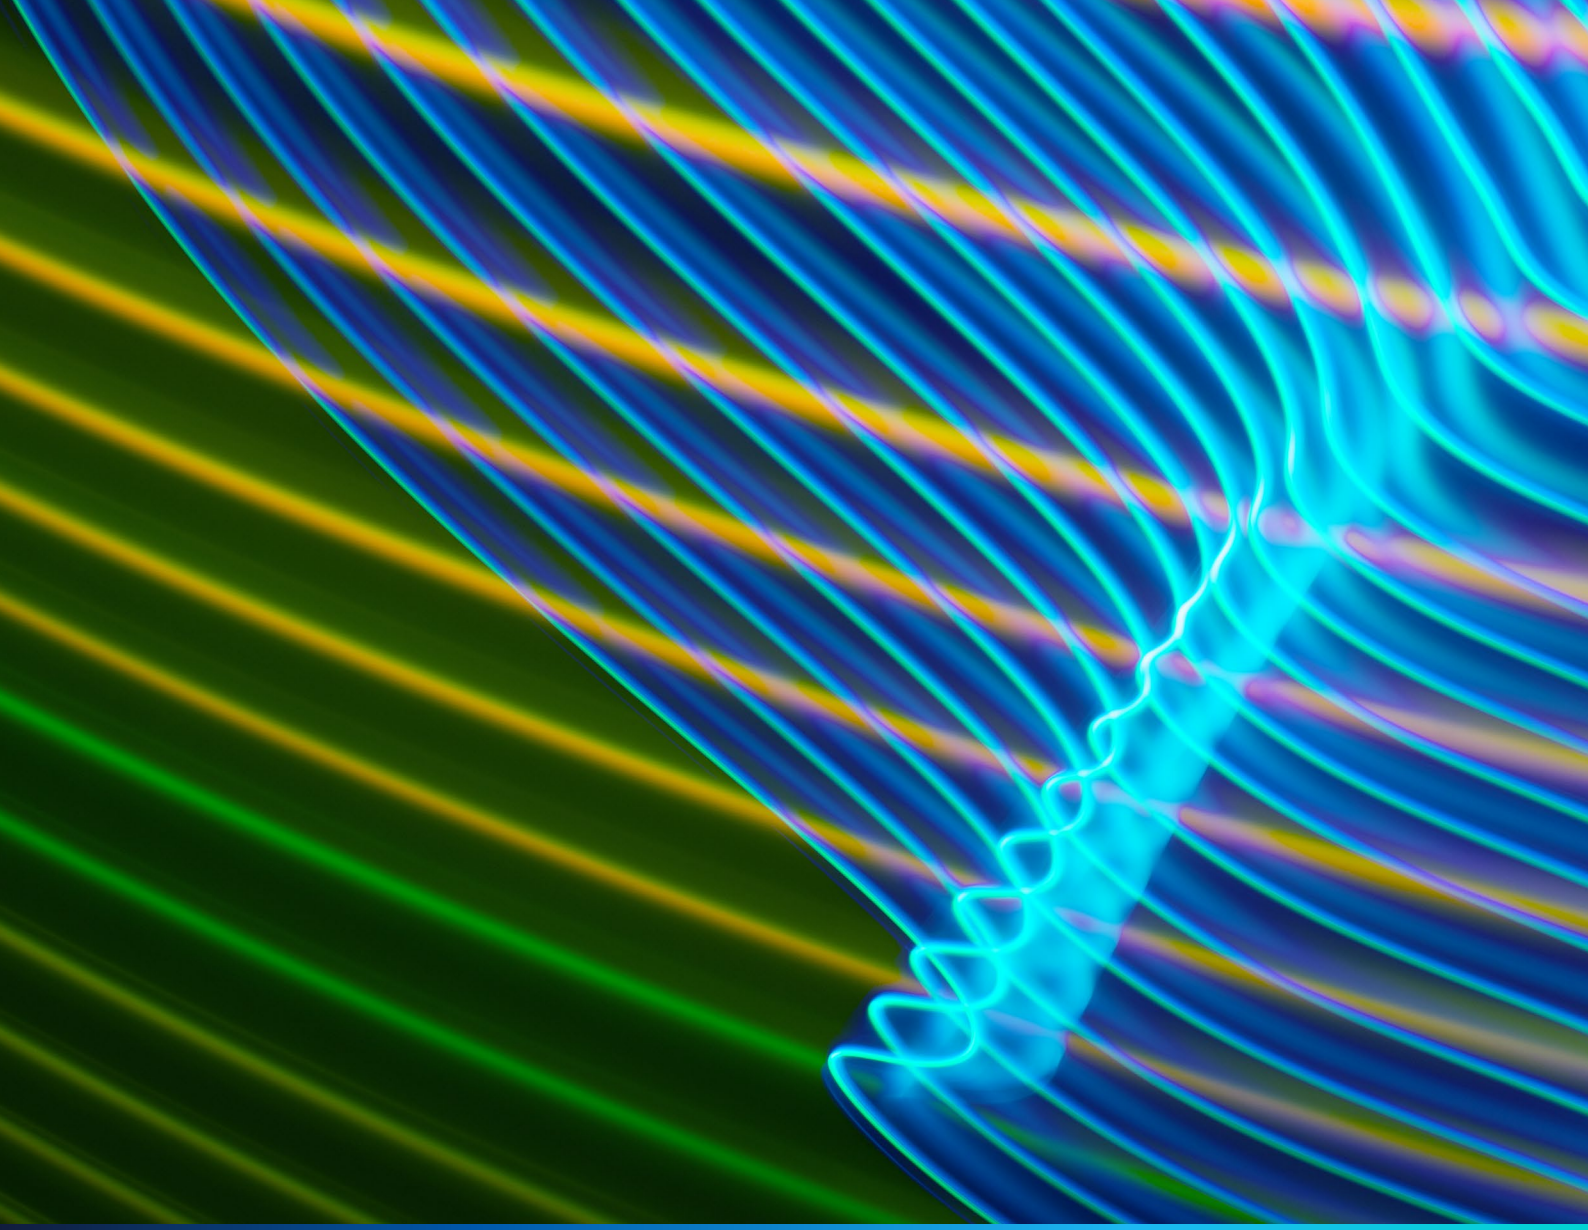

Advancing Data & AI Innovation in *Connected Healthcare*

Capgemini

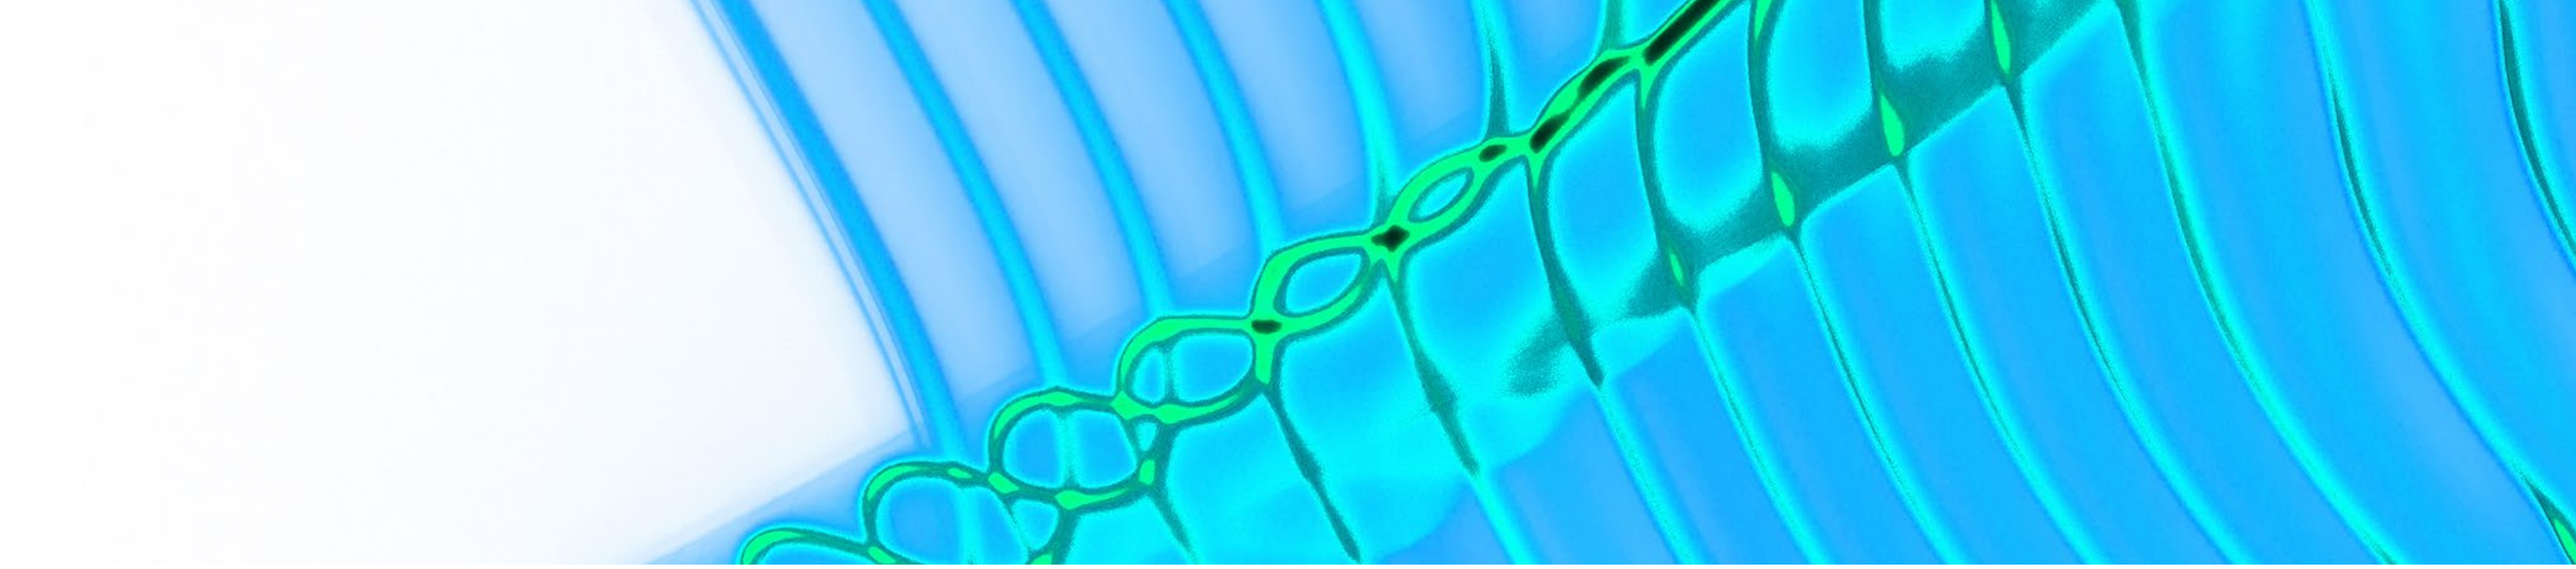

Executive Summary

Healthcare systems across Europe are under unprecedented pressure: rising costs, growing workforce shortages, and increasing expectations for better patient outcomes are forcing hospitals, insurers, and MedTech companies to rethink how care is delivered. In response, the MedTech industry has introduced a wave of innovation, from advanced remote monitoring devices to digital health platforms and cloud enabled diagnostics, creating the foundation for connected, data driven care. Yet, despite this momentum, the full value of connected health remains largely unrealized. Technical fragmentation, regulatory complexity, data silos, and the absence of coordinated operational models continue to slow adoption.

No single stakeholder can address these barriers alone. Hospitals often lack the resources or incentives to develop analytics and AI capabilities in house and increasingly rely on interoperable, governance enabled platforms. Insurers are exploring ways to adopt and operationalize disease specific algorithms but require guidance to ensure trustworthy, explainable, and compliant deployment. MedTech companies, meanwhile, must evolve beyond device centric B2B2C (business-to-business-to-consumer) approaches toward integrated service models that connect devices, data, and clinical intelligence. Regulators contribute by defining the guardrails for safe and ethical AI through policies such as the EU AI Act and the European Health Data Space (EHDS). However, the operationalization of these frameworks remains a work in progress.

These challenges, while significant, are solvable. This Point of View outlines a practical approach to overcoming them by focusing on three pillars:

1. **Interoperable, high quality data foundations** that connect devices, clinical workflows, and operational systems across the healthcare ecosystem.

2. **Trustworthy, governed AI** that is explainable, bias controlled, continuously monitored, and aligned with the needs of clinicians and care organizations.

3. **Cross ecosystem collaboration models** that bring together medical experts, data scientists, engineers, and regulatory stakeholders to design and scale connected care solutions.

Early steps, which includes the EU AI Act, EHDS, and first adopters of federated architectures, which demonstrate that the shift is underway. But scaling connected health requires a coordinated effort to align incentives, share capabilities, and adopt technologies that can bridge long standing gaps between devices, data, and clinical decision making.

Capgemini's role is to help orchestrate this transformation: by designing interoperable data platforms, delivering regulatory aligned AI, and enabling the connected care ecosystem to collaborate with confidence. With the right architecture, governance, and cross functional expertise in place, connected health can move from aspiration to reality and unlocking better outcomes for patients, more sustainable operations for providers, and a clearer path to value for MedTech and insurers alike.

Contents

Who should read this report and why?

