

DECENTRALIZED AI IN HEALTHCARE: FROM THEORY TO PRACTICE

A CASE STUDY-BASED ANALYSIS

Date:

NOVEMBER 2025

Produced by:

Carmen Cucul
Antonio Pesqueira
Thomas Egelhof

Leilei Tang
Natalia Sofia
Stephanie Fuchs

On behalf of:

Ethereos HealthData
Foundation

Crypto Valley Association
(Sustainability Working Group)



TABLE OF CONTENT

Executive Summary	03
Glossary of Terms	16
About the Authors	21
Methodology	25
Chapter 1: Setting the Stage	27
Chapter 2: Foundational Framework for Decentralization in Healthcare	59
Chapter 3: Decentralized Data Collection and Processing	81
Chapter 4: Decentralized Model Development and Training	109
Chapter 5: Decentralized Validation and Deployment	134
Chapter 6: Future Directions and Strategic Recommendations	156
Closing Remarks	175
References	176

EXECUTIVE SUMMARY



OVERVIEW

Healthcare stands at a pivotal inflection point. By 2030, the World Health Organization projects a shortage of 10 million healthcare workers globally, while half of doctors and nurses already report burnout. Healthcare costs continue to rise unsustainably, data silos fragment care delivery, and traditional centralized AI approaches face increasing barriers around privacy, regulatory compliance, and equitable access. Against this backdrop, decentralized AI (dAI) emerges not as a replacement for centralized systems but as a complementary approach—addressing specific challenges where distributed architectures provide decisive advantages while recognizing that centralized solutions may remain superior for other functions.


THE CURRENT STATE

As of 2025, decentralized AI in healthcare remains in its infancy. While the technological foundations have matured and promising implementations demonstrate viability, the field is characterized by pilot projects, experimental deployments, and emerging standards rather than widespread operational adoption. Many of the case studies documented in this report represent early-stage implementations—proving concepts, establishing governance frameworks, and navigating regulatory uncertainties. Yet this nascent stage should not obscure the trajectory: convergent forces across technology, regulation, economics, society are creating unprecedented momentum, positioning 2025 as a pivotal inflection point from which more mature implementations will emerge through 2030.

This report examines decentralized AI in healthcare as it transitions from theoretical promise to practical implementation. Drawing on systematic literature review, 10 detailed case studies spanning five continents, expert interviews with healthcare leaders and patient advocates, and analysis of emerging regulatory frameworks, we document how organizations are strategically deploying federated learning, blockchain infrastructure, edge computing, and decentralized governance where these approaches deliver clear benefits over traditional centralized architectures.

ABOUT THIS REPORT

This business-oriented report targets healthcare executives, technology leaders, policymakers, and researchers who operate at the intersection of AI and healthcare strategy. We assume readers have foundational understanding of conventional AI concepts (machine learning basics, neural networks, clinical AI applications) but are neither complete beginners requiring introductory explanations nor deep technical experts needing mathematical proofs. While we maintain academic rigor through systematic literature review and peer consultation, the structure follows business logic designed for strategic decision-making. Chapter 2 presents more technical content covering terminology essential for understanding subsequent case studies and future developments—readers should have prior familiarity with standard AI terminology to fully benefit from this foundational chapter.



THE CONVERGENT FORCES DRIVING DECENTRALIZATION

Four interconnected forces have reached critical mass in 2025, creating unprecedented momentum for dAI adoption:



01 TECHNOLOGICAL MATURATION

Federated learning has evolved from theoretical concept to operational implementation, supported by enterprise-ready frameworks (MONAI, NVIDIA FLARE), edge computing proliferation enabling hospital-grade local processing, blockchain infrastructure offering substantial cost reductions versus traditional cloud providers, and standardized interoperability advances through FHIR APIs.

02 REGULATORY MANDATES

Some privacy laws now mandate decentralization rather than merely encouraging it. GDPR's global influence, China's Personal Information Protection Law (PIPL), the European Health Data Space, and healthcare-specific protections raise legal barriers to traditional centralized data aggregation that decentralized approaches navigate by design.

03 ECONOMIC IMPERATIVES

Unsustainable healthcare spending, workforce crises requiring AI augmentation, competitive collaboration needs where individual capabilities cannot match collective intelligence benefits, and vendor lock-in concerns are driving organizations toward shared, distributed solutions.

04 SOCIETAL SHIFTS

Post-pandemic privacy awareness, AI explainability demands from clinicians, patient empowerment movements seeking data ownership, and democratic governance expectations converge on decentralization's core value proposition—preserving individual control while enabling collective benefit.

GLOBAL LEADERSHIP PATTERNS

Regional analysis reveals complementary rather than competing approaches:

Europe leads privacy-first collaborative frameworks (e.g. MELLODDY Consortium in federated learning)

North America drives platform-driven commercial integration (e.g. Mayo Clinic Platform)

Asia-Pacific implements sovereign AI solutions (e.g. China's DeepSeek hospital deployments)

Switzerland pioneers a unique "third way" combining national sovereignty with open-source transparency (Apertus LLM)

Africa advances inclusive federated learning addressing resource constraints and health equity

This distributed innovation landscape creates a diverse global ecosystem where different regions contribute specialized capabilities while addressing shared healthcare challenges.

EVIDENCE BASE AND RESEARCH METHODOLOGY

Our systematic literature review following PRISMA 2020 guidelines analyzed 87 peer-reviewed studies published between January 2023 and March 2025, with 13 papers selected for in-depth qualitative analysis. The review reveals three technology architectures at different maturity levels:

1. **Federated Learning (Most Mature):** The vanguard architecture for privacy-preserving clinical compute, with consistent application in medical imaging and electronic health record analysis. Primary challenge: statistical heterogeneity of non-identically distributed data, addressed through advanced algorithms like FedProx and Scaffold.
2. **Blockchain (Operational Infrastructure):** Functions as trust-enabling infrastructure providing tamper-proof audit trails, optimal for clinical trials management, digital identity, and multi-stakeholder data exchange where auditability and integrity are paramount.
3. **Decentralized Autonomous Organizations (Most Transformative, Nascent):** Sociotechnical frameworks for reimagining collaboration through blockchain-based governance, applied to civic medical data trusts and decentralized science (DeSci) funding, though socio-economic and legal frameworks remain in infancy.

THE THREE PILLARS FRAMEWORK

Decentralized AI in healthcare operates across three interconnected architectural pillars that work synergistically throughout the AI development lifecycle:



PILLAR 1: DECENTRALIZED DATA MANAGEMENT

Patient data remains within institutional boundaries rather than aggregating in central repositories. Technologies include federated learning (training models across distributed datasets without centralizing data), privacy-preserving computation (analyzing data without exposing raw information), and local data sovereignty (institutional control over sensitive health information). This pillar addresses regulatory compliance, competitive collaboration, and trust requirements.

PILLAR 2: DECENTRALIZED COMPUTE AND CONTROL

AI processing occurs across distributed infrastructure rather than centralized cloud platforms. This encompasses edge computing (local model deployment within healthcare facilities), distributed GPU networks (decentralized computation marketplaces), and autonomous agents (systems operating independently across institutional boundaries). This pillar addresses cost efficiency, latency requirements, and infrastructure resilience.

PILLAR 3: DECENTRALIZED GOVERNANCE

Decision-making authority is distributed among stakeholders rather than being concentrated in single entities. Mechanisms include blockchain-based voting (transparent, auditable governance), multi-stakeholder consortia (collaborative decision frameworks), and community-driven research prioritization (democratic resource allocation). This pillar addresses accountability, transparency, and stakeholder participation.

These pillars manifest throughout the AI model lifecycle—from initial data collection through training to validation and deployment. The following chapters (3-5) are structured around this lifecycle, demonstrating how decentralization emerges at each development stage. Each case study explicitly identifies which pillars it exhibits, recognizing that implementations may emphasize one pillar strongly or integrate all three in sophisticated hybrid architectures.

CASE STUDIES ACROSS THE AI LIFECYCLE

The report documents 10 primary case studies structured around the AI model development lifecycle (data collection → model training → validation and deployment), with each implementation exhibiting one or more of the three pillars depending on its specific architecture and objectives. Beyond these detailed analyses, we include bonus resources highlighting additional initiatives that offer valuable insights into emerging patterns and potential future directions—though these may represent earlier development stages, more centralized current operations, or enabling infrastructure rather than complete implementations at the maturity level of our case studies.

CHAPTER 3: DECENTRALIZED DATA COLLECTION & PROCESSING

This chapter examines how organizations collect and process health data while maintaining privacy, sovereignty, and regulatory compliance through decentralized approaches.

MELLODDY Consortium (Europe) - Ten competing pharmaceutical companies collaborate on drug discovery through federated learning, processing 1.3+ billion compound-assay data points across proprietary libraries without sharing raw data. Three-year implementation required sophisticated governance frameworks defining intellectual property rights and conflict resolution alongside technical architecture. Exhibits: Pillar 1 (Data Management) & Pillar 3 (Governance)

Mayo Clinic Platform "Discover" and "Connect" (USA) - "Data Behind Glass" architecture enables 81 developers to analyze de-identified data from 45+ million patient lives without raw data ever leaving Mayo's infrastructure. FHIR-based federated queries across 25+ partner organizations demonstrate how competitive academic medical centers can collaborate while maintaining institutional data control. Exhibits: Pillar 1 (Data Management)

SEOVE - Cadeia de Cuidados (Brazil) - Federated care coordination platform for domestic violence survivors and elderly care in resource-constrained settings. Demonstrates that decentralized approaches can succeed despite limited infrastructure by leveraging open-source frameworks and strategic partnerships while maintaining LGPD compliance and enhanced confidentiality for vulnerable populations. Exhibits: Pillar 1 (Data Management) & Pillar 3 (Governance)

Acoer Clinical Trials Platform (USA) - Blockchain-based "Cryptographic Data Mesh" on Hedera Hashgraph provides immutable audit trails for clinical trial data provenance, tamper-proof consent tracking, and transparent AI model governance. Enables federated analysis across trial sites without centralizing sensitive patient data. Exhibits: Pillar 1 (Data Management) & Pillar 3 (Governance)

KEY CHAPTER INSIGHTS

Competitive collaboration requires governance investment comparable to technical development. Organizations consistently require 2-3 years developing legal frameworks, partnership agreements, and conflict resolution mechanisms before achieving operational federated learning—technical architecture alone proves insufficient.

CHAPTER 4: DECENTRALIZED MODEL DEVELOPMENT&TRAINING

This chapter explores how AI models are trained, refined, and optimized using decentralized computational infrastructure and collaborative frameworks.

DeepSeek R1/VL (China) - Open-source sovereign AI foundation models deployed across 100+ Chinese healthcare institutions, enabling state-of-the-art reasoning and vision-language capabilities with complete local data control. Demonstrates that geopolitical drivers (data sovereignty, supply chain independence) can be as powerful as technical considerations in driving dAI adoption. Exhibits: Pillar 1 (Data Management) & Pillar 2 (Computation)

Akash Network and Aethir (Global) - Decentralized GPU marketplaces offering up to 85% cost savings versus traditional cloud providers for AI training workloads. Akash provides community-driven computation with transparent pricing; Aethir delivers enterprise-grade infrastructure for high-performance requirements. Demonstrates complementary rather than replacement relationship with centralized cloud. Exhibits: Pillar 2 (Computation)

Bittensor (Global) - Open marketplace where AI models are collaboratively trained through incentivized contribution, with healthcare applications in medical imaging and drug discovery. Validators assess model quality; miners contribute computation/data; cryptocurrency rewards align incentives for continuous improvement without central data custody. Exhibits: Pillars 2 (Computation) & 3 (Governance)

ASI Alliance (Singapore/UK) - Merges autonomous agents, data marketplaces, and computation infrastructure for decentralized healthcare workflows. Specialized agents demonstrate superhuman diagnostic performance in pathology; platform coordinates multi-domain medical reasoning while maintaining transparent provenance. Exhibits: All 3 Pillars

KEY CHAPTER INSIGHTS

Decentralized computation networks complement rather than replace centralized cloud infrastructure. Hybrid architectures emerge as a sustainable pattern—centralized cloud for mission-critical clinical systems requiring maximum reliability; decentralized computation for research, development, and training workloads prioritizing cost efficiency and data locality.

CHAPTER 5: DECENTRALIZED VALIDATION AND DEPLOYMENT

This chapter examines how AI systems transition from development to production environments, emphasizing governance models essential for scaling within existing healthcare infrastructure.

VitaDAO (Switzerland) - Decentralized autonomous organization governing AI-driven longevity research through blockchain-based decision-making. Over 10,000 community members direct \$4.1M in funding across 22+ projects via transparent token-weighted voting. Demonstrates transformative governance potential while highlighting need for healthcare-specific accountability frameworks. Exhibits: All 3 Pillars

Changi General Hospital's AimSG Platform (Singapore) - National-scale AI deployment for medical imaging across Singapore's public healthcare system. AI models for chest X-ray triage operate locally within hospitals while validation, accreditation, and oversight coordinate at national level—exemplifying how decentralized technical architecture coexists with appropriate centralized governance. Exhibits: Pillar 1 (Data Management) & Pillar 2 (Computation)

KEY CHAPTER INSIGHTS

Production deployment requires hybrid architectures balancing decentralization with selective centralization. Successful implementations combine local institutional autonomy (data sovereignty, clinical decisions) with coordinated national governance (quality monitoring, regulatory compliance, interoperability standards) rather than pursuing ideological purity.

KEY FINDINGS ACROSS CASE STUDIES

01

Hybrid Architectures Outperform Pure Decentralization

Organizations succeed by combining local autonomy (data sovereignty, institutional decision-making) with selective centralization (quality monitoring, regulatory compliance, interoperability standards) rather than pursuing ideological purity.

02

Governance Investment Precedes Technical Success

The timeline for operational federated learning consistently spans 2-3 years, with governance framework development consuming comparable resources to technical implementation. Organizations that underinvest in legal structures, partnership agreements, and conflict resolution mechanisms face deployment delays regardless of technical readiness.

03

Organizational Readiness Constrains Adoption

The gap between technical possibility and organizational capability emerges consistently across implementations. Success through 2030 depends as much on change management, training, and cultural transformation as on technology development.

04

Certain Domains Particularly Benefit from Decentralized Approaches

Mental health innovation (privacy-reducing stigma), rare disease research (uniting geographically dispersed populations), vulnerable population care (enhanced confidentiality), and cross-border collaboration (navigating conflicting regulations) demonstrate distinctive advantages that justify higher implementation complexity.

05

Economic Models Remain Experimental

While technical viability is demonstrated, sustainable economic models require further development. Subscription-based platform access (Mayo Clinic), tokenized incentive mechanisms (VitaDAO), and cost-sharing consortia (MELLODDY) represent early experiments rather than proven templates.

FUTURE TRAJECTORY THROUGH 2030

The report identifies five domains where emerging trends will shape decentralized AI's evolution over the next five years, analyzing which forces create tailwinds accelerating adoption versus headwinds introducing resistance and complexity.

TECHNOLOGY TRENDS

Represent strong tailwinds for dAI adoption. Sovereign AI implementations driven by geopolitical considerations are emerging as one of the most powerful accelerants, with nation-states like Switzerland and China developing indigenous AI capabilities where data remains within national boundaries. Strategic deployment of decentralized compute infrastructure demonstrates how geopolitical constraints are already accelerating alternative architectures. Agentic AI systems enabling autonomous medical decision support are rapidly maturing, with blockchain-enabled platforms demonstrating compelling healthcare applications across diagnostics, drug discovery, and resource optimization. Multimodal federated learning integrating imaging, genomics, and electronic health records is expanding analytical capabilities. Edge AI advancement continues reducing cloud dependency, with hospital-grade processing power enabling sophisticated local deployment without compromising data locality requirements.

DATA & INFRASTRUCTURE TRENDS

Present a mixed picture of opportunities and challenges. Open-source momentum through frameworks like MONAI and Flower.ai accelerates innovation and reduces barriers to entry, though concerns about quality assurance and enterprise support persist for production deployments. Decentralized computation networks including Akash and Render offer substantial cost advantages—up to 85% savings versus traditional cloud providers—but introduce reliability questions and coordination complexity that healthcare's risk-averse culture finds challenging. The critical challenge of semantic interoperability remains inadequately addressed, requiring sustained multi-stakeholder investment that no single organization can justify individually. Standards development for healthcare-specific federated learning protocols progresses slowly, creating friction as early adopters develop incompatible implementations that will prove difficult to harmonize later.

GOVERNANCE AND ECONOMIC TRENDS

Generate moderate tailwinds alongside persistent headwinds. DAO governance models are maturing beyond pure experimentation, with organizations like VitaDAO demonstrating that blockchain-based decision-making can coordinate thousands of stakeholders transparently. However, healthcare's unique accountability requirements—regulatory mandates for clinical oversight, liability frameworks demanding identifiable decision-makers, patient safety provisions requiring rapid intervention—create tensions with pure decentralized governance that require healthcare-specific adaptations. Tokenized incentive mechanisms enabling compensation for data contributors and model developers show promise but face regulatory classification uncertainty. Will health data tokens be treated as securities?

How do tax implications affect cross-border research collaborations? These questions lack definitive answers, creating hesitancy among institutional adopters. Perhaps most significantly, institutional resistance to power redistribution creates adoption friction—centralized healthcare systems and dominant technology vendors often resist architectures that would diminish their control and market position.

SUSTAINABILITY TRENDS

Present complex tradeoffs. Energy efficiency gains from edge computing and distributed processing offer meaningful improvements over centralized data center models, reducing both costs and environmental impact. However, blockchain computational overhead—particularly for proof-of-work consensus mechanisms—can offset these gains, requiring careful architectural choices toward more efficient alternatives like proof-of-stake. Broader AI scaling limitations are forcing architectural innovation toward efficiency regardless of centralization approach, as the exponential growth in model size and training costs proves unsustainable. The digital divide presents a critical equity concern: will decentralized AI democratize access by reducing infrastructure requirements, or will it create new disparities between institutions that can deploy sophisticated federated systems and those that cannot? Addressing this requires explicit equity-focused design rather than assuming decentralization inherently improves access.

REGULATORY TRENDS

Represent perhaps the greatest source of uncertainty shaping the 2025-2030 trajectory. The European Health Data Space creates precedent for federated infrastructure supporting cross-border health data sharing while maintaining national sovereignty—potentially serving as a template that other regions adopt or adapt. The EU AI Act establishes risk-based compliance frameworks that, while designed for centralized systems, contain provisions accommodating decentralized approaches for lower-risk applications. However, critical liability framework gaps for autonomous distributed systems persist, creating adoption hesitancy among healthcare organizations unwilling to deploy systems where accountability remains legally ambiguous. If a federated AI model produces a harmful clinical decision, which institution bears responsibility—the one that contributed training data, the one that hosted the final inference, or the consortium that governed development? These questions lack clear legal answers in most jurisdictions. International coordination challenges compound these issues, as federated networks increasingly operate across borders with conflicting regulatory requirements. The emergence of regulatory sandboxes and pilot programs in several countries offers pathways for experimentation, but the transition from sandbox to scaled deployment remains uncertain—will regulators extrapolate from successful pilots to create enabling frameworks, or will they maintain restrictive default positions requiring case-by-case approval?

STRATEGIC RECOMMENDATIONS BY STAKEHOLDER GROUP

1

FOR HEALTHCARE DELIVERY ORGANIZATIONS

Begin controlled experimentation now (2025-2027) while maintaining hybrid centralized-decentralized operations. Focus on: identifying specific use cases where data sovereignty provides decisive regulatory advantages; joining established federated networks rather than building infrastructure independently; investing comparably in governance frameworks and technical implementation; measuring success through strategic positioning rather than only operational metrics; collaborating through professional societies for liability framework advocacy.

Critical Questions: Which clinical domains justify decentralized complexity versus centralized efficiency? How do you balance local autonomy demands with system-wide standardization requirements? What partnership selection criteria distinguish genuine collaborators from those exploiting asymmetric data contributions?

2

FOR PATIENT GROUPS AND ADVOCACY ORGANIZATIONS

Prioritize data governance participation over pure technology adoption. Focus on: joining DAO governance structures to influence research priorities; demanding transparent, auditable AI decision-making; supporting policy frameworks enabling patient data ownership; educating communities about privacy-sovereignty tradeoffs; advocating for equity-focused design in decentralized systems.

Critical Questions: How do you ensure governance participation doesn't tokenize patient voices while maintaining efficiency? What mechanisms verify that decentralization delivers tangible privacy benefits beyond marketing claims?

3

FOR TECHNOLOGY DEVELOPERS AND VENDORS

Recognize that healthcare requirements differ fundamentally from consumer applications. Focus on: building interoperability by design rather than proprietary lock-in; developing healthcare-specific governance templates addressing liability and oversight; creating graduated deployment pathways from pilot to production; establishing transparent security and model provenance; designing for resource-constrained environments alongside high-capability settings.

Critical Questions: How do you monetize open-source federated frameworks while maintaining community trust? What technical abstractions hide complexity without sacrificing necessary clinical control?

4

FOR RESEARCH INSTITUTIONS AND FUNDERS

Accelerate standards development and capability building. Focus on: establishing multi-institutional federated research networks; funding interoperability standards consortia; supporting comparative effectiveness research on centralized versus decentralized approaches; training next-generation workforce in distributed AI systems; creating regulatory sandboxes for decentralized AI experimentation.

Critical Questions: How do you balance open science principles with institutional intellectual property concerns? What funding mechanisms incentivize collaborative standards development that no single institution can justify individually?

5

FOR POLICYMAKERS AND REGULATORS

Develop adaptive frameworks accommodating innovation while ensuring safety. Focus on: establishing approval pathways for decentralized systems in 2025-2027 to inform scaled frameworks by 2028-2030; creating graduated liability guidance rather than waiting for comprehensive frameworks; facilitating international regulatory coordination as federated networks cross borders; designing regulatory sandboxes enabling experimentation without premature technology lock-in; developing monitoring frameworks providing visibility without undermining privacy benefits.

Critical Questions: How do you regulate distributed systems where traditional assumptions about control and accountability no longer apply? What mechanisms transition successful pilots to scaled deployment while filtering unsuccessful approaches?

CONCLUSION: AN EVOLUTIONARY, COMPLEMENTARY PATH FORWARD

This report documents a field at a pivotal moment. The technological capabilities have matured, regulatory frameworks increasingly accommodate federated approaches, and economic pressures make purely centralized AI architectures insufficient for many healthcare challenges. Yet significant barriers around liability frameworks, interoperability standards, and organizational readiness will determine whether decentralized AI fulfills its potential or remains confined to narrow applications.

We envision an evolutionary rather than revolutionary trajectory through 2030, with decentralized and centralized approaches coexisting as complementary architectures rather than competing paradigms. Healthcare will witness strategic integration of decentralized components where they provide decisive advantages:

- **Data sovereignty** addressing regulatory requirements where centralized aggregation violates privacy laws
- **Federated learning** enabling multi-institutional collaboration where competitive or ethical factors prevent data sharing
- **Edge computing** supporting real-time clinical decisions where latency or connectivity constraints exist
- **DAO governance** coordinating research consortia where stakeholder participation improves legitimacy and innovation

Simultaneously, **centralized elements** will persist and remain superior for functions requiring:

- **Unified oversight** for quality assurance and safety monitoring
- **Standardized compliance** with regulatory frameworks demanding clear accountability
- **Rapid coordination** in time-critical clinical scenarios
- **Integrated service ecosystems** where orchestration complexity outweighs decentralization benefits
- **Clear liability frameworks** where distributed responsibility creates unacceptable legal uncertainty

The outcome depends not on technological determinism but on strategic decisions diverse stakeholders make over the next 3-5 years. Organizations that experiment strategically now while building governance capabilities position themselves to shape emerging standards rather than inheriting externally imposed frameworks. Those that delay experimentation until "the technology matures" risk discovering in 2028-2030 that dominant architectures, business models, and regulatory templates were established by competitors and collaborators who moved earlier—yet those who pursue decentralization ideologically without recognizing centralized solutions' enduring advantages risk costly misalignment between architecture and actual requirements.

Success requires sophisticated judgment about which specific functions genuinely benefit from decentralization versus which improve through centralization—a capability built through experimentation rather than theoretical analysis. The transformation toward more strategically architected healthcare AI, deploying each approach where it delivers optimal results, is underway. This report provides the foundational understanding and strategic frameworks to navigate that transformation effectively.



GLOSSARY OF TERMS

This glossary provides definitions for key technical terms, acronyms, and concepts used throughout the report. Terms are organized by category for easier reference.

CORE AI CONCEPTS

ARTIFICIAL INTELLIGENCE (AI)

Computer systems capable of performing tasks that typically require human intelligence, including learning from data, recognizing patterns, making decisions, and generating predictions or recommendations.

FOUNDATION MODELS

Large-scale AI systems trained on broad, diverse datasets that can be adapted to multiple specific tasks through fine-tuning. Examples include models for text, images, audio, video, and multimodal combinations.

LARGE LANGUAGE MODELS (LLMs)

A subset of foundation models specifically trained on vast amounts of text data, focused on understanding and generating human language. Examples include GPT, Claude, and DeepSeek.

MULTIMODAL MODELS

AI systems that can process and integrate multiple types of data simultaneously (text, images, genomics, sensor data, etc.), enabling comprehensive analysis across different information sources.

MULTI-TASK LEARNING

Training AI models to perform multiple related tasks simultaneously, allowing different institutions to contribute diverse expertise without centralizing sensitive information.

MACHINE LEARNING (ML)

A subset of AI where systems learn patterns from data and improve their performance over time without being explicitly programmed for every scenario.

DECENTRALIZED AI ARCHITECTURES

DECENTRALIZED AI (DAI)

AI systems where control, computation, and governance are distributed across multiple entities rather than concentrated in a single centralized authority. Data remains at its source while models and algorithms move to the data.

FEDERATED LEARNING (FL)

A machine learning approach that trains AI models across decentralized locations without exchanging raw data. The model visits each institution to train locally, then returns with learned insights while original datasets remain behind.

HORIZONTAL FEDERATED LEARNING

Federated learning where participating institutions have similar data structures (e.g., all have chest X-rays) but different patient populations.

VERTICAL FEDERATED LEARNING

Federated learning where institutions have different data types about overlapping populations (e.g., hospital has clinical records, insurer has claims data, pharmacy has medication history for the same patients).

MULTI-CHAIN FEDERATED LEARNING

Federated learning coordinated across different blockchain infrastructures, allowing institutions on different technical platforms to collaborate.

DISTRIBUTED AI

AI computation distributed across multiple locations or devices, which may or may not preserve data sovereignty (broader term than federated learning).

SOVEREIGN AI

Nation-specific AI development that maintains local control over both algorithms and training data, ensuring data remains within national boundaries for strategic and regulatory reasons.

COMPUTING INFRASTRUCTURE

EDGE COMPUTING

Processing data locally at or near the source (hospitals, clinics, medical devices) rather than sending it to centralized cloud servers, reducing latency and maintaining data locality.

EDGE AI

AI models deployed and executed on edge infrastructure, enabling real-time insights without cloud dependency or data transmission.

CLOUD COMPUTING

Computing services (storage, processing, databases) delivered over the internet through centralized data centers operated by providers like AWS, Google Cloud, or Microsoft Azure.

DECENTRALIZED COMPUTE NETWORKS

Blockchain-based marketplaces where computational resources (GPUs, CPUs) are distributed across independent providers rather than centralized cloud platforms. Examples include Akash Network and Aethir.

GPU (GRAPHICS PROCESSING UNIT)

Specialized hardware originally designed for rendering graphics but now essential for AI model training and inference due to parallel processing capabilities.

PRIVACY-ENHANCING TECHNOLOGIES

PRIVACY-ENHANCING TECHNOLOGIES (PETS)

Technologies that enable computation and analysis on sensitive data without exposing underlying information to unauthorized parties.

DATA BEHIND GLASS®

An architecture where external parties can analyze data within secure containers, but the data never leaves institutional boundaries. Pioneered by Mayo Clinic Platform.

HOMOMORPHIC ENCRYPTION

Cryptographic technique allowing mathematical operations to be performed on encrypted data without ever decrypting it, enabling secure computation on sensitive medical information.

DIFFERENTIAL PRIVACY

Mathematical technique that adds carefully calibrated noise to datasets or query results to protect individual privacy while preserving aggregate statistical insights for AI models.

ZERO-KNOWLEDGE PROOFS (ZKPS)

Cryptographic methods that verify patient credentials or medical facts without revealing the actual underlying data.

SECURE MULTI-PARTY COMPUTATION (SMC)

Cryptographic protocols enabling multiple institutions to jointly compute results from their combined data while keeping each party's individual inputs private.

SYNTHETIC DATA

Artificially generated datasets that statistically resemble real patient data, used for AI training without exposing actual patient information.



BLOCKCHAIN & DLTS

BLOCKCHAIN

A distributed ledger technology that records transactions in tamper-proof, chronological blocks linked through cryptography, providing transparent and immutable audit trails.

DISTRIBUTED LEDGER TECHNOLOGY (DLT)

Broader category of technologies (including blockchain) where data is recorded across multiple locations simultaneously, with no central administrator.

SMART CONTRACTS

Self-executing contracts with terms written directly in code, automatically enforcing agreements when predefined conditions are met. Used for data access governance and automated transactions.

DECENTRALIZED AUTONOMOUS ORGANIZATION (DAO)

Blockchain-based governance structure where decision-making authority distributes among stakeholders through transparent, token-weighted voting mechanisms rather than centralized management.

TOKENIZATION

Representing assets, rights, or data access permissions as digital tokens on a blockchain, enabling fractional ownership, transparent tracking, and programmable incentives.

NON-FUNGIBLE TOKEN (NFT)

Unique digital token representing ownership or authenticity of a specific asset (e.g., intellectual property rights, individual medication batches in supply chains).

GOVERNANCE & IDENTITY

SELF-SOVEREIGN IDENTITY (SSI)

Identity systems where individuals control their own digital identities through blockchain-based credentials, enabling selective disclosure and granular permission management without central authority.

DECENTRALIZED IDENTIFIERS (DIDS)

Blockchain-based digital identifiers that give patients control over their healthcare identities, enabling verifiable credentials without relying on centralized identity providers.

DATA SOVEREIGNTY

The principle that data remains under the control and jurisdiction of the entity or nation that generated it, with legal authority to determine how it's accessed and used.

PATIENT DATA SOVEREIGNTY

The right of patients to maintain control over their personal health information, determining who accesses it, for what purposes, and under what conditions.

DECENTRALIZED SCIENCE (DESCI)

DECENTRALIZED SCIENCE (DESCI)

Movement applying blockchain, DAOs, and decentralized technologies to scientific research, enabling transparent funding, open collaboration, and community governance of research priorities.

RESEARCH DAO

Decentralized autonomous organization focused on funding, coordinating, or governing scientific research through blockchain-based governance mechanisms.

HEALTHCARE DATA STANDARDS

FHIR (FAST HEALTHCARE INTEROPERABILITY RESOURCES)

Modern standard for electronic health information exchange providing APIs for real-time data access and standardized data structures, enabling seamless communication across healthcare systems.

HL7 (HEALTH LEVEL SEVEN)

Earlier healthcare data exchange standard still widely used in legacy systems, focusing on clinical and administrative data messaging between healthcare applications.

DICOM (DIGITAL IMAGING AND COMMUNICATIONS IN MEDICINE)

International standard for medical imaging data, ensuring compatibility across imaging equipment and systems.

PHI (PROTECTED HEALTH INFORMATION)

Any individually identifiable health information, including medical records, test results, and billing information, protected under regulations like HIPAA.

EHR (ELECTRONIC HEALTH RECORD)

Digital version of a patient's complete medical history, maintained by healthcare providers and containing clinical data, diagnoses, medications, treatment plans, and test results.

HEALTHCARE DATA REGULATIONS

HIPAA (HEALTH INSURANCE PORTABILITY AND ACCOUNTABILITY ACT)

U.S. federal law establishing privacy and security standards for protecting patient health information.

GDPR (GENERAL DATA PROTECTION REGULATION)

European Union regulation governing data protection and privacy, requiring explicit consent for data processing and granting individuals rights over their personal data.

PIPL (PERSONAL INFORMATION PROTECTION LAW)

China's comprehensive data protection law regulating collection, storage, and processing of personal information, including health data.

LGPD (LEI GERAL DE PROTEÇÃO DE DADOS)

Brazil's General Data Protection Law, similar to GDPR, protecting personal data privacy and establishing requirements for data processing.

EHDS (EUROPEAN HEALTH DATA SPACE)

EU initiative establishing frameworks for cross-border health data sharing while maintaining national sovereignty, supporting federated research infrastructure.