This report is intended for senior leaders across the healthcare ecosystem who are shaping the next wave of AI-enabled connected health. It provides evidence-based insights, practical considerations, and strategic recommendations for organizations seeking to unlock value while navigating regulatory and technical complexity.

It is particularly relevant for:

- Hospital executives and clinical operations leaders looking to improve patient outcomes, streamline care pathways, and scale digital innovation responsibly.
- Payer and insurance decision-makers evaluating how AI-driven connected health can strengthen prevention, enhance member engagement, and optimize cost and risk structures.
- Public-sector leaders and policymakers responsible for funding, regulating, and enabling digital health transformation at system level.
- Chief Data Officers, Chief Information Officers, and digital health strategists accountable for governance, interoperability, and the trustworthy deployment of AI at scale.
- Contract Research Organizations (CROs) seeking to evolve their role in decentralized trials, device enabled data collection, and AI supported evidence generation.
- CEOs and founders of small, privately owned MedTech companies navigating the shift from device centric models to integrated, data driven, service oriented solutions.

For these stakeholders, this report offers a clear view of where value can be realized today, what barriers must be addressed to move forward towards how organizations can act decisively to accelerate adoption while ensuring compliance, trust, and safety.

Disclaimer – Terminology used in this report

To ensure clarity throughout this report, we distinguish explicitly between three closely related domains in the health ecosystem:

Life Sciences
Refers to pharmaceuticals, biotechnology, and biomedical innovation, including the research, development, and commercialization of therapeutics, biologics, and scientific technologies.

Healthcare
Describes the broader aggregation of all sectors that provide goods and services aimed at maintaining, improving, or restoring human health, spanning hospitals, payers, public health organizations, service providers, and care delivery systems.

MedTech
Encompasses medical devices, diagnostic systems, and digital health technologies that generate, process, or act upon clinical data.

01

Introduction

Healthcare is undergoing a profound transformation. As hospitals and health systems face demanding pressure to reduce costs, improve patient outcomes, and address critical workforce shortages, the promise of connected, data-driven, and remotely managed care has never been more compelling. In this environment, the promise of connected health including continuous patient monitoring, early detection of health deterioration , and more proactive intervention pathways, is increasingly persuasive. The ability to monitor patients beyond hospital walls, particularly in the critical days following surgery, offers a path to earlier discharges, optimized resource utilization, improved health and enhanced patient satisfaction.

MedTech companies have responded with a wave of innovation, developing advanced remote monitoring devices, digital platforms and cloud-enabled diagnostic tools, that enable continuous patient data collection and now generate vast streams of physiological and behavioral data and disease-related data.