CLINICAL & OPERATIONAL TERMS

CLINICAL DECISION SUPPORT SYSTEM (CDSS)

Software providing healthcare professionals with patient-specific assessments or recommendations to aid clinical decision-making.

POINT-OF-CARE DIAGNOSTICS

Medical testing performed at or near the site of patient care, providing rapid results that can inform immediate treatment decisions.

INTEROPERABILITY

Ability of different healthcare information systems, devices, and applications to access, exchange, and cooperatively use data in a coordinated manner.

DATA MINIMIZATION

Privacy principle requiring that only data necessary for a specific purpose be collected and processed, reducing privacy risks and compliance requirements.

AUDIT TRAIL

Chronological record documenting the sequence of activities affecting data or system operations, essential for accountability and regulatory compliance.

EMERGING TECH & CONCEPTS

AGENTIC AI

AI systems capable of autonomous action, goal-directed behavior, and decision-making with minimal human intervention, often coordinated through multi-agent frameworks.

MULTI-AGENT SYSTEMS

Networks of autonomous AI agents that interact, coordinate, and collaborate to achieve individual or collective goals.

RETRIEVAL-AUGMENTED GENERATION (RAG)

AI technique combining large language models with real-time information retrieval, enabling models to access current data and provide more accurate, verifiable responses.

EXPLAINABLE AI (XAI)

AI systems designed to provide transparent, interpretable explanations for their decisions and predictions, essential for clinical trust and regulatory compliance.

MOVE-TO-EARN

Gamification mechanism rewarding users with digital tokens for physical activity or healthy behaviors, aligning economic incentives with preventive healthcare.

KEY ACRONYMS & ABBREVIATIONS

API - Application Programming Interface
DAO - Decentralized Autonomous Organization
DeSci - Decentralized Science
DID - Decentralized Identifier
DLT - Distributed Ledger Technology
EHR - Electronic Health Record
EHDS - European Health Data Space
FHIR - Fast Healthcare Interoperability Resources
FL - Federated Learning
GDPR - General Data Protection Regulation
GPU - Graphics Processing Unit

HIPAA - Health Insurance Portability and Accountability Act
IoT - Internet of Things
LLM - Large Language Model
ML - Machine Learning
NFT - Non-Fungible Token
PET - Privacy-Enhancing Technology
PIPL - Personal Information Protection Law
RAG - Retrieval-Augmented Generation
SMC - Secure Multi-Party Computation
SSI - Self-Sovereign Identity
ZKP - Zero-Knowledge Proof

*This glossary reflects terminology as used throughout the report. For detailed technical explanations and healthcare-specific applications, please refer to Chapter 2.

ABOUT THE AUTHORS

SUPPORTING ORGANISATIONS:



Ethereos HealthData Foundation

Ethereos HealthData Foundation (EHF) is a US-incorporated nonprofit organization with a Swiss subsidiary, dedicated to enabling individuals globally to own, share, and benefit from their health data. Its goal is to foster a shift towards a more equitable, inclusive, and personalized healthcare future. EHF operates at the intersection of emerging technologies and patient empowerment, focusing on privacy-enhancing technologies, decentralized data management, and self-sovereign identity solutions.

The foundation actively engages in thought leadership through research initiatives, including its 2024 collaborative report with the Crypto Valley Association on "Decentralized Health Data Management: An overview of solutions empowering individuals to own, share and benefit from their health data". Additionally, EHF organizes educational events bringing together healthcare providers, technology developers, and patient advocates to advance the field of decentralized health data management.

Through its team of consultants, researchers, and Young Ambassador program, EHF bridges the gap between cutting-edge technology and practical healthcare applications, advocating for systems where patients maintain control over their medical records while enabling secure, seamless data sharing with healthcare providers. The foundation is also developing an Impact Studio program designed to accelerate the development of decentralized health data management solutions.

Website: <https://etheroshealthdata.org/>



Crypto Valley Association (CVA) is a leading blockchain and crypto ecosystem in Switzerland, representing more than 250 corporations and 900 individual members. It has been fostering growth, collaboration and integrity in the global blockchain economy since 2017. CVA operates as an independent association headquartered in Zug, Switzerland—the heart of "Crypto Valley"—serving as a hub for blockchain innovation, regulatory advocacy, and ecosystem development.

The association organizes flagship events including the annual Crypto Valley Conference in Rotkreuz and the Banking Symposium in Lugano, which bring together global leaders in blockchain, finance, and technology. CVA actively produces thought leadership through its working groups, publishing influential research including the annual CVA Research Journal (featuring topics such as "Risk Management in Web3" in 2023 and "DeFi in Traditional Finance" in 2022), the "Next Step: Sustainable AI" 2023 white paper exploring blockchain's role in addressing AI's ethical and environmental challenges, as well as regulatory position papers on critical topics such as crypto staking services.

The association connects entrepreneurs, investors, technology developers, legal experts, and policymakers to advance responsible blockchain adoption across industries. Its members benefit from networking opportunities, access to regulatory expertise, participation in policy discussions with Swiss and international authorities, and visibility within one of the world's most established blockchain communities. CVA has been instrumental in establishing Switzerland as a global leader in blockchain innovation while maintaining high standards for compliance, ethics, and technical excellence.

Website: <https://cryptovalley.swiss/>

CO-AUTHORS (ALPHABETICAL ORDER):



ANTONIO PESQUEIRA

Antonio is a specialist in the field of healthcare and pharmaceutical operations, with a focus on digital transformation and strategic technology application. With his experience of over 15 years in the industry, including areas such as supply chain, commercial excellence and project management. Antonio's research contributions are focused on Generative AI, Agent-Based Modelling, and Blockchain, exploring their application in enhancing operational efficiency and compliance within complex healthcare supply chains. His work has been published in respected journals, including Elsevier's Journal of Medical Systems and Springer Nature's Software Quality Journal. He currently serves as an Editorial Review Board Member for the International Journal of Artificial Intelligence in Medicine and Healthcare (IJAIMH) and contributes to INESC-INOV and as a reviewer for Springer Nature's Discover Artificial Intelligence and Health Economics.



CARMEN CUCUL

Carmen Cucul is a Blockchain & AI Healthcare Consultant with over 15 years of professional experience at the intersection of healthcare and emerging technologies, holding roles in digital health, customer excellence, learning & development, and strategic marketing in the pharmaceutical industry. Carmen holds an MBA from INSEAD (Singapore) and has contributed to thought leadership in the space, co-authoring the CVA's "Next Step: Sustainable AI" white paper. She co-leads CVA's Sustainability Working Group and serves as a pro-bono consultant for EHF. Carmen's passion for AI was ignited by coordinating the scale-up of several AI algorithms in the pharmaceutical industry. More recently, she has been conducting her own experiments with decentralized computation for AI, exploring how distributed networks can make AI more accessible, transparent, and sustainable while maintaining strong data privacy.



LEILEI TANG

LeiLei Tang is a Medical Affairs consultant based in Switzerland. She holds a PhD in hematology from Radboud University Nijmegen and has over 10 years of experience in pharmaceutical and biotech industries, with expertise in rare diseases, vaccines, and medical devices. Her career spans strategic leadership roles at Takeda, where she served as Global Medical Lead in the Vaccine Business Unit managing dengue vaccine launch preparation, and at companies including Baxalta/Shire (now part of Takeda) where she led the AHEAD study—the largest real-world evidence cohort in hemophilia involving 23 countries. LeiLei is skilled in and passionate about digital innovation in healthcare, with a particular focus on AI-powered digital solutions for improving patient outcomes and streamlining medical affairs operations. She also volunteers as a pro-bono consultant for EHF.



NATALIA SOFIA

Natalia Sofia, is a pro-bono Data Ecosystem Lead at Ethers HealthData Foundation. She is a pharmacist with dual master's degrees in Drug Sciences and Toxicology from the University of Basel and specialized training in Digital Acquisition of Big Data from Columbia University. She began her career in retail pharmacy, building expertise in patient engagement and personalized care, before moving into strategic roles, where she led digital initiatives in clinical trials, data management, and patient access. She also contributed to the PharmaLedger project, gaining hands-on experience in applying blockchain to healthcare. Currently, as pro-bono Ecosystem Lead at EHF, Natalia leverages her expertise in blockchain and AI to drive innovation in the Web3 space, while actively supporting technology start-ups and blockchain communities.



STEPHANIE FUCHS

Stephanie Fuchs is a Swiss and Liechtenstein tax advisor and CEO of Stéphanie Fuchs Consulting. With over a decade of experience in tax advisory, she specializes in Swiss and Liechtenstein tax law for corporates and individuals, with particular expertise in blockchain and emerging technologies, cross-border structuring, and DAO setups. Stephanie holds an LL.M. in International Tax Law from the University of Zurich, an M.A. in Accounting and Finance from the University of St. Gallen, and a CAS in Applied Information Technology from ETH Zurich. Since 2017, she has been deeply engaged in the crypto industry and serves as a founding member of the Swiss Blockchain Federation's tax work group. She is a member of the Crypto Valley Association (CVA), Women in Web3 Switzerland (WIW3CH), and serves as an ambassador for sustainability in Web3.



THOMAS EGELHOF

Thomas Egelhof, is the Chief Radiologist and Member of the Management Board at Merian Iselin Clinic for Orthopaedics and Surgery in Basel, Switzerland. With more than 30 years of experience in healthcare, Thomas combines deep expertise in clinical practice and patient care with a strong passion for innovation. He is particularly fascinated by the potential of blockchain and AI to shape the future of medicine. As a bridge between traditional medicine and digital technologies, he actively contributes to the point where experience meets the future. His extensive background includes work and leadership roles at prominent Swiss healthcare institutions including University Hospital Basel and Merian Iselin Clinic Basel, where he developed expertise in interventional and diagnostic radiology and health care management while maintaining a focus on integrating emerging technologies into clinical workflows. Thomas is also an active member of CVA.

CONTRIBUTORS:

This report has been enriched by diverse perspectives from exceptional professionals across the healthcare and technology ecosystem. We extend sincere gratitude to six expert contributors who generously shared their insights through interviews featured in Chapter 1: Livio Francescucci (ICT Lead Strategic Projects, Hirslanden Group, Switzerland), Annie Axelle (Head of Programmatic Partnerships, AfyaRekod, Kenya), Richard Zhong (USA), Lars Münter (Denmark), Ülkü Cibik (MLL Legal AG, Switzerland), and Andreia de Bem Machado (Postdoctoral Researcher, Federal University of Santa Catarina, Brazil). Each will be introduced more extensively in their respective interview sections.

We also acknowledge the valuable contributions to the development of this report by Tjasa Zajc (Digital Health Expert, Host of Faces of Digital Health, Business Developer at Better), Michele Soavi (COO/Chief Sustainability Officer at ImpactScope) and Minahil Riaz (Graphic Designer). Special appreciation goes to Michele Soavi and colleagues from EHF/CVA who provided thoughtful peer review that significantly strengthened this report. Their commitment exemplifies the collaborative spirit essential to this emerging field.

DISCLAIMERS:

The views and opinions expressed in this interview are solely those of the experts and do not necessarily reflect the official position or policies of their employers. Due to confidentiality obligations, no proprietary or internal company information has been disclosed in this report. All insights shared are based on general professional experience in the relevant field and publicly available knowledge.

Inclusion of projects as either "case studies" or "bonus resources" does not constitute endorsement or recommendation by the authors or affiliated organizations. Case studies represent more mature implementations with documented outcomes, while bonus resources highlight emerging approaches at earlier development stages that carry inherent risks. References to digital tokens are purely descriptive and educational; they should not be construed as financial advice. Readers should conduct independent due diligence before engaging with decentralized AI platforms or digital assets.

METHODOLOGY:

This report, albeit of business nature, employs a rigorous mixed-methods approach combining systematic literature review, desk research, case study analysis, and expert interviews to provide comprehensive insights into dAI implementation in healthcare. The methodology is designed to bridge theoretical foundations with practical implementation realities in this rapidly evolving field, ensuring both academic rigor and actionable guidance for practitioners.

Our systematic literature review follows PRISMA 2020 guidelines, analyzing peer-reviewed studies from PubMed, Scopus, and Web of Science published between January 2023 and March 2025. The review, which will be published independently from this report, synthesizes research on federated learning, blockchain, DAOs, agentic AI, and other decentralized approaches to address three core research questions:

01

What are the primary decentralized AI architectures and use cases in healthcare?

02

What is the current implementation maturity across different settings?

03

What are the most significant benefits, challenges, best practices?

Using a multi-layered conceptual framework, data extraction systematically captures study characteristics, technical solutions, governance models, quantitative outcomes, and implementation challenges. Bibliometric analysis using VOSviewer complements the qualitative synthesis by mapping research trends, intellectual networks, and temporal evolution of key concepts within the dAI ecosystem.

Case studies were systematically selected to ensure geographic diversity (Europe, North America, Asia-Pacific, emerging markets), varied healthcare applications (imaging, drug discovery, clinical trials, operations), and different implementation maturity levels. Each case employs a standardized analytical framework examining: organization profile and challenge, technical solution architecture, business implementation approach, quantitative and qualitative outcomes, and key insights.

Complementary desk research analyzed industry reports, regulatory frameworks (GDPR, HIPAA, PIPL), technical documentation from open-source platforms (MONAI, Flower), and market intelligence to capture developments not yet reflected in peer-reviewed literature due to publication lag.

Six expert interviews with technical leaders, clinical stakeholders, legal experts, entrepreneurs, academics, and patient advocates provide practitioner perspectives and real-world insights unavailable in published sources. These semi-structured interviews explore firsthand experiences with decentralized AI, identify underexplored applications, and surface critical implementation considerations across different functional roles and geographic contexts.

To ensure quality and accuracy, the report underwent a peer review process by subject matter experts across technology, healthcare, legal, and policy domains. Multiple review cycles incorporating feedback from both internal contributors and external reviewers strengthened the technical accuracy, practical relevance, and balanced perspective of the final work.

CHAPTER 1: SETTING THE STAGE



Foreword by the author: Carmen Cucul, Blockchain & AI Healthcare Consultant, Ethers HealthData Foundation

My interest in AI's sustainability challenges began in 2023 while co-authoring the "[Next Step: Sustainable AI](#)" report for the Crypto Valley Association. That work explored how blockchain technology could address AI's ethical and environmental concerns through decentralized data ownership, privacy-preserving computation, and transparent governance. The RESTART framework we developed—emphasizing Restrainability, Effectiveness, Security, Transparency, Accessibility, Representativity, and Trust—revealed that decentralization was not merely a technical feature but a fundamental requirement for AI to serve healthcare's diverse global needs in a more sustainable fashion.

Two years later, the imperative for decentralization seems to have crystallized. Data centers' exponential energy consumption, the sovereign AI movement recognizing health data as strategic national assets, the maturation of federated learning from concept to implementation, and regulatory frameworks like the European Health Data Space have created unprecedented momentum.

My professional experience coordinating AI-driven disease prediction algorithms long before ChatGPT's emergence taught me critical lessons: privacy protection unlocks patient participation, democratized data access sparks local innovation, and collaborative approaches uncover patterns that centralized, siloed methods miss. My recent experiments with blockchain-based platforms for decentralized computation have reinforced these insights, demonstrating how distributed networks can make AI more accessible, transparent, and resilient, despite their early stage of development.

This chapter examines the convergent forces—regulatory, technological, economic, and societal—driving this paradigm shift. Through expert interviews spanning IT leaders, patient advocates, entrepreneurs, and legal experts across diverse geographies, we capture the multifaceted reality of implementing decentralized AI. This work represents our collective effort to reimagine AI systems that are more equitable, sustainable, and ultimately more effective at improving global health outcomes.

NOTE: The views expressed are those of the author and do not necessarily reflect those of the employer, Ethers HealthData Foundation.

1.1 CHAPTER OVERVIEW

Chapter 1 establishes the foundational context for understanding decentralized AI in healthcare by examining the critical shift from centralized to distributed approaches. We begin by analyzing the healthcare AI landscape, exploring why traditional centralized models—despite their computational power—face insurmountable limitations around data privacy, regulatory compliance, institutional silos, and equitable access. The chapter articulates how decentralization offers pathways to more robust, privacy-preserving, and collaborative AI systems that respect data sovereignty while enabling powerful model development across diverse patient populations.

The chapter maps the convergent forces driving this paradigm shift across four interconnected domains: technological maturation (federated learning frameworks, edge computing, blockchain infrastructure), regulatory evolution (GDPR, HIPAA, emerging AI regulations, European Health Data Space), economic pressures (unsustainable healthcare costs, workforce shortages, competitive collaboration imperatives), and societal expectations (patient data ownership, algorithmic transparency, equitable access demands). These forces are examined not as isolated trends but as synergistic enablers that have reached critical mass, making 2025 a pivotal moment for dAI adoption in healthcare settings worldwide.

To ground these concepts in practical reality, we present a comprehensive analysis of global leadership patterns and healthcare application domains where decentralized AI demonstrates measurable impact. Regional innovation landscapes reveal distinctive approaches: Europe's privacy first collaborative frameworks, North America's platform-driven commercial integration, Asia Pacific's sovereign AI implementations, Switzerland's unique "third way" combining national sovereignty with open-source transparency, and Africa's inclusive federated learning initiatives addressing resource constraints. Across medical imaging, drug discovery, clinical trials, remote patient monitoring, and specialized care networks, we document where decentralized approaches are moving from pilot projects to operational deployments.

This report is designed for healthcare decision-makers, AI product managers, health technology investors, and policymakers with intermediate AI familiarity—what we characterize as a 4/10 knowledge level. We assume readers understand foundational concepts like "machine learning," "data privacy," and "cloud computing" but may need structured frameworks for understanding decentralized approaches. Chapter 2 provides essential technical definitions and architectural foundations that inform the case studies in Chapters 3-5. Some sections, particularly in Chapter 2, delve into technical specifications necessary for evaluating implementation feasibility. Readers seeking lighter overviews may focus on chapter summaries and case study outcomes, while those requiring deeper technical understanding will find comprehensive architectural details throughout.

Our systematic literature review provides rigorous evidence mapping the decentralized AI ecosystem. Analyzing 87 peer-reviewed studies published between 2023-2025, we identified three technology pillars at different maturity levels: federated learning as the most mature architecture for privacy-preserving computation, blockchain functioning as trust-enabling infrastructure, and DAOs representing transformative yet nascent governance frameworks. The synthesis reveals key benefits driving adoption (enhanced privacy, broken silos, patient sovereignty) alongside persistent challenges (technical complexity, regulatory uncertainty, governance risks), while documenting best practices and regional innovation patterns that inform the case studies presented throughout this report.

The chapter concludes with expert interviews featuring six practitioners representing diverse functional perspectives—technical leadership, entrepreneurship, patient advocacy, public policy, legal expertise, and academic research—across multiple geographies. These conversations surface critical implementation considerations, underexplored opportunities, and practical challenges that complement the report's systematic literature review and case study analysis. Together, these components provide you with both the conceptual framework and real-world context necessary to understand decentralized AI's transformative potential and implementation realities in healthcare.

1.2 INTRODUCTION TO DECENTRALIZED AI IN HEALTHCARE

Healthcare stands at a critical juncture. The promise of artificial intelligence to transform patient care, accelerate research, and optimize clinical operations is immense—yet traditional centralized AI approaches face fundamental limitations that threaten to undermine this potential. Data privacy concerns, regulatory constraints, institutional silos, and the inherent complexity of healthcare information have become significant barriers to AI adoption. However, 2025 marks a pivotal year where these challenges are driving rapid innovation in decentralized approaches.

Healthcare was among the first industries to embrace artificial intelligence, driven by acute needs for improved diagnostic accuracy, operational efficiency, and personalized treatment approaches. From early rule-based expert systems in the 1970s to today's sophisticated deep learning models, healthcare has consistently been at the forefront of AI application. This evolution has accelerated dramatically: according to McKinsey (1), by the fourth quarter of 2024, 85% of healthcare leaders were exploring or had already adopted generative AI capabilities. Ambient listening (i.e. the passive monitoring and interpretation of conversations, images or sounds within an environment using AI) has recently become almost "table stakes" in healthcare settings.

Yet this rapid adoption has highlighted a crucial challenge: the "black box" nature of many AI systems makes them problematic for clinical contexts where transparency and explainability are essential. Clinicians and patients rightfully demand to understand how an AI system arrived at a specific diagnostic recommendation or treatment plan—a requirement that many centralized, monolithic AI approaches struggle to satisfy.

Decentralized AI represents a paradigm shift that fundamentally redistributes control, computation, and governance away from centralized entities toward distributed networks of stakeholders. Rather than aggregating sensitive patient data into centralized repositories—a practice increasingly restricted by regulations like GDPR, HIPAA, and China's PIPL—dAI operates on one of more of the following core principles: data remains within its original boundaries, computational processing occurs across distributed nodes, and decision-making authority is shared among network participants. Approaches such as federated learning (training AI models across distributed datasets without centralizing data), blockchain-based computation (dAI model computing on blockchain infrastructure), multi-agent local orchestrations (local coordination of autonomous agents) or DAOs (governance structures coordinating stakeholders through blockchain-based voting) enable AI systems to harness collective intelligence while preserving institutional autonomy, regulatory compliance, and stakeholder control.

The stakes could not be higher. Healthcare systems worldwide face unprecedented pressures:

- Aging population increasing demand for services
- Workforce shortages projected to reach 10 million health workers globally by 2030, according to WHO (2)
- Rising costs making current healthcare delivery models unsustainable
- Health inequities persisting across and within countries
- Administrative burden overwhelming clinicians, with studies showing up to half of doctors and nurses reporting burnout (3)
- Infrastructure strain as traditional centralized approaches struggle to scale effectively

Decentralized AI addresses these challenges by enabling:

- Privacy-preserving collaboration across institutions, regions, and borders
- Equitable access to AI capabilities, regardless of an organization's size or resources
- Localized control that respects cultural, regulatory, and institutional differences
- Reduced data duplication and fragmentation through federated approaches
- Enhanced model robustness by learning from diverse, real-world datasets
- Infrastructure resilience through distributed architectures that reduce single points of failure

However, decentralization also introduces meaningful challenges that organizations and data providers must navigate:

- Technical complexity requiring specialized expertise and infrastructure investment
- Coordination overhead across participating institutions with varying capabilities
- Performance considerations when balancing privacy preservation with model effectiveness
- Governance frameworks establishing sustainable collaborative relationships
- Resource disparities affecting equitable participation across different organization types
- Change management addressing clinical workflow integration and staff training
- Digital illiteracy preventing many individuals (especially elderly) to contribute their data

Despite these challenges, the significance of dAI extends beyond technical innovation. By fundamentally rethinking how healthcare data is shared, analyzed, and leveraged for AI, decentralization creates pathways to more equitable and sustainable healthcare systems. It enables powerful AI models to benefit from diverse patient populations while respecting data sovereignty 24 and privacy rights—bridging the gap between AI's theoretical potential and practical implementation in healthcare's complex landscape.

This report examines how decentralized AI is moving from theoretical promise to practical application in healthcare, through detailed case studies across Europe, North America, and Asia. By highlighting both successes and challenges in real-world implementations, we provide actionable insights for healthcare providers, technology developers, research institutions, and policymakers navigating this transformative approach.

REFERENCES

01

Pardo Martin, C., Lamb, J., Dahab, A., Jones, J. & Bhasker, S. (2025) 'Generative AI in healthcare: Current trends and future outlook', McKinsey & Company, 26 March. Available at: <https://www.mckinsey.com/industries/healthcare/our-insights/generative-ai-in-healthcare-current-trends-and-future-outlook>

02

Boniol, M. et al. (2022). "The global health workforce stock and distribution in 2020 and 2030: a threat to equity and 'universal' health coverage?" *BMJ Global Health*, 7(6):e009316

03

Chateau Health & Wellness (2025) 'Addressing Healthcare Worker Burnout in 2025'. Available at: <https://www.chateaufrecovery.com/addressing-healthcare-worker-burnout-2025>

1.3 CONVERGENT FORCES DRIVING AI DECENTRALIZATION

The current wave of AI decentralization in healthcare is not an isolated technological trend but rather the result of multiple convergent forces across technology, society, law, and industry that have reached a critical inflection point in 2025. Understanding these enabling factors provides essential context for why decentralized approaches are gaining rapid adoption, albeit still being in early stages of development.

TECHNOLOGICAL MATURATION: INFRASTRUCTURE READY FOR DECENTRALIZATION

Recent technological breakthroughs have finally made decentralized AI practically feasible at scale. The maturation of federated learning frameworks, edge computing capabilities, and blockchain infrastructure (see Section 2.2 for technical details) has created the technical foundation necessary for healthcare organizations to participate in collaborative AI development without sacrificing control or security.

Key technological enablers include:

- Edge computing proliferation: Hospital-grade processing power enabling local AI deployment without cloud dependency
- Standardized federated learning frameworks: Open-source platforms like MONAI and mature commercial solutions reducing implementation complexity
- Blockchain infrastructure maturity: Enterprise-ready decentralized computation networks offering substantial cost reductions compared to traditional cloud providers
- Interoperability advances: FHIR standards and API-driven integrations enabling seamless connection between heterogeneous healthcare systems

REGULATORY AND LEGAL ACCELERATION: PRIVACY LAWS MANDATING DECENTRALIZATION

The global proliferation of stringent data protection regulations has transformed decentralization from a technical preference into a legal necessity. Healthcare organizations worldwide must now navigate complex, jurisdiction-specific requirements that often prohibit or severely restrict traditional centralized data aggregation approaches.

Regulatory drivers include:

- GDPR's global influence: European data protection standards adopted worldwide, requiring data minimization and purpose limitation
- National sovereignty requirements: China's PIPL, India's data localization laws, and similar regulations mandating local data control
- Healthcare-specific protections: Enhanced requirements for medical data handling creating additional barriers to centralized approaches

SOCIETAL SHIFTS: TRUST, TRANSPARENCY, AND PATIENT EMPOWERMENT

Societal expectations around data control, algorithmic transparency, and patient autonomy have fundamentally shifted following high-profile data breaches and AI bias incidents. Patients and healthcare professionals increasingly demand systems that preserve individual control while enabling collective benefit—precisely the value proposition that decentralized AI delivers.

Social catalysts include:

- Post-pandemic privacy awareness: COVID-19 contact tracing and vaccine passport debates heightening public sensitivity to health data usage
- AI explainability demands: Healthcare professionals requiring transparent, auditable AI decisions that centralized "black box" models cannot provide
- Patient empowerment movements: Growing expectations for individual control over health data and treatment decisions
- Democratic governance expectations: Communities seeking participatory control over AI systems affecting their health outcomes

HEALTHCARE INDUSTRY PRESSURES: ECONOMIC AND OPERATIONAL NECESSITY

The healthcare industry faces unprecedented economic and operational pressures that make collaborative AI approaches not just attractive but essential for survival. Resource constraints, workforce shortages, and competitive dynamics are driving organizations toward shared solutions that would have been unthinkable under traditional business models.

Industry transformation factors include:

- Workforce crisis response: WHO's projected 10 million healthcare worker shortage by 2030 forcing automation and AI augmentation strategies
- Cost containment imperatives: Unsustainable healthcare spending growth requiring shared infrastructure and collaborative efficiency gains
- Competitive collaboration needs: Organizations recognizing that individual AI capabilities cannot match collective intelligence benefits
- Vendor lock-in concerns: Healthcare systems seeking alternatives to dependence on major cloud providers for critical AI infrastructure

While we recognize this is still the early days of AI decentralization—with many implementations remaining in pilot phases and adoption varying significantly across regions and organization types—the signs are encouraging. The case studies examined throughout this report demonstrate how forward-thinking organizations are capitalizing on this convergence to build competitive advantages while addressing fundamental healthcare challenges through collaborative, privacy-preserving AI systems.



1.4 GLOBAL LEADERSHIP IN DECENTRALIZED AI FOR HEALTHCARE

Understanding the current landscape of decentralized AI implementation requires examining both the geographic distribution of innovation and the specific healthcare applications where these approaches are demonstrating measurable impact. This dual perspective reveals how different regions are leveraging their unique strengths and regulatory environments to advance decentralized AI, while highlighting the healthcare domains where collaborative, privacy-preserving approaches are proving most effective.

REGIONAL LEADERSHIP: DISTINCTIVE APPROACHES TO DECENTRALIZED AI

The global landscape of decentralized AI in healthcare reveals distinct regional innovation patterns, each shaped by unique regulatory environments, technological capabilities, and healthcare system structures. Europe leads in privacy-preserving collaborative frameworks driven by GDPR compliance, North America focuses on scalable commercial platforms with sustainable business models, Asia Pacific emphasizes sovereign AI implementations within national boundaries, while emerging regions like Africa pioneer inclusive approaches addressing resource constraints. Switzerland has carved out a distinctive "third way" through fully open sovereign AI development, demonstrating how smaller nations can lead through transparency and international collaboration.

Europe: Privacy-First Collaborative Frameworks

Europe has emerged as the global leader in privacy-preserving healthcare AI, driven by GDPR, EU AI Act requirements and strong research networks. The European approach emphasizes cross-border collaboration within strict data protection frameworks, exemplified by the MELLODDY Consortium uniting 10 pharmaceutical companies, the European Health Data Space enabling data sharing across 27 EU member states, the MONAI open-source framework supporting 80+ research institutions globally, and the Personal Health Train initiative connecting 30+ European health repositories through "train-to-data" methodologies.

North America: Platform-Driven Innovation and Commercial Integration

North American implementations focus on scalable platform development and sustainable business models, combining advanced technical capabilities with commercial viability, rapid clinical integration and lighter oversight for most AI models. Notable developments include Mayo Clinic Platform's "Data Behind Glass" model connecting 25+ organizations for federated analysis of 30+ million patient records, Care.ai's edge computing deployment across 200+ hospitals using federated learning for operational AI improvement, Acoer's decentralized clinical trials platform transforming research through distributed data collection, and academic networks creating federated research capabilities across major medical institutions.

Asia-Pacific: Sovereign AI and Scale Implementation

Asian markets, particularly China and Singapore, emphasize national sovereignty and large scale deployment within domestic frameworks, creating distinctive models that prioritize local control alongside collaborative capabilities. Major implementations include China's DeepSeek deployment across 100+ tertiary hospitals representing the world's largest sovereign AI healthcare implementation, Singapore's AimSG Platform enabling cross institutional collaboration within sovereign boundaries, and Changi General Hospital's smart hospital approach combining AI, robotics and automation.

Switzerland: Open Sovereign AI Pioneer

Switzerland has pioneered a unique "third way" approach through the Swiss AI Initiative, releasing in 2025 Apertus, a fully open-source large language model trained on the "Alps" super computationr involving over 800 researchers across 10+ academic institutions with 20 million GPU hours annually. This represents the world's largest open science effort for AI foundation models, featuring full transparency with source code, model weights, and training data publicly available under Apache 2.0 license, multilingual capability supporting over 1,000 languages, 100% carbon-neutral training on nationally controlled infrastructure, and built-in compliance with Swiss data protection laws and EU AI Act requirements.

Emerging African Initiatives: Inclusive Federated Learning

Africa is developing innovative approaches to inclusive federated learning, with projects enabling collaboration across borders while accommodating resource constraints. Notable initiatives include fetal ultrasound screening programs across Algeria, Ghana, Egypt, Malawi, and Uganda using edge devices like Raspberry Pi, resource-adaptive frameworks designed for heterogeneous hardware capabilities, cross-border model training for improved population representation, mobile health integration leveraging widespread device adoption, and capacity building programs for minimally trained clinical workers to utilize AI assisted diagnostic tools.

HEALTHCARE APPLICATION LEADERSHIP: WHERE DECENTRALIZED AI DELIVERS IMPACT

Across healthcare domains, decentralized AI is demonstrating measurable impact in areas where traditional centralized approaches face fundamental limitations. Medical imaging leads in maturity with proven federated learning implementations for diagnostic collaboration, while remote patient monitoring leverages distributed sensor networks for continuous care management. Drug discovery and clinical research benefit from competitive collaboration models that enable knowledge sharing while protecting proprietary information, operational AI optimizes workflows across institutional networks, and specialized care networks address rare diseases requiring large patient cohorts that no single institution possesses.

MEDICAL IMAGING: COLLABORATIVE DIAGNOSTICS WITHOUT DATA SHARING

Medical imaging represents the most mature application area for decentralized AI, with proven implementations spanning brain imaging networks for tumor segmentation and Parkinson's disease detection, COVID-19 diagnostic collaboration through federated chest X ray analysis across global networks, cancer screening programs using federated breast imaging and dermatological AI systems, and specialized imaging for rare diseases requiring large, diverse datasets without data centralization.

OPERATIONAL AI: WORKFLOW OPTIMIZATION THROUGH DISTRIBUTED INTELLIGENCE

Healthcare operations are benefiting from federated learning approaches that optimize workflows, resource allocation, and administrative processes across institutional networks, including collaborative supply chain demand forecasting and inventory management, federated staffing and scheduling models predicting patient volume and resource needs, quality improvement initiatives sharing best practices without exposing institutional performance data, and administrative automation for collaborative development of billing, coding, and regulatory compliance tools.

DRUG DISCOVERY AND CLINICAL RESEARCH: COMPETITIVE COLLABORATION

Pharmaceutical research is leveraging decentralized approaches to accelerate discovery while protecting proprietary information, demonstrated by MELLODDY's training of predictive models on over 1 billion chemical activity labels across competing companies, decentralized clinical trials through platforms like Acoer enabling patient recruitment without centralized information management, federated biomarker discovery analyzing genomic and proteomic data across research institutions, and personalized medicine development through collaborative training on patient response data while maintaining privacy.

REMOTE PATIENT MONITORING: DISTRIBUTED INTELLIGENCE FOR CONTINUOUS CARE

Medical imaging represents the most mature application area for decentralized AI, with proven implementations spanning brain imaging networks for tumor segmentation and Parkinson's disease detection, COVID-19 diagnostic collaboration through federated chest X ray analysis across global networks, cancer screening programs using federated breast imaging and dermatological AI systems, and specialized imaging for rare diseases requiring large, diverse datasets without data centralization.

SPECIALIZED CARE NETWORKS: RARE DISEASES AND COMPLEX CONDITIONS

Decentralized AI is proving particularly valuable for rare diseases and complex conditions where individual institutions lack sufficient patient volumes for robust model development, enabling federated rare disease diagnosis through collaborative pattern recognition, specialized pediatric care networks training models across children's hospitals, multi institutional cancer care coordination for treatment planning and outcome prediction, and genetic counseling networks conducting federated analysis of genomic variants and family histories while maintaining patient privacy.

The convergence of regional innovation strengths with specific healthcare application needs is creating a dynamic global ecosystem where different geographic areas contribute unique capabilities while addressing shared healthcare challenges. This distributed innovation model suggests that the future of healthcare AI will be characterized not by a single dominant approach, but by a diverse ecosystem of specialized, collaborative, and sovereign solutions tailored to specific regional needs and healthcare applications.

1.5 SYSTEMATIC LITERATURE REVIEW: MAPPING THE DECENTRALIZED AI ECOSYSTEM

To establish a rigorous, evidence base for this business-orientated report, which is complementary to the case study analyses and expert interviews, we conducted a systematic literature review following PRISMA 2020 guidelines, analyzing peer-reviewed studies from PubMed, Scopus, and Web of Science published between January 2023 and March 2025.

Led by Antonio Pesqueira, this comprehensive review—which has been submitted for publishing as a separate academic paper—examined 87 eligible studies, with 13 papers selected for in-depth qualitative analysis based on architectural representativeness, implementation clarity, and intellectual contribution (see table 1 below). The analysis synthesizes current research on federated learning, blockchain integration, decentralized autonomous organizations (DAOs), and agentic AI to address 3 main objectives:

- map the technological landscape of decentralized AI initiatives in healthcare
- assess implementation maturity of these initiatives
- identify key benefits, challenges, and best practices shaping decentralized AI adoption in healthcare

These research objectives and findings will be revisited throughout the report: Chapter 2 addresses the architectural question, Chapters 3-5 demonstrate maturity levels through case studies, and Chapter 6 synthesizes benefits, challenges, and best practices into strategic recommendations.

THREE TECHNOLOGY PILLARS WITH VARYING MATURITY

The review reveals three distinct technology architectures with different levels of development:

- map the technological landscape of decentralized AI initiatives in healthcare
- assess implementation maturity of these initiatives
- identify key benefits, challenges, and best practices shaping decentralized AI adoption in healthcare

FEDERATED LEARNING

The most mature approach, serving as the vanguard architecture for privacy-preserving clinical computation with consistent application in medical imaging and electronic health record analysis. Bibliometric analyses confirm its centrality at the nexus of interoperability, edge computing, and explainable AI. Key challenge: statistical heterogeneity of non-IID (non-identically distributed) data, being addressed through advanced algorithms like FedProx and Scaffold.

BLOCKCHAIN

Functions primarily as trust-enabling infrastructure rather than an AI training mechanism, providing tamper-proof ledgers where auditability and integrity are paramount. Optimal applications: operational and administrative domains, particularly clinical trials management. Keyword clustering tightly links blockchain to security, digital identity challenges, and trust deficits in multi-stakeholder health data exchange.

DAO

Decentralized Autonomous Organisations (DAOs) - The most transformative yet nascent architecture, representing a sociotechnical framework for reimagining collaboration. Applications include governing civic medical data trusts and funding decentralized science (DeSci) initiatives. However, socio-economic and legal frameworks necessary for functional healthcare DAOs remain in their infancy.

BENEFITS DRIVING ADOPTION VS. IMPLEMENTATION CHALLENGES

Our synthesis reveals a clear dichotomy between pull factors and barriers. The primary benefits driving adoption include enhanced privacy through local data processing, trust and security via transparent audit trails, patient data sovereignty and control, the breaking down of institutional data silos, and more generalizable models trained on diverse datasets. These advantages address fundamental problems: privacy regulations preventing traditional data aggregation, institutional silos limiting collaborative research, and centralized systems that concentrate control with attendant misuse risks.

Conversely, implementation faces substantial challenges spanning technical, organizational, and governance domains. Infrastructural limitations and integration complexity with existing hospital systems present immediate hurdles, while regulatory uncertainties and the evolving legal landscape—particularly GDPR's "Right to be Forgotten" provisions—create compliance concerns. Blockchain implementations struggle with scalability issues and high computational costs. DAOs face the most complex socio-technical challenges: governance capture by large token-holders, low member participation (voter apathy), unclear legal status across jurisdictions, and smart contract vulnerabilities. The temporal evolution in our bibliometric analysis shows research shifting from defining broad challenges (2019-2022) to investigating specific implementation issues, with pronounced emphasis on data privacy and integrating generative AI into decentralized frameworks.

BEST PRACTICES BY TECHNOLOGY PILLAR

Beyond identifying benefits and challenges, our literature review provides actionable implementation guidance derived from documented case studies and expert experience. These best practices reflect the current state of knowledge for organizations embarking on dAI implementations, with recommendations tailored to the specific characteristics and maturity levels of each technology pillar. Some of our learnings are reflected below under the form of a checklist.

TECHNOLOGY PILLAR	IMPLEMENTATION BEST PRACTICES
Federated Learning	<ul style="list-style-type: none"> □ Use formal implementation checklists □ Integrate privacy-enhancing technologies (Differential Privacy, Secure Multi-Party Computation) □ Benchmark on diverse, real-world datasets □ Deploy advanced algorithms for non-IID data challenges
Blockchain Applications	<ul style="list-style-type: none"> □ Start with non-critical applications like consent tracking before clinical decision support □ Ensure interoperability with existing EMR systems □ Address regulatory frameworks proactively (GDPR compliance)
DAO Governance	<ul style="list-style-type: none"> □ Implement delegative democracy (liquid democracy) rather than simple token voting □ Use hybrid models: community participation + expert review councils □ Establish clear "constitutions" (legal agreements or smart contract encoded rules) □ Support on-chain processes with robust off-chain community discussion

REGIONAL INNOVATION PATTERNS

Our geographic distribution analysis examined citation patterns, institutional affiliations, and documented implementations across the literature corpus to map how different regions approach decentralized AI in healthcare. This analysis reveals that innovation is not uniformly distributed but rather shaped by distinctive regulatory environments, technological capabilities, and healthcare system structures. The resulting patterns demonstrate that decentralized AI development is evolving along multiple parallel trajectories, each leveraging unique regional strengths:

- Europe: Privacy-preserving collaborative frameworks driven by GDPR (MELLODDY consortium, European Health Data Space)
- North America: Platform-driven innovation with commercial viability (Mayo Clinic Platform, Care.ai deployments)
- Asia-Pacific: Large-scale sovereign AI implementations (China's DeepSeek across 100+ hospitals, Singapore's national programs)
- Switzerland: "Third way" model—Apertus sovereign LLM that is nationally controlled yet globally accessible, reconciling data sovereignty with collaborative innovation
- Emerging markets: Inclusive approaches addressing resource constraints (African federated learning initiatives)

This distributed innovation landscape suggests the future of healthcare AI will be characterized not by a single dominant approach but by a diverse ecosystem of specialized, collaborative, and sovereign solutions tailored to specific regional needs and healthcare applications.



TABLE 1: ANALYSIS OF BENEFITS, CHALLENGES AND RISKS

AUTHOR, YEAR	REPORTED BENEFITS	REPORTED CHALLENGES & RISKS
Teo et al. (2024) (Review)	Data Privacy: Enables model training without sharing raw patient data. Access to Diverse Data: Breaks down data silos, leading to more generalizable and robust models.	Statistical Heterogeneity (Non-IID data), high communication costs, system heterogeneity (hardware/software variance), vulnerability to model poisoning attacks.
Myrzashova et al. (2023) (Review)	Enhanced Security & Trust: Blockchain provides an immutable audit trail of model updates and data access. Incentivization: Smart contracts can automate rewards for data contributors.	Scalability issues of blockchain, high computational/energy cost of PoW, and integration complexity between FL and blockchain systems.
Al-Marridi et al. (2024)	Decentralized Decision-Making: Agents learn optimal policies without a central controller. Security & Auditability: Blockchain records all actions and decisions transparently.	Simulation may not capture real world complexity. High setup complexity for MARL systems.
Puri et al. (2024)	Improved data security for sensitive IoT data, enhanced privacy for remote monitoring, and created a trusted ecosystem for data sharing.	Interoperability of diverse IoT devices, data quality from consumer-grade sensors, and managing network latency.
Li et al. (2025) (Review)	Privacy preservation is the primary benefit. Also cites improved fairness and reduced bias by including data from diverse populations.	"Data-centric" challenges (non-IID, missing data) and "Model-centric" challenges (communication bottlenecks, security). Highlights "pitfalls" like data leakage risks.
Castro et al. (2024) (Review)	Enhanced Data Integrity & Transparency: Immutable record of trial data and protocol adherence. Patient Empowerment: Patients can control access to their data via smart contracts.	Interoperability with existing hospital EMR systems, regulatory uncertainty (e.g., FDA/EMA acceptance), and data privacy concerns if not designed correctly (Right to be Forgotten).
Cunningham et al. (2024)	Democratic Control: Gives citizens a direct voice in how their data is used. Transparency: All governance decisions are on a public ledger.	Legal & Regulatory Uncertainty: DAOs lack a clear legal status. Scalability of Governance: Ensuring efficient decision-making with many participants.

Savioz et al. (2025)	Access to Capital: Opens new, global funding avenues. Community Ownership: Allows patients/researchers to own a stake in the research they support.	Valuation of IP: Difficulty in accurately valuing early-stage research. Regulatory Compliance (securities law).
Saurabh et al. (2024)	Automation of Governance: Reduces administrative overhead. Censorship Resistance: Decisions cannot be easily overturned by a central party.	Plutocracy: Risk of governance being dominated by large token holders. Smart Contract Vulnerabilities.
Ellinger et al. (2024)	Resilience & Decentralization: Has operated for years without central control, allowing community engagement.	Centralization Risks from early token distributions. High Complexity for new users. Regulatory Scrutiny.
Hue et al. (2024)	Provides a structured language for describing DAOs, making them more understandable to enterprise architects and IT managers.	Highlights a cultural and methodological gap between formal enterprise structures and the fluid nature of DAOs.
Farry (2025)	Global & Inclusive Governance: Allows members worldwide to participate. Transparent Funding: Donors can track fund usage on chain.	Member Apathy: Ensuring sustained participation from a non-financially motivated membership base. Technical Barriers for users.
Ly & Shojaei (2025)	Transparent & Efficient Management: Automating tasks like maintenance requests and access control. Shared Ownership of physical assets.	Legal Integration: Connecting on chain ownership to real-world legal titles. IoT Security: Ensuring the integrity of sensor data.

Source: The Decentralized AI Ecosystem in Healthcare: Systematic Review of Technologies, Governance, and Implementation, Antonio Pesqueira et al, 2025 (submitted for publication)

1.6 EXPERT INTERVIEWS

To complement our systematic literature review and case study analysis, we convened a virtual roundtable of six practitioners who bring diverse perspectives from across the decentralized AI ecosystem. These experts represent different functional roles, geographic contexts, and stages of implementation—from entrepreneurial ventures pioneering patient data sovereignty in Africa to established healthcare IT leaders navigating regulatory complexities in Switzerland, from young researchers exploring AI safety at leading universities to seasoned policy advocates shaping civil society engagement with health data governance.



Livio Francescucci

ICT Lead Strategic Projects,
Hirslan Group (Switzerland)



Annie Axelle

Head of Programmatic
Partnerships, AfyaRekod (Kenya)



Andreia de Bem Machado

Postdoctoral Researcher, Federal
University of Santa Catarina
(Brazil)



Richard Zhong

Master's Student, NYU Tandon
School of Engineering; EHF
(USA)



Lars Münter

Public Policy and Patient Advocacy
Expert (Denmark)



Ülkü Cibik

Legal Expert, MLL Legal AG
(Switzerland)

Through semi-structured interviews, these experts share firsthand experiences with decentralized AI technologies, identify underexplored applications within their respective contexts, and surface critical implementation considerations—from technical integration challenges to governance frameworks, from regulatory compliance to cultural change management. Their diverse voices provide the practical wisdom and contextual nuance that complement this report's research foundation, offering you actionable insights grounded in real-world implementation across varied healthcare systems and organizational contexts.



INTERVIEW #1

Interview with Livio Francescucci, ICT Lead Strategic Projects, Hirslanden Group (Switzerland)

Interviewed by: Carmen Cucul

From your unique professional and personal background, what firsthand experience do you have with decentralized AI technologies in healthcare? What drives your interest in this emerging field?

Livio: Over the past few years working in IT and most recently in healthcare IT, I've seen how difficult it can be to implement innovative technologies like decentralized AI — not because the ideas aren't promising, but because integration and interoperability are still major roadblocks in many real-world settings. Of course, given the challenges, this makes the technological space even more interesting. Personally, I believe in a future where patients can control their own data — where their entire medical history is securely encrypted and stored on a blockchain, and access is granted only with their private key. Whether they're visiting a hospital, a clinic, or a specialist abroad, their records could be instantly available — but always under their control. It's an ambitious vision, but one that could change the way we think about data ownership and access in healthcare.

Based on your functional expertise and regional healthcare landscape, which specific healthcare applications do you believe would benefit most from decentralized AI approaches— particularly those that may be underexplored or emerging?

Livio: One clear opportunity I see is in enabling cross-site research and learning without centralizing data. Right now, we often have multiple research institutions requesting similar types of data — for tumor boards, quality reviews, population studies — but the process is slow and fragmented, and sometimes not even feasible. A decentralized, blockchain-based approach could allow us to share anonymized insights securely, while keeping the data protected and distributed. This would be especially useful in Switzerland, where the healthcare landscape is very diverse and decentralized by nature. If we could break down the technical and governance barriers, we could unlock a lot of untapped potential.

What are the most critical considerations, challenges, or requirements that stakeholders in your position must address to successfully implement decentralized AI in healthcare settings?

Livio: Even if the technology is ready, the real challenge is often elsewhere. Legal and regulatory frameworks still lag behind what's needed to adopt decentralized models — and in some cantons in Switzerland, the use of cloud services for storing patient data is still highly restricted or not clearly allowed. That alone can make it nearly impossible to move forward. On top of that, we also face cultural resistance. In Switzerland, medical practices and IT maturity vary widely between regions, and even between hospitals. Clinicians are often trained with different systems, tools, and workflows — which makes standardization incredibly difficult. All of this contributes to a landscape that's not very friendly to new, cross-cutting technologies — at least not yet.

Livio Francescucci serves as ICT Lead Strategic Projects at Hirslanden Group, Switzerland's largest private hospital network. An experienced technology executive with a background spanning software engineering leadership at organizations like EF Education First and Ferrari's racing team, Livio brings deep expertise in building engineering organizations and navigating the practical challenges of implementing innovative technologies within complex healthcare environments. His insights reflect the realities of integrating decentralized AI within Switzerland's highly decentralized and regulated healthcare landscape. LinkedIn profile: <https://www.linkedin.com/in/livio-francescucci-b379042/>

Disclaimer: The views and opinions expressed in this interview are solely those of the interviewee and do not reflect the official position or policies of their employer. Due to confidentiality obligations, no proprietary or internal company information has been disclosed. All insights shared are based on general professional experience in the relevant field and publicly available knowledge.

INTERVIEW #2

Interview with Annie Axelle, Head of Programmatic Partnerships, AfyaRekod (Kenya)

Interviewed by: Carmen Cucul

From your unique professional and personal background, what firsthand experience do you have with decentralized AI technologies in healthcare? What drives your interest in this emerging field?

Annie: My journey with AfyaRekod started because of lack of mobility of health-data across systems, countries and various ecosystems. It has been rooted in addressing one of Africa's biggest healthcare challenges: fragmented, inaccessible, and often unreliable patient health records. Professionally, I have worked at the intersection of AI, blockchain, and health systems (myrekod.com) designing platforms that empower patients to own their data while enabling providers and governments to leverage AI for predictive and preventive care.

First-hand, I have seen how decentralised AI can bridge gaps where centralised systems have failed particularly in contexts with poor infrastructure, low trust in institutions, and rising demand for patient-centric care. This means giving the patient access to their health-record and also allow them to be central-decentralized and become the collectors and mobility framework of a connected health data ecosystem. Personally, my motivation comes from the belief that health data sovereignty is a human right. Families should not lose their medical history when they move from one hospital to another or across borders. Decentralized, AI-driven healthcare allows us to give patients true control while improving outcomes at scale.

Based on your functional expertise and regional healthcare landscape, which specific healthcare applications do you believe would benefit most from decentralized AI approaches— particularly those that may be underexplored or emerging?

Annie: From the African healthcare landscape, I see few high-impact applications for decentralised AI that remain under-explored:

- Chronic Disease Management: AI-driven decentralised patient registries for conditions like diabetes, hypertension, and HIV. With a blockchain for trust and immutability, patients can maintain lifelong care records accessible anywhere, enabling continuity of care and predictive insights.
- Maternal & Child Health: Decentralized AI platforms can track maternal journeys and neonatal care, aggregating data across fragmented facilities to detect risks early (e.g., pre eclampsia or malnutrition).
- Cross-Border Health Portability: In regions like East Africa, patients often migrate for work. A decentralised health wallet powered by blockchain ensures their health history is portable across hospitals, insurers, and even countries.
- Community-Driven Public Health Intelligence: Local clinics and NGOs can plug into federated AI models without ceding raw data. This helps governments detect outbreaks or patterns (malaria, pandemics, environmental health risks) in real time without compromising privacy.

What are the most critical considerations, challenges, or requirements that stakeholders in your position must address to successfully implement decentralized AI in healthcare settings?

Annie: To implement decentralized AI in healthcare successfully, several aspects must be addressed:

Firstly, data transformation/digitization & standardization. Many facilities still operate on paper or fragmented digital systems. Building decentralised AI requires digitization, data transformation and building mobility trust networks. The human challenge and lack of knowledge has been a big problem in this space.

Secondly, what about trust, privacy & regulation? Blockchain gives transparency, but adoption depends on patient trust and regulatory alignment. Stakeholders must balance data sovereignty with AI innovation, governments, hospitals, and communities must all be part of governance frameworks.

Thirdly, infrastructure & accessibility. Decentralized systems can run lightweight on mobile devices, but rural connectivity and affordability are still barriers across Africa. Solutions must be designed for low-resource environments and function offline with periodic synchronisation.

Fourthly, let's not forget about human capacity & adoption. Clinicians, nurses, and administrators need training to interact with AI and blockchain tools. Success is not just about building tech but embedding it into day-to-day workflows, as we have experienced firsthand at AfyaRekod.

And lastly, funding and long term approach are not to be underestimated, either. For long-term adoption, decentralised AI platforms must have R&D funding to experiment and also create truly tested networks as well as align with financial incentives of insurers, governments, and donors. For example, predictive insights can reduce claims costs for insurers, making them natural partners.

Annie Axelle is Head of Programmatic Partnerships at AfyaRekod, a Kenyan health-tech startup that has built a blockchain-driven Universal Patient Portal serving over 150,000 users across Kenya, Nigeria, South Africa, Tanzania, and Zambia. Founded in 2019 as an Adanian Labs venture, AfyaRekod addresses Africa's critical challenge of fragmented, inaccessible patient health records by empowering patients to own and control their medical data while enabling AI-driven predictive and preventive care. Annie's perspective illuminates how decentralized approaches can solve fundamental infrastructure gaps in emerging markets while advancing patient-centric care models. LinkedIn profile: <https://www.linkedin.com/in/annie-axelle-b-71b488184/>

Disclaimer: The views and opinions expressed in this interview are solely those of the interviewee and do not reflect the official position or policies of their employer. Due to confidentiality obligations, no proprietary or internal company information has been disclosed. All insights shared are based on general professional experience in the relevant field and publicly available knowledge.

INTERVIEW #3



Interview with: Andreia de Bem Machado, Postdoctoral Researcher, Department of Knowledge Engineering Federal University of Santa Catarina Research Center for Intelligence, Management and Technology for Innovation - IGTI (Brazil)

Interviewed by: Antonio Pesqueira

From your unique professional and personal background, what firsthand experience do you have with decentralized AI technologies in healthcare? What drives your interest in this emerging field?

Andreia: My background integrates engineering and knowledge management, innovation, and public policy, with academic work and educational leadership. In recent years, I have focused on the intersections of AI governance, data privacy, and health/life sciences, publishing articles and chapters that engage directly with distributed architectures and decentralized technologies. My involvement in research centers such as IGTI reinforces an interdisciplinary approach to IT innovation applied to management and services, including healthcare.

This experience—combined with leadership roles and projects on human rights and urban technologies—strengthens my commitment to decentralized AI solutions that preserve data sovereignty, auditability, and transparency. In sum, my experience combines scholarly output and applied projects in privacy, governance, blockchain, and distributed systems design in the context of health and life sciences, motivated by social impact, patient safety, and ethics in the sector’s digital transformation.

Based on your functional expertise and regional healthcare landscape, which specific healthcare applications do you believe would benefit most from decentralized AI approaches — particularly those that may be underexplored or emerging?

Andreia: Drawing on my academic experience and committee work in Health Management in Santa Catarina, Brazil, I see the greatest gains with decentralized AI in:

- Pharmaceutical traceability and cold chain (blockchain + IoT at the edge) to reduce diversion, counterfeiting, and stockouts;
- Data governance in mental health and primary care, with granular consent, self-sovereign identity, and audit trails;
- Federated learning for diagnostic support (imaging, risk triage) among mid-size hospitals, preserving privacy;
- Home monitoring with embedded/edge AI for chronic conditions and ASD, prioritizing data protection and interoperability;
- Verifiable credentials for qualification and continuing education of healthcare professionals.

These applications align with my work on blockchain/AI and sustainability, and with committee experience in pharmaceutical assistance, medication transport, and litigation regarding access—indicating regulatory and logistical challenges that decentralization can mitigate. They favor continuous auditability, cost reduction, and secure scalability across heterogeneous networks.

What are the most critical considerations, challenges, or requirements that stakeholders in your position must address to successfully implement decentralized AI in healthcare settings?

Andreia: Implementing decentralized AI in healthcare requires addressing governance, technical, and adoption dimensions simultaneously. Critical points in my view are:

- Regulatory and ethics compliance: LGPD/ANPD, ANVISA, CEP/IRB; preparation of DPIA, legal bases, granular consent, and algorithmic risk management (bias, explainability, accountability).
- Interoperability and data quality: HL7 FHIR/openEHR, clinical taxonomies, metadata catalogs and stewardship; data use agreements (DUAs)
- Architecture and security: cloud-to-edge, federated learning, end-to-end encryption, key management/PKI, zero trust, decentralized identity (DID/VC), and audit trails
- Clinical MLOps: versioning, drift monitoring, post-deployment revalidation, and independent audits

- Procurement and sustainability: public purchasing with interoperability/portability criteria (avoiding lock-in), TCO/ROI, and consortial governance models
- Care integration: workflow redesign, patient safety, multicenter pilots with clinical outcomes, cost-effectiveness, and equity
- Capacity building and cultural change: ongoing training for clinical/managerial teams, data literacy, and engagement mechanisms with users and local councils.

Andreia de Bem Machado is a Postdoctoral Researcher at the Federal University of Santa Catarina's Research Center for Intelligence, Management and Technology for Innovation (IGTI) in Brazil. Her interdisciplinary work integrates engineering, knowledge management, innovation, and public policy, with recent focus on AI governance, data privacy, and distributed architectures in health and life sciences. Andreia's academic research provides theoretical grounding and comparative international perspectives on how decentralized technologies can preserve data sovereignty, auditability, and transparency while enabling collaborative innovation. LinkedIn profile: <https://www.linkedin.com/in/andreiadebem/>

Disclaimer: The views and opinions expressed in this interview are solely those of the interviewee and do not reflect the official position or policies of their employer. Due to confidentiality obligations, no proprietary or internal company information has been disclosed. All insights shared are based on general professional experience in the relevant field and publicly available knowledge.

INTERVIEW #4



Interview with Richard Zhong, Master's Student NYU and Young Ambassador, Ethers HealthData Foundation (USA)

Interviewed by: Carmen Cucul

From your unique academic and personal background, what firsthand experience do you have with decentralized AI technologies in healthcare? What drives your interest in this emerging field?

Richard: As a master's student researching in the AI field, I am interested in seeing the new computational capabilities of AI models being used to accelerate healthcare research. Whether in drug development or in population health policies, AI can assist researchers in handling massive amounts of health data and extracting insights, ultimately getting patients higher quality care quicker. While I was in the Young Ambassador program at Ethers HealthData Foundation, I had the opportunity to research the problem of unused data silos of patient health data and the challenges preventing its use in healthcare research.

Based on your functional expertise and regional healthcare landscape, which specific healthcare applications do you believe would benefit most from decentralized AI approaches— particularly those that may be underexplored or emerging?

Richard: Decentralized AI has the potential to address critical gaps in healthcare, particularly in areas where privacy and collaboration are essential. Rare disease diagnosis stands out as a promising application to me: federated learning allows hospitals to train models on local patient data without sharing raw information, enabling more accurate and timely diagnoses across dispersed populations. Decentralized clinical trials are another underexplored area, where blockchain can ensure data integrity and patient consent while federated models improve inclusivity by analyzing diverse, real-world data securely. Patients themselves can benefit from decentralized identities, giving them greater control over how their information is shared and used.

What are the most critical considerations, challenges, or requirements that stakeholders in your position must address to successfully implement decentralized AI in healthcare settings?

Richard: A critical factor for implementing decentralized AI in healthcare is ensuring that data privacy is not only protected but also perceived as trustworthy by the public. While techniques like federated learning and differential privacy reduce the need to centralize sensitive health data, the black-box nature of machine learning models introduces risk: information encoded in model weights can, in theory, be reverse-engineered to reveal patient details. Robust safeguards must be developed to guarantee that this leakage cannot occur, alongside clear regulatory standards that hold systems accountable. Equally important is building public confidence. Without widespread trust that decentralized AI protects individual privacy, patients (especially those from my generation) may hesitate to share their data, leaving models underpowered and ineffective.

Richard Zhong is a Master's student at NYU Tandon School of Engineering and volunteer with Ethers HealthData Foundation, focusing on improving transformer-based models' perception and reasoning capabilities while researching data privacy issues in healthcare transitions with American University. Representing the next generation of AI researchers, Richard is deeply engaged with AI safety, alignment to human values, and healthcare AI's transformative potential. His perspective as a young ambassador bridges technical AI research with practical healthcare applications and ethical considerations. LinkedIn Profile: <https://www.linkedin.com/in/richard-zhong/>

Disclaimer: The views and opinions expressed in this interview are solely those of the interviewee and do not reflect the official position or policies of their school. All insights shared are based on relevant professional and academic experience of the interviewee.

INTERVIEW #5



Interview with Lars Münter, Public Policy and Patient Advocacy Expert (Denmark)

Interviewed by: Carmen Cucul

From your unique professional and personal background, what firsthand experience do you have with decentralized AI technologies in healthcare? What drives your interest in this emerging field?

Lars: From a civil society perspective, citizens should care about health data in the same way that they care about how their tax money is being spent—this is a common good so not worrying about it would be negligent. My primary interest lies in ensuring that when we deploy AI, that the entire purpose should be helping people stay healthy. If we are using AI to optimize knee surgeries by 50%, we are approaching this incorrectly. The objective should be to reduce the need for knee surgeries by 100%. We live in a world shaped by generations of suboptimal health policy decisions, and we now have sophisticated tools at our disposal to address some of those systemic design flaws—if we choose to leverage decentralized approaches effectively. The fundamental question is whether we will use AI to fix underlying determinants of health or merely optimize flawed processes.

Based on your functional expertise and regional healthcare landscape, which specific healthcare applications do you believe would benefit most from decentralized AI approaches— particularly those that may be underexplored or emerging?

Lars: Decentralization fundamentally improves equity and innovation by democratizing access to AI benefits. Centralized systems inherently concentrate control, creating risks of narrow perspectives and limited stakeholder representation. Decentralized systems foster collaborative innovation by engaging more diverse contributors across different contexts. A promising development is the emergence of multilingual AI models, like the new Swiss LLM that incorporates diverse languages—it makes little sense to deploy advanced technology while constraining it to analyze data through such a narrow lens like one language alone. The healthcare AI ecosystem should emulate successful decentralized models - with strong joint standards - like the global financial system, ensuring seamless interoperability while preventing linguistic and cultural bias from undermining system effectiveness.

What are the most critical considerations, challenges, or requirements that stakeholders in your position must address to successfully implement decentralized AI in healthcare settings?

Lars: The most critical challenge is not the AI bias—it is the human bias embedded in prompt engineering. As we decentralize AI systems, we increase the risk of introducing varied biases unless we simultaneously improve AI literacy across all stakeholders. Everyone involved should develop sophisticated prompting skills, not only to achieve higher quality outputs but also to contribute more diverse, high-quality insights that AI models can learn from real-world applications. Another factor is interoperability, which remains absolutely essential—without effective system communication, decentralized approaches cannot realize their potential benefits. We also face the substantial risk of creating fragmented information ecosystems where conflicting (or straightaway fabricated) narratives undermine shared understanding, potentially contaminating datasets with unreliable information. Lastly, the fundamental challenge lies not in technological limitations but in stakeholders' willingness to act decisively on AI-generated insights—many systems could identify root causes of inefficiencies and inequities, yet organizational and political commitment to implement these findings often remains insufficient.

Lars Münter is a European leader in health literacy and wellbeing. As International Director of the Nordic Wellbeing Academy, he advances practical initiatives and systems change, including serving as Communications Lead for the Nordic Health 2030 Movement and active member of the European Health Futures Forum. For over 20 years, he has built cross-border platforms in Denmark, the Nordics, and Europe to promote self-care, empowerment, and citizen action. By bridging communication and science, Lars fosters knowledge sharing among stakeholders and sectors, embedding corporate social responsibility and collaborative action at the heart of health, innovation, and sustainable transformation. LinkedIn profile: <https://www.linkedin.com/in/lars-muenter/>

Disclaimer: The views and opinions expressed in this interview are solely those of the interviewee and do not reflect the official position or policies of their employer. Due to confidentiality obligations, no proprietary or internal company information has been disclosed. All insights shared are based on general professional experience in the relevant field and publicly available knowledge.



INTERVIEW #6

Interview with Ülkü Cibik, MLL Legal AG (Switzerland)

Interviewed by: Stephanie Fuchs

From your unique academic and personal background, what firsthand experience do you have with decentralized AI technologies in healthcare? What drives your interest in this emerging field?

Ülkü: From my perspective the emergence of decentralized AI technologies in healthcare represents a fascinating convergence of innovation, governance, and risk management. What resonates most with me is how decentralized architectures can potentially reframe questions of accountability, data ownership, and value distribution, topics that are deeply embedded in my daily legal practice.

In traditional corporate and financial structures, value flows and liability chains are clearly defined through centralized entities. Decentralized AI, however, challenges this model by distributing both decision-making and data processing across networks of participants. In the healthcare context, this creates new opportunities for patient empowerment and data sovereignty, but it also complicates the allocation of legal responsibility and oversight.

Personally, I find this shift compelling because it requires rethinking existing regulatory and contractual frameworks, similar to how structured finance once pushed the boundaries of legal innovation in capital markets. The rise of decentralized AI invites legal professionals to engage in shaping governance models that balance technological autonomy with the rule of law, patient protection, and ethical accountability. It also raises questions about sustainable design and the potential of decentralized infrastructure to reduce systemic risk in sensitive sectors like healthcare.