Yet despite this technological progress, the industry faces a structural bottleneck: data without insight does not improve care. The challenge is no longer access to data, but the ability to transform heterogeneous, incomplete, and often siloed information into clinically meaningful intelligence.

This transformation depends on close collaboration across the healthcare ecosystem, providers, MedTech organizations, insurers, and regulators, each of whom contributes essential context, data, or domain knowledge. While MedTech companies innovate rapidly, they do not always have access to operational or longitudinal clinical data. Hospitals, meanwhile, face regulatory, technical, and governance constraints that limit data sharing and interoperability. As a result, the organizations generating the most clinically relevant data are not always the ones building the AI models that rely on it.

However, a significant bottleneck remains: the ability of hospitals and healthcare providers to harness this data effectively. The challenge lies not in the availability of technology, but in the capacity to transform raw data into actionable insights. This requires seamless collaboration between medical professionals, who understand the clinical context, AI experts and data scientists, who can build robust, explainable algorithms and models. Addressing this bottleneck requires not only improved access to integrated clinical data but also the use of trustworthy AI, which are AI systems that are explainable, reliable, bias-controlled, and aligned with clinical and regulatory expectations. Without these safeguards, even advanced models cannot be adopted with confidence by clinicians or integrated into critical care processes.

Unlocking the value of connected health requires a shift toward ecosystem-wide collaboration. While hospitals and insurers generate much of the relevant health data, most organizations lack the capacity or intention to build AI and analytics themselves and rely instead on interoperable, governance-enabled platforms.

To make connected care work at scale, insurers, providers, MedTech companies and regulators must collaborate far more closely, aligning on shared data infrastructures, selecting the right devices, adopting disease-specific algorithms, and ensuring trustworthy, explainable AI across all touchpoints. Only through this connected effort can AI become a true clinical capability rather than an isolated feature attached to a device.

01

To make this possible, several foundational technical challenges must be addressed:

- Clean, high-quality and structured data to ensure reliable model performance
- Interoperability across devices, (Electronic Health Records EHRs), and monitoring systems to break clinical silos
- Secure architectures that protect sensitive patient information
- Explainable models that clinicians can understand, trust, and act upon
- Traceable, auditable AI outputs aligned with regulatory requirements (EU AI Act, EHDS, Medical Devices Regulation (MDR), GxP)

Connected health is not only about transmitting signals from devices; it is about bringing data together in a governed, interoperable architecture that enables real-time analytics and on-the-fly decision support. To unlock this value, healthcare providers need an analytics engine capable of identifying clinical patterns, supporting assisted diagnosis, and enabling proactive interventions, which are powered by AI models that are continuously monitored, bias-controlled, and adapted to local population needs.

This paper introduces the technical and organizational challenges that stand in the way of realizing the full potential of connected healthcare to help the healthcare ecosystem move beyond device innovation and toward hospital-led, data-driven care. We outline how interoperable data platforms, trustworthy AI, and cross-industry engineering practices can enable hospitals to derive actionable insight from connected data, while enabling MedTech companies, insurers, and public health institutions to collaborate around a shared goal: sustainable, patient-centric, intelligent healthcare delivery.

For AI in healthcare to move from pilot to practice, explainability, trust, and engineering-led delivery must come together. But none of this is possible without a resilient data foundation. Before advanced AI models can support clinical decisions, the underlying clinical data must be clean, connected, interoperable, and governed. The next section examines this critical prerequisite in detail, how primary clinical data can be transformed into reliable secondary insights through high-quality, interoperable data infrastructure, and why this forms the backbone of any scalable AI-driven healthcare solution.