Based on your functional expertise and regional healthcare landscape, which specific healthcare applications do you believe would benefit most from decentralized AI approaches—particularly those that may be underexplored or emerging?

Ülkü: From a Swiss legal and regulatory perspective, several areas in healthcare appear particularly well suited to decentralized AI approaches. First, medical data management and interoperability present a significant opportunity. Decentralized AI systems, operating on blockchain or distributed ledger infrastructures, could enable secure and privacy-preserving sharing of patient data between healthcare institutions, both within Switzerland and across borders. This is particularly relevant in light of the Swiss Federal Act on Data Protection (FADP), which imposes strict requirements on data processing, localization, and informed consent. To comply with these obligations and with guidance from the Federal Data Protection and Information Commissioner (FDPIC), decentralized systems must implement robust anonymization and pseudonymization techniques to protect patient identities while enabling data utility.

Second, clinical research and pharmaceutical development in Switzerland could benefit substantially from decentralized AI frameworks. Through federated learning, research institutions and hospitals could contribute to model training without transferring raw data, thus maintaining control over sensitive information. This approach aligns with the ethical and legal standards of the Human Research Act (HRA) and oversight by swissethics and cantonal ethics committees. Given that both Swiss and European data protection regimes limit centralized data pooling, decentralized models offer a legally compliant way to accelerate drug discovery and innovation while respecting participants' privacy and consent rights.

Third, insurance and healthcare financing could evolve through the introduction of smart contract-based solutions that automate claims processing, verification, and fraud detection using decentralized AI analytics. However, this raises complex legal issues regarding the validity and enforceability of automated contracts, the allocation of liability for algorithmic or autonomous decisions, and the classification of such systems under the Insurance Supervision Act (ISA) and the supervision of the Swiss Financial Market Supervisory Authority (FINMA).

From a broader regulatory standpoint, the intersection between medical device law, AI governance, and decentralized architectures requires further examination. Switzerland has largely aligned its Medical Devices Ordinance (MedDO) with the EU Medical Device Regulation (MDR) and closely follows the ongoing development of the EU AI Act. However, both frameworks are built on assumptions of centralized accountability and traceability. Decentralized AI systems challenge these assumptions and call for adaptive legal thinking around concepts such as shared accountability, auditability in distributed environments, and the definition of a "provider" in decentralized contexts.

What are the most critical considerations, challenges, or requirements that stakeholders in your position must address to successfully implement decentralized AI in healthcare settings?

Ülkü: To enable responsible and compliant implementation of decentralized AI systems in healthcare, several key legal, ethical, and governance factors must be addressed.

First, there must be clarity on liability allocation. Decentralized systems inherently distribute functions among multiple actors, developers, data contributors, validators, and users, making it difficult to assign responsibility for harms or errors. Legal frameworks should consider hybrid models of collective accountability, possibly using multi-party agreements or DAO-style governance mechanisms that define decision-making rights and obligations.

Second, data governance and patient consent require careful design. Decentralized AI must comply with GDPR and related health data regulations, which demand clear identification of data controllers and processors. Novel concepts such as self-sovereign identity and on-chain consent management could help operationalize compliance, but they also require legal validation and interoperability with existing eIDAS and health data frameworks.

Third, ethical transparency and auditability are crucial. In healthcare, algorithmic explainability is not only a technical issue but a legal and ethical necessity. Decentralized AI models must incorporate verifiable audit trails, potentially via blockchain, allowing regulators and patients to trace data provenance and model updates.

Fourth, governance and sustainability must be integrated from the outset. Governance models should ensure inclusiveness, accountability, and resilience, while also considering the environmental footprint of decentralized infrastructures. Aligning decentralized AI systems with ESG principles could help reinforce public trust and attract responsible investment.

Finally, collaboration between legal, technical, and medical experts will be critical. Decentralized AI challenges siloed thinking, its responsible implementation depends on interdisciplinary governance frameworks that combine legal certainty with technical robustness and ethical reflection.

In conclusion, decentralized AI in healthcare holds immense transformative potential - but realizing it responsibly requires a sophisticated balance between innovation and regulation. As a legal professional, I see my role as helping to build that bridge: translating legal principles into flexible, technology-aware governance models that uphold both human rights and market integrity.

Ülkü Cibik is Counsel at MLL Legal AG in Zurich, one of Switzerland's leading full-service law firms. Ülkü is a member of MLL Legal's Banking & Finance and Corporate and M&A team, where she advises clients on a wide range of finance transactions, capital markets matters, and regulatory issues as well as on corporate and M&A transactions. She has particular expertise in structured products and derivatives, advising major domestic and international financial institutions, issuers, and service providers on all aspects of product structuring, issuance, listing, and distribution. Her expertise also spans, syndicated financing, Islamic finance and project finance. As a recognised expert in her field, Ülkü serves as MLL Legal's delegate to the Swiss Structured Products Association (SSPA). Personally mandated by the association, she acts as legal advisor and supports its Legal & Regulations division in regulatory matters. She also contributes her expertise as a long-standing member of the expert jury for the Swiss Derivative Awards - a key annual event in Switzerland's structured products sector. Ülkü established and leads the firm's Türkiye Desk. In this capacity, she advises Turkish clients, companies, and investors on a regular basis with regard to their Swiss-Turkey related operations and transactions. LinkedIn profile: <https://www.linkedin.com/in/%C3%BCIk%C3%BC-cibik-58307761/>

Disclaimer: The views and opinions expressed in this interview are solely those of the interviewee and do not reflect the official position or policies of their employer. Due to confidentiality obligations, no proprietary or internal company information has been disclosed. All insights shared are based on general professional experience in the relevant field and publicly available knowledge.

1.7 CHAPTER INSIGHTS AND CONCLUSIONS

This chapter examines the foundational context for decentralized AI in healthcare by mapping the convergent forces, regional innovation patterns, and evidence base that establish 2025 as a pivotal inflection point for this paradigm shift. Through systematic literature review, global landscape analysis, and expert consultations, several critical insights emerge that frame the opportunities and challenges ahead. Importantly, this analysis acknowledges that decentralized AI in healthcare remains an emerging field—while momentum is accelerating rapidly, much of the current research exists at the conceptual or proof-of-concept stage rather than large-scale operational deployment.

Key Finding 1: Convergence of Four Critical Forces Creates Unprecedented Momentum-

The analysis reveals that decentralized AI's advancement stems from the simultaneous maturation of four interconnected domains. Technologically, federated learning evolves from theoretical concept to operational implementation, supported by edge computing capabilities and blockchain infrastructure. Regulatory frameworks—particularly GDPR, HIPAA, and the European Health Data Space—codify privacy requirements that centralized approaches cannot satisfy. Economic pressures including unsustainable healthcare costs, workforce shortages reaching crisis levels (with half of doctors and nurses reporting burnout), and competitive collaboration imperatives converge with rising societal expectations around patient data ownership, algorithmic transparency, and equitable access. These forces are not isolated trends but synergistic enablers that collectively make decentralization not merely advantageous but necessary.

Key Finding 2: Regional Innovation Patterns Reveal Complementary Rather Than Competing Approaches

Geographic analysis demonstrates that distinct regional approaches are creating a diverse global ecosystem rather than competing for dominance. Europe's privacy-first collaborative frameworks leverage GDPR compliance as competitive advantage; North America's platform-driven commercial integration focuses on scalability and interoperability; Asia-Pacific's sovereign AI implementations prioritize national data control; Switzerland carves a unique "third way" combining sovereignty with open-source transparency; and Africa's inclusive federated learning initiatives address resource constraints through distributed collaboration. These patterns suggest that successful global implementation will require orchestrating these complementary strengths rather than imposing uniform solutions.

Key Finding 3: Evidence Synthesis Identifies Significant Maturity Gaps Across Technology Pillars

The systematic literature review of 87 peer-reviewed studies (2023-2025) reveals significant variation in technological maturity across the three foundational pillars, with a pronounced "concept-to-practice" gap throughout. Federated learning emerges as the most mature architecture with documented implementations moving from pilots to operational deployment in medical imaging, drug discovery, and clinical trials—yet even here, a significant proportion of research remains at the proof-of-concept stage. Blockchain functions effectively as trust-enabling infrastructure for data integrity, access control, and audit trails, though interoperability challenges and integration with legacy systems persist as major barriers. DAOs represent the least mature pillar—transformative in potential for healthcare governance but nascent in practical implementation, with governance complexity emerging as a formidable barrier requiring interdisciplinary expertise spanning economics, law, ethics, and organizational behavior. This maturity assessment underscores that while the value proposition of decentralized AI is widely accepted, the evidence base for real-world clinical and economic impact remains to be established.

PRAGMATIC EXPERIMENTATION AND INTERDISCIPLINARY COLLABORATION DEFINE THE PATH FORWARD

Expert interviews with practitioners across technical leadership, clinical operations, legal compliance, entrepreneurship, patient advocacy, and academic research consistently emphasize that technological sophistication alone does not guarantee success in this emerging field. Implementations succeed when organizations: start with well-defined problems where decentralization offers undeniable advantage; demonstrate value at smaller scale before pursuing integration complexity; invest in change management and stakeholder alignment; address governance frameworks proactively rather than reactively; and assemble interdisciplinary teams capable of navigating technical, regulatory, economic, and ethical dimensions simultaneously. Given the field's nascent state, the path forward requires pragmatic experimentation, transparent documentation of both successes and failures, and sustained commitment to collaborative problem solving across institutional and national boundaries. The transition from concept to practice demands patience, realistic expectations about implementation timelines, and recognition that early adopters are pioneering solutions that will shape the field's future trajectory.

CHAPTER 2: FOUNDATIONAL FRAMEWORK FOR DECENTRALIZATION IN HEALTHCARE



Foreword by the author: Antonio Pesqueira, IGI Global Editorial Review Board Member, ISCTE-IUL, University of Lisbon Researcher

The Dawn of Collaborative Intelligence: Charting a Human-Centered Future for AI in Medicine

For years, the promise of AI in medicine has felt like a distant horizon—powerful, yet just out of reach. The reason has been a fundamental conflict: the insatiable appetite of AI for data has clashed with our non-negotiable right to privacy and the deep-seated trust that must exist between patients and caregivers. The old paradigm, which demanded we pool our most sensitive health information into centralized vaults, was a barrier not of technology, but of philosophy. This chapter introduces a new philosophy. As our systematic literature review reveals, a profound paradigm shift is underway—one that stops trying to move the data to the algorithm and instead brings the algorithm securely to the data. This is the world of decentralized AI (dAI), and it paves the way for a future built not on AI as a replacement, but as a collaborator. It is a future of hybrid intelligence, where human intuition and machine precision merge to create something greater than the sum of its parts.

The concepts within this chapter are the building blocks for these new collaborative spaces. They are the language we will use to build a more trusted, human-centric healthcare system.

We will explore Foundation Models and LLMs not as abstract tools, but as future digital colleagues, working securely alongside clinicians within a hospital's walls to help draft notes, synthesize research, and augment decision-making, all without ever compromising patient confidentiality. We will see Multimodal Models as a new kind of diagnostic lens, one that allows a clinical team to see a truly holistic view of a patient for the first time, weaving together the scattered threads of images, lab results, and genomic data that were previously trapped in institutional silos. Finally, we will understand Federated Learning as the "digital handshake" that enables this new era of cooperation, allowing institutions across the globe to collaborate on curing rare diseases or predicting pandemics, all while honoring the promise of privacy made to every single patient.

The future we are charting is one where a doctor's expertise is amplified by an AI that can scan a million similar cases in a second. It is a future where a "virtual tumor board" of specialized AI agents works in concert with human radiologists, pathologists, and oncologists to devise the best possible treatment plan. And it is a future where patients themselves can become active participants in the governance of research, ensuring that the path of scientific progress is guided by the very communities it aims to serve.

Mastering these concepts is the first step toward building this future—a future where technology does not dehumanize medicine, but gives us the tools to make it more intelligent, more collaborative, and profoundly more human.

*The views expressed are those of the author and do not necessarily reflect those of the employer, ISCTE-University of Lisbon.

2.1 CHAPTER OVERVIEW

Chapter 2 provides the essential conceptual foundation for understanding decentralized AI in healthcare by establishing a shared vocabulary and analytical framework. Building on the convergent forces identified in Chapter 1, this chapter moves from abstract principles to concrete definitions, mapping the key technologies, stakeholders, regulatory requirements, and architectural patterns that enable decentralized approaches. We begin by systematically defining the core technical concepts—from foundation models and federated learning to DAOs and decentralized science—that form the building blocks of decentralized healthcare AI systems examined throughout subsequent chapters.

The chapter then analyzes the healthcare AI ecosystem as a network of evolving stakeholder relationships, recognizing that decentralization fundamentally transforms traditional vendor-buyer dynamics into multi-stakeholder collaborations. Healthcare providers transition from passive technology adopters to active participants in federated networks; patients evolve from data subjects to governance participants with meaningful control over their health information; and technology developers shift from product vendors to platform orchestrators enabling collaborative innovation. Understanding these shifting roles illuminates why decentralized approaches can succeed where centralized models face insurmountable barriers around trust, privacy, and institutional autonomy.

Recognizing that technical capability means nothing without regulatory compliance, we map the complex global regulatory landscape governing healthcare AI—from the EU's comprehensive privacy-first frameworks and the United States' sector-specific approach to China's data sovereignty requirements and emerging international coordination mechanisms. This analysis demonstrates how regulatory requirements increasingly accommodate and even encourage decentralized architectures that keep data under local control while enabling collaborative learning.

The chapter concludes by introducing the three foundational pillars of decentralized AI—decentralized data management, decentralized computation and control, and decentralized governance—explaining how these pillars work synergistically across the AI development lifecycle examined in Chapters 3-5.

Together, these conceptual building blocks equip you to understand the practical implementations, technical architectures, and strategic implications explored throughout the remainder of the report.

2.2 KEY TERMS AND APPROACHES IN DECENTRALIZED AI

Understanding decentralized AI requires mastering key concepts that build on the convergent forces outlined in Chapter 1. These terms form the foundation for the case studies detailed in subsequent chapters of this report.

2.2.1 FOUNDATION MODELS AND LARGE LANGUAGE MODELS (LLMS)

Foundation Models are AI systems trained on broad data that adapt to specific tasks through fine tuning. They include various modalities such as text, images, audio, video, multimodal combinations. Large Language Models (LLMs) are a subset of foundation models, specifically trained on text data and focused on language understanding and generation. In healthcare, foundation models and LLMs now provide clinical documentation assistance, medical literature analysis, patient education, and diagnostic support. Recent advances in reasoning capabilities and multimodal integration have expanded their applications significantly. Decentralization enables:

- Local deployment within hospital infrastructure
- Sovereign AI implementations maintaining national control over sensitive health data
- Federated prompt engineering for collaborative development across institutions while preserving privacy

Examples: DeepSeek's deployment across 100+ Chinese hospitals, Switzerland's open-source sovereign LLM project Apertus, Mayo Clinic's federated clinical documentation systems.

2.2.2 MULTIMODAL MODELS AND MULTI-TASK LEARNING

Multimodal models process text, images, genomics, and sensor data simultaneously, enabling comprehensive healthcare analysis. Multi-task learning allows different institutions to contribute diverse medical expertise without centralizing sensitive information, addressing challenges like modal alignment across institutions and managing varying data quality while balancing computational demands with model complexity.

Examples: European MONAI implementations combining radiology and pathology, African federated ultrasound screening across multiple countries.

2.2.3 FEDERATED LEARNING

Federated Learning (FL) trains AI models across decentralized locations without exchanging raw data. FL sends the AI model to each institution where it trains locally, then returns with learned insights while leaving original datasets behind. Key approaches address different data scenarios: horizontal FL works when hospitals have similar data structures (e.g. x-rays) but from different patient populations, vertical FL applies when organisations have different data types about the same patient (e.g. hospital has clinical records, insurer has claims data, pharmacy has medication history), and multi-chain federated learning uses blockchain coordination across different technical infrastructures to allow for data sharing. Healthcare implementations span multi-hospital rare disease detection, cross-border population health research, and pharmaceutical collaborations.

Examples: MELLODDY consortium across 10 pharmaceutical companies, Singapore's AimSG Platform for cross-institutional collaboration, Owkin's FL network for diagnostics.

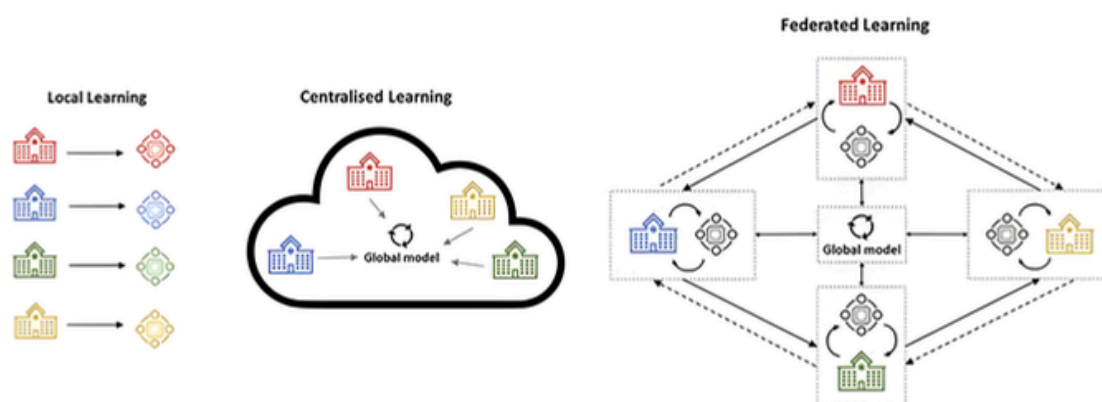


Figure 1. Federated learning in comparison to local learning and centralized learning

Source: Teo, J.T.H., Lim, J.L.W., Lee, J.Y., et al. (2024) 'Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture', Cell Reports Medicine, 5, 101419

2.2.4 EDGE COMPUTING & EDGE AI

Edge Computing refers to processing data locally at or near the source (hospitals, clinics, medical devices) rather than sending it to centralized cloud servers. Edge AI specifically describes AI models deployed and executed on this edge infrastructure, enabling real-time insights without cloud dependency or data transmission. In healthcare, edge computing addresses several critical challenges such as latency reduction, data sovereignty, bandwidth optimization and resilience.

Examples: Care.ai's edge computing across 200+ hospitals, SEOVE's local processing in Brazilian care facilities, Changi General Hospital's institutional AI deployment, DeepSeek's hospital-based inference across China, Google's AI Edge.

2.2.5 DECENTRALIZED/DISTRIBUTED COMPUTATION

Decentralized computation distributes AI processing across multiple locations without central coordination, enabling cost-effective access to computational resources while maintaining data sovereignty. This approach leverages marketplace dynamics to reduce costs and democratize access to high-performance computing for healthcare AI development.

Key platforms include Akash Network (permissionless cloud marketplace with potential for 85% cost reductions), Aethir (enterprise-grade GPU infrastructure with 425,000+ containers), Render Network and io.net (distributed GPU computing for AI training). These platforms enable healthcare organizations to access computational resources without vendor lock-in while supporting federated learning implementations.

Examples: Solve.Care's patient data sovereignty on Akash Network, enterprise AI training on Aethir's distributed infrastructure, medical imaging processing on decentralized GPU networks.

2.2.6 AGENTIC AI AND ORCHESTRATION

Agentic AI refers to AI systems that can autonomously plan and execute complex sequences of actions to achieve goals. In healthcare, this includes:

- Clinical workflow agents: Automating routine clinical tasks
- Research agents: Synthesizing literature and generating hypotheses
- Patient-facing agents: Supporting self-management and care navigation

Decentralized approaches to agentic AI include:

- Local-first agents: Operating primarily on local data with minimal external dependencies
- Federated agents: Collaborating across institutions while maintaining data locality
- Orchestrated agent teams: Multiple specialized agents coordinating on complex healthcare tasks, such as virtual tumor boards where different AI specialists (radiology, pathology, oncology) collaborate on patient cases while keeping data distributed across institutional boundaries

This orchestration enables sophisticated clinical decision support that mirrors human multidisciplinary teams while preserving privacy and institutional autonomy.

Examples: curewise.ai's virtual tumor board simulation (even if not decentralized), VitaDAO's Aubrey AI research advisor, Tsinghua's 42 AI doctor agents across 21 medical departments (also not decentralized, but a good example of agentic AI with clinical benefits).

2.2.7 DECENTRALIZED PHYSICAL AI (DEPAI)

Decentralized Physical AI (dePAI) combines digital intelligence with physical-world interactions through distributed networks of medical devices, robotic systems, and autonomous healthcare infrastructure. Unlike centrally controlled systems, dePAI enables coordinated but independent operation of medical equipment while maintaining patient safety and privacy.

Applications include autonomous medical device coordination (surgical robots communicating without central servers), smart hospital infrastructure (HVAC, lighting, security systems responding to patient needs), distributed diagnostic equipment (imaging devices sharing insights across networks), and IoT medical sensor networks (patient monitoring systems operating independently while contributing to collective intelligence).

Examples: Aethir's decentralized GPU cloud for intelligent robots, IoTeX connecting isolated machine networks through blockchain and enabling AI agent coordination.

INFO BOX: Readers familiar with Web3 may recognize "DePIN" (Decentralized Physical Infrastructure Networks)—projects like Akash Network, Render, and Filecoin that provide decentralized computation, storage, or network resources. While some DePIN projects appear in this report (Akash, Aethir in Chapter 4), we focus on their role in dePAI (Decentralized Physical AI)—architectures where AI interacts with physical medical devices, robotic systems, and autonomous healthcare infrastructure in distributed ways.

KEY DISTINCTION:

DePIN = Infrastructure layer (decentralized GPUs, storage, bandwidth)

dePAI = AI application layer (AI controlling physical devices using decentralized coordination)

DePIN provides the foundation; dePAI describes how healthcare AI leverages that infrastructure to coordinate physical-world systems while preserving privacy and institutional autonomy.

2.2.8 DECENTRALIZED AUTONOMOUS ORGANIZATIONS (DAOS)

DAOs are governance structures that operate through blockchain-based smart contracts rather than traditional hierarchical management. Unlike centralized decision-making, DAOs enable distributed stakeholders to propose, vote on, and automatically execute decisions through programmable rules. Token holders participate in governance by voting on proposals, with decisions automatically implemented by smart contracts, creating transparent, community-driven organizations without traditional executives or boards.

In healthcare, DAOs enable collaborative decision-making across diverse stakeholders (patients, clinicians, researchers), transparent resource allocation for research and development with all transactions publicly recorded, and patient-controlled data governance where communities collectively determine data usage policies. This decentralized approach ensures no single entity controls critical decisions while maintaining accountability through immutable voting records.

Examples: VitaDAO's \$7M+ longevity research governance, Molecule DAO's pharmaceutical IP-NFT management, patient data collectives in European health initiatives.

2.2.9 DECENTRALIZED SCIENCE (DESCI)

DeSci reimagines scientific research by replacing traditional gatekeepers (journals, funding bodies, institutions) with decentralized technologies and community governance. Instead of researchers depending on centralized institutions for funding approval and publication, DeSci uses blockchain based DAOs for transparent research funding, tokenizes intellectual property through IP-NFTs to enable community ownership of discoveries, and creates open, verifiable research processes. This approach democratizes access to research funding and ensures scientific discoveries benefit broader communities rather than only institutional shareholders.

DeSci intersects with decentralized AI in healthcare through:

- Federated model development across independent research organizations
- Collaborative training datasets governed by scientific communities rather than corporations
- Open model architectures with transparent documentation and validation
- Democratized access to advanced AI capabilities for researchers globally

Examples: LabDAO's open wetlab protocols, SpectruthDAO's work in PTSD treatment and trauma healing, African genomics research cooperatives such as H3Africa.

2.2.10 ESSENTIAL PRIVACY-ENHANCING TECHNOLOGIES

Privacy-enhancing technologies enable computation and analysis on sensitive health data without exposing the underlying information to unauthorized parties, which is an essential enabler of decentralized AI models. Data Behind Glass® allows external analysis within secure containers where data never leaves institutional boundaries, homomorphic encryption performs mathematical operations on encrypted medical data without ever decrypting it, zero-knowledge proofs verify patient credentials or medical facts without revealing the actual data, differential privacy adds carefully calibrated mathematical noise to protect individual privacy while preserving aggregate insights for AI models, and secure multi-party computation enables multiple institutions to jointly analyze data while keeping each party's information private.

Examples: Mayo Clinic Platform's Data Behind Glass implementation, European hospital networks using homomorphic encryption, Swiss health identity verification systems (eIDs).

2.2.11 INFRASTRUCTURE AND IDENTITY CONCEPTS

Infrastructure and identity technologies provide the foundational building blocks for decentralized healthcare AI systems. These concepts address where computation happens, how data is created and managed, and how individuals control their digital presence in healthcare systems. Sovereign AI refers to nation-specific AI development maintaining local control over algorithms and training data, synthetic data creates artificially generated datasets that statistically resemble real patient data for training without privacy risks, decentralized identifiers (DIDs) give patients self-sovereign control over their digital healthcare identities through blockchain-based identity systems, and federated data validation ensures data quality and consistency across distributed institutions through collaborative verification processes.

Examples: China's sovereign DeepSeek implementations, Netherlands' Personal Health Train synthetic data initiatives, Hedera Hashgraph's or Internet Computer Protocol's DID wallets.

2.2.12 HEALTHCARE INTEROPERABILITY STANDARDS

Fast Healthcare Interoperability Resources (FHIR) is the modern standard for electronic health information exchange, providing APIs for real-time data access and standardized data structures that enable seamless communication across heterogeneous healthcare systems. FHIR evolved from earlier Health Level Seven (HL7) standards - which are still widely used in legacy systems, addressing their limitations through web-native architecture. FHIR's API-driven architecture naturally supports federated approaches by allowing institutions to expose controlled data access without centralizing storage, enabling federated learning frameworks to query distributed patient cohorts, harmonize data formats for collaborative AI training, and maintain regulatory audit trails—providing the "common language" essential for privacy-preserving collaboration across the healthcare ecosystem.

Examples: Mayo Clinic Platform's FHIR-based federated queries, European Health Data Space FHIR implementation, MONAI framework's FHIR integration for federated medical imaging.

While not exhaustive, understanding these concepts provides the foundation for appreciating how decentralized AI can be effectively implemented in healthcare's complex environment.

2.3 THE HEALTHCARE AI ECOSYSTEM AND STAKEHOLDERS

Healthcare's decentralized AI ecosystem involves complex interactions between diverse stakeholders who traditionally operated in silos but are now becoming interconnected through shared AI development, federated governance, and collaborative data utilization. Unlike traditional healthcare technology adoption where vendors sell solutions to buyers, decentralized AI creates multi-stakeholder networks where participants simultaneously contribute resources (data, computation, expertise) and share benefits (improved models, cost reductions, enhanced capabilities).

This ecosystem transformation reflects the convergent forces outlined in Chapter 1, as regulatory requirements, technological capabilities, and economic pressures drive new forms of collaboration that transcend traditional organizational boundaries while respecting institutional autonomy and patient privacy.

2.3.1 CORE STAKEHOLDERS AND THEIR EVOLVING ROLES

Infrastructure and identity technologies provide the foundational building blocks for decentralized healthcare AI systems. These concepts address where computation happens, how data is created and managed, and how individuals control their digital presence in healthcare systems. Sovereign AI refers to nation-specific AI development maintaining local control over algorithms and training data, synthetic data creates artificially generated datasets that statistically resemble real patient data for training without privacy risks, decentralized identifiers (DIDs) give patients self-sovereign control over their digital healthcare identities through blockchain-based identity systems, and federated data validation ensures data quality and consistency across distributed institutions through collaborative verification processes.

Healthcare Providers function as both data generators and active participants in federated AI networks, transitioning from passive technology adopters to collaborative partners in AI development. They maintain primary custody of clinical data while participating in distributed training networks, deploy AI models locally to preserve data sovereignty, and contribute clinical expertise to validate and improve shared AI systems. Their evolving role encompasses data stewardship, federated validation partnership, and local AI deployment while maintaining patient care responsibilities.

Examples: Mayo Clinic Platform's "Data Behind Glass®" architecture, Chinese hospital networks implementing DeepSeek's federated learning, SEOVE's protective care coordination in Brazil.

Patients and Communities are transitioning from passive data subjects to active governance participants who control their health information, vote on AI development priorities, and directly benefit from personalized AI healthcare systems. Modern frameworks enable patients to manage granular consent for data usage, participate in research governance through community voting mechanisms, and interact with AI agents that respect their privacy preferences while providing personalized care support.

Examples: Patient-controlled data vaults, community voting on AI development priorities, VitaDAO's patient governance participation.

Payers and Insurers are evolving from traditional claims processors to active participants in federated health analytics, contributing population health data for AI model development while gaining insights for risk management and care optimization. They participate in federated learning networks to improve outcome predictions, cost forecasting, and fraud detection while maintaining competitive advantages through data privacy preservation. Their role expands to include federated analytics partnership, population health insights contribution, and collaborative risk modeling.

Examples: Blue Cross Blue Shield federated population health analytics, German statutory health insurers' collaborative AI development, Singapore's Ministry of Health national AI programs.

Technology Developers now build decentralized ecosystems rather than centralized platforms, focusing on enabling federated collaboration, providing privacy-preserving infrastructure, and creating tools that respect institutional autonomy while enabling collective intelligence. Their role shifts from software vendors to ecosystem enablers who provide technical infrastructure for distributed AI development and deployment.

Examples: MONAI's open-source federated framework, Owkin Connect's pharmaceutical research platform, Akash Network's or Aethir's decentralized cloud computing.

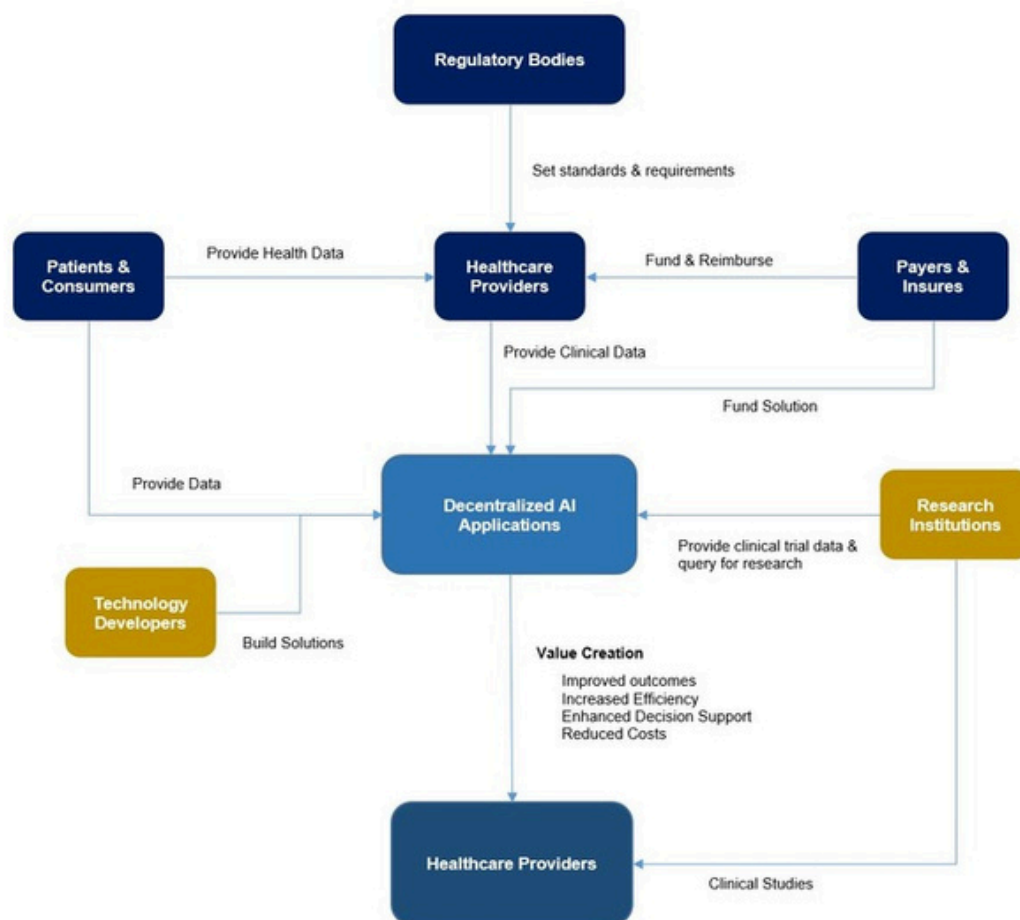
Regulatory Bodies are developing adaptive frameworks that accommodate federated systems while maintaining safety standards, creating mechanisms for cross-jurisdictional coordination, and establishing governance models that can evolve with technological capabilities. They serve as framework developers who enable innovation while ensuring patient protection across distributed AI systems.