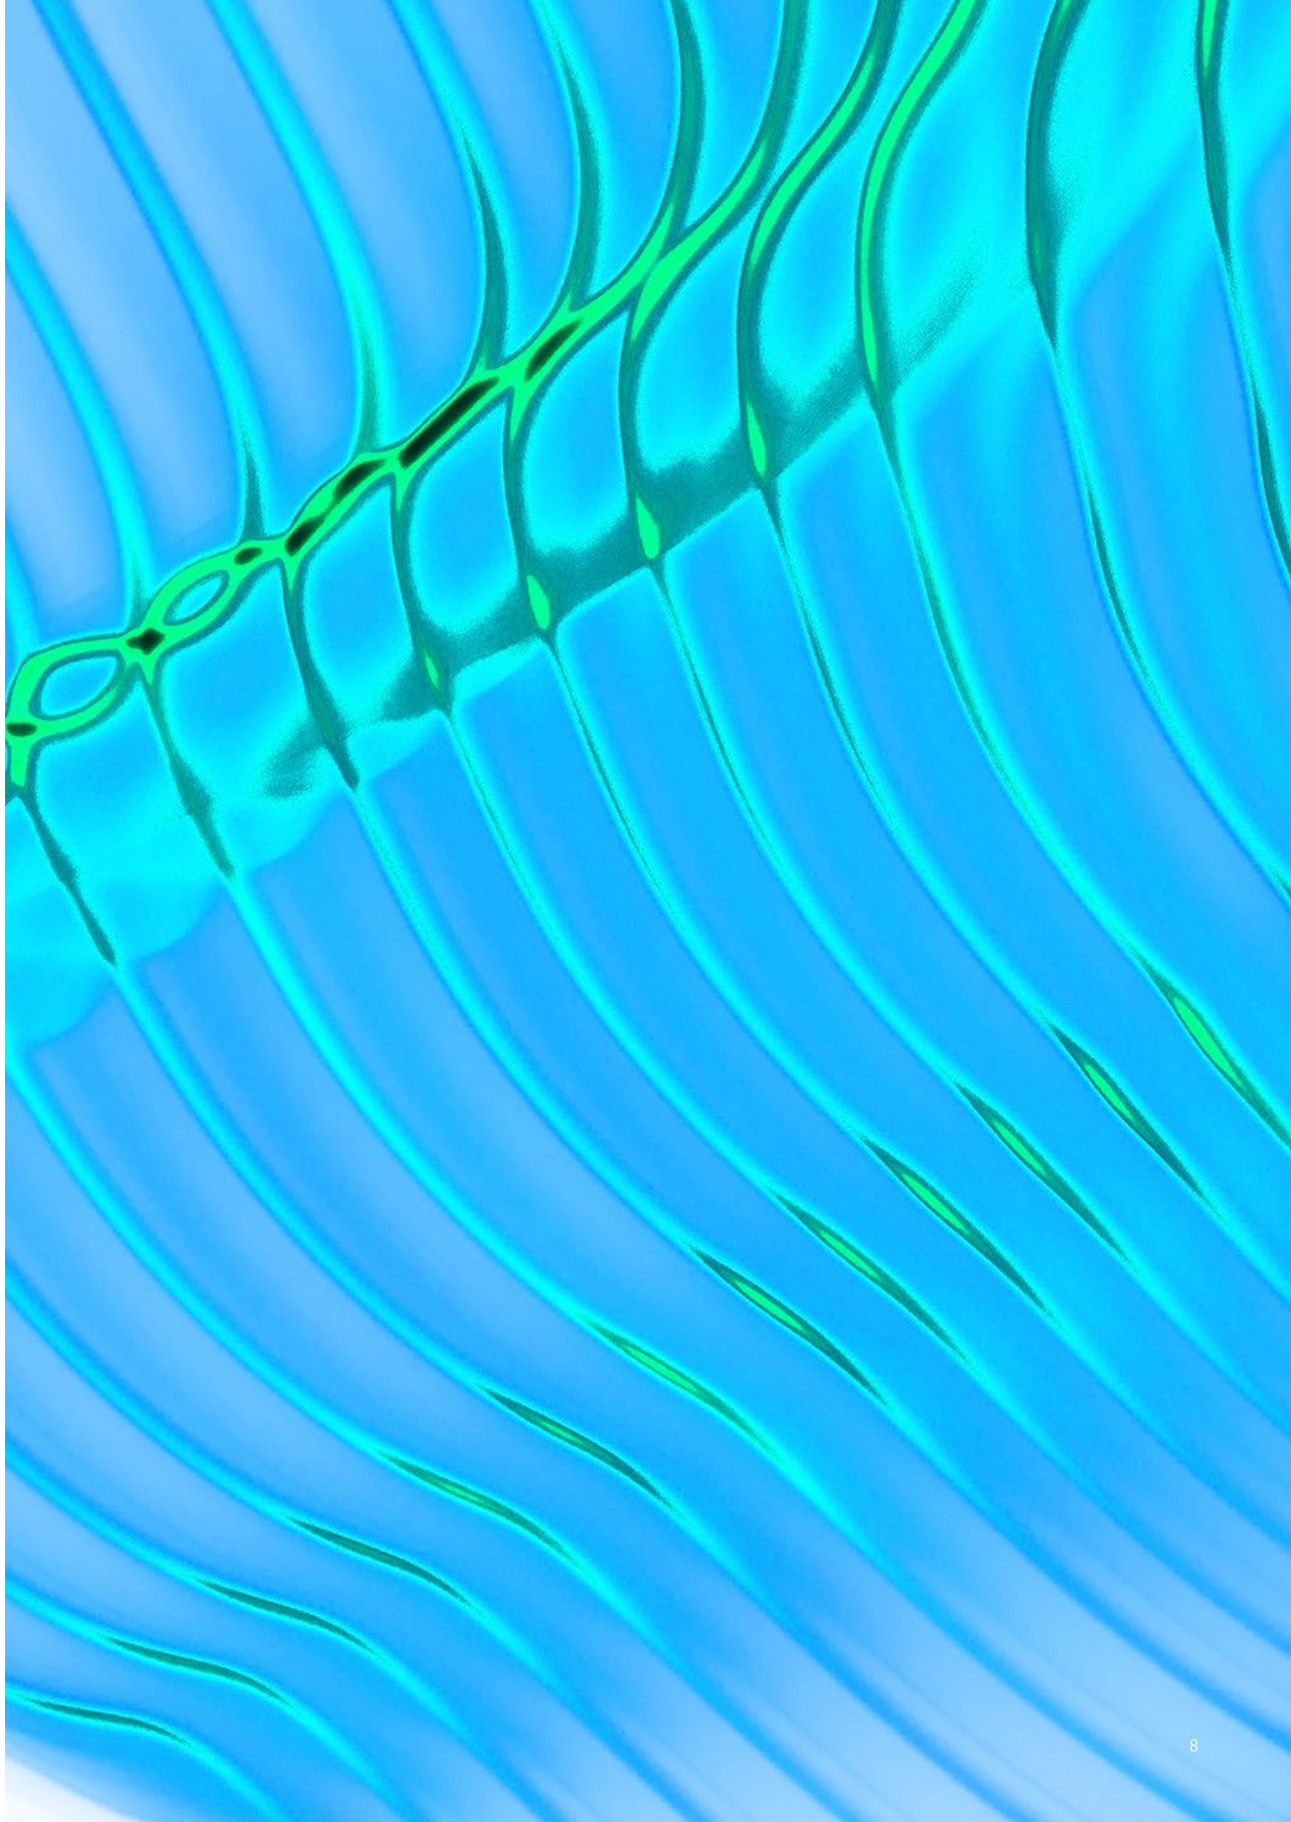

02

The data interoperability challenge

Hospitals and medical technology providers face mounting pressure to deliver integrated, data-driven care. However, interoperability barriers and data silos continue to undermine clinical and operational efficiency across European healthcare systems.

The Operational Reality

Disconnected systems: EHRs , Laboratory Information Management Systems (LIMS), medical devices, imaging platforms, and patient monitoring tools are often built on legacy infrastructure and lack interoperability, causing them to operate independently and preventing real-time data access at the point of care. Clinicians waste valuable time navigating multiple interfaces while critical patient information remains scattered across platforms.

1. **Clinical inefficiency:** Fragmented patient data forces clinicians to manually aggregate information across disconnected systems which leads to preventable harms like duplicate testing, delayed diagnoses, and a higher risk of patient safety incidents.

2. **AI deployment barriers:** Machine learning algorithms and clinical decision support tools depend on comprehensive, real-time data integration. Fragmented data architecture prevents AI from scaling beyond isolated use cases, keeping hospitals stuck in Proof-of-Concept (PoC) stage where promising pilots never achieve enterprise deployment.

Integration Complexity

1. **Legacy infrastructure:** Non-standard data formats and aging systems demand **bespoke interfaces** for each integration, creating unsustainable technical debt and preventing **scalable interoperability** across hospital networks.

2. **GxP validation constraints:** GxP-validated systems (certified EHRs, LIMS, Clinical Trial Management Systems (CTMS), quality management platforms) must maintain strict compliance with ALCOA+ (Attributable, Legible, Contemporaneous, Original, Accurate, Complete, Consistent, Enduring, and Available) data integrity principles. Any data extraction or analytical use requires extensive validation documentation and governance oversight, significantly increasing integration timelines and resource requirements compared to non-regulated systems.

3. **Semantic heterogeneity:** Inconsistent data standards, proprietary APIs, and varied terminology across vendors create **integration friction**, delaying deployment and increasing costs.

Business Impact

These technical challenges translate directly to operational consequences:

- **Increased operational costs** from manual workflows and system redundancy.
- **Extended implementation timelines** for new technologies and AI tools.
- **Limited scalability** of digital health initiatives across hospital groups.

These challenges highlight the need for a structured, scalable approach to data readiness. The following section outlines how Capgemini addresses this through proven solutions in the pharmaceutical, BioTech and MedTech sector, solutions that are directly applicable to healthcare.

2.1. Capgemini’s approach to build interoperable data ecosystems

Access alone is insufficient; data must be analytics ready. The 2025 Healthcare Data Quality Report[1] found that “limited time and resources” were cited as the primary barrier by 70% of respondents, rather than a specific measure of time spent by data scientists on data cleaning. This highlights the persistent challenge of converting fragmented, inconsistent data into structured, usable formats.

Capgemini has addressed this challenge at scale in the healthcare, pharmaceutical and life sciences sectors, delivering enterprise-grade data platforms that transform legacy, siloed data into trusted, interoperable assets. Our approach is to build on a modular, sovereign cloud -native architecture that integrates ingestion, transformation, governance, and analytics enablement. Our solutioning begins with establishing a strong data foundation. We identify and connect to diverse data sources-ranging from ERP, CRM, and R&D systems to manufacturing and supply chain platforms. We implement reusable ingestion patterns (e.g. delta, snapshot, streaming) and apply standardised transformations to align data with common models. This ensures consistency and prepares data for downstream use. We then embed data trust capabilities directly into the platform. This includes data quality checks, lineage tracking, metadata cataloguing, and governance aligned with regulatory frameworks such as GxP.

Our teams define Architecture and Governance Decision Records (ADRs/GDRs), deploy foundational components (storage, compute, orchestration), and implement automated validation and quality dashboards. These capabilities ensure that data is not only available but also reliable, and compliant.

Our success is evidenced by large-scale programs in the pharmaceutical and MedTech sector. For a global pharmaceutical and animal health company, we served as a strategic partner in delivering a Global Data Platform Program. Over four years, our multidisciplinary team, comprising data architects, engineers, GxP experts, and data scientists, migrated more than 20 use cases across clinical, manufacturing, and supply chain domains onto a scalable, sovereign cloud-native platform built on state-of-the-art systems. The solution included modular ingestion pipelines, standardized (cleaned and curated) data, integrated data quality and lineage controls, and automated orchestration workflows. This enabled significant performance improvements. This strategic partnership enabled the client to take the critical decision in migrating on-premises platforms to cloud infrastructure, achieving scalability and eliminating performance bottlenecks.

In another engagement, we delivered a sovereign cloud native enterprise data platform for a leading biotechnology company to support end to end CSRD (Corporate Sustainability Reporting Directive) reporting. The solution consolidates ESG (Environmental, Social, and Governance) data from across the organization into a single, reliable foundation, ensuring consistency, traceability, and audit readiness. Curated, domain specific data across Environmental, Social, and Procurement areas improves the accuracy of regulatory reporting and enhances internal transparency. Once loaded into internal databases and reporting tools, this curated data enables analysis and visualization, while integration with a third party CSRD solution supports automated compliance reporting from data collection through to insight generation.

What makes our approach successful is our ability to deliver end-to-end, from architecture and ingestion to governance and insight activation. Our teams combine deep domain knowledge with technical expertise in cloud platforms, data engineering, and compliance. We leverage accelerators, reusable templates, and a structured working plan to reduce implementation time and ensure consistency. While these solutions were delivered in pharma, the underlying challenges-data fragmentation, quality assurance, regulatory compliance, and the need for real-time analytics-are equally relevant in healthcare provider settings. The architectural patterns, governance models, and automation frameworks we’ve implemented are directly transferable to hospitals, health systems, and public health agencies.

At DMEA 2026 in Berlin, we showcased a production-grade demonstrator illustrating how fragmented health data can be converted into scalable, AI-ready intelligence. Built on open interoperability standards such as HL7 FHIR and designed with sovereign-by-design and patient-controlled data governance at its core, the platform unifies continuous wearable signals, longitudinal clinical records, and contextual health information into a single, validated digital patient representation. More than a concept, our demonstrator operates in real time, showcasing a standardized integration layer that makes health data immediately consumable by advanced analytics and GenAI use cases. It enables healthcare stakeholders to experience how interoperable, cloud-native architectures can power preventive care, accelerate clinical insight, and unlock cross-ecosystem collaboration. By turning data silos into a connected intelligence fabric, it demonstrates Capgemini’s capability to engineer the foundations of future-proof, AI-driven healthcare systems where smarter, trusted data directly translates into better outcomes.

In summary, Capgemini brings a proven, industrialised approach to solving the data quality and transformation challenge. Our success in pharma demonstrates our ability to build trusted, scalable data ecosystems capabilities that are essential for healthcare organisations seeking to unlock the full potential of AI, improve patient outcomes, and meet evolving regulatory demands. Building on this foundation, it becomes clear that data integration alone is not enough. The next frontier lies in operationalising AI at scale.

[1] Clinical Architecture, 2025 Healthcare Data Quality Report, June 2025, https://clinicalarchitecture.com/2025-data-quality-report/.

03

Why integration isn’t enough to scale AI in healthcare

AI adoption in healthcare has reached a critical juncture: the transition from pilots to practice.

While adoption rates are rising with 71% of U.S. hospitals using predictive AI integrated with electronic health records as of 2024, and approximately 50% projected to adopt generative AI by end of 2025[2] and Europe’s AI healthcare market projected to reach USD 2.6 billion by 2030 with a compound annual growth rate of 35.9%[3] , a substantial implementation gap persists. 42% of enterprises run AI in production and 40% are piloting, yet 88% of PoCs fail, 95% due to data issues, and only 20% achieve ROI despite 77% exploring AI[4]. This gap exists despite mounting systemic pressures: the WHO projects healthcare workforce shortage of 4.1 million workers in Europe by 2030 (varying by source and scope)[5] , while regulatory frameworks including the EU AI Act (which entered into force on 1 August 2024)[6] and the EHDS regulation (which entered into force on 26 March 2025)[7] create both compliance obligations and interoperability opportunities.

Hospitals have built robust data infrastructures, but integration alone isn’t enough. The following sub sections explore why AI fails without addressing four critical challenges: governance, bias, drift, data sharing, and explainability.

[2] HealthIT.gov, “Hospital Trends in the Use, Evaluation, and Governance of Predictive AI, 2023-2024” (2024); Becker’s Hospital Review, “Half of US hospitals to adopt generative AI by end of 2025, study finds” (December 2025)

[3] Grand View Research, “Europe AI In Healthcare Market Size & Outlook, 2030,” Horizon Databook (2024),

[4] Available at: https://www.improving.com/thoughts/ai-strategy-and-roadmap-assessment/ (Accessed: 20 February 2026)

[5] European Parliament (2025). Healthcare sector: addressing labour shortages and working conditions. Available at: www.europarl.europa.eu (Accessed: 20 February 2026).

[6] European Commission (2024) AI Act | Shaping Europe’s digital future. Available at: https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai (Accessed: 20 February 2026)

[7] European Commission (2024) European Health Data Space. (Accessed: 27 February 2026).

3.1. Governance landscape: FDA, EU AI Act and EHDS

Regulatory frameworks are undergoing significant changes as AI systems increasingly play key roles in safety-critical healthcare tasks, ranging from clinical decision support to automated diagnostics. The shift is clear: regulators now mandate continuous monitoring, bias mitigation, and transparent explainability as core requirements, not optional best practices. Two pivotal frameworks are driving this change: the FDA's (Food and Drug Administration) Predetermined Change Control Plan (PCCP)[8] and the EU's regulatory frameworks including the EU AI Act and EHDS.

FDA PCCP:

The FDA's PCCP framework, finalized in August 2025, enables manufacturers of AI-enabled devices to describe planned modifications, the methodology to develop and validate those changes, and an assessment of their impact within a single marketing submission. The FDA reviews the PCCP as part of the initial 510(k), De Novo, or PMA submission, then allows manufacturers to implement each described modification without additional marketing submissions. This approach supports iterative improvement of AI-enabled devices while maintaining reasonable assurance of safety and effectiveness, fundamentally transforming the regulatory pathway from static approvals to adaptive oversight. PCCP doesn't eliminate regulatory burden; it shifts it from episodic submissions to continuous, automated monitoring.

The EU AI Act:

The EU AI Act, which entered into force in August 2024, establishes legally binding requirements for high-risk AI systems including medical devices. Article 10 explicitly mandates data governance practices "with a view of possible biases" that could affect health, safety, or fundamental rights. This makes the bias detection and mitigation a legal obligation, not a recommendation. Article 14 requires human oversight through transparent, interpretable outputs that enable meaningful human supervision, directly mandating explainability.

EHDS:

Entered into force in March 2025 with main provisions applying key parts by March 2029. The EHDS Regulation aims to establish a common framework for the use and exchange of electronic health data across the EU. It creates dual frameworks: primary use enabling patients to access and share electronic health records across borders for healthcare delivery, and secondary use providing legal pathways for researchers and companies to access anonymized health data for research, innovation, policy-making, and regulatory activities. It mandates interoperable health data exchange across EU member states through standardized formats and infrastructure, creating the technical foundation that makes AI Act compliance achievable at scale. EHDS facilitates secure, federated access to a wide range of health data from different countries without aggregating patient records. This enables organizations to validate AI models across various European populations, helping detect bias in demographic groups, monitor drift at multiple sites, and collaboratively improve models and all while staying compliant with GDPR regulations. This infrastructure transforms AI Act's fairness and monitoring mandates from theoretical requirements into operationally feasible capabilities.

The regulations make it mandatory to demonstrate that your AI is fair for all groups and adapts over time. Although compliance requirements may seem intricate, organizations that commit to comprehensive bias detection, ongoing monitoring, and the implementation of explainable AI will be well-positioned to foster trust among regulators, clinicians, and patients. However, meeting these obligations first requires a clear understanding of where bias originates, how it manifests across clinical settings, and what practical strategies exist to address it. The following section explains the bias problem and outlines the path forward.

[8] U.S. Food and Drug Administration. (2025, August 18). Marketing submission recommendations for a predetermined change control plan for artificial intelligence-enabled device software functions: Guidance for industry and Food and Drug Administration staff. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence

3.2. The bias problem and path forward

The challenge:

As AI systems scale across clinical settings, algorithmic bias has emerged as a critical patient safety and compliance risk. Bias in healthcare AI refers to systematic errors that lead to unequal treatment across demographic groups. These disparities often stem from historical inequities embedded in training data, underrepresentation of certain populations, or flawed proxy variables such as healthcare costs or race-adjusted metrics. Bias manifests in multiple forms:

1. **Historical bias:** Reflects entrenched societal disparities. For example, the widely used 'Get With The Guidelines-Heart Failure (GWTG-HF) Risk Score'[9] historically assigned higher risk points to "non-Black" patients, leading to under-treatment of Black and Latinx individuals in emergency departments.
2. **Representation bias:** Arises when certain groups are missing or underrepresented in training data. A diabetic retinopathy model trained predominantly on light-skinned patients, for instance, performs poorly on darker skin tones.
3. **Automation and feedback loop bias:** Occurs when clinicians over-rely on AI outputs (automation bias) or when biased predictions are used to train future models, reinforcing disparities.