Examples: EU AI Act federated learning accommodations, FDA guidance on distributed AI validation, regulatory sandboxes for decentralized AI testing in EU, Singapore etc.

Research Institutions coordinate federated networks that enable multi-institutional collaboration while developing methodological foundations for decentralized AI, serving as network orchestrators who bridge academic research with practical healthcare implementation across institutional boundaries.

Examples: The Distributed AI Research Institute (DAIR), European MELLODDY consortium participation, The Institute for Decentralized AI (IDAI).

HEALTHCARE AI ECOSYSTEM DIAGRAM



2.3.2 COMMUNITY-DRIVEN PLATFORMS FOR DECENTRALIZED AI IN HEALTHCARE

The healthcare AI landscape is increasingly shaped by community-driven platforms that enable collaborative development, shared governance, and open innovation. Unlike traditional vendor client relationships, these platforms create ecosystems where diverse stakeholders contribute expertise, resources, and governance while sharing collective benefits. This community-centric approach addresses healthcare's unique requirements for transparency, inclusivity, and distributed control while accelerating innovation through collaborative intelligence.

OPEN-SOURCE COLLABORATIVE FRAMEWORKS

Community-driven open-source platforms provide the foundational infrastructure for decentralized healthcare AI development, enabling global collaboration while maintaining clinical relevance and regulatory compliance.

HUGGING FACE

Hugging Face represents the largest AI community platform with 350,000+ shared models and datasets, including specialized healthcare repositories like BioMed-Transformers. The platform's Healthcare initiative supports 120+ healthcare-specific models and facilitates federated fine-tuning through community-contributed infrastructure. Their ecosystem demonstrates how open collaboration can scale globally while maintaining quality through peer review and community validation.

MONAI

MONAI exemplifies successful community governance in medical AI, developed through collaboration between NVIDIA, King's College London, and a global community of 15,000+ contributors across 80+ research institutions. The platform's technical steering committee includes representatives from academic institutions, technology companies, and clinical organizations, ensuring decisions reflect diverse stakeholder needs while maintaining scientific rigor.

FLOWER

Flower provides federated learning infrastructure with implementations across 40+ medical research institutions, developed through community contributions that address real-world healthcare deployment challenges. The platform's community-driven development model ensures compatibility with diverse healthcare IT environments while maintaining privacy preserving capabilities.

DECENTRALIZED GOVERNANCE AND RESEARCH COMMUNITIES

Community-governed platforms demonstrate how stakeholders can collectively direct AI development priorities, funding allocation, and research directions through transparent, participatory decision-making processes.

VITADAO

VitaDAO pioneered community-governed healthcare research with over \$7M in research assets managed through token-weighted voting by 2,000+ community members. The platform enables patients, researchers, clinicians, and investors to collectively determine research priorities, evaluate project proposals, and share ownership of resulting intellectual property through IP-NFTs. This model demonstrates how decentralized governance can accelerate specialized research while maintaining scientific standards.

LABDAO

LabDAO coordinates open wetlab infrastructure across 50+ independent laboratories through community governance, enabling collaborative protocol development and resource sharing. The platform shows how physical research infrastructure can be coordinated through decentralized mechanisms while maintaining quality and safety standards.

MOLECULE DAO

Molecule DAO facilitates over \$4M in research investments through community-governed IP-NFTs, connecting researchers with funding while enabling shared ownership of discoveries. The platform demonstrates how traditional research funding can be democratized while maintaining accountability and scientific rigor.

PATIENT-CENTERED COMMUNITY PLATFORMS

Emerging platforms enable patient communities to collectively control their health data, participate in research governance, and share benefits from AI systems developed using their information.

PATIENT DATA

Patient data cooperatives allow communities to pool health information while maintaining collective control over usage rights, research participation, and benefit distribution. These platforms demonstrate how patients can transition from passive data subjects to active research partners with meaningful governance participation.

CITIZEN SCIENCE

Citizen science networks enable patient communities to contribute to research design, validation, and priority-setting while maintaining control over their participation and data contributions. These platforms show how community involvement can improve research relevance and accelerate discovery.

2.3.3 STAKEHOLDER INTERACTION PATTERNS

The evolution toward decentralized AI has created new interaction patterns that differ significantly from traditional healthcare technology adoption:

- **Multi-Stakeholder Governance:** Traditional buyer-seller relationships are evolving into collaborative partnerships where patients, providers, researchers, and technology companies share governance responsibilities and economic benefits.
- **Federated Innovation Networks:** Rather than isolated R&D efforts, innovation increasingly occurs through federated networks where participants contribute data, expertise, and resources while maintaining institutional independence.
- **Community-Driven Development:** Open-source principles and community governance are becoming standard in healthcare AI, with platforms like MONAI demonstrating how collaborative development can achieve both technical excellence and clinical relevance.
- **Patient-Centric Value Chains:** New models are emerging where patients can control their data, participate in governance, and share in the economic benefits of AI systems developed using their information.

These evolving relationships create both opportunities and challenges for implementing decentralized AI, requiring careful attention to incentive alignment, governance frameworks, and sustainable economic models that benefit all participants in the healthcare ecosystem.

2.4 REGULATORY CONTEXT AND DATA GOVERNANCE REQUIREMENTS

The regulatory landscape for healthcare AI has evolved rapidly as governments worldwide grapple with balancing innovation enablement against patient protection. Understanding this complex, multi-jurisdictional environment is essential for designing compliant decentralized AI systems, particularly as the regulatory frameworks themselves are increasingly recognizing and accommodating decentralized approaches.

KEY REGULATORY FRAMEWORKS

EUROPEAN UNION: COMPREHENSIVE PRIVACY-FIRST APPROACH

The EU has established the world's most comprehensive regulatory framework for healthcare AI, building on strong privacy foundations while enabling innovation through structured approaches to decentralization.

General Data Protection Regulation (GDPR):
Establishes strict requirements for processing personal health data, including purpose limitation, data minimization, and explicit consent requirements that often necessitate decentralized approaches

EU AI Act: Risk-based regulatory approach with specific provisions for high-risk healthcare AI systems, including requirements for transparency, human oversight, and algorithmic auditing

European Health Data Space (EHDS):
Framework enabling secure cross-border health data sharing while maintaining patient control and national sovereignty

Medical Device Regulation (MDR): Governs AI systems qualifying as medical devices, with specific pathways for software-based diagnostics and treatment recommendation systems

The EU framework explicitly recognizes federated learning and decentralized approaches as preferred methods for compliance with privacy requirements while enabling cross-border research collaboration.

UNITED STATES: INNOVATION-FOCUSED WITH EMERGING FEDERAL COORDINATION

The US approach emphasizes innovation enablement while building regulatory frameworks that can accommodate both centralized and decentralized approaches.

Health Insurance Portability and Accountability Act (HIPAA): Governs protected health information with Security and Privacy Rules that influence AI system design

Food and Drug Administration (FDA): Evolving frameworks for AI/ML-based medical devices, including Software as a Medical Device (SaMD) pathways and emerging guidance on continuous learning systems

21st Century Cures Act: Provisions for health data interoperability and information blocking that support federated approaches

Executive Orders on AI: Recent federal guidance emphasizing the need for safe, secure, and trustworthy AI development across government agencies

California SB 53: Pioneering state-level framework that imposes stringent requirements only on "frontier AI models" (those exceeding specific computation thresholds and capabilities) while maintaining light-touch oversight for most healthcare AI applications.

The fragmented US regulatory landscape—combining federal frameworks (HIPAA, FDA), state innovations (California SB 53), and industry self-regulation—creates both challenges and opportunities for decentralized AI. Organizations must navigate multiple jurisdictions while state level experimentation may yield best practices that inform future federal coordination.

CHINA: SOVEREIGN AI WITH CONTROLLED COLLABORATION

China's regulatory approach emphasizes national data sovereignty while enabling domestic collaboration and controlled international partnerships.

Personal Information Protection Law (PIPL): Comprehensive data protection framework with specific provisions for health data requiring local processing and control

Cybersecurity Law: Requirements for critical information infrastructure, including major healthcare systems and AI platforms

Data Security Law: Classification system for data with heightened protection for important health information

National Medical Products Administration (NMPA): Guidelines for AI medical devices with expedited pathways for domestic innovations

China's framework specifically supports domestic federated learning initiatives while restricting cross-border data sharing, leading to the development of sovereign AI capabilities like the DeepSeek implementations across Chinese hospital networks.

OTHER KEY JURISDICTIONS

SWITZERLAND

Crypto-friendly regulatory environment enabling DAO governance structures while maintaining strict healthcare privacy protections

SINGAPORE

National AI governance framework with healthcare-specific guidelines supporting both domestic and international federated research

CANADA

Provincial health data regulations creating opportunities for federated approaches across jurisdictional boundaries

BRAZIL

LGPD data protection framework requiring local data processing that supports federated learning implementations

HOW REGULATIONS INFLUENCE DECENTRALIZATION APPROACHES

Modern regulatory frameworks increasingly recognize decentralized AI as preferred approaches for addressing fundamental regulatory challenges rather than merely compliant alternatives to centralized systems.

DATA SOVEREIGNTY AND LOCALIZATION

Regulatory requirements for data localization create natural incentives for decentralized approaches. Cross-border restrictions like PIPL and national data localization laws make traditional centralized AI development legally impossible in many contexts, while federated compliance allows adherence to local regulations while enabling international collaboration. Decentralized architectures naturally address national and institutional sovereignty requirements by keeping data within jurisdictional boundaries.

SECURITY AND RESILIENCE ADVANTAGES

Regulatory frameworks increasingly recognize security benefits of decentralized approaches. Distributed architectures eliminate single points of failure that create attractive targets for malicious actors, while federated systems maintain functionality even if individual nodes experience security incidents. Collaborative security through shared threat detection across federated networks enhances overall system security.

CONSENT AND ALGORITHMIC TRANSPARENCY

Privacy regulations emphasizing specific consent and purpose limitation align well with decentralized approaches. Federated systems maintain alignment with original consent parameters by keeping data within its initial context, while emerging frameworks allow patients to modify consent preferences in real-time across federated networks. Regulatory requirements for AI transparency are increasingly accommodated through distributed auditing, where multiple institutions can independently validate AI systems without centralizing sensitive data, and collaborative oversight enabling shared responsibility for AI system behavior.

EMERGING REGULATORY TRENDS AND IMPLEMENTATION STRATEGIES

ADAPTIVE AND COORDINATED FRAMEWORKS

Regulators are developing flexible approaches that evolve with technological capabilities through regulatory sandboxes for testing innovative approaches, continuous monitoring using real-world evidence, and increased stakeholder engagement including patients and providers in regulatory processes. Cross-jurisdictional coordination mechanisms include mutual recognition agreements, international standards development for federated healthcare AI, and diplomatic protocols enabling cross-border collaboration.

PATIENT-CENTRIC REGULATION

Regulatory frameworks increasingly emphasize patient rights including data portability rights enabling participation in federated networks, algorithmic rights to understand and challenge AI driven decisions, and recognition of patient communities as stakeholders in AI system governance.

IMPLEMENTATION STRATEGIES

Successful decentralized AI implementations employ design for multiple jurisdictions with modular components adapting to different requirements, privacy-by-design architecture embedding protections at the system level, collaborative governance frameworks satisfying regulatory requirements while enabling decentralized decision-making, continuous compliance monitoring tracking regulatory changes, and active stakeholder engagement involving regulators, patients, and providers in system design.

Organizations must establish robust monitoring and adaptation capabilities to maintain compliance as requirements evolve while contributing to regulatory development through active participation in policy discussions and pilot programs.

2.5 THE THREE PILLARS OF DECENTRALIZED AI

Decentralized AI in healthcare rests on three fundamental pillars that work together to create systems that are more secure, collaborative, and responsive to stakeholder needs than traditional centralized approaches. These pillars—decentralized governance, decentralized data management, and decentralized computation—address the core challenges identified in Chapter 1 while enabling the innovative implementations detailed in subsequent chapters.

01

DECENTRALIZED
DATA MANAGEMENT

02

DECENTRALIZED
COMPUTE & CONTROL

03

DECENTRALIZED
GOVERNANCE

2.5.1 DECENTRALIZED DATA MANAGEMENT

Decentralized data management keeps sensitive health information distributed across multiple locations under local control while enabling collaborative analysis. This approach resolves the tension between needing large, diverse datasets for effective AI training and requirements for privacy, security, and institutional autonomy.

Local Data Sovereignty maintains institutional control through data remaining within originating institutions with strict access controls, edge storage for real-time processing, and "Data Behind Glass®" approaches where external analysis occurs within secure containers without data movement.

Blockchain Integration provides infrastructure for decentralized health data management through immutable audit trails, smart contract automation for data access governance, and patient controlled digital identities enabling granular sharing control.

Cryptographic Protection enables secure collaboration through homomorphic encryption allowing computation on encrypted data without decryption, secure multi-party computation for joint analysis while keeping inputs private, zero-knowledge proofs verifying data validity without revealing information, and differential privacy adding calibrated noise to protect individual privacy.

Examples: Mayo Clinic Platform's "Data Behind Glass®", Filecoin Health's cryptographic storage, Internet Computer Protocol's patient-controlled health records, Hedera Hashgraph's healthcare identity management, the IPFS file sharing and storage protocol.

2.5.2 DECENTRALIZED COMPUTATION

Decentralized computation distributes AI processing across multiple locations, enabling collaborative model development without centralizing sensitive data while leveraging collective computational resources. This transforms how AI models are trained, validated, and deployed in healthcare settings.

Federated Learning Architectures include horizontal federated learning training models across institutions with similar data structures but different patients, vertical federated learning combining different data types from institutions serving overlapping populations, and multi-task federated learning simultaneously training for multiple related tasks to improve efficiency.

Distributed Infrastructure leverages collective resources through decentralized cloud platforms like Akash Network offering 85% cost reductions, enterprise-grade GPU infrastructure like Aethir with 425,000+ containers globally, and edge computing networks enabling local AI deployment while maintaining federated capabilities.

Advanced Coordination supports asynchronous federated learning enabling participation by institutions with different computational capabilities, hierarchical coordination accommodating regional and national structures, and resource optimization balancing privacy protection with computational efficiency.

Examples: DeepSeek's multi-task federated learning across 100+ Chinese hospitals, NVIDIA Clara's federated medical imaging, Akash Network's decentralized cloud computing for AI.

2.5.3 DECENTRALIZED GOVERNANCE AND CONTROL

Decentralized governance shifts decision-making authority from centralized entities to distributed stakeholders, creating more equitable, transparent, and resilient systems. Rather than single organizations controlling AI development and deployment, multiple parties including patients, clinicians, researchers, and technology providers share governance responsibilities through structured frameworks.

Multi-Stakeholder Decision-Making enables meaningful participation through consensus mechanisms for reaching decisions across competing interests, expertise-weighted voting that accounts for domain knowledge while maintaining democratic principles, and stake-based participation where influence reflects contribution while protecting minority interests.

DAO-Based Healthcare Governance represents sophisticated evolution in AI oversight through blockchain-enabled transparency and automated enforcement. DAOs enable multi-stakeholder representation in AI governance decisions, create immutable records of all governance activities, and implement programmable compliance with healthcare-specific requirements like mandatory clinical review periods or automatic bias detection.

Economic Models align incentives through contribution-based value-sharing, allocating benefits based on data quality or computational resources contributed, token-based governance using digital assets for network participation, and reputation systems tracking participant reliability over time.

Examples: MELLODDY's pharmaceutical consortium agreements, VitaDAO's community research funding, patient data collectives with democratic governance showing measurably improved stakeholder engagement.

These three pillars—decentralized data management, decentralized computation and control, and decentralized governance—work synergistically to create healthcare AI systems that respect institutional boundaries and regulatory requirements while enabling powerful collaborative capabilities. You will encounter these pillars interwoven throughout the entire AI model development lifecycle, from initial data collection through model training to validation and real world deployment.

The following chapters (3-5) are structured around this AI lifecycle rather than individual pillars, demonstrating how decentralization manifests at each stage: Chapter 3 examines decentralized data collection and processing (primarily showcasing data management and governance through cases like MELLODDY, Mayo Clinic Platform, and SEOVE); Chapter 4 explores decentralized model development and training (emphasizing computation alongside governance and data management through DeepSeek, Akash Network, Bittensor, and the ASI Alliance); and Chapter 5 investigates decentralized validation and deployment (where all three pillars converge in production environments like VitaDAO's research governance and Changi General Hospital's national imaging platform).

In each case study, we explicitly identify which of the three pillars are exhibited, helping you recognize how these foundational concepts translate into practical healthcare AI solutions. Some implementations emphasize a single pillar strongly, while others integrate all three in sophisticated ways, demonstrating both the theoretical framework and its practical manifestation across the AI development lifecycle.

2.6 CHAPTER INSIGHTS AND CONCLUSIONS

This chapter establishes the conceptual foundation for understanding decentralized AI in healthcare by defining key technologies, mapping stakeholder ecosystems, analyzing regulatory frameworks, and introducing the three architectural pillars. Several critical findings emerge that shape how organizations approach decentralization in practice.

Key Finding 1: Terminology Confusion Creates Implementation Barriers in an Emerging Field

The rapid evolution of decentralized AI has produced overlapping and sometimes conflicting terminology that creates confusion during deployment. "Federated learning" versus "syndicated learning," "edge computing" versus "cloud computing," "decentralized" versus "distributed"—these distinctions matter operationally but lack standardization. Organizations frequently misalign on technical requirements because stakeholders attach different meanings to the same terms. Some vendors market "federated" solutions that actually centralize model aggregation, while others use "edge AI" to describe cloud-adjacent processing rather than true device-level computation. This terminological ambiguity becomes particularly problematic when negotiating consortium agreements or regulatory compliance documentation, where precise technical definitions carry legal weight.

Key Finding 2: Decentralized Models Fundamentally Alter Stakeholder Roles, Incentives, and Power Dynamics

Stakeholder relationships in decentralized AI differ markedly from traditional centralized development. In centralized models, healthcare providers act as customers purchasing AI systems, patients serve as passive data subjects, and technology vendors retain full control over algorithms. Decentralization inverts these dynamics: providers become active participants contributing data and computational resources while maintaining sovereignty; patients transition to governance participants with voting rights on data usage and potential economic benefit-sharing; technology companies shift from product vendors to platform operators enabling collaboration. These transformed roles create new incentive structures—providers gain collective intelligence without losing competitive advantages, patients control data while accessing personalized AI—but also introduce coordination complexity requiring explicit governance frameworks and trust mechanisms that don't rely on centralized authority.

Key Finding 3: Regulatory Maturity Varies Dramatically, With Some Frameworks Actively Enabling Decentralization While Others Lag

Global regulatory analysis reveals stark disparities in how jurisdictions accommodate decentralized approaches. The European Union demonstrates the most advanced supportive framework: GDPR's data minimization principle explicitly favors federated architectures; the AI Act provides clear validation pathways; and the European Health Data Space promotes cross-border federated research while maintaining national data sovereignty. The United States presents a fragmented landscape—HIPAA accommodates decentralized architectures but provides no specific guidance, creating implementation uncertainty. China emphasizes data sovereignty, effectively mandating federated approaches for health data but lacking detailed technical standards. Many jurisdictions including most of Latin America, Africa, and Southeast Asia have yet to develop any AI-specific healthcare regulations, creating legal vacuums where decentralized approaches exist in regulatory gray zones.

DECENTRALIZATION OPERATES ON A CONTINUUM ACROSS THE AI LIFECYCLE, WITH FULL DECENTRALIZATION NEITHER POSSIBLE NOR DESIRABLE

Decentralization manifests at varying degrees across different AI development stages, with no implementation achieving complete decentralization. During data collection, MELLODDY maintains strong decentralization—pharmaceutical companies' data never leaves institutional boundaries—while Mayo Clinic's "Data Behind Glass" uses centralized secure containers with decentralized access policies. In model training, DeepSeek deploys local inference but coordinates updates through centralized repositories; Bittensor decentralizes model validation across nodes while maintaining centralized subnet governance. For deployment, Changi General Hospital decentralizes inference across clinical sites while centralizing quality monitoring at the national level. This continuum reflects practical realities: some centralization improves efficiency (centralized model registries reduce duplication), ensures safety (centralized monitoring catches deployment errors faster), or satisfies legal requirements (centralized entities can sign contracts and accept liability). Successful implementations explicitly identify which components require decentralization to achieve core objectives—typically data privacy, institutional autonomy, or collaborative governance—versus which benefit from selective centralization for operational efficiency, regulatory compliance, or economic sustainability.

CHAPTER 3: DECENTRALIZED DATA COLLECTION AND PROCESSING



Foreword by: LeiLei Tang, Medical Affairs Consultant, Ethers HealthData Foundation

AI-driven decentralized data collection and processing is revolutionizing global healthcare by enabling secure, privacy-compliant analysis of medical data across borders and institutions. Leveraging advanced frameworks such as federated learning, blockchain-based integrity validation, and privacy-enhancing technologies like homomorphic encryption, this approach allows diverse medical entities—from hospitals in China to research centers in Europe and North America—to collaboratively develop robust AI models without centralizing sensitive information. These systems adhere to stringent international regulations, including GDPR, HIPAA, and China's PIPL, while supporting applications ranging from diagnostic assistance and personalized treatment to drug discovery and pandemic response.

Recent implementations highlight the scalability and adaptability of decentralized AI. China's integration of platforms like DeepSeek in hospital settings demonstrates how large language models can enhance clinical decision-making and operational efficiency without compromising data locality. Multimodal federated learning further expands these capabilities by incorporating diverse data types—imaging, genomics, and electronic health records—into unified analytical frameworks. These advancements are forging a new paradigm in medical research and care delivery: one that prioritizes ethical data use, global collaboration, and equitable access to AI-powered health innovations.

As a modern medical affairs professional working across pharmaceuticals, biotech, and digital innovation, I have the privilege of learning and utilizing advanced AI tools. Except for LLMs like ChatGPT, DeepSeek, and Grok that are broadly used, specialized GenAI platforms are also introduced to the medical field such as Jenni.ai and Litmaps for literature reviews and KOL mapping. These tools significantly enhance efficiency and are powerful catalysts for novel insights. We are indeed both privileged and challenged to be building careers in this transformative AI era.

The views expressed are those of the author and do not necessarily reflect those of the employer.

3.1 CHAPTER OVERVIEW

Chapter 3 examines decentralized data collection and processing through four primary case studies and two bonus resources spanning pharmaceutical drug discovery, multi-institutional research, vulnerable population care, and clinical trial management. MELLODDY Consortium demonstrates federated learning across 10 competing pharmaceutical companies training AI models on proprietary compound libraries without sharing raw data. Mayo Clinic Platform's "Discover" and "Connect" services enable multi-institutional research collaboration through "Data Behind Glass" architecture where external researchers analyze data within secure containers. SEOVE's Brazilian elderly care network coordinates services across satellite facilities using federated learning to address healthcare provider shortages while protecting vulnerable populations. Acoer's platform applies blockchain technology to clinical trials, creating immutable audit trails for data integrity and AI model transparency.

Each case study follows a consistent structure to facilitate comparison and practical assessment: Organization & Challenge establishes institutional context and the specific problem requiring decentralized approaches; The Solution describes the technical architecture, governance frameworks, and implementation approach; Business Implementation details timeline, stakeholder coordination, and operational deployment; and Impact & Outcomes presents quantitative results and qualitative benefits. This standardized format enables readers to quickly identify which decentralization pillars each implementation exhibits (data management, computation, governance), assess solution maturity (pilot vs. production scale), and evaluate applicability to their own organizational context and constraints.

Beyond the primary case studies, the chapter includes "bonus resources"—initiatives we found relevant and insightful but that don't fully qualify as case studies. These may be enabling infrastructure rather than business implementations, proof-of-concepts still maturing at the time of writing, or systems with potential for decentralization but currently operating in centralized mode. For this chapter: Project MONAI provides open-source infrastructure for medical imaging AI with federated learning support used by thousands of institutions globally, while KnowS demonstrates AI powered medical knowledge platforms for pharmaceutical literature synthesis and translation currently operating centrally but offering insights into future decentralized knowledge curation possibilities.

3.2 CASE STUDY: MELLODDY CONSORTIUM (EU)

WORLD'S LARGEST FEDERATED LEARNING COLLABORATION FOR DRUG DISCOVERY IN PHARMACEUTICAL INDUSTRY

Documented by: Carmen Cucul, Blockchain & AI Healthcare Consultant, Ethers HealthData Foundation

CASE STUDY AT-A-GLANCE: MELLODDY CONSORTIUM (EU)

<p>MATURITY LEVEL: Production (completed pilot, framework replicable)</p> <p>GEOGRAPHY: Europe (10 pharma companies)</p> <p>REPLICABILITY: High (model applicable to other industries)</p>	
<p>THE CHALLENGE:</p> <p>Ten competing pharmaceutical companies needed to train AI on proprietary compound libraries without sharing sensitive intellectual property or violating competitive boundaries.</p>	<p>THE SOLUTION:</p> <p>Federated learning architecture enabling local AI training across company boundaries with cryptographic privacy guarantees and consortium governance framework.</p> <ul style="list-style-type: none"> ✓ Pillar 1: Decentralized Data Management ✓ Pillar 3: Decentralized Governance
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none"> • 10 pharmaceutical companies collaborated successfully • 1.3B+ compound-assay data points analyzed privately • 3-year governance framework development for IP rights 	<p>LESSONS LEARNED:</p> <p>Technical architecture alone is insufficient - competitive collaboration requires equally sophisticated governance frameworks defining IP rights and contribution rules.</p>

1. ORGANIZATION & CHALLENGE

Organization Profile: MELLODDY (Machine Learning Ledger Orchestration for Drug Discovery) is an Innovative Medicines Initiative-funded consortium comprising 10 major pharmaceutical companies: Amgen, Astellas, AstraZeneca, Bayer, Boehringer Ingelheim, GSK, Institut De Recherches Servier, Janssen Pharmaceutica NV, Merck KGaA, and Novartis. The consortium also includes seven technical partners: Budapesti Muszaki Es Gazdasagtudomanyi Egyetem (BME), Iktos, Kubermatic, KU Leuven, NVIDIA, Owkin, and Substra Foundation. This unprecedented collaboration brings together the world's largest pharmaceutical companies to jointly train predictive models on an order of magnitude larger dataset than any previous ML experiment in drug discovery.

The Problem: Drug discovery faced fundamental challenges that individual companies could not solve alone:

- Data Volume Limitations: Each pharmaceutical company's individual datasets, while valuable, were insufficient to train robust AI models capable of predicting complex biological interactions across diverse molecular structures
- Competitive Data Hoarding: Companies possessed vast proprietary chemical compound libraries (collectively more than 10 million annotated small molecules and 1 billion activity labels) but could not share this data due to competitive and intellectual property concerns
- "Coopetition" Challenge: Need to create new paradigm of collaborative research in low trust environment where competitors have mutual interest in building predictive models while protecting private research, data, and commercial advantages
- Technical Barriers: No existing platform could enable secure, large-scale collaboration between competing industrial partners while maintaining data privacy and regulatory compliance

Why Decentralization?: Traditional centralized approaches requiring data pooling were impossible due to competitive constraints and intellectual property protection. MELLODDY pushes beyond the privacy requirements of traditional federated learning, because the pharmaceutical companies are investigating different experimental assays, which must remain private in that they cannot be disclosed to each other. The consortium needed a solution that would facilitate "coopetition" where competitors could benefit from collective intelligence without compromising proprietary advantages.

Key Stakeholders: The consortium balances diverse interests across pharmaceutical giants (each protecting multi-billion dollar R&D investments), academic research institutions (BME, KU Leuven), technology providers (NVIDIA for GPU optimization, Owkin for federated learning platform, Kubermatic for infrastructure), open-source foundations (Substra), and EU regulatory bodies overseeing the Innovative Medicines Initiative funding.

2. THE SOLUTION

MELLODDY developed a secure federated learning platform built on Owkin Connect technology that enables competing pharmaceutical companies to collaboratively train AI models without sharing proprietary data. The platform allows each partner to register their datasets locally within their own infrastructure, then uses cryptographic gradient aggregation to combine model updates across all participants, creating enhanced global models that benefit from collective knowledge while maintaining zero data exposure. Deployed on AWS multi-account architecture with Kubernetes clusters, the system supports multi-task learning across diverse biological targets and assays, enabling companies to benefit from learnings across different experimental approaches without revealing specific compound structures, assay details, or experimental results to competitors.

Decentralization Approach:

- ✓ Decentralized Data: Each pharmaceutical partner maintains complete control over proprietary datasets within their own local infrastructure
- ✓ Decentralized Computation: Multi-task federated learning across distributed nodes with secure gradient aggregation
- ✓ Decentralized Governance: Distributed rights and permissions administered across 10 competing pharmaceutical companies with formal consortium agreements

Technical Architecture - MELLODDY Platform: Built on Owkin Connect technology with specialized enhancements for pharmaceutical collaboration:

- Secure Registration: Partners securely register their proprietary datasets in their own local instance of the distributed platform, which allows the private models to learn from the aggregated knowledge of all partners, without sharing private data
- Multi-Task Learning: The platform enables simultaneous training across multiple biological targets and assays, allowing each company to benefit from learnings across diverse experimental approaches without revealing specific assay details
- Cryptographic Gradient Aggregation: The federated model was trained on the platform by aggregating the gradients of all contributing partners in a cryptographic, secure way following each training iteration
- AWS Multi-Account Architecture: Deployed on Amazon Web Services multi-account architecture running Kubernetes clusters in private subnets, ensuring enterprise-grade security and scalability

Federated Training Process: The MELLODDY approach implements sophisticated collaborative learning:

- Local Model Initialization: Each pharmaceutical partner receives the same base model architecture for local training on their proprietary compound libraries
- Distributed Training Cycles: Partners train models locally on their datasets, generating gradients and model parameters without exposing underlying chemical structures or biological activity data
- Secure Aggregation: Cryptographic methods combine model updates across all partners, creating an enhanced global model that benefits from collective knowledge
- Iterative Improvement: The enhanced model returns to each partner for additional training cycles, progressively improving predictive performance

Privacy & Security Architecture: MELLODDY implements the most stringent security requirements in federated learning:

- Triple-Layer Protection: External security audits, individual partner IT team verification, and cryptographic aggregation methods
- Zero Data Exposure: No raw data, compound structures, or experimental results ever leave partner premises
- Assay Privacy: Advanced privacy preservation ensures that even the types of biological assays being investigated remain confidential between competitors
- Regulatory Compliance: Framework designed to meet GDPR, various national data protection laws, and pharmaceutical industry regulatory requirements

3. BUSINESS IMPLEMENTATION

Implementation Timeline: MELLODDY's three-year journey demonstrated the evolution from concept to operational reality:

- 2019-2020: Platform design, development, and extensive security auditing across all 10 pharmaceutical partners
- September 2020: Achievement of year one objective with first successful federated learning run using the secure platform - a technical triumph and key milestone
- 2021: Year two objective achievement with significant model improvement through second federated run at scale using improved and re-audited platform
- 2022: Project completion with demonstrated superior predictive performance compared to standalone models and preparation for industry scaling

Governance Model: MELLODDY pioneered complex multi-stakeholder governance for competitive collaboration:

- Consortium Structure: Formal agreements defining roles, responsibilities, and intellectual property sharing across 10 competing pharmaceutical companies
- Technical Partnership Division: Research partners (BME, Iktos, NVIDIA) focused on ML implementation and privacy optimization, while operational partners (Owkin, Kubermatic, KU Leuven, Substra Foundation) developed platform code
- Distributed Decision-Making: Organizational roles codified as different rights and permissions on the platform, administered in a decentralized way across all partners
- IP Protection Framework: Agreements ensuring that collective model improvements benefit all partners while protecting individual company proprietary advantages

Economic Framework: The consortium established sustainable funding and value-sharing models:

- EU IMI Funding: Significant public investment from Innovative Medicines Initiative providing neutral foundation for competitive collaboration
- Shared Infrastructure Costs: Distributed AWS infrastructure costs and computational resources across consortium members
- Collective Value Creation: Each partner gains access to predictive models trained on datasets orders of magnitude larger than individual company holdings
- Competitive Advantage Preservation: Framework ensures that collaboration enhances rather than compromises individual company R&D capabilities

Change Management: MELLODDY required unprecedented coordination across competitive organizations, with extensive project management, legal framework development, and technical integration efforts spanning multiple countries and regulatory jurisdictions.