The solution:

To address these risks, frameworks like (Justice, Equity, Fairness, and Anti-bias) **JustEFAB**[10] and **EquiLense**[11] ensure both ethical and individual fairness. By applying these Fairness Matrices, hospitals can verify that screening opportunities are equally distributed and that the model is not learning spurious correlations based on patient metadata.

1. **JustEFAB Framework:** This framework distinguishes between "fair" biological heterogeneity (clinical need) and "unfair" performance differences caused by unequal access to care. For example, it identified that the Sybil lung cancer model's AUROC gap (0.88 for women vs. 0.81 for men) was an unfair bias that clinical confounders could not explain.
2. **EquiLense:** A post-hoc, model-agnostic auditing tool that ensures individual fairness by ensuring "like cases are treated alike." It utilizes clinical similarity matching and the Mean Predicted Probability Difference (MPPD) metric to quantify inconsistencies between similar patients across demographic groups.

Besides compliance (discussed in section 3.1), fairness is essential for trust and effective care. Hospitals using biased algorithms may risk patient outcomes and lose clinician support, resulting in wasted resources. Organizations that develop fair and transparent AI not only fulfill legal and ethical responsibilities but also enhance patient care and protect public trust in digital health innovation. However, building trust is not a one-time achievement. Even the most transparent and well-calibrated AI systems can degrade over time if they are not actively maintained. This brings us to a critical operational challenge in clinical AI deployment: ensuring sustained performance in the face of changing data environments.

[9] Vyas DA, Eisenstein LG, Jones DS. Hidden in Plain Sight - Reconsidering the Use of Race Correction in Clinical Algorithms. N Engl J Med. 2020 Aug 27;383(9):874-882. doi: 10.1056/NEJMms2004740. Epub 2020 Jun 17. PMID: 32853499.

[10] Mccradden, M. D., et al. (2023). What's fair is… fair? Presenting JustEFAB, an ethical framework for operationalizing medical ethics and social justice in the integration of clinical machine learning: JustEFAB. In Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency (FAccT '23), 1505–1519. https://doi.org/10.1145/3593013.3594096.

[11] Xu, J., & Strohmer, T. (2025). Integrating Group and Individual Fairness in Clinical AI: A Post-Hoc, Model-Agnostic Framework for Fairness Auditing. medRxiv. doi.org.

3.3. Combating AI aging

The challenge:

Clinical AI systems are not "set and forget" tools; they are susceptible to performance (temporal drift) and across different settings (site-based or spatial drift), often referred to as **"AI aging"**.

For instance, in 2020[12], a high-performing emergency admission model began to deteriorate. Originally trained on data from before the pandemic, it was an expert at predicting caseloads, until the world changed. As COVID-19 altered patient demographics and care patterns, the model's assumptions were no longer grounded in reality. The "mix" of patients shifted away from mild cases toward acute ones, causing the model's accuracy (AUROC) to decline from 0.856 to 0.826. This was not a coding error, but a textbook case of **Temporal Drift,** where a model's effectiveness reduces substantially as the data landscape evolves over time. This aging process is driven by the volatile nature of healthcare and is frequently fuelled by operational changes such as how clinicians document data or demographic shifts in patient populations. It leads to **Covariate/ Data Drift**, where input distributions change (e.g., the transition from ICD-9 to ICD-10). However, if the relationship between symptoms and outcomes changes it is called as **Concept Drift.** The challenge is magnified when moving models between different set-ups, which means that a model trained at one hospital can suffer from **Site-Based Drift** immediately upon deployment at a new site with different lab protocols.

The solution:

To address this, hospitals are adopting **continuous monitoring** strategies. For instance, using **Information Geometric Temporal (IGT) projections[13]**, clinical leads can visualize shifts in data distribution like a "weather map," flagging variations before they impact patients. It serves as an **early indicator** of performance variations (covariate and concept shifts) that directly correlate with the decline of models.

However, model adaptation can't overcome the limitations of fragmented source data. The next challenge is enabling secure, collaborative learning across institutional boundaries.

3.4. Secure multi-institutional AI through federated learning

The challenge:

Healthcare data is often described as the "new oil" for AI, yet it remains locked in silos due to privacy laws, institutional policies, and competitive concerns. No single hospital, not even a national health system, holds enough data to train robust AI models. This is especially true for rare diseases or diverse populations. For example, a paediatric oncologist may see only a dozen cases of a rare childhood cancer each year-far too few to train a reliable model. While thousands of such cases exist globally, data protection regulations such as GDPR and HIPAA (Health Insurance Portability and Accountability Act) prevent centralised pooling. This fragmentation limits progress. AI models trained on isolated datasets often fail to generalize.

The solution:

A promising solution is Federated Learning (FL), a privacy-preserving approach that enables collaborative model training without sharing raw data. In FL, an initial model is distributed to participating hospitals, where it is trained locally on each institution's data. Only the updated model parameters without patient data are sent back to a central server, which aggregates them to refine the global model. This process repeats over several rounds, allowing the model to learn from distributed data while maintaining compliance with privacy regulations. FL is now moving from research into real-world deployment. Platforms such as NVIDIA FLARE[14] and open-source tools like Flower[15] are simplifying implementation. In Europe, the EHDS is embracing FL to enable cross-border health research while preserving data sovereignty. Challenges remain. Hospitals differ in data formats, coding standards, and patient demographics. Algorithms like FedProx[16] address these issues by weighting updates appropriately.

In summary, FL is dissolving the data-sharing deadlock. It empowers hospitals to collaboratively train AI models for rare diseases, pandemics, and other complex conditions without compromising privacy. This innovation is turning regulatory constraints into a catalyst for secure, scalable AI in healthcare. The next challenge is to make AI decisions understandable and actionable at the point of patient care.

3.5. Breaking the AI trust barrier through explainability

The challenge:

For instance, a nurse receives an alert: "Patient 402: 91% risk of acute kidney injury." She checks and the patient is resting comfortably, with vitals within normal ranges. The system offers no explanation, just a number. Is this a genuine early warning, or another false alarm driven by noisy data? She might dismiss the alert as it has no rationale to guide her. Hours later, the patient's condition deteriorates. It is difficult for clinicians to act on AI-generated predictions without understanding the reasoning behind them. When an algorithm declares "91% risk" without context, it forces clinicians into an impossible position to either trust a black-box system that may be detecting spurious correlations or ignore a potentially life-saving signal.

The core issue is explainability not accuracy. Most AI models are built to identify patterns, not causes. They can flag statistical associations but cannot answer the questions clinicians care about most: "Why is this happening?" and "What will happen if we intervene?" AI requires causal reasoning and explainability to be reliable and trusted in clinical settings; prediction alone is insufficient. It must clarify its outputs to support clinical decisions and earn a role in patient care.

The solution:

To bridge this gap, AI must evolve into a transparent partner using three core methodological approaches:

- **Feature Attribution and Local Interpretability:** Instead of a raw score, frameworks like SHAP and LIME decompose predictions into specific clinical drivers. For example, by showing that a creatinine spike or a drop in urine output triggered the alert, the AI provides a rationale the nurse can immediately validate against the patient's physical state, reducing "alarm fatigue."
- **Visual Grounding and Generative Counterfactuals:** For imaging, tools like Grad-CAM highlight pathological regions with heatmaps. More advanced Diffusion Models take this further by generating "synthetic healthy" versions of a patient's scan. Visualizing the difference between the pathology and a healthy baseline allows clinicians to intuitively "see" the reasoning behind a diagnosis.
- **Uncertainty Calibration and Causal Reasoning:** Reliable systems must quantify their own limits. Using Bayesian Uncertainty, the system avoids making "forced guesses" on poor-quality data, instead flagging images or vitals as "ungradable". Furthermore, by moving from simple patterns to Causal Inference, the AI becomes a treatment simulator. It doesn't just flag a risk; it helps clinicians visualize the future by predicting how specific interventions, such as a dosage change, will statistically improve the patient's recovery path.

In summary, explainability is the catalyst that transforms a statistical prediction into a clinical decision. Without these transparent frameworks, even the most accurate model remains a liability in a high-stakes environment where "because the computer said so" is an unacceptable medical rationale. By shifting the focus from raw accuracy to causal transparency, healthcare providers can finally overcome the 80% failure rate of AI implementations. Moving from a "black-box" oracle to an interpretable assistant is more than a technical advancement; it is essential for earning clinical trust and, in the end, achieving better patient outcomes.

The next section discusses how Capgemini's cross-industry expertise and internal practices address challenges in Healthcare AI, with examples of use cases delivered.

[12]Duckworth, C., Chmiel, F.P., Burns, D.K. et al. Using explainable machine learning to characterize data drift and detect emergent health risks for emergency department admissions during COVID-19. Sci Rep 11, 23017 (2021). https://doi.org/10.1038/s41598-021-02481-y

[13]Fernández-Narro D, Ferri P, et al. 'Unsupervised Characterization of Temporal Dataset Shifts as an Early Indicator of AI Performance Variations: Evaluation Study Using the Medical Information Mart for Intensive Care-IV Dataset; JMIR Med Inform 2025;13:e78309; https://doi.org/10.2196/78309

[14]https://developer.nvidia.com/flare (Accessed: 26 February 2026)

[15]https://flower.ai/ (Accessed: 26 February 2026)

[16]Li, T., Sahu, A.K., Zaheer, M., Sanjabi, M., Talwalkar, A. and Smith, V. (2020) 'Federated Optimization in Heterogeneous Networks', Proceedings of Machine Learning and Systems (MLSys), 2, pp. 429–450. Available at: proceedings.mlsys.org (Accessed: 26 February 2026).

04

Capgemini's cross-industry experience

All of the challenges described above are indeed solvable, as demonstrated by Capgemini's extensive experience in other highly regulated industries. The lessons learned and best practices developed in these sectors can be directly applied to address similar issues in healthcare and MedTech. Capgemini's teams have successfully implemented a range of AI and data solutions in such environments, emphasizing data integration, regulatory compliance, reliability, and security. The following section highlights how Capgemini's methodologies enable the deployment of compliant AI solutions in healthcare, supported by cross-industry use cases that showcase the application of advanced AI under real-world constraints. To address the unique challenges of healthcare AI, Capgemini has established a comprehensive framework Augmented Engineering approach designed to ensure the delivery of safe, reliable, and high-quality AI solutions.

Capgemini's Augmented Engineering (Hybrid AI) approach is an engineering discipline in which AI is integrated into engineering activities to change how engineering work is carried out. It combines Generative AI with other types of AI and engineering methods to enhance engineers' capabilities and enable more efficient and intelligent solutions. It is not limited to generative AI alone. Engineering work is described as multimodal, involving far more than text, including structured data, visuals, and domain specific artefacts. Augmented Engineering reflects this reality by embedding AI into engineering processes rather than using it as a standalone tool.

This approach is beneficial for constrained environments that require particular care. For instance, high levels of standardization and regulation, strong emphasis on quality, conformity, reliability, and correctness and a need for precision, where incorrect outputs can have significant consequences. Hybrid AI combines the Statistical and machine learning models, which are powerful but unreliable on their own, with structured, rule based, or knowledge driven systems, which provide determinism and accuracy. This combination provides flexibility while maintaining correctness, which is described as essential for engineering contexts where trust and precision matter.

In healthcare, where patient safety, clinical impact, and system reliability are important, this approach provides a disciplined way to apply AI without undermining engineering rigor. It supports the use of AI in environments where decisions and outputs must be dependable, traceable, and fit within established clinical and engineering processes.

Proven expertise in AI value chain for several industries:

Clinical Research Optimization – Pharmaceutical Industry: Capgemini developed a sophisticated data-driven platform for a global pharmaceutical company that optimizes the selection of clinical trial sites and countries. By combining internal and external data into predictive models, the platform accelerates clinical trial planning and enables researchers and project managers to make informed, strategic decisions. The solution simplifies identifying promising trial locations with real-time insights for thorough performance evaluation and proactive risk management throughout the clinical trial process. This demonstrates our capability to deploy AI that directly addresses pharmaceutical R&D's most critical bottlenecks such as trial site selection and patient recruitment where delays cost millions and impact time-to-market.

Clinical Monitoring & Decision Support – Medical Device Industry: Capgemini developed a clinically validated AI-powered system for continuous non-invasive blood pressure monitoring. The solution combines advanced signal processing, automated data annotation, and privacy-preserving techniques to provide continuous hemodynamic monitoring that matches invasive arterial line accuracy while reducing patient risk-now positioned for medical device regulatory approval. This exemplifies our capability to deliver AI systems that meet both clinical validation standards and regulatory pathways to market.

Regulatory Intelligence & Compliance Automation – MedTech Industry: For a medical technology platform migration, Capgemini deployed a generative AI-based knowledge assistant that organizes hundreds of legacy SDLC documents, technical specifications, and quality records into a unified knowledge ontology. Engineering and regulatory teams retrieve specifications, test protocols, and risk documentation with improved speed and accuracy, accelerating onboarding while maintaining the traceability required for medical device submissions. This solution demonstrates our expertise in embedding AI within regulated quality management systems.

Life Sciences R&D Acceleration – Crop Sciences & Agrochemical Industry: In collaboration with a leading crop sciences organization, Capgemini evaluated advanced generative AI platforms for molecular design to explore how AI capabilities from drug discovery can accelerate next-generation agrochemical development. We delivered a first platform prototype enabling computational scientists to test in silico bioengineering approaches, demonstrating 30% faster model deployment and 20% reduction in computational requirements, paving the way for productive solutions that compress early-stage research timelines. This showcases our capability to implement cutting-edge generative AI platforms for research acceleration across life sciences domains where speed and accuracy are competitive differentiators.

Intelligent Manufacturing – Energy & Industrial Equipment: For an energy sector manufacturing operation, Capgemini deployed Hybrid AI combining computer vision and automation to create production plans for custom engineered components. The solution tackles challenges like fragmented data and changing manufacturing processes by supporting planners with an autonomous system that drafts initial plans, demonstrating AI's ability to enhance planning accuracy while maintaining full traceability in regulated manufacturing environments.

Our delivery model: From concept to sustained value

Capgemini delivers AI transformation through a full-lifecycle model that addresses the complete journey from research to production:

- **Data-Driven Decision Intelligence:** Our solutions combine diverse data sources into predictive models that enable strategic decision-making in complex, high-stakes environments such as clinical trial planning, site selection, and resource allocation.
- **Workflow-Native Integration:** Our solutions embed directly into existing clinical and enterprise systems (EHR, LIMS, MES, QMS), ensuring AI enhances existing workflows rather than creating parallel processes.
- **Explainability & Validation:** We employ co-design with clinical and research users and implement interpretable AI architectures that provide transparent reasoning building trust through validation, not opacity.
- **Regulatory-First Architecture:** Our solutions integrate EU AI Act conformity requirements, MDR quality management, and EHDS data governance.
- **Scalable Deployment Infrastructure:** We provide flexible engagement models from platform prototypes to fully productive solutions, with MLOps (Machine Learning Operations) frameworks and domain-specific accelerators that compress time-to-value.

Capgemini brings a strong foundation in solving complex, regulated challenges across healthcare and life science sector through close collaboration with industry experts. Our experience in delivering AI-enabled solutions positions us to support high-impact use cases such as early warning systems for clinical deterioration, diagnostic support in medical imaging, operational optimisation for care delivery, and remote monitoring via wearable data. We combine technical depth with domain insight to co-develop scalable, compliant, and impactful AI solutions and help clients to deploy secure cloud environments.