4. IMPACT & OUTCOMES

Quantitative Results:

- **Dataset Scale:** Collaborative training on more than 10 million annotated small molecules and 1 billion activity labels - unprecedented scale in pharmaceutical ML
- **Computing Performance:** Platform processed 713,796 GB of data transfer and 912,778 EC2 hours during second federated run, demonstrating enterprise scalability
- **Task Completion:** MELLODDY achieved more than 100,000 ML tasks representing more than 40,000 tests across three years of operation
- **Partner Network:** Successfully coordinated collaboration across 10 major pharmaceutical companies and 7 technical partners spanning multiple countries

Qualitative Benefits:

- **Industry Transformation:** MELLODDY project demonstrated that drug discovery research using a federated model outperforms standalone models, establishing new paradigm for pharmaceutical collaboration
- **Technical Breakthrough:** World's first federated learning experiment in drug discovery performed at industrial scale between competitive partners, proving feasibility of "coopetition" in highly competitive industries
- **Model Performance:** Demonstrated greater predictive performance of models used to improve drug discovery process, with promising observations for domain applicability across diverse biological targets
- **Regulatory Validation:** Platform passed extensive security audits and regulatory review, establishing framework for future pharmaceutical collaborations under strict privacy and IP protection requirements
- **Open Innovation:** Project catalyzed release of underlying Substra technology as open source through Linux Foundation, democratizing federated learning capabilities for broader pharmaceutical research community
- **Knowledge Transfer:** Successful demonstration of federated learning principles applicable across other competitive industries requiring secure collaboration on sensitive data

Key Insights Summary:

- **"Coopetition" model validation:** Successfully demonstrated that competing pharmaceutical companies can collaborate on foundational AI research while protecting proprietary advantages through careful governance and technical architecture
- **Neutral oversight essential:** EU funding and third-party project coordination proved critical for building trust among competitors in highly competitive industry
- **Scale drives value:** Collective dataset 10x larger than individual company holdings delivered superior model performance, validating the collaborative approach for all participants

REFERENCES:

- Innovative Health Initiative (2025) MELLODDY Project Factsheet. Available at: <https://www.ih.europa.eu/projects-results/project-factsheets/melloddy>.
- Kubermatic (2024) 'MELLODDY: Turning pharma competition into "coopetition"'. Available at: <https://www.kubermatic.com/blog/melloddy-turning-pharma-competition-into-coopetition/>
- Oldenhof, M., et al. (2024) 'Industry-scale orchestrated federated learning for drug discovery', Proceedings of the AAAI Conference on Artificial Intelligence, 37(13), pp. 15576-15584. doi: 10.1609/aaai.v37i13.26847.
- Owkin (2024) 'MELLODDY project hits its year one goal: First secure platform for federated learning in drug discovery'. Available at: <https://www.owkin.com/newsfeed/melloddy-project-meets-its-year-one-objective-deployment-of-the-worlds-first-secure-platform-for-multi-task-federated-learning-in-drug-discovery>.

3.3 CASE STUDY: MAYO CLINIC PLATFORM “DISCOVER” AND “CONNECT” (US)

MULTI-INSTITUTION COLLABORATION TO GENERATE DE-IDENTIFIED DATASETS FOR RESEARCH

Documented by: Carmen Cucul, Blockchain & AI Healthcare Consultant, Ethers HealthData Foundation

CASE STUDY AT-A-GLANCE: MAYO CLINIC PLATFORM (US)

<p>MATURITY LEVEL: Production (operational commercial service) GEOGRAPHY: North America (expanding globally) REPLICABILITY: High (platform model, not custom builds)</p>	
<p>THE CHALLENGE: Healthcare institutions needed to enable external researchers to analyze sensitive patient data without moving data outside institutional security boundaries.</p>	<p>THE SOLUTION: "Data Behind Glass" architecture where researchers analyze data within secure containers; data never leaves Mayo's infrastructure. FHIR-based federated queries across 25+ partner organizations.</p> <p>✅ Pillar 1: Decentralized Data Management</p>
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none"> • 30M+ patient records accessible for federated analysis • 25+ healthcare organizations connected • Zero data breaches while enabling 100+ research projects • Commercial model: researchers pay for computation access 	<p>LESSONS LEARNED: "Data Behind Glass" demonstrates that controlled access within institutional boundaries can enable collaboration without compromising data sovereignty or security.</p>

1. ORGANIZATION & CHALLENGE

Organization Profile: Mayo Clinic is a nonprofit academic medical center with major campuses in Rochester, Minnesota; Phoenix/Scottsdale, Arizona; and Jacksonville, Florida. The organization sees approximately 1.3 million unique patients annually from all 50 states and 138 countries, and operates as one of the largest integrated healthcare delivery systems in the United States with over 70 hospitals and clinics across Minnesota, Iowa, and Wisconsin. As a world-renowned institution known for treating complex and rare conditions, Mayo Clinic has accumulated vast amounts of diverse clinical data across multiple specialties, with the Mayo Clinic Platform launching in 2019 as a strategic initiative to improve healthcare through insights and knowledge derived from data.

The Problem: Healthcare innovation faced a fundamental bottleneck that required a new approach beyond traditional data sharing. Key challenges included:

- Access barriers: Solution developers typically need several years and millions of dollars to bring a digital health solution to the point of care, with limited access to high-quality, diverse clinical datasets for AI model development
- Privacy and control concerns: Current methods for leveraging data analytics require data to be shared with a central location, like a startup with a novel algorithm. This forces data generators to give away most of the control they held by being the sole possessors of that data
- Fragmented innovation ecosystem: Startups lacked clinical expertise while established healthcare systems struggled with technology integration
- Global collaboration barriers: The diversity of languages presented a significant barrier to consolidating clinical data on a global scale

Why Decentralization?: Loss of control and concerns about data privacy are some of the disincentives keeping many health care companies and organizations from sharing their data, and it is precisely the problem that federated learning was created to solve. Traditional centralized approaches created "virtual treasure chests of data, more enticing to hackers than the individual data sources." Mayo Clinic needed an approach that could unlock the value of its clinical data while maintaining complete control and ensuring patient privacy.

Key Stakeholders: Primary participants include digital health companies and startups, pharmaceutical companies conducting research, medical device manufacturers seeking real-world evidence, academic researchers, and Mayo Clinic's internal innovation teams. The platform now connects eight of the world's leading health systems across three continents through Mayo Clinic Platform_Connect, including Hospital Israelita Albert Einstein (Brazil), Sheba Medical Center (Israel), University Health Network (Canada), Seoul National University Hospital, SingHealth, UC Davis Health, Mercy, and Mayo Clinic.

2. THE SOLUTION

Platform Architecture: Mayo Clinic Platform operates through two complementary systems with different data hosting models:

- Mayo Clinic Platform_Discover: Provides access to Mayo Clinic's own de-identified dataset containing 13.6M+ patient records (accumulated over time) hosted in Mayo's secure cloud environment - a centralized approach where external developers access data within Mayo's infrastructure
- Mayo Clinic Platform_Connect: A distributed data network where partner data never leaves each institution's local environment - Mayo Platform intermediates secure access between partner clouds without hosting any partner data

Decentralization Approach:

- ☒ Decentralized Data: Data remains within each institution's infrastructure boundaries using the proprietary "Data Behind Glass®" model
- ☒ Decentralized Computation: Federated learning approach where algorithms travel to data rather than data traveling to algorithms
- ☐ Decentralized Governance: Currently centralized governance with Mayo Clinic maintaining primary control, though individual Connect partners retain control over their data nodes

Technical Architecture - "Data Behind Glass®" Model: Based on Mayo Clinic's comprehensive white paper (2023), this approach implements different architectures for each service:

- For Discover: Mayo hosts de-identified data from its own 70+ care sites in a private cloud container within Google Cloud infrastructure, where external parties can access data for analytics but data never leaves Mayo's container
- For Connect: True federated architecture where de-identified clinical data from each partner never leave the partner's local environment, with encryption and de-identification performed locally by each partner using industry-accepted statistical methods
- Secure communication: Partners communicate via cryptographic methods and cloud enabled federated tools, with Mayo Platform providing network governance but not data hosting for Connect partners

Privacy & Security: The platform implements multiple layers of protection as detailed in their technical documentation:

- All data is de-identified and Mayo Clinic Platform customers cannot see or interact with identifiable data at any point
- Contractual prohibitions against co-mingling Mayo Clinic Platform data with any other data sources
- Strict identity and access management controls with code repository whitelisting
- Cryptographic methods enabling secure, blind communication between data nodes during model validation

3. BUSINESS IMPLEMENTATION

- Implementation Timeline: Mayo Clinic Platform launched in 2019 as an "innovation factory," with the Clinical Data Analytics Platform announced as the first venture in 2020 featuring nference as the initial partner to accelerate drug discovery and development. The Solutions Studio program launched in 2024 to provide comprehensive support from concept to clinical integration. By 2023, 28 developers were using the platform with potential impact on 45 million lives, growing to 81 developers and 56 million patient lives by 2025.
- Governance Model: Mayo Clinic maintains centralized control while establishing collaborative frameworks with partners. Decision-making prioritizes patient privacy and clinical outcomes over commercial considerations, with formal partnership agreements defining roles, responsibilities, and data usage rights.

- **Economic Framework:** The platform operates on a subscription-based model where partners pay for access to de-identified data and development tools. Depending on the opportunity, Mayo Clinic Platform may consider investment or new company creation. Value creation is shared through improved clinical outcomes, accelerated time-to-market for innovations, and potential equity arrangements in successful ventures.
- **Change Management:** Success requires addressing the core business challenges of healthcare providers, such as improving margins, enhancing staff retention, and optimizing workflow efficiency. The platform provides extensive clinical guidance and workflow integration support, with about 40 companies participating in accelerator programs that include clinical expertise and real-world testing opportunities.

4. IMPACT & OUTCOMES

Quantitative Results:

- **Data scale:** Mayo Clinic Platform_Discover contains 13.6M+ de-identified patient records (accumulated over time), including 5.8B+ images, 2.72B+ lab test results, 10.1M pathology reports, and 698M clinical notes
- **Global network growth:** Mayo Clinic Platform_Connect expanded from initial partnership with Mercy (2022) to eight leading health systems across three continents by 2024, representing four of the top 11 hospitals in the world
- **Developer engagement:** 28 developers using the platform in 2023 with potential impact on 45 million lives, growing to 81 developers and 56 million patient lives by 2025
- **Partnership acceleration:** Mayo Clinic Platform_Accelerate program has graduated multiple companies with successful funding rounds (e.g., Acorai's \$4.5M series seed round)

Qualitative Benefits:

- **Enhanced global collaboration:** Using aggregated, de-identified clinical data to generate patterns to pinpoint disease earlier and identify the best treatment options, including representation from diverse global regions and languages
- **Improved model performance:** Access to diverse, high-quality datasets enables development of more robust and generalizable AI models that overcome model bias and create more diverse treatment recommendations
- **Accelerated innovation:** The combination of data, clinical expertise, and validation infrastructure significantly reduces traditional development timelines from years to months
- **Addressing health inequities:** Platform enables inclusion of previously underrepresented populations, such as Latin American populations in global studies with cutting-edge technology and artificial intelligence
- **Sustainability indicators:** Continuous expansion with new global health system partners, established governance frameworks, and proven commercial viability through subscription based revenue streams
- **Real-world validation:** Mayo Clinic and Mercy reached their first major milestone in data collaboration in 2024, demonstrating cross-institutional analysis capabilities for discovering new diagnostic and treatment approaches

KEY INSIGHTS SUMMARY

- Federated vs. Syndicated: Mayo Clinic Platform employs true federated learning architecture where algorithms travel to data rather than data being centralized, distinguishing it from syndicated data approaches that aggregate information in central repositories
- "Data Behind Glass®" innovation: The proprietary model successfully balances data accessibility for innovation with institutional control and patient privacy, creating a replicable framework for healthcare data collaboration
- Global federation success: Demonstration that federated learning can scale across different continents, languages, and regulatory environments while maintaining local data sovereignty and control

REFERENCES:

- HealthSystemCIO (2023) 'Data behind glass: Exploring a federated model for data management', 11 October. Available at: <https://healthsystemcio.com/2023/10/11/data-behind-glass-exploring-a-federated-model-for-data-management/>
- Healthcare IT News (2025) 'Exploring Mayo Clinic's digital health innovation platform'. Available at: <https://www.techtarget.com/healthtechnanalytics/feature/Exploring-Mayo-Clinics-digital-health-innovation-platform>
- Mayo Clinic Platform (2023) Mayo Clinic Platform keeps data behind glass: White paper. Available at: https://www.marco.health/content/files/2023/05/Platform_DBG-White Paper.pdf
- Mayo Clinic Platform (2025) Mayo Clinic Platform Connect. Available at: <https://www.mayoclinicplatform.org/mayo-clinic-platform-connect/>
- Mayo Clinic Platform (2025) Mayo Clinic Platform Discover. Available at: <https://www.mayoclinicplatform.org/discover/>
- NCBI (2025) 'Mayo-Google partnership: Sharing health data'. Available at: <https://www.ncbi.nlm.nih.gov/books/NBK594445/>

3.4 CASE STUDY: SEOVE – CADEIA DE CUIDADOS DECENTRALIZED CARE PLATFORM (BR)

INTEGRATED FEDERATED LEARNING, AI RESOURCE ALLOCATION, AND BLOCKCHAIN SUPPLY CHAIN FOR ELDERLY CARE

Documented by: Antonio Pesqueira, IGI Global Editorial Review Board Member, ISCTE-IUL, University of Lisbon Researcher

CASE STUDY AT-A-GLANCE: SEOVE CADEIA DE CUIDADOS (BR)

<p>MATURITY LEVEL: Pilot (18-month implementation, Jan 2025- Jul 2026)</p> <p>GEOGRAPHY: Brazil (Florianópolis region, expanding south)</p> <p>REPLICABILITY: Medium (requires consortium + municipal buy-in)</p>	
<p>THE CHALLENGE:</p> <p>Brazilian elderly care network serving vulnerable populations (domestic violence victims) needed coordinated care while protecting both medical privacy and personal financial data from exploitative family members.</p>	<p>THE SOLUTION:</p> <p>Integrated platform combining federated learning for health predictions, edge computing for real-time analytics, and blockchain for medication supply chain transparency—all keeping data local per Brazil's LGPD requirements.</p> <ul style="list-style-type: none">✓ Pillar 1: Decentralized Data Management✓ Pillar 2: Decentralized Computation (Edge AI)✓ Pillar 3: Decentralized Governance (Consortium)
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none">• 20% projected reduction in hospital admissions • Dual privacy protection (medical + financial data)• Medication inventory automation for single-nurse staff• Municipal partnership model for sustainable funding	<p>LESSONS LEARNED:</p> <p>Operational challenges (supply chain) must be solved before clinical AI can succeed. Blockchain-based inventory became the foundation enabling predictive health interventions.</p>

1. ORGANIZATION & CHALLENGE

Organization Profile: SEOVE (Sociedade Espírita Obreiros da Vida Eterna, or Spiritist Society of Workers for Eternal Life) is a Brazilian philanthropic institution founded on February 10, 1972, dedicated to improving quality of life for the elderly and children through comprehensive health, well-being, and community support rooted in Spiritist doctrine principles. Based in Florianópolis serving nearly 1 million people, SEOVE operates multiple satellite care homes coordinated through a regional center. The organization serves vulnerable populations including elderly women who are victims of domestic violence, as well as children and families in situations of social vulnerability, providing medical care, on-site pharmacy and dental services, spiritual guidance, cultural activities for trauma processing, and family relationship counseling. Community programs include education initiatives, food distribution, and cultural activities fostering mutual support. Originally publicly funded, SEOVE transitioned to private funding and municipal sponsorship as public resources became limited.

The Problem: SEOVE confronted four fundamental and interdependent barriers limiting coordinated, protective care across its network:

- **Healthcare Resource Scarcity:** The Florianópolis region faces severe shortages of medical doctors and nurses, making it difficult to connect satellite care homes with regional healthcare networks when residents require hospital care or on-site medical visits
- **Pervasive Operational Inefficiencies:** Day-to-day supply management is highly manual and burdensome. SEOVE's single nurse must physically travel to public health centers to acquire medications, navigate complex bureaucracy, and wait in long lines. Manual spreadsheet based inventory creates high risk of error, unaccounted stock losses ("quebras"), and medication expiration, diverting skilled staff from patient care to administrative logistics.
- **Dual Privacy Threats:** Beyond medical privacy mandated by Brazil's Lei Geral de Proteção de Dados (LGPD), residents face financial fraud attempts from family members, requiring protection of both health and personal financial information with any centralized data storage creating unacceptable risk to this vulnerable population
- **Data Fragmentation:** Lack of integration across care homes created information silos that impeded effective care coordination and best practice sharing, while Brazil's LGPD regulations strictly limited traditional centralized data approaches, preventing patient information aggregation for system-wide analysis

Why Decentralization?: Traditional centralized approaches were impossible due to LGPD compliance requirements and the need to protect vulnerable populations from both medical privacy breaches and potential financial exploitation. SEOVE needed a solution enabling collective intelligence and care coordination while keeping sensitive data within each facility's local infrastructure.

Key Stakeholders: The ecosystem includes vulnerable elderly residents and families in situations of social vulnerability, care home staff needing clinical decision support, regional healthcare providers with limited availability, the municipality of Florianópolis providing sponsorship, Brazilian regulatory bodies overseeing LGPD compliance, and research collaborators across southern Brazil.

2. THE SOLUTION

The "Cadeia de Cuidados" (CareChain) platform is a comprehensive decentralized care coordination system integrating three complementary pillars to address SEOVE's interconnected challenges. While this case study provides an overview of all three functionalities to demonstrate the integrated architecture, subsequent sections will focus specifically on the federated learning and AI components (Pillars 1 and 2). The platform employs federated learning to train AI models locally across care homes without centralizing sensitive data, combined with edge computing for real-time analytics at the point of care and blockchain technology for transparent supply chain management. The platform delivers predictive analytics for early health interventions, decision support for tailored care plans, automated medication inventory management with traceable provenance, and enhanced communication among healthcare teams while ensuring data sovereignty for each facility and reducing privacy risks by keeping all data local.

Decentralization Approach:

- ✓ Decentralized Data: Patient health and personal protection data remains within each facility's local infrastructure, with medication and supply chain data recorded on blockchain
- ✓ Decentralized Computation: Federated learning enables local model training with secure aggregation, while edge computing provides real-time analytics
- ✓ Decentralized Governance: Consortium model where satellite care homes co-manage the platform while the regional center coordinates services

Technical Architecture: SEOVE's Cadeia de Cuidados platform integrates three decentralized technologies to create a comprehensive protective care ecosystem:

Pillar 1: Predictive Health Interventions (Federated Learning)

- Federated Learning Framework: Local AI models train at each satellite care home using facility-specific data, with only encrypted parameters shared for central aggregation, enabling collective learning while keeping medical and personal protection data local
- Predictive Analytics: Utilizes federated learning to analyze distributed health data and identify residents at high risk for adverse events (falls, hospital readmission), enabling early interventions projected to reduce hospital admissions by 20%
- Chronic Disease Management: Platform focuses initial pilot on chronic disease patterns, establishing baseline metrics across the care network

Pillar 2: AI-Assisted Resource Allocation

- Clinical Decision Support Integration: AI algorithms identify and allocate scarce healthcare resources, helping care homes connect with available medical professionals across the Florianópolis healthcare network
- Intelligent Matching System: Creates automated connections between care facilities and available doctors/nurses to address the regional provider shortage, expected to increase staff efficiency by 15%

Pillar 3: Automated and Traceable Supply Chain Management (Blockchain)

- Blockchain-Enabled Traceability: Each medication batch is represented as a unique Non Fungible Token (NFT) on the Cardano blockchain using Aiken smart contracts, creating transparent and immutable audit trails from manufacturer to patient

- Automated Inventory Management: Digital system with automated alerts for low stock levels and approaching expiration dates, replacing manual spreadsheets and eliminating time-consuming physical travel to health centers
- Online Requisition System: Staff can submit medication orders digitally to hospitals or public health system (SUS), with real-time status tracking replacing in-person queuing and uncertainty
- Product Authenticity Verification: Smart contracts verify medication integrity at each supply chain step, addressing staff concerns about counterfeit products through traceable origin documentation

Cross-Pillar Integration:

- Edge Computing for Real-Time Protection: On-site devices process real-time data for immediate insights, including vital sign monitoring, fall detection, and security alerts for potential family-based financial threats - supporting both clinical care (Pillar 1) and resident safety
- Regional Knowledge Sharing: Federated approach enables institutions outside São Paulo and Florianópolis to contribute clinical insights while benefiting from SEOVE's protective care expertise

Privacy & Security Architecture: Multiple protection layers designed for LGPD compliance and personal safety:

- Dual Privacy Protection: Secure handling of both medical information and personal/financial data targeted by exploitative family members
- Data Localization: All sensitive information remains within facility infrastructure with comprehensive access controls
- Encrypted Model Sharing: Only mathematically transformed parameters shared between facilities, preventing reconstruction of individual resident information
- Blockchain Transparency with Privacy: Supply chain transactions provide verifiable audit trails while maintaining resident anonymity through cryptographic techniques

3. BUSINESS IMPLEMENTATION

- Implementation Timeline: Phased deployment strategy validating effectiveness in protective care settings, with design-thinking methodology emphasizing stakeholder empathy and iterative refinement:

Phase 1 (January 2025 - July 2026): 18-month pilot coordinated by Dr. Andreia de Bem Machado, SEOVE's Treasury Board Member and Financial Delegated Responsible, focusing on chronic disease management and establishing baseline metrics. Phase includes deep stakeholder engagement workshops with clinical staff, administrators, residents, and families to map complete workflows and pain points. Highest priority identified: internal medication stock control, which becomes foundation for MVP development. Deliverables include stakeholder needs assessment, technical specifications, UI/UX wireframes designed for mobile accessibility and simplicity, and consortium governance charter.

Phase 2 (Mid-2026): Expansion to all SEOVE facilities with enhanced fall detection, full medication management integration across blockchain and AI systems, medication management, and security monitoring

Phase 3 (2026-2027): Scaling to partner organizations across southern Brazil and exploring regional collaboration

- **Governance Model:** Consortium-based governance adapted for protective non-profit healthcare delivery, with regional coordination through the central Florianópolis facility and municipal partnership for sustainable funding. Smart contracts on the blockchain layer can programmatically enforce consortium rules such as data usage agreements and partner responsibilities, creating transparent and auditable governance.
- **Economic Framework:** Sustainable funding model combining AI-driven resource optimization, municipal support, and shared protection costs across participating facilities. Blockchain supply chain reduces medication waste and fraud, while federated learning eliminates centralized data storage costs, contributing to long-term economic viability.

4. IMPACT & OUTCOMES

Quantitative Results (Expected based on pilot implementation):

- **Clinical Outcomes:** Projected 20% reduction in hospital admissions through predictive analytics and improved healthcare resource access
- **Operational Efficiency:** Expected 15% increase in staff efficiency through AI-assisted resource allocation and healthcare provider matching, plus significant time savings from eliminating physical travel for medication procurement
- **Care Coordination:** Anticipated 30% faster care plan updates through real-time data processing and automated regional healthcare connections
- **Protection Enhancement:** Improved security protocols for medical privacy and financial protection against family exploitation, plus reduced medication waste and stockouts through blockchain-enabled inventory management

Qualitative Benefits:

- **Enhanced Protective Care:** Federated learning enables protective care strategy sharing while maintaining strict resident confidentiality
- **Healthcare Resource Access:** AI-driven matching systems overcome regional provider shortages by efficiently connecting care homes with available professionals
- **Comprehensive Service Integration:** Platform supports medical care coordination plus spiritual guidance, cultural programs, and family relationship management
- **Supply Chain Integrity:** Blockchain traceability provides verifiable product authenticity and complete provenance documentation, addressing staff concerns about counterfeit medications and creating accountability from supplier to resident
- **Regional Knowledge Expansion:** Federated approach enables systematic sharing of protective elderly care innovations across southern Brazil
- **Vulnerable Population Empowerment:** Demonstrates that advanced AI and blockchain can enhance care for society's most vulnerable while maintaining highest privacy standards

Sustainability Indicators: Municipal commitment, regional scalability interest, and academic validation through 18-month research pilot.

KEY INSIGHTS SUMMARY

- **Protective Care Innovation:** Successfully demonstrates that decentralized AI integrated with blockchain supply chain management can enhance care coordination for vulnerable populations while providing additional layers of personal and financial protection beyond traditional medical privacy
- **Resource Optimization in Scarcity:** Proves that federated learning and AI-driven resource allocation can help overcome healthcare provider shortages by efficiently matching available professionals with care needs
- **Holistic Service Integration:** Platform supports comprehensive protective services beyond medical care, demonstrating how multiple decentralized technologies (FL, edge computing, blockchain) can coordinate complex social support systems
- **Regulatory and Social Compliance:** Successful LGPD-compliant technology implementation that addresses social protection needs, creating a model for privacy-preserving AI requiring enhanced personal security
- **Operational Foundation Enables Clinical AI:** Recognition that supply chain inefficiencies must be solved before clinical AI can succeed - blockchain-based inventory management (Pillar 3) provides the operational stability necessary for predictive health interventions (Pillar 1) to deliver impact
- **Regional Collaboration Potential:** Federated approach enables knowledge sharing across institutions serving vulnerable populations while maintaining strict confidentiality

REFERENCES AND ADDITIONAL RESOURCES:

- SEOVE (2025) Sociedade Espírita Obreiros da Vida Eterna. Available at: <https://seove.org.br/>
- Brazil (2018) Lei Geral de Proteção de Dados Pessoais (LGPD), Lei nº 13.709. Available at: https://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/l13709.htm
- Interview with Dr. Andreia de Bem Machado, Postdoctoral Researcher, Federal University of Santa Catarina, SEOVE's Treasury Board Member and Financial Delegated Responsible
- Interviews with Andréia Carla Tonin, SEOVE Technical Manager and clinical staff

3.5 CASE STUDY: ACOER – DECENTRALIZED CLINICAL TRIALS PLATFORM (US)

TRANSFORMING CLINICAL RESEARCH THROUGH BLOCKCHAIN-ENABLED DATA INTEGRITY AND AI TRANSPARENCY

Documented by: Natalia Sofia, Data Ecosystem Lead, Ethers HealthData FoundationAdvisor

CASE STUDY AT-A-GLANCE: ACOER CLINICAL TRIALS PLATFORM (US)

MATURITY LEVEL: Production (multiple trials using platform)

GEOGRAPHY: United States (expanding internationally)

REPLICABILITY: High (platform-as-a-service model)

THE CHALLENGE:

Clinical trials suffer from fragmented data management, lack of transparency in data provenance, and trust deficits between sponsors, sites, and regulators.

THE SOLUTION:

Blockchain-based "Cryptographic Data Mesh" on Hedera Hashgraph providing immutable audit trails, tamper-proof consent tracking, and transparent AI model governance.

✓ Pillar 1: Decentralized Data Management (Blockchain)

✓ Pillar 3: Decentralized Governance (Smart contracts)

KEY OUTCOMES:

- Immutable consent and data access logs
- Real-time trial monitoring for sponsors/regulators
- Reduced fraud through tamper-proof audit trails
- Federated AI analysis without centralizing trial data

LESSONS LEARNED:

Blockchain excels in operational/administrative domains (consent, audit trails) rather than AI training itself. Creates trust infrastructure enabling federated analytics.

1. ORGANIZATION & CHALLENGE

Organization Profile Acoer is a U.S.-based healthcare technology company pioneering the use of blockchain, web3, and privacy-preserving AI for trusted health data management. The company's flagship innovation is the Cryptographic Data Mesh®, a decentralized infrastructure that ensures data provenance, tamper proof auditability, and transparent consent tracking. Anchored on the Hedera Hashgraph network, Acoer's products—including the Health Data Explorer and Cancer Trials Explorer—support clinical trial integrity, regulatory compliance, and data-driven decision-making.

The Problem: Traditional clinical trial processes are slow, costly, and heavily burdened by fragmented data management. Patient data is siloed across institutions, data sharing agreements are cumbersome, and compliance with HIPAA, GDPR, and other frameworks creates significant overhead. Moreover, patients and regulators often lack visibility into how data is used, eroding trust. These inefficiencies directly impact the speed of drug development, trial reliability, and patient safety.

Why Decentralization?: Acoer's approach leverages decentralization to address these systemic challenges. Rather than centralizing trial data, Acoer ensures that it remains in the custody of data owners—patients, hospitals, or sponsors—while cryptographic proofs of access, consent, and use are immutably logged on Hedera. This architecture prevents single points of failure, supports federated AI analysis, and enables transparent, verifiable collaboration between multiple stakeholders.

Key Stakeholders:

- Patients contributing data and participating in clinical trials.
- Hospitals and research institutions generating and managing trial datasets.
- Pharmaceutical companies and sponsors seeking faster, compliant trials.
- Regulators requiring auditable, privacy-compliant processes.
- Technology partners such as Hedera and Andromeda for ledger infrastructure and AI onboarding.

2. THE SOLUTION

- Platform Architecture: The Cryptographic Data Mesh® underpins Acoer's decentralized clinical trials platform. It enables distributed data nodes to interact via cryptographically signed transactions recorded on Hedera. Each consent, data submission, or access request is verifiable in real time. The Cancer Trials Explorer acts as a transparent dashboard, allowing stakeholders to monitor trial progress, data provenance, and compliance status.
- Decentralization Approach: Data never leaves its original repository. Instead, AI models and analytics tools are deployed through privacy-preserving methods, such as federated learning or distributed computation, to generate insights without exposing raw patient records. Every event—whether a patient's consent, a model update, or a sponsor's data query—is recorded immutably, creating trust across the ecosystem.

Decentralization Approach:

- ✓ Decentralized Data: Clinical trial data remains in its original location under local control, with the Cryptographic Data Mesh enabling federated search and analysis across distributed sources through cryptographic proofing without data movement.
- ✓ Decentralized Computation: AI models are trained locally at each institution using federated learning, with only model parameters (not raw patient data) shared and aggregated to create global models.
- ✓ Decentralized Governance: Smart contracts on Hedera Hashgraph automate and enforce data access controls, consent management, and compliance requirements through immutable, auditable records while maintaining institutional data ownership.

Technical Architecture

- Blockchain Layer: Hedera Hashgraph provides fast, low-cost, immutable logging.
- Data Mesh Layer: Distributed nodes across hospitals and sponsors, cryptographically linked.
- Application Layer: Health Data Explorer & Cancer Trials Explorer for visualization, dashboards, and AI insights.
- Integration: Partnerships with Andromeda's Pulsar streamline AI onboarding and orchestration for research teams.

Privacy & Security

Acoer embeds compliance into its system architecture. Patient data remains under institutional or individual control, never being centralized or shared without explicit, blockchain-recorded consent. This ensures HIPAA and GDPR compliance while reducing the risk of data breaches.

3. BUSINESS IMPLEMENTATION

- Governance Model Acoer operates as platform steward, providing the infrastructure and tools, while trial sponsors and institutions maintain data ownership and compliance responsibilities. Governance is enforced through smart contracts and cryptographic proofs, ensuring automated, auditable enforcement of consent and regulatory requirements.
- Economic Framework Decentralization reduces trial costs by minimizing reconciliation overhead, reducing compliance risks, and streamlining monitoring. Pharmaceutical sponsors benefit from accelerated timelines and greater transparency, while hospitals gain easier onboarding and regulatory relief. Patients benefit from clearer consent processes and assurances of data integrity. The open, vendor-neutral design also prevents lock-in, supporting a competitive and scalable ecosystem.
- Change Management To support adoption, Acoer emphasizes onboarding support, integration workshops, and iterative stakeholder feedback. Tools such as Andromeda's Pulsar simplify AI deployment for non-technical users. Transparency dashboards improve engagement with patients and regulators, reinforcing trust throughout the trial process.

4. IMPACT & OUTCOMES

Acoer's Cryptographic Data Mesh® has demonstrated real-world application through partnerships with Hedera Hashgraph for blockchain infrastructure and Andromeda for AI orchestration. The Health Data Explorer deployment in the Jefferson County Coroner/Medical Examiner's Office via the CDC MDI Connect project replaced manual data exchange with secure, FHIR-based interoperability. The company's Cancer Trials Explorer and Clinical Trials Explorer aggregate federal clinical trial data, demonstrating capability to handle large-scale healthcare datasets while maintaining data provenance and transparency.

The platform addresses critical research challenges through blockchain-enabled audit trails, automated consent tracking via RightsHash/ConsentHash, and federated data access without centralization. By keeping patient data under local control while enabling cryptographically verified collaboration, the architecture reduces reconciliation overhead, streamlines compliance monitoring, and enhances stakeholder trust. The vendor-neutral design positions the platform as infrastructure for multi-institutional collaboration, with growing adoption through partnerships with cancer research organizations and enterprise blockchain networks.