05

Unlocking business value by overcoming AI & compliance challenges in connected health

As connected health matures across Europe, the integration of AI into remote monitoring and connected healthcare is no longer a futuristic ambition, it is a strategic imperative. But this promise remains unrealized unless we overcome the dual barriers of regulatory complexity and technical challenges. Solving these challenges is not merely a compliance exercise, rather it is the key to unlocking transformative value for patients, providers, payers, and MedTech innovators.

The value of overcoming regulatory and technical barriers

Solving the challenge we described in the chapter above enables:

- **Clinically meaningful AI at scale:** When AI systems are explainable, interoperable, and compliant with frameworks like the FDA, EU AI Act, MDR and GxP regulations, they can move from pilot to practice. This means earlier detection of deterioration, personalized interventions, and reduced hospitalizations, directly improving patient outcomes.

- **Trustworthy, equitable care:** Regulatory alignment ensures that AI systems are fair, auditable, and safe. Technical readiness such as data quality, interoperability, and explainability will ensure that clinicians can trust and act on AI insights, reducing disparities and improving care equity.

- **Operational efficiency and cost savings:** AI-enabled remote monitoring reduces unnecessary admissions, enables earlier discharges, and supports proactive care. Hospitals and insurers benefit from reduced costs and improved resource utilization.

- **Accelerated innovation and market access:** A harmonized regulatory environment (e.g., DiGA in Germany, EHDS across the EU) and scalable technical infrastructure reduce time-to-market for MedTech companies, enabling faster deployment of AI-powered solutions.

A collective responsibility

Although we have discussed a range of strategies to tackle technical challenges, it is important to clarify which stakeholders are responsible for unlocking the process and how this should be achieved. Delivering this value is not the responsibility of one stakeholder. It is a shared mandate, where Hospitals, MedTech companies, insurers, and regulators need to recognize that no single party can solve this alone. Each must invest, align, and collaborate to build a connected health ecosystem that is not only technically feasible but clinically impactful and sustainable.

- Hospitals must modernize their data infrastructure and embed AI into clinical workflows, but they cannot do so without support from technology partners who understand both healthcare and data engineering.

- MedTech device companies must go beyond device innovation to co-develop AI solutions with providers and patients, but they need access to real-world data and regulatory clarity to do so safely and effectively.

- Insurers must fund what works and incentivize adoption, but they need robust evidence and interoperable data to assess value and outcomes.

- Regulators must create clear, harmonized pathways, but they need input from industry and clinical stakeholders to ensure frameworks are practical and innovation friendly.

Europe has made noteworthy strides in addressing AI's regulatory and technical hurdles in connected health: strong frameworks and early successes (like DiGA's reimbursements, the AI Act, EHDS groundwork, and hospitals' data investments) prove the value and feasibility of AI-driven care. However, critical initiatives are still missing or nascent. These include integration and trust at the provider level, data-sharing partnerships, broad payer incentives, and unified guidance – all of which are works in progress. Crucially, no single sector can finish this alone. The full value of AI in connected health will only be delivered to patients when all stakeholders share responsibility and expertise: hospitals, MedTech firms, insurers, and regulators must invest together, align their efforts, and leverage expert Data & AI capabilities to build systems that are interoperable, compliant, and clinically effective. This collective push, which is supported by teams with deep data engineering and clinical AI know-how, is what will transform promising pilots into routine, life-improving care at scale.

At the heart of this transformation is the need for a trusted partner who can bridge the gap between clinical expertise and data science. Capgemini brings a unique combination of scientific rigor, engineering excellence, and scalable AI delivery. With a proven track record across industries such as automotive and pharmaceuticals, we are ideally positioned to help healthcare organizations navigate this complex landscape and implement connected care solutions that are secure, reliable, and impactful.

The role of data and AI experts

At the centre of the ongoing transformation in healthcare is the expertise found within data and AI engineering. Building AI systems that are explainable, compliant, and scalable goes well beyond merely developing algorithms. Success in this area relies on:

- Deep understanding of clinical workflows and regulatory constraints.
- Reliable data platforms that support interoperability, traceability, and maintain high quality.
- AI frameworks designed to be transparent, minimise bias, and allow for continuous oversight.
- Teams that combine clinical, technical, and operational expertise across disciplines.

Without these skills and resources, even the most promising connected health initiatives risk not advancing beyond early trial stages. When these capabilities are in place, it becomes possible to develop AI that is trusted by clinicians, beneficial to patients, and sustainable for health systems.

This shift represents both a technical challenge and a broader strategic opportunity. Through collaborative investment and the inclusion of varied expertise, connected health can progress from an aspiration to a practical solution, providing measurable benefits to patients and the wider healthcare community.

06

Conclusion & outlook

The healthcare system is undergoing a structural shift. Devices alone cannot deliver continuous, proactive care; AI alone cannot function without high-quality, governed clinical data; and financial incentives alone cannot drive sustainable change. Real transformation emerges only when hospitals, insurers, and MedTech companies align around a shared, interoperable data and AI foundation.

Healthcare is rapidly moving toward a more connected, data-driven ecosystem, and several developments will shape how hospitals, insurers, and MedTech organizations operate in the coming years:

- **Seamless Data Ecosystems:** The EHDS will enable real-time, cross-border data exchange, creating the foundation for collaboration and secondary use of clinical data.
- **AI-as-Assistant:** Automation and intelligent support systems will take over routine tasks, reducing workload and enhancing clinical workflows.
- **Personalised Care:** Predictive analytics and advanced modeling will enable earlier interventions and more individualized treatment pathways.
- **Ecosystem Collaboration:** Modular, interoperable IT architectures will allow hospitals, MedTech companies, and insurers to integrate best-of-breed AI tools without vendor lock-in.
- **Continuous Learning:** AI systems will evolve continuously through MLOps, monitoring, and governance, ensuring safe and reliable performance over time.

Capgemini supports organizations in navigating these shifts by solving the underlying Data & AI technical challenges, from building interoperable data platforms and ensuring data quality, to developing explainable, traceable AI and operationalizing it safely at scale. With deep clinical, engineering, and regulatory expertise, we help hospitals, insurers, and MedTech companies turn these industry developments into tangible value.

By combining trustworthy AI, interoperable data, and engineering-led delivery, we create a foundation that is secure, scalable, and regulatory-ready. This allows all ecosystem actors to deploy clinically meaningful AI faster, at lower risk, and with greater impact.

The future of connected health is collaborative. It is built on shared data, shared intelligence, and shared responsibility.

As the industry enters the era of the European Health Data Space and the EU AI Act, now is the moment to establish the architectures, partnerships, and governance models that will define the next decade of healthcare innovation.

Healthcare leaders, across hospitals, insurers, and MedTech companies, now have the opportunity to shape a system where data empowers clinicians, AI enhances human judgment, and patients receive care that is more personalized, more proactive, and more equitable.

Together, we can build an intelligent, patient-centric healthcare ecosystem that is ready for the challenges of today and tomorrow.

07

About the authors

Meghna Goyal

Data Scientist – Applied AI & Generative AI
Capgemini Engineering

meghna.goyal@capgemini.com

Meghna Goyal works across AI R&D, solution architecture, and the delivery of data-driven AI solutions spanning Generative AI, Agentic AI, Natural Language Processing (NLP), and Knowledge Graphs. Her work focuses on providing technical leadership to translate applied AI research into safe, trustworthy, and production-ready systems deployed across Life Sciences and other regulated industries.

Franziska Wolff

Lead for Scientific & Data Driven Innovation
Capgemini Engineering

franziska-elisabeth.wolff@capgemini.com

Franziska Wolff drives AI powered and data driven transformation across industries, helping organizations accelerate innovation and build scalable, future ready R&D capabilities. Her work combines advanced analytics, scientific intelligence, and quantum enabled simulation to unlock new insights, shorten discovery cycles, and enable digital transformation at scale.

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

Make it real.
www.capgemini.com