KEY INSIGHTS SUMMARY

The Acoer case illustrates how decentralized architectures can transform clinical trial operations. Key insights include:

- Blockchain-enabled provenance provides immutable audit trails that support trust-building and compliance verification across stakeholders.
- Decentralized data meshes enable collaboration without compromising privacy by keeping data under local control while allowing federated access.
- Vendor-neutral infrastructure reduces lock-in and supports scalability through open architecture design compatible with existing healthcare systems.
- Federated AI capabilities allow model training across institutions without centralizing sensitive patient data, preserving privacy while enabling collaborative insights.
- Transparency dashboards make data provenance and consent tracking visible to patients, researchers, and regulators, fostering ecosystem engagement.

REFERENCES:

- Acoer (2024) 'Federated learning and cryptographic data mesh: Technical explanation of privacy-preserving AI implementation', Acoer Blog, July. Available at: <https://www.acoer.com/news/blog/federated-learning>
- Acoer (2025) Cancer Trials Explorer: Clinical trials data aggregation and visualization platform. Available at: <https://cancertrials.acoer.com/>
- Acoer (2025) Health Data Explorer Platform: Medicolegal death investigation and public health reporting. Available at: <https://www.healthdataexplorer.io/>
- Hedera (2021) 'Acoer's RightsHash™ builds on Hedera to pioneer decentralized management and protection of user's rights', Hedera Blog, July. Available at: <https://hedera.com/blog/acoers-rightshash-builds-on-hedera-to-pioneer-decentralized-management-and-protection-of-users-rights>
- Hedera (2024) 'Acoer use case: Cryptographic data mesh technology and Hedera integration', November. Available at: <https://hedera.com/users/acoer>

3.6 BONUS RESOURCE: PROJECT MONAI – OPEN-SOURCE DATA EXCHANGE AND STANDARDS (GLOBAL)

Documented by: Thomas Egelhof, Chief Radiologist, Member of the Management Board, Merian Iselin Clinic for Orthopaedics and Surgery, Basel, Switzerland

Project MONAI (Medical Open Network for AI) is a key advancement in decentralized healthcare AI infrastructure. Started in 2019 as a partnership between NVIDIA and King's College London, it later grew to include the National Institutes of Health and many academic medical centers. Today, MONAI is the most widely used open-source framework for medical imaging AI. It has over 5.5 million downloads, is used by leading healthcare institutions worldwide, and has been cited in over 3,000 research papers. MONAI shows how open-source tools can support decentralized innovation while upholding high technical and clinical safety standards.

OPEN-SOURCE FOUNDATION AND COMMUNITY GOVERNANCE

MONAI uses the Apache 2.0 license, allowing free access for both research and commercial use. Its governance combines expert oversight with community input through working groups focused on areas like federated learning, ethical AI, clinical deployment, and specific medical applications. An advisory board of experts from top academic and industry organizations guides the project's direction via regular meetings and joint decisions. However, final choices stay with these expert groups rather than using distributed systems like blockchain voting.

The framework emphasizes modularity and standards, built on PyTorch with tools tailored for medical imaging tasks such as segmentation, classification, registration, and image generation. MONAI covers the full AI development process: MONAI Label uses AI to assist with annotation, cutting labeling time by up to 75%; MONAI Core provides efficient training tools with smart caching and GPU speed-ups that shorten training from days to hours; MONAI Deploy supports clinical use through containerized apps; and the MONAI Model Zoo offers over 40 pre-trained models for various imaging types.

FEDERATED LEARNING AND DECENTRALIZED DATA MANAGEMENT

MONAI's strongest decentralized feature is its support for federated learning, which lets healthcare organizations train AI models together without sharing sensitive patient data. It works smoothly with platforms like NVIDIA FLARE, Intel's OpenFL, and Substra, giving users options for setup. The Federated Tumor Segmentation (FeTS) project highlights this on a large scale—71 institutions from six continents used MONAI and OpenFL to train brain tumor segmentation models. This improved accuracy, even on new data, while keeping all patient information private. Each site trains locally on its own data, sharing only model updates—not the data itself—for combining into better global models.

For clinical use, MONAI Application Packages (MAPs) provide standardized, containerized apps that follow medical standards like DICOM and FHIR. These run on-site in hospitals, so patient data stays secure behind firewalls while supporting advanced AI for diagnosis and treatment. Mayo Clinic's Center for Augmented Intelligence in Imaging shows this in action, integrating MONAI-based AI into radiology workflows for faster diagnostics and better efficiency. The partnership with Siemens Healthineers' Digital Marketplace, announced in late 2024, extends this further by making Mayo Clinic's AI tools available to thousands of institutions globally with easy, no-code setup.

THE INFRASTRUCTURE FOR DECENTRALIZED HEALTHCARE AI

MONAI stands out as a core infrastructure for decentralized healthcare AI. While its governance is centralized in expert groups to ensure quality and safety, the open-source license and federated learning tools create a standard base for distributed work. Recent updates, including generative AI extensions in 2023 and new releases through 2024, enhance its capabilities for multimodal models and clinical deployment. MONAI proves a key point about healthcare AI: success often mixes centralized tech foundations with decentralized data handling and model building. This balances clinical reliability and safety with privacy and collaboration benefits. MONAI acts as the "foundation" for global institutions to create their own decentralized AI solutions, keeping data control and meeting regulations.

REFERENCES:

- Cardoso, M.J. et al. (2022) 'MONAI: An open-source framework for deep learning in healthcare', arXiv preprint arXiv:2211.02701. Available at: <https://arxiv.org/abs/2211.02701>
- Intel (2025) OpenFL Documentation. Available at: <https://openfl.readthedocs.io/en/latest/>
- MONAI (2025a) MONAI Project. Available at: <https://monai.io/>
- MONAI (2025b) Mayo Clinic MONAI Implementation. Available at: <https://monai.io/mayo-case-study.html>
- MONAI (2025c) MONAI GitHub Repository. Available at: <https://github.com/ProjectMONAI/MONAI>
- NVIDIA (2025a) MONAI Deploy. Available at: https://docs.monai.io/projects/monai_deploy/en/stable/
- Patel, S. et al. (2022) 'The federated tumor segmentation (FeTS) tool', Frontiers in Neuroinformatics. Available at: <https://www.frontiersin.org/journals/neuroinformatics>

3.7 BONUS RESOURCE: KNOWS – AI POWERED MEDICAL KNOWLEDGE PLATFORM (CHINA)

Documented by: LeiLei Tang, Medical Affairs Consultant, Ethers HealthData Foundation

KnowS represents an innovative approach to medical knowledge translation and scientific communication, developed by Shanghai Null Hypothesis Information Technology Co., Ltd. (上海零假设信息科技有限公司).

As Steven Yuwen, its founder, explains "We have developed the first medical AI agent called KnowS. It was initially developed in 2019 as a WeChat mini program, then in 2024 it further evolved as a web-based Enterprise Medical AI agent." The system has expanded from serving domestic pharmaceutical companies to international healthcare organizations with its English version enabling global uptake.

The platform operates with a simple, powerful trigger: a scientific question. KnowS responds by conducting real-time searches of peer-reviewed, full-text publications and generating a concise literature summary. Users can then command the AI to produce diverse tailored outputs including data-driven clinical guidance, customized educational content for Key Opinion Leader (KOL) engagement, accessible popular science essays for patient education, and structured academic reports with professionally formatted presentation slides. This multi-format capability addresses a critical gap in medical communications—the need to translate complex scientific findings into appropriate formats for different audiences. With the integration of natural language processing and medical domain expertise, KnowS enables rapid knowledge extraction and translation that can significantly accelerate the academic engagement process.

Currently KnowS operates as a fundamentally centralized AI system, with all processing occurring within its infrastructure and Shanghai Null Hypothesis maintaining primary control. Despite this, KnowS offers compelling implications for future decentralized scientific knowledge platforms. Future iterations could leverage federated learning where institutions contribute to knowledge curation while maintaining control over proprietary insights—potentially scaling KnowS's capabilities across the global medical research ecosystem while ensuring equitable access and institutional independence.

REFERENCES:

- MedKnowS (2025) MedKnowS Official Website. Available at: <https://www.medknows.com/>
- Interview with Steven Yuwen, Founder of KnowS. Interviewed by LeiLei Tang for Decentralized AI in Healthcare Report

3.8 CHAPTER INSIGHTS AND CONCLUSIONS

This chapter examines decentralized data collection and processing through implementations spanning competitive pharmaceutical collaboration, multi-institutional research, vulnerable population care, and clinical trial integrity. Several critical findings guide organizations approaching decentralized AI at the data collection stage.

Key Finding 1: Competitive Collaboration Requires Technical Architecture Plus Governance Frameworks

MELLODDY's 10-pharmaceutical-company consortium demonstrates that federated learning prevents raw data exposure, but technical architecture alone proves insufficient. The consortium required three years developing governance frameworks defining intellectual property rights, data contribution expectations, and conflict resolution mechanisms before achieving operational federated learning. Mayo Clinic Platform similarly invested heavily in legal frameworks and institutional coordination alongside technical implementation. Organizations pursuing competitive collaboration must allocate comparable resources to governance design as to technical development, recognizing that trust mechanisms require explicit negotiation rather than emerging spontaneously from privacy-preserving technology.

Key Finding 2: Decentralization Enables Previously Impossible Scale and Diversity

Decentralized approaches don't merely replicate centralized capabilities with better privacy—they enable a fundamentally different scale, which is impossible through data aggregation. MELLODDY's federated learning accessed compound libraries orders of magnitude larger than any individual company possesses, improving model performance precisely because decentralization unlocked proprietary datasets that centralization could never aggregate. SEOVE's network demonstrates similar dynamics at a smaller scale, where federated learning enables knowledge sharing that regulatory and ethical constraints prevent from centralizing. Organizations should evaluate decentralization as strategic capability enabling access to datasets that competitive, regulatory, or ethical factors make inaccessible to centralized approaches.

KEY FINDING 3: RESOURCE-CONSTRAINED ENVIRONMENTS CAN IMPLEMENT FEDERATED SOLUTIONS

SEOVE's implementation in resource-limited Brazilian healthcare challenges assumptions that decentralized AI requires extensive infrastructure or expertise. The organization successfully deployed federated learning despite severe provider shortages and limited technical capacity by leveraging open-source frameworks, external partnerships, and phased implementation. Decentralization accessibility depends more on problem-solution fit and partnership models than absolute resource availability. Resource-constrained organizations should identify strategic partners, open-source tools, and phased approaches aligned with available capabilities.

PRACTICAL IMPLICATIONS FOR STAKEHOLDER GROUPS

Healthcare institutions should identify collaborators with complementary datasets, aligned incentives, and compatible regulations before selecting technical architecture. Medtech startups should position federated learning as enabling previously impossible collaborations rather than merely improving compliance. Patient advocacy groups should demand transparency into which institutions participate in federated networks, what governance protects patient interests, and how benefits flow back to communities. Regulators should assess governance frameworks, contribution transparency, and benefit distribution alongside privacy and security criteria, recognizing that decentralization shifts risk from centralized breaches to distributed governance failures requiring different oversight.

CHAPTER 4: DECENTRALIZED MODEL DEVELOPMENT AND TRAINING

Foreword by the author: Natalia Sofia, Data Ecosystem Lead, Ethers HealthData Foundation

My path into decentralized systems started on the blockchain side of the house—designing tokenized data-sharing pilots and privacy-preserving consent flows long before “sovereign AI” was a headline. That work taught me two durable lessons: first, sensitive data rarely needs to move if intelligence can; second, governance is the real product. Reading this chapter through that lens, the common thread across all cases isn’t just new models or cheaper GPUs—it’s architecture that lets institutions keep control of data, policy, and risk while still compounding shared learning.

DeepSeek’s sovereign deployments show what happens when open models meet strict data-locality rules: hospitals can run state-of-the-art reasoning and VLM capabilities without exporting a single pixel or token of PHI. Akash and Aethir tackle the adjacent bottleneck— computation scarcity—by turning GPUs into a marketplace rather than a gate. Bittensor extends the idea further, aligning incentives so that models improve collaboratively without central custody of data. And the ASI Alliance ties agents, data markets, and computers into workflows that look a lot like real clinical operations—routing tasks, checking constraints, and recording provenance.

From a blockchain practitioner’s standpoint, the novelty isn’t “crypto in healthcare”; it’s credible accountability at scale.

On-chain audit trails, open licensing, and stake-based curation give us tools to answer the questions clinicians actually ask: Who touched the data? Why did the model decide this? What happens if it fails? The cases in this chapter don't claim perfection, but they demonstrate workable patterns: keep data local, push models to the edge, price resources in open markets, and make governance legible.

If we hold to those principles, decentralized AI can lower costs and variance without lowering standards. It won't replace clinical judgment; it will make it safer to apply. That is the promise I've seen in blockchain, now maturing in healthcare AI: not maximal decentralization for its own sake, but the minimum decentralization necessary to preserve trust where it matters most.

The views expressed are those of the author and do not necessarily reflect those of the employer.

4.1 CHAPTER OVERVIEW

Chapter 4 examines decentralized model development and training through four primary case studies spanning sovereign AI deployment, decentralized cloud computing, collaborative machine learning networks, and autonomous agent platforms.

- DeepSeek R1/VL demonstrates how open-source foundation models enable hospitals to run state-of-the-art reasoning locally while maintaining data sovereignty across 100+ Chinese institutions.
- Akash Network and Aethir provide decentralized GPU infrastructure—Akash offering community-driven computation marketplaces at up to 85% cost savings, while Aethir delivers enterprise-grade resources for high-performance AI workloads.
- Bittensor creates an open marketplace where models are trained collaboratively through incentivized contribution, with applications in medical imaging and drug discovery.
- The ASI Alliance merges autonomous agents, data marketplaces, and computation infrastructure to enable decentralized healthcare workflows, with specialized models demonstrating superhuman performance in multi-domain diagnostics.

Each case study follows the structure established in Chapter 3: Organization & Challenge, The Solution, Business Implementation, and Impact & Outcomes. This format enables readers to identify which decentralization pillars each implementation exhibits (Chapter 4 emphasizing decentralized computation and control alongside data management and governance), evaluate solution maturity, and assess applicability to their organizational constraints around computation costs, data sovereignty, and regulatory compliance.

The chapter addresses a critical bottleneck: the concentration of computational resources and model development expertise in centralized platforms. Decentralized approaches democratize access to advanced AI capabilities while enabling institutions to maintain control over sensitive data and participate meaningfully in model development. However, decentralization introduces new risks including coordination complexity, quality assurance across distributed nodes, and governance challenges that differ fundamentally from centralized oversight. The cases demonstrate how decentralization transforms AI training economics—turning GPUs into open markets, incentivizing collaborative improvement, and enabling sovereign deployment—while maintaining medical-grade technical sophistication and addressing these emerging challenges.

4.2 CASE STUDY: DEEPSEEK R1/VL – SOVEREIGN AI FOR MEDICAL APPLICATIONS (CHINA)

LARGE-SCALE DEPLOYMENT OF OPEN-SOURCE FOUNDATION MODELS ACROSS CHINESE HEALTHCARE INSTITUTIONS

Documented by: LeiLei Tang, Medical Consultant, Ethers HealthData Foundation

CASE STUDY AT-A-GLANCE: DEEPSEEK R1/VL (CHINA)

<p>MATURITY LEVEL: Production (operational at scale)</p> <p>GEOGRAPHY: China (100+ hospitals nationwide)</p> <p>REPLICABILITY: High (open-source model, replicable approach)</p>	
<p>THE CHALLENGE:</p> <p>Chinese hospitals needed state-of-the-art AI reasoning capabilities while maintaining data sovereignty and compliance with national data localization requirements.</p>	<p>THE SOLUTION:</p> <p>Open-source sovereign AI foundation models deployed locally across 100+ hospitals, enabling advanced reasoning and multimodal medical analysis without cloud dependency.</p> <ul style="list-style-type: none">✅ Pillar 1: Decentralized Data Management (Local storage)✅ Pillar 2: Decentralized Computation (Local inference)
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none">• 100+ tertiary hospitals running DeepSeek locally• Superhuman performance on medical licensing exams• Multimodal capabilities (text, images, clinical data)• Complete data sovereignty within Chinese borders	<p>LESSONS LEARNED:</p> <p>Sovereign AI demonstrates that nation-scale decentralization is viable when open-source models enable local deployment without sacrificing performance or requiring centralized cloud infrastructure.</p>

1. ORGANIZATION & CHALLENGE

Organization Profile: DeepSeek is a Chinese AI research company founded in 2023 by High-Flyer Capital Management, developing open-source large language models and vision-language models. The company operates under China's AI sovereignty framework, creating domestic alternatives to Western AI systems. DeepSeek has released multiple model variants including DeepSeek-R1 (a 671 billion parameter reasoning model) and DeepSeek-VL (a multimodal vision-language model), both distributed under permissive open-source licenses. By early 2025, nearly 100 Chinese hospitals have announced or initiated local deployment of DeepSeek models, representing one of the world's largest sovereign AI healthcare implementations.

The Problem: China's healthcare system faced critical challenges that centralized AI solutions could not adequately address:

- **Data Sovereignty Requirements:** China's Personal Information Protection Law (PIPL) mandates that sensitive medical data remain within national boundaries, making dependence on foreign AI platforms legally problematic
- **Regional Healthcare Disparities:** Significant gaps in medical expertise between urban centers (Shanghai, Beijing) and rural areas, with limited access to specialist radiologists and diagnostic capabilities
- **Infrastructure Heterogeneity:** Vast differences in hospital IT capabilities, from advanced tertiary hospitals with sophisticated systems to rural clinics with basic infrastructure
- **Diagnostic Accuracy Variations:** Inconsistent diagnostic quality across regions due to varying levels of medical expertise and equipment standards

Why Decentralization?: Traditional centralized AI approaches were incompatible with China's regulatory environment and healthcare infrastructure diversity. DeepSeek's open-source, locally deployable approach enables hospitals to maintain data sovereignty while accessing advanced AI capabilities, ensuring compliance with PIPL while addressing regional healthcare disparities through standardized AI assistance.

Key Stakeholders: The ecosystem includes the Chinese Ministry of Health coordinating national AI strategy, nearly 100 participating hospitals ranging from tertiary medical centers to rural clinics (as reported by local media outlets), DeepSeek as the core technology provider, local system integrators facilitating hospital deployments, and medical AI companies such as Schule Informatics developing specialized applications built on DeepSeek's foundation models.

2. THE SOLUTION

DeepSeek's solution implements a sovereign AI approach where open-source foundation models are deployed locally within Chinese hospital networks, enabling advanced medical AI capabilities while maintaining complete data sovereignty. The platform combines DeepSeek-R1's 671-billion parameter reasoning capabilities for complex diagnostic tasks with DeepSeek-VL's multimodal vision-language processing for medical imaging analysis, both operating entirely within institutional boundaries to ensure PIPL compliance and data security.

Decentralization Approach:

- ✓ Decentralized Data: All patient data remains within local hospital infrastructure with no external data sharing
- ✓ Decentralized Computation: AI processing occurs locally on hospital servers or approved domestic cloud platforms
- ✓ Decentralized Governance: Individual hospitals control model deployment, customization, and clinical integration decisions

Technical Architecture - Sovereign AI Implementation: DeepSeek's approach implements comprehensive local AI deployment designed for China's regulatory and infrastructure environment:

- Open-Source Foundation Models: DeepSeek-R1 (671B parameters) for complex reasoning tasks and DeepSeek-VL for multimodal medical imaging analysis, both released under MIT license enabling local customization and deployment
- Local Deployment Options: Models deployable on hospital-owned servers, approved domestic cloud platforms (Alibaba Cloud), or hybrid configurations based on institutional capabilities and security requirements
- Hospital-Specific Calibration: Open-source framework enables fine-tuning to local disease patterns, imaging equipment specifications, and clinical workflows, with different model sizes (32B, 70B, 671B parameters) matched to hospital computational capacity
- Medical Integration Protocols: Seamless integration with existing Hospital Information Systems (HIS), Picture Archiving and Communication Systems (PACS), and Electronic Medical Record (EMR) systems through standardized APIs and protocols

Privacy & Security Architecture: Comprehensive sovereignty protection ensuring complete data localization:

- PIPL Compliance: All processing occurs within Chinese territorial boundaries with no international data transmission
- Local Data Processing: Patient information never leaves hospital networks, with AI analysis performed entirely on local infrastructure
- Audit Trail Capabilities: Complete traceability of AI decision-making processes for regulatory compliance and clinical oversight
- Differential Privacy Options: Additional privacy protection layers available for institutions requiring enhanced anonymization

3. BUSINESS IMPLEMENTATION

- Implementation Timeline: DeepSeek's healthcare deployment represents rapid sovereign AI adoption across China's medical system:
 - 2023: DeepSeek founded with initial model development and research partnerships
 - Late 2024: DeepSeek-R1 and VL models released open-source, first hospital pilots initiated
 - November 2024: Shanghai Eastern Hospital launches Med-Go clinical testing using DeepSeek-R1

- Early 2025: Nearly 100 hospitals announce DeepSeek deployment plans across multiple medical specialties
- Mid-2025: Documented implementations include diagnostic support, medical imaging analysis, and administrative automation
- Governance Model: Distributed governance where individual hospitals maintain autonomy over AI deployment while participating in informal knowledge-sharing networks. The Chinese government provides regulatory guidance through health ministry directives, while DeepSeek maintains open source development with community contributions from participating institutions.
- Economic Framework: Open-source licensing reduces AI access costs dramatically compared to proprietary alternatives, while hospitals invest in local infrastructure and customization. Value creation occurs through improved diagnostic efficiency, reduced specialist consultation needs, and enhanced care standardization across regional facilities.
- Change Management: Implementation requires significant clinical workflow integration, staff training on AI-assisted diagnostics, and establishment of human-AI collaboration protocols. Success depends on addressing physician concerns about AI reliability while demonstrating clear clinical value.

4. IMPACT & OUTCOMES

Quantitative Results:

- Scale of Deployment: Nearly 100 Chinese hospitals have announced or initiated DeepSeek implementations across multiple medical specialties, as reported by local media outlets
- Clinical Validation: DeepSeek-R1 achieved superior diagnostic reasoning compared to other LLMs, with average Likert scores of 3.61 vs 3.22 (ChatGPT) and 3.13 (Llama 3.1-405B) on medical tasks ($p < 0.005$). However it reported lower performance in imaging report summarization tasks (Tordjman, M. et al. (2025))
- Diagnostic Accuracy: Huashan Hospital achieved >95.2% accuracy in lung nodule differentiation using multimodal DeepSeek integration
- Efficiency Improvements: Med-Go integration at Shanghai Eastern Hospital reported >10% improvement in diagnostic accuracy, with more pronounced gains in complex cases

Qualitative Benefits:

- Healthcare Equity: Standardized AI capabilities reduce diagnostic disparities between urban and rural healthcare facilities by providing consistent specialist-level analysis
- Data Sovereignty: Complete compliance with Chinese data protection requirements while enabling advanced AI capabilities, proving sovereign AI viability in regulated industries
- Clinical Workflow Enhancement: Integration with existing hospital systems (PACS, EMR) streamlines radiologist workflows and reduces interpretation time for routine cases
- Medical Education: AI-assisted diagnostics provide learning opportunities for medical professionals in underserved regions, improving overall system capability
- Innovation Acceleration: Open-source approach enables rapid customization and improvement by local medical institutions, fostering indigenous AI development

KEY INSIGHTS SUMMARY

- **Sovereign AI Feasibility:** Successfully demonstrated that open-source foundation models can provide competitive medical AI capabilities while maintaining complete national data sovereignty, proving alternatives to Western AI platforms are viable for sensitive applications
- **Scale-Driven Network Effects:** Deployment across nearly 100 hospitals creates knowledge sharing opportunities and collective improvement in AI model performance, showing how distributed implementation can generate network benefits without centralized data aggregation
- **Infrastructure Adaptability:** DeepSeek's modular approach successfully accommodates China's diverse hospital IT capabilities, from advanced tertiary centers to resource constrained rural facilities, demonstrating scalable AI deployment across heterogeneous healthcare infrastructure

REFERENCES AND ADDITIONAL RESOURCES:

- Gradient Flow (2025) 'DeepSeek in Action: Practical AI Applications Transforming Chinese Healthcare', 7 March. Available at: <https://gradientflow.com/deepseek-in-action-practical-ai-applications-transforming-chinese-healthcare/>
- Healthcare IT News (2025) 'Chinese health players begin integrating DeepSeek'. Available at: <https://www.healthcareitnews.com/news/asia/chinese-health-players-begin-integrating-deepseek>
- PandaYoo Analysis (2025) 'DeepSeek Diagnosis: Is AI Finally Revolutionizing Healthcare in China?' Available at: <https://pandayoo.com/post/deepseek-diagnosis-is-ai-finally-revolutionizing-healthcare-in-china/>
- Tordjman, M., Liu, Z., Yuce, M., Forghani, R., Tang, A., Kadoury, S. and Duron, T. (2025) 'Comparative benchmarking of the DeepSeek large language model on medical tasks and clinical reasoning', Nature Medicine. doi: 10.1038/s41591-025-03726-3.

4.3 CASE STUDY: AKASH NETWORK AND AETHIR FOR DECENTRALIZED CLOUD COMPUTING (GLOBAL)

COMPLEMENTARY DECENTRALIZED GPU CLOUD PLATFORMS ENABLING SCALABLE AI MODEL DEVELOPMENT AND TRAINING

Documented by: Carmen Cucul, Blockchain & AI Healthcare Consultant, Ethers HealthData Foundation

CASE STUDY AT-A-GLANCE: AKASH NETWORK & AETHIR (GLOBAL)

<p>MATURITY LEVEL: Production (both operational) GEOGRAPHY: Global (distributed infrastructure) REPLICABILITY: High (marketplace model scales globally)</p>	
<p>THE CHALLENGE: Healthcare AI developers face prohibitive cloud computing costs and vendor lock in, limiting access to GPU resources needed for training and inference.</p>	<p>THE SOLUTION: Decentralized computation marketplaces: Akash (permissionless community model, 85% cost reduction) and Aethir (enterprise-grade GPU infrastructure, 425K+ containers globally).</p> <ul style="list-style-type: none"> ✓ Pillar 2: Decentralized Computation ✓ Pillar 3: Decentralized Governance (Akash: token-based)
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none"> • Up to 85% cost reduction vs traditional cloud (Akash) • Enterprise-grade GPU access without vendor lock-in • Democratized access to high-performance computation 	<p>LESSONS LEARNED: Decentralized computation addresses both cost barriers and vendor lock-in concerns. Different models (permissionless vs enterprise-managed) serve different market segments.</p>

Note: This case study examines two complementary decentralized cloud computing platforms rather than a single implementation. Akash Network represents a fully permissionless, community-driven approach supporting diverse AI applications and development across multiple industries, while Aethir provides enterprise-grade decentralized GPU infrastructure optimized for large-scale, high performance AI workloads primarily in gaming and AI sectors (currently). While both platforms serve various industries, this analysis focuses on their existing and potential future applications in healthcare.

1. ORGANIZATION & CHALLENGE

Organization Profile:

Akash Network is a fully open-source, permissionless decentralized cloud computing platform built on the Cosmos blockchain, founded in 2019 and operated by Overclock Labs. The network enables anyone worldwide to contribute computation resources (CPU, GPU, storage) to a global marketplace, creating what they call the "Supercloud"—a censorship-resistant alternative to traditional cloud providers. With over 58 million GB hours of memory leased globally and prices up to 85% lower than traditional cloud services, Akash has processed significant healthcare workloads including partnerships with Solve.Care for patient data sovereignty.

Aethir is an enterprise-focused decentralized GPU cloud platform launched in 2021, specifically designed for AI and high-performance computing workloads across gaming, AI development, and enterprise applications. The platform aggregates GPU resources from enterprises, data centers, miners, and retail providers into a unified network supporting 425,000+ GPU containers globally. With over \$125 million in annual recurring revenue and partnerships with NVIDIA for H200 and B200 GPU access, Aethir primarily serves gaming studios, AI companies, and enterprise clients requiring intensive computational power. Healthcare applications represent a potential growth area rather than a current primary focus.

The Problem: Healthcare AI development faces critical infrastructure barriers limiting innovation and accessibility:

- **Computational Resource Scarcity:** AI development across all sectors, including healthcare, requires substantial GPU resources for training complex models, but access to high performance computing is dominated by expensive centralized providers
- **Cost Barriers:** Traditional cloud GPU pricing creates prohibitive costs for research institutions, startups, and smaller organizations across all sectors developing AI applications, with healthcare particularly affected due to its resource-constrained environment
- **Geographic Limitations:** Centralized cloud infrastructure concentrates computation resources in specific regions, creating latency issues for real-time healthcare applications and limiting global access to advanced AI capabilities
- **Vendor Lock-in:** Dependence on major cloud providers (AWS, Google Cloud, Azure) creates strategic vulnerabilities for healthcare organizations and limits flexibility in model deployment and data management
- **Privacy and Sovereignty Concerns:** Data sovereignty requirements across regulated industries, particularly healthcare, often conflict with centralized cloud architectures, especially for cross-border collaborations and sensitive data protection

Why Decentralization?: Decentralized cloud computing addresses healthcare AI infrastructure challenges by democratizing access to computational resources, reducing costs through marketplace competition, enabling geographic distribution of computation closer to data sources, and providing censorship-resistant infrastructure that supports data sovereignty requirements. For healthcare specifically, this approach enables federated learning implementations, privacy-preserving AI development, and equitable access to advanced computational capabilities regardless of institutional size or location.

Key Stakeholders: The ecosystem includes AI researchers and developers across multiple industries, medical institutions exploring federated learning applications, biotech startups developing AI-driven solutions, GPU owners (including hospitals, research centers, data centers) monetizing underutilized resources, and organizations seeking to democratize AI capabilities.

2. THE SOLUTION

Both platforms implement decentralized cloud computing but target different segments of the healthcare AI market through complementary approaches:

- Akash Network's Community-Driven Model operates as a fully permissionless marketplace where anyone can contribute computing resources through a reverse auction mechanism. Organizations deploy AI workloads by specifying requirements in deployment files, receiving competitive bids from global providers, and paying with \$AKT tokens. The platform's Solve.Care partnership demonstrates concrete healthcare applications, enabling patient-controlled health data through Care.Nodes deployed on Akash's distributed infrastructure. This approach particularly benefits smaller AI projects across all sectors, with potential applications in federated learning pilot programs and research initiatives requiring flexible, cost-effective computation access.
- Aethir's Enterprise-Grade Framework aggregates high-performance GPUs from institutional providers into a managed decentralized network optimized for large-scale AI workloads primarily serving gaming and AI companies. While the platform's current focus is on gaming and general AI applications future developments might involve healthcare enterprises in need to access enterprise grade GPU clusters (including NVIDIA H100s, H200s, and B200s) through standardized APIs with 24/7 support.

Decentralization Approach:

Akash Network:

- ☒ Decentralized Data: Supports patient data sovereignty through Care.Nodes and local data processing
- ☒ Decentralized Computation: Fully permissionless computation marketplace with global provider participation
- ☒ Decentralized Governance: Community-governed through \$AKT token holders and open-source development

Aethir:

- ☒ Decentralized Data: Enterprise-controlled data processing with edge computing capabilities (potential healthcare applications)
- ☒ Decentralized Computation: Managed network of distributed GPU resources with institutional providers
- ☐ Decentralized Governance: Centralized platform management with \$ATH token utilities for access and rewards

Technical Architecture:

- **Akash Network Architecture:** Built on Kubernetes and the Cosmos blockchain, Akash operates as a permissionless marketplace where users define application requirements through Service Definition Language (SDL) and receive competitive bids from global providers through reverse auctions. All transactions are recorded on-chain for transparency, while the distributed provider network enables edge computing capabilities that support data sovereignty requirements and real-time applications across various industries, with potential healthcare applications benefiting from up to 85% cost reductions compared to traditional cloud providers.
- **Aethir Enterprise Infrastructure:** Aethir aggregates enterprise-grade GPU resources (including NVIDIA H100s, H200s, and B200s) through a proprietary orchestration layer that manages containerized workloads across 425,000+ GPU containers in 93+ global locations. The platform provides standardized APIs for enterprise integration, managed resource allocation with 24/7 support, and edge computing capabilities that bring computation closer to data sources, enabling scalable deployment from single nodes to 4,000-GPU clusters while maintaining enterprise-grade security and performance guarantees primarily for gaming and AI companies, with potential applications across regulated industries including healthcare.

3. BUSINESS IMPLEMENTATION

- **Implementation Timeline:**

Akash Network Multi-Industry Growth:

- 2019-2021: Platform development and initial AI/blockchain research deployments across various sectors
- 2022: Solve.Care partnership enabling patient data sovereignty through Care.Nodes
- 2023-2024: Expansion to diverse AI workloads including machine learning, gaming, and research applications
- 2025: Growing ecosystem of AI projects across multiple industries leveraging decentralized computation

Aethir Enterprise Scaling:

- 2021-2022: Platform launch focused on gaming and enterprise GPU aggregation
- 2023: Major partnerships with NVIDIA, expansion to 425,000+ GPU containers serving primarily gaming and AI companies
- 2024: \$125M ARR achievement through gaming and enterprise clients, H200 GPU integration for advanced AI workloads
- 2025: B200 Blackwell architecture support, continued focus on gaming and general enterprise applications

● Governance Models:

Akash: Community governance through \$AKT token holders voting on protocol upgrades, network parameters, and resource allocation decisions. Open-source development enables transparent platform evolution driven by community needs.

Aethir: Managed platform governance with \$ATH token providing network access, staking rewards, and quality assurance mechanisms. Enterprise partnerships guide platform development priorities.

4. IMPACT & OUTCOMES

Quantitative Results:

Akash Network:

- Global Reach: Over 58 million GB hours of computation resources leased across diverse AI and blockchain applications in multiple industries
- Cost Reduction: Up to 85% cost savings compared to traditional cloud providers, making AI accessible to resource-constrained organizations across all sectors
- Network Growth: Active global provider network with deployments across multiple continents and industries
- Verified Healthcare Partnership: Solve.Care integration enabling patient data sovereignty

Aethir:

- Enterprise Scale: 425,000+ GPU containers deployed globally, supporting large-scale AI initiatives primarily in gaming and enterprise sectors
- Revenue Performance: \$125+ million annual recurring revenue, demonstrating sustainable business model primarily serving gaming studios and AI companies
- Hardware Leadership: First decentralized platform offering NVIDIA H200 and B200 GPUs for cutting-edge AI applications across various industries
- Global Infrastructure: 93+ locations worldwide enabling low-latency applications across multiple sectors (with potential for healthcare applications)

Qualitative Benefits:

- Democratized AI Access: Both platforms enable smaller organizations and research institutions to access advanced AI capabilities previously limited to well-funded entities. For healthcare, this could foster innovation in underserved markets, though widespread adoption remains limited currently.
- Enhanced Data Sovereignty: Decentralized computation infrastructure supports data sovereignty requirements across regulated industries through local data processing. Healthcare applications like the Solve.Care partnership demonstrate concrete implementation of patient data control.

- **Innovation Acceleration:** Reduced infrastructure costs enable broader AI experimentation across sectors. Healthcare could benefit through accelerated medical imaging, drug discovery, and personalized medicine development, though these represent potential rather than documented implementations.
- **Infrastructure Resilience:** Distributed networks provide redundancy for critical applications. For healthcare, this could ensure care continuity, though extensive deployment validation is still needed.

KEY INSIGHTS SUMMARY

- **Complementary Decentralization Models:** Akash Network's fully permissionless approach and Aethir's enterprise-managed framework demonstrate different valid paths to decentralizing cloud infrastructure, each serving distinct market segments across multiple industries while maintaining core decentralization benefits.
- **Cost-Driven Adoption Potential:** Significant cost reductions (up to 85% for Akash, competitive enterprise rates for Aethir) could prove essential for healthcare AI adoption, enabling organizations to redirect resources from infrastructure to research and patient care.
- **Industry-Agnostic Infrastructure with Healthcare Applications:** Both platforms address broad infrastructure needs including data sovereignty, regulatory compliance, real-time processing capabilities, and global accessibility. These features align well with healthcare requirements, paving the way for potential increased adoption in the future.

REFERENCES AND ADDITIONAL RESOURCES:

- Aethir (2025) Aethir Official Website. Available at: <https://aethir.com/>
- Aethir Ecosystem (2025) 'Aethir, GAIB, and GMI Cloud launch first decentralized AI computation powered by H200 GPUs'. Available at: <https://ecosystem.aethir.com/blog/posts/aethir-gaib-and-gmi-cloud-launch-first-decentralized-ai-compute-powered-by-h200-gpus>
- Akash Network (2025) Akash Network Official Website. Available at: <https://akash.network/>
- Akash Network (2025) 'Akash Network and Solve.Care bring true patient data ownership to the healthcare industry'. Available at: <https://akash.network/blog/akash-network-and-solve-care-bring-true-patient-data-ownership-to-the-healthcare-industry/>

Disclaimer: References to the \$AKT and \$ATH token are descriptive only and should not be construed as financial advice. Readers should conduct independent due diligence before engaging with any blockchain platform or digital asset.

4.4 CASE STUDY: BITTENSOR, A DECENTRALIZED NETWORK FOR SHARING AI MODELS (GLOBAL)

DECENTRALIZED MACHINE LEARNING NETWORK INCENTIVIZING COLLABORATIVE AI FOR HEALTHCARE DIAGNOSTICS AND BIOTECH RESEARCH

Documented by: Thomas Egelhof, Chief Radiologist, Member of the Management Board, Merian Iselin Clinic for Orthopaedics and Surgery

CASE STUDY AT-A-GLANCE: BITTENSOR (GLOBAL)

<p>MATURITY LEVEL: Production (network operational; healthcare apps early-stage)</p> <p>GEOGRAPHY: Global (distributed network)</p> <p>REPLICABILITY: Medium (requires blockchain/token expertise)</p>	
<p>THE CHALLENGE:</p> <p>Centralized AI development concentrates expertise and computational power, limiting innovation and creating single points of failure.</p>	<p>THE SOLUTION:</p> <p>Open marketplace where AI models are collaboratively trained through incentivized contribution. Validators assess model quality; miners contribute computation/data; tokens reward valuable contributions.</p> <ul style="list-style-type: none">✓ Pillar 2: Decentralized Computation✓ Pillar 3: Decentralized Governance (Token-based)
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none">• 32 specialized AI subnets (including medical imaging)• Incentivized model improvement through token rewards• No single entity controls the AI model• Healthcare applications in drug discovery emerging	<p>LESSONS LEARNED:</p> <p>Token-based incentives can coordinate global AI development without central control, but healthcare applications require clinical validation and regulatory compliance frameworks.</p>

NOTE: Bittensor covers a broad range of AI applications, from text generation to vision models. This article focuses on its uses in healthcare and biotech to show how a rewarded, decentralized machine learning system can support secure data sharing, faster drug discovery, and privacy-protected diagnostics.

1. ORGANIZATION & CHALLENGE

Organization Profile: Bittensor is a global decentralized machine learning protocol founded in 2021 by the Opentensor Foundation, led by Canadian engineers Jacob Steeves and Ala Shaabana. It has no traditional headquarters and relies on community-driven development. The project raised no formal venture capital funding at inception, instead growing through fair-launch mining. As of October 2025, the network has expanded to over 128 active subnets, supported by thousands of miners, validators, and developers who hold \$TAO tokens. Bittensor creates an open marketplace for AI, where models are trained collaboratively and rewarded based on their value. While it supports many applications, it increasingly focuses on healthcare through subnets for cancer detection, protein folding, and federated learning.

The Problem: Traditional healthcare AI faces major barriers that slow innovation and limit access:

- **Data Silos and Privacy Concerns:** Medical data is scattered across providers. Centralized systems risk data breaches and offer few rewards for secure sharing in research or diagnostics.
- **Inefficient Model Training:** Tasks like drug discovery or image analysis need costly, proprietary hardware. This leads to slow progress, biases, and exclusion of smaller researchers or clinics.
- **Limited Collaboration:** Centralized AI platforms give power to big tech or pharmaceutical companies. This restricts open access to models for personalized medicine or disease prediction, while issues like bias continue without clear oversight.
- **Scalability and Cost Hurdles:** Biotech simulations, such as protein folding, demand high computing power. Regulatory challenges widen gaps in global healthcare innovation.

Why Decentralization? Bittensor decentralizes AI by building a rewarded marketplace. Miners contribute models through subnets, validators check quality, and users access AI without permission. This promotes ethical, team-based healthcare AI. Data owners keep control, biases are reduced through spread-out validation, and rewards open up participation in advances like early cancer detection.

Key Stakeholders: The system involves healthcare researchers (such as biotech labs training models), AI developers creating medical subnets, hospitals and clinics using diagnostics, patients sharing anonymized data, pharmaceutical companies, and regulators. It operates in a global framework that works with or without cryptocurrency.

2. THE SOLUTION

Bittensor's protocol builds a blockchain-based marketplace where nodes, called neurons, create specialized subnets for AI tasks. Miners train and submit models using smart contracts, validators rank them for rewards, and outputs are tokenized for ownership and sharing. In healthcare, this enables privacy-focused workflows, like federated cancer detection or protein analysis. Researchers, providers, and patients can work together on AI solutions while keeping data control and earning \$TAO.

Decentralization Approach:

- ✓ Decentralized Data: Healthcare datasets are shared and tokenized across subnets, allowing secure, rewarded access without central holders.
- ✓ Decentralized Computation: Model training spreads across global miners using everyday hardware, supporting federated learning for sensitive medical data.
- ✓ Decentralized Governance: Network updates and subnet registrations use \$TAO-weighted voting. Validators automate quality checks for ethical AI in health uses.

Technical Architecture: Network Governance Process: Bittensor uses a proof-of-intelligence consensus designed for AI in sensitive fields like healthcare:

- Subnet Proposal: Developers register new subnets (for example, for medical imaging) through the DAO. They specify tasks, rewards, and validation rules. Validators first screen them to ensure they fit network ethics and healthcare needs.
- Model Submission: Miners submit trained models or computations via the subnet's chain. They use tools like federated learning for privacy, avoiding raw data sharing in biotech tasks.
- Validation Process: Validators, who stake \$TAO, evaluate submissions with Yuma Consensus. They score based on accuracy, novelty, and value—for instance, diagnostic precision in cancer models—using field-specific benchmarks.
- Multi-Stage Consensus: \$TAO holders vote on broad governance, like protocol upgrades, through on-chain proposals. Subnet-specific decisions are weighted by stake. Penalties like slashing prevent bad actions.
- Reward Execution: Smart contracts on the Polkadot-derived chain distribute \$TAO emissions (1 \$TAO per block, with halvings over time) based on rankings. This includes milestones for long-term projects like drug discovery simulations.
- Tokenization Framework: Valuable models and datasets are tokenized as assets. This lets the community govern licensing, royalties, and updates while providing traceability for medical rules.
- Continuous Monitoring: Validators and community dashboards track performance, such as model bias in diagnostics. Underperforming nodes are automatically removed.

AI-Integrated Healthcare Portfolio: Bittensor hosts focused subnets that show decentralized AI in healthcare:

- Cancer-AI Subnet: A subnet for cancer detection using open-source tools to analyze imaging data. It supports early diagnosis with models from global miners, targeting over 88% accuracy in spotting anomalies in medical scans.
- FLock Subnet (SN96): A federated learning platform for healthcare AI. It allows model training on spread-out patient data without centralizing it, with uses in diagnostics and personalized medicine. It integrates with tools like Alibaba Cloud's Qwen.
- Protein Folding Subnet (Mainframe/SN25): Bittensor's first decentralized science subnet, operated by Macrocosmos in partnership with Rowan Scientific. Their open-source Egret-1 AI model can achieve quantum-level accuracy at significantly reduced computational costs.
- Protein Folding Subnet: Focuses on biotech simulations for drug discovery. It uses AI to predict molecular structures faster than traditional methods, showing about 25% better accuracy for research on longevity and diseases.

- NATIX AI Subnet: Applies physical AI to healthcare-related tasks, like monitoring medical devices for anomalies. It started with uses like detecting litter or road signs but is expanding into biotech imaging.
- General Medical AI Initiatives: New subnets for multi-modal tasks like genomics and radiology. They use Bittensor's marketplace for full diagnostics and treatment planning with community-backed models.

Additional Healthcare Subnets: Beyond these core applications, Bittensor hosts several other healthcare-focused subnets including BetterTherapy (Subnet 102), a mental health platform providing AI-driven therapy through digital "twin" doctors; Dippy Roleplay (Subnet 11), addressing loneliness and mental health with over 4 million users; NOVA (Subnet 68) by Metanova, a decentralized drug discovery platform that has screened 4.8 million molecules across 7,000 protein targets; and a general healthcare subnet for disease diagnosis using medical imaging analysis.

3. BUSINESS IMPLEMENTATION

Bittensor launched its mainnet (Kusanagi) in 2021, migrated to Nakamoto later that year, and to Finney in 2023. Subnet growth sped up in 2024, reaching over 128 active by October 2025. Key 2025 milestones include subnet expansion (such as FLock launch and Cancer-AI pilots), biotech integrations, and the first halving planned for December, scaling to over 150 use cases.

- Governance Models: Governance is decentralized through \$TAO staking and voting, with no central control. Validators stay neutral across subnets, focusing on ethics in healthcare like bias checks. Community proposals balance innovation with regulatory needs.
- Economic Framework: \$TAO emissions (7,200 per day until the December 2025 halving, then 3,600) reward miners and validators. In healthcare, fees from model access (like diagnostic queries) go to the treasury. Deflationary burns and staking returns fund growth, such as new biotech subnets.

4. IMPACT & OUTCOMES

As of October 2025, Bittensor's \$TAO has a market cap of about \$3.1 billion, with a circulating supply of around 10 million (of a 21 million max) and daily volume near \$100 million. Healthcare subnets like Cancer-AI have processed thousands of datasets, boosting detection accuracy by 20-25%. Federated models, trained on distributed records, account for about 10% of network activity. Mainframe/SN25 has completed over 400,000 protein folding tasks since June 2024, with 30 active validators conducting more than 3,000 simulations simultaneously.

- Qualitative Benefits: Rewarded governance builds trust in AI diagnostics through openness and lower biases. Subnets reduce barriers for small clinics, enabling 30% faster protein simulations and hybrid research. Cross-subnet links speed personalized medicine, helping underserved areas. The Mainframe subnet's partnership with Rowan Scientific demonstrates concrete commercial adoption, with the Egret-1 models providing pharmaceutical researchers with DFT-level accuracy at multiple order-of-magnitude speedups. This enables smaller research labs and biotech startups to access computational capabilities previously limited to organizations with expensive supercomputer infrastructure.

- AI-Specific Achievements: Subnets like Protein Folding reach superhuman efficiency in biotech tasks. Federated systems like FLock support self-improving models, with recognition for collaborative, privacy-focused AI in healthcare. The Egret-1 neural network potentials developed through the Mainframe subnet equal or exceed the accuracy of routinely employed quantum-chemical methods on standard tasks including torsional scans, conformer ranking, and geometry optimization.

KEY INSIGHTS SUMMARY

- Incentivized Subnet Efficacy: Specialized, rewarded subnets (such as for cancer detection) outperform centralized training. They speed adoption in regulated fields without sacrificing privacy. Egret-1 models developed through the Mainframe subnet achieve quantum chemical accuracy on drug discovery tasks, enabling rational design of complex therapeutics like macrocyclic drugs.
- Ethical AI Marketplace: Token economics align rewards for secure healthcare data teamwork, breaking silos better than proprietary systems.
- Scalable Biotech Integration: Decentralized computing enables flexible, global workflows from diagnostics to discovery, fostering fair innovation in healthcare AI. Real-world partnerships like Rowan Scientific's integration with Bittensor validate the commercial viability of decentralized infrastructure for pharmaceutical research.

REFERENCES:

- BetterTherapy (2025) 'BetterTherapy Subnet: AI-driven mental health platform'. Available at: <https://subnetalpha.ai/subnet/bettertherapy/>
- Bittensor (2025) Bittensor Official Website. Available at: <https://bittensor.com/>
- Bittensor Blog (2025) 'TAO token economy explained'. Available at: <https://blog.bittensor.com/tao-token-economy-explained-17a3a90cd44e>
- CryptoSlate (2025) 'NATIX launches decentralized AI subnet on Bittensor to advance autonomous driving and physical AI'. Available at: <https://cryptoslate.com/press-releases/natix-launches-decentralized-ai-subnet-on-bittensor-to-advance-autonomous-driving-and-physical-ai/>
- FLock (2025) 'Overview of federated healthcare AI on Bittensor'. Available at: <https://substack.com/home/post/p-165074870>
- NOVA (Metanova) (2025) 'Drug Discovery: Decentralized molecular screening platform'. Available at: <https://subnetalpha.ai/subnet/nova/>
- Rowan Scientific (2025) 'Partnering with Macrococosmos: Collaboration for next-generation neural network potential development'. Available at: <https://rowansci.com/blog/partnering-with-macrococosmos>
- Rowan Scientific (2025) 'Egret-1: Pretrained neural network potentials for bioorganic simulation'. Available at: <https://www.rowansci.com/publications/egret-1-pretrained-neural-network-potentials>
- Safe Scan AI (2025) Cancer-AI Subnet GitHub: Cancer detection tools. Available at: <https://github.com/safe-scan-ai/cancer-ai>

Disclaimer: References to the \$TAO token are descriptive only and should not be construed as financial advice. Readers should conduct independent due diligence before engaging with any blockchain platform or digital asset.

4.5 CASE STUDY: ARTIFICIAL SUPERINTELLIGENCE ALLIANCE (ASI), A PLATFORM FOR AUTONOMOUS AI AGENTS AND DATA EXCHANGE (SG/UK)

DECENTRALIZED AI INFRASTRUCTURE ALLIANCE ENABLING AGENT-BASED HEALTHCARE DIAGNOSTICS AND RESOURCE OPTIMIZATION

Documented by: Thomas Egelhof, Chief Radiologist, Member of the Management Board, Merian Iselin Clinic for Orthopaedics and Surgery

CASE STUDY AT-A-GLANCE: ARTIFICIAL SUPERINTELLIGENCE ALLIANCE (ASI) (SG.UK)

<p>MATURITY LEVEL: Early Production (infrastructure live; healthcare apps emerging)</p> <p>GEOGRAPHY: Singapore/UK headquarters, global network</p> <p>REPLICABILITY: Medium (requires sophisticated integration)</p>	
<p>THE CHALLENGE:</p> <p>Healthcare needs coordinated AI agents that can autonomously discover data, access computation, and collaborate across institutional boundaries without central orchestration.</p>	<p>THE SOLUTION:</p> <p>Merged infrastructure combining autonomous agents (Fetch.ai), data marketplaces (Ocean Protocol), and computation (SingularityNET) enabling decentralized AI workflows with specialized medical models achieving superhuman diagnostic performance.</p> <ul style="list-style-type: none"> ✓ Pillar 1: Decentralized Data (Data marketplaces) ✓ Pillar 2: Decentralized Computation (Agentic AI) ✓ Pillar 3: Decentralized Governance (Token coordination)
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none"> • Autonomous agent coordination without central control • Data marketplace enabling privacy-preserving exchange • Integrated computation + data + agent infrastructure 	<p>LESSONS LEARNED:</p> <p>Future healthcare AI requires integrating all three pillars: autonomous agents discovering/using decentralized data and computation. Early stage but demonstrates architectural vision.</p>

NOTE: The Artificial Superintelligence Alliance (ASI) offers decentralized AI tools across sectors like finance and robotics. This case study highlights its healthcare applications to show how the merged platforms support privacy-protected diagnostics, faster drug discovery, and efficient medical workflows.

1. ORGANIZATION & CHALLENGE

The Artificial Superintelligence Alliance (ASI) is a Singapore-based decentralized AI collective formed in 2024 through the merger of Fetch.ai (founded in 2017 in Cambridge, UK), SingularityNET, Ocean Protocol, and CUDOS (joined following a community vote in September 2024, with the merger completed in November 2024). Backed by over \$100 million from investors like Binance Labs and Bitmain, it has built a community of thousands of developers and token holders. ASI advances decentralized AI by combining autonomous agents, data marketplaces, AI services, and computation infrastructure to challenge centralized AI dominance. While it applies to many industries, ASI increasingly focuses on healthcare through initiatives like ASI: Train for specialized models in pathology and agentic systems for patient care.

The Problem: Conventional healthcare AI faces deep-rooted barriers that limit scalability and fairness:

- Data Fragmentation and Privacy Vulnerabilities: Medical records are isolated in institutions. Centralized AI systems are vulnerable to breaches and lack rewards for secure, collaborative sharing.
- Inefficient Diagnostics and Resource Management: Manual processes in triage, imaging analysis, and supply chains lead to delays, errors, and poor resource allocation, increasing costs and harming patient outcomes.
- Restricted Access to Advanced AI: Proprietary platforms block small clinics and researchers from using models for tumor detection or personalized medicine, leading to biases and innovation gaps.
- Ethical and Regulatory Hurdles: Unclear oversight in AI deployment can cause inaccurate diagnoses, compatibility issues, and unequal access to tools like early disease prediction.

Why Decentralization? ASI decentralizes healthcare AI by enabling open data sharing, autonomous agents for flexible tasks, and token-based development. This creates ethical models where providers keep data control, reduce biases through distributed training, and reward participation to open up advances in diagnostics and treatment.

Key Stakeholders: The system includes hospitals and clinics deploying agents, AI developers building medical models, patients sharing anonymized data, pharmaceutical researchers, and regulators, all within supportive frameworks in Singapore and the UK.

2. THE SOLUTION

ASI combines autonomous agents (from Fetch.ai), AI marketplaces (SingularityNET), data tokenization (Ocean Protocol), and distributed computation (CUDOS) into a unified ecosystem. Users launch agents via smart contracts for tasks like federated diagnostics or resource optimization, with outputs tokenized for shared ownership. In healthcare, this supports privacy-focused workflows, such as tumor detection with specialized models or agent-driven triage. Providers, researchers, and patients can collaborate while maintaining data sovereignty.

Decentralization Approach:

- ✓ Decentralized Data: Healthcare datasets are tokenized and shared through Ocean's marketplace, enabling secure access without central controllers.
- ✓ Decentralized Computation: Model training is spread across CUDOS' network, supporting federated learning for sensitive medical data.
- ✓ Decentralized Governance: Proposals and updates use \$FET token voting (with a pending migration to ASI ticker as of October 2025), where agents automate ethical checks for health applications.

Technical Architecture:

Alliance Governance Process: ASI uses a modular governance system designed for AI in regulated areas like healthcare:

- Proposal Submission: Developers or medical groups submit ideas via the DAO portal. They outline goals (such as a pathology agent), resources (data and computation), and impacts (like faster cancer detection). Initial screening ensures fit with ethics and health priorities.
- Expert Review Process: Proposals advance to councils with specialists (for example, medical AI ethicists) who assess feasibility, compliance (like HIPAA or GDPR), and potential using simulations such as ASI-1 Mini.
- Multi-Stage Voting: \$FET holders vote in phases—strategic (health focus), allocation (grants), and execution (deployment)—with staking-based weighting and measures to prevent concentration.
- Smart Contract Execution: Funding is released through ASI Chain (a Cosmos-based blockchain), triggered by milestones (such as model validation after training), with cross chain support via IBC.
- Tokenization Framework: Models and datasets are tokenized as assets, enabling community control over licensing and royalties while providing traceability for health regulations.
- Continuous Monitoring: Agents and dashboards track performance (for instance, diagnostic accuracy), alerting to issues like biases.

AI-Integrated Healthcare Portfolio: ASI builds targeted AI projects that demonstrate decentralized health applications:

- ASI: Pathology Tumor Detection: Focuses on lymph node metastases in breast cancer using convolutional neural networks on decentralized datasets. It achieves high precision in identifying metastases while protecting privacy, with clinical pilots deployed in 2025.
- ASI: Brain Tumor Detection: An agentic project for early brain tumor segmentation from imaging. It uses federated learning to process scans autonomously, aiming to integrate into hospital workflows for 20-30% faster diagnostics.
- ASI: Prostate Cancer Detection: A specialized model for prostate imaging analysis. It represents decentralized AI for oncology, outperforming traditional methods in accuracy and allowing community-backed updates.

- ASI: GMAI (General Medical AI): Seeks to exceed benchmarks in areas like radiology and genomics, using multi-agent simulations for full diagnostics and treatment planning.
- Autonomous Healthcare Agents: Multi-agent systems improve patient care and resources, such as predicting bed needs from electronic health records and wearables with 80% accuracy in UK pilots. This extends to robotics in healthcare via the Cortex model, which supports adaptive tasks in hospitals.
- Biotech Expansion: Planned models under ASI: Train for drug discovery and longevity, using brain-inspired architectures to speed compound identification in biotechnology.

3. BUSINESS IMPLEMENTATION

Fetch.ai's agent mainnet launched in 2020 and evolved into the ASI merger in 2024. In 2025, key developments include the ASI-1 Mini launch, expansions in ASI: Train (such as tumor models in Q2), and biotech roadmaps, scaling to over 150 use cases with healthcare integrations.

- Governance Models: Projects have specific autonomy (like Fetch.ai agents) under a unified council. \$FET voting ensures decisions are health-neutral but ethics-focused, balancing innovation with safety in medical deployments.
- Economic Framework: Fees from data and computation usage, plus model inference, support the treasury. In healthcare, royalties from licensed diagnostics (projected over \$10 million from models like Cortex) fund token burns and grants, reinvesting in areas like biotech.

4. IMPACT & OUTCOMES

As of October 2025, ASI's \$FET has a market cap of about \$1.41 billion, with a supply of around 2.5 billion tokens and daily volume near \$69 million. Healthcare projects like ASI: Train have processed thousands of images, improving detection accuracy by 15-25%. Federated models, trained on millions of records, make up about 20% of alliance activity.

- Qualitative Benefits: Tokenized governance fosters trust in AI diagnostics through openness and bias reduction. Agents reduce administrative burdens (such as 30% faster triage), supporting hybrid human-AI care. Cross-project links speed personalized medicine, helping underserved regions.
- AI-Specific Achievements: Specialized models like GMAI show superhuman reasoning in multi domain diagnostics. Agentic frameworks enable self-improving systems, with Cortex advancing robotics in healthcare for seamless collaboration.

KEY INSIGHTS SUMMARY

- Specialized Model Efficacy: Decentralized, field-focused AI (like oncology detection) outperforms general models, speeding clinical adoption without privacy risks.
- Ethical Decentralization: Governance aligns rewards for secure health data sharing, cutting silos and biases better than centralized systems.
- Scalable Health Integration: Agentic and federated methods create efficient, adaptable workflows from diagnostics to resource management, promoting fair global healthcare innovation.

REFERENCES:

- ASI Alliance (2025) ASI Alliance Official Website. Available at: <https://superintelligence.io/>
- ASI Alliance (2025) 'Announcement of lymph node metastasis model', X (formerly Twitter), 17 October. Available at: https://x.com/ASI_Alliance/status/1915439728658198904
- Bittensor Blog (2025) 'TAO token economy explained'. Available at: <https://blog.bittensor.com/tao-token-economy-explained-17a3a90cd44e>
- CoinBureau (2025) 'ASI Alliance review: Comprehensive alliance overview'. Available at: <https://coinbureau.com/review/asi-alliance-review/>
- Fetch.ai (2025) 'ASI Alliance revolutionizes AI earning with ASI Train: Unlock the power of DeSci models'. Available at: <https://fetch.ai/blog/asi-alliance-revolutionizes-ai-earning-with-asi-train-unlock-the-power-of-de-sci-models>
- Fetch.ai (2025) Press & Media Kit: Overview of autonomous agents in healthcare systems. Available at: <https://fetch.ai/press-media>

Disclaimer: References to the \$FET token are descriptive only and should not be construed as financial advice. Readers should conduct independent due diligence before engaging with any blockchain platform or digital asset.

4.6 CHAPTER INSIGHTS AND CONCLUSIONS

This chapter examined decentralized model development and training through implementations spanning sovereign AI deployment, decentralized GPU infrastructure, collaborative machine learning networks, and autonomous agent platforms. Three critical findings emerge that shape the future trajectory of healthcare AI development in terms of decentralized model development and training.

Key Finding 1: Sovereign AI Deployment Demonstrates That Large-Scale Decentralization Is Achievable When Aligned With National Priorities

DeepSeek's deployment across 100+ Chinese healthcare institutions proves that decentralized AI can scale nationally when it aligns with sovereign data governance priorities. China's strategic push for AI self-sufficiency, combined with strict data localization requirements, created conditions where decentralized, locally-deployed foundation models became not just technically feasible but strategically essential. The success demonstrates that open-source models with local deployment capabilities can deliver state-of-the-art reasoning and vision-language performance while maintaining complete institutional control over sensitive health data. This finding suggests that geopolitical drivers—data sovereignty concerns, supply chain independence, and regulatory frameworks—may be as powerful catalysts for decentralized AI adoption as technical or economic considerations are.

Key Finding 2: Decentralized Compute Networks Are Emerging as Complements Rather Than Replacements for Centralized Cloud Infrastructure

Akash Network and Aethir demonstrate that decentralized GPU marketplaces address specific limitations of centralized cloud providers—offering significant cost reductions, reducing single-point of-failure cybersecurity risks, minimizing data transmission for privacy-sensitive workloads, and enabling more sustainable computation through distributed energy consumption. However, these platforms complement rather than replace AWS, Google Cloud, and Azure, in our view. Centralized providers maintain decisive advantages in reliability guarantees, integrated service ecosystems, regulatory compliance certifications, and enterprise support that production healthcare systems require. An emerging sustainable pattern shows hybrid architectures: centralized cloud for mission critical clinical systems requiring maximum reliability, and decentralized computation for research, development, and training workloads where cost efficiency and data locality outweigh the coordination complexity of distributed infrastructure.

Key Finding 3: Blockchain-Enabled Agentic AI Activity Surges With Promising Healthcare Pilots, Though Scalability Depends on Resolving Interoperability and Legal Frameworks

Blockchain-enabled agentic AI has experienced remarkable growth, with initiatives like Bittensor and the ASI Alliance demonstrating compelling healthcare applications across diagnostics, drug discovery, and resource optimization. The ASI Alliance's specialized healthcare agents achieving superhuman diagnostic performance in pathology, combined with Bittensor's incentivized collaborative model development across medical imaging and biotech research, illustrate the genuine potential of this approach. Token-based governance enables global participation in research priorities, smart contract-based traceability creates unprecedented accountability for AI decisions, and cryptocurrency rewards align incentives for continuous model improvement. Multiple healthcare-focused DAOs and agent networks have emerged from pilot to operational status, attracting significant investment and institutional partnerships. However, translating this momentum into mainstream healthcare adoption at scale depends critically on the speed at which the ecosystem resolves two fundamental challenges: interoperability standards enabling agents to operate seamlessly across different blockchain networks and integrate with traditional healthcare IT systems, and legal frameworks establishing clear accountability, liability, and ownership structures for decentralized governance models where no single entity controls decision-making.

PRACTICAL IMPLICATIONS FOR STAKEHOLDER GROUPS

The transition toward decentralized model development requires stakeholders to adopt hybrid strategies rather than binary choices between centralized and decentralized approaches. Healthcare organizations should architect computation infrastructures combining centralized reliability for production clinical systems with decentralized alternatives for research and training workloads, while institutions in data sovereignty-focused jurisdictions must prioritize local deployment capabilities. Patient advocacy groups should demand transparency about where models run and who controls health data, supporting sovereign approaches that maintain institutional accountability while enabling collaborative learning. MedTech developers must build solutions compatible with both centralized and decentralized infrastructure to maximize market reach, focusing sovereign AI offerings on privacy-sensitive applications and data localization markets, while leveraging the emerging opportunities of blockchain-enabled agent platforms when and where applicable. Research institutions should leverage decentralized computation for cost-effective experimentation while actively shaping interoperability standards and governance frameworks that will allow blockchain-based collaborative learning to scale beyond proof-of-concept. Policy makers face the most complex challenge: creating regulatory environments that incentivize sovereign AI development for national security without fragmenting global healthcare innovation, establishing legal frameworks for autonomous agent accountability and DAO governance before widespread clinical adoption, and supporting infrastructure diversity through technology-neutral regulations rather than mandating specific architectural approaches.

CHAPTER 5: DECENTRALIZED VALIDATION AND DEPLOYMENT

Foreword by the author: Stephanie Fuchs, International Tax Advisor, Owner of Stephanie Fuchs Consulting

Decentralized artificial intelligence offers hospitals the possibility to strengthen diagnostics while maintaining responsibility for their own systems, as the deployment of the AimSG platform at Changi General Hospital illustrates. The hospital integrates AI models into its own infrastructure, with data processing and triage carried out locally, while validation, accreditation and oversight are coordinated at the national level. This arrangement combines the flexibility of institutional autonomy with the consistency of shared standards, creating a structure that is both adaptable and reliable. It demonstrates how AI can be embedded into clinical workflows without undermining accountability or patient safety.

When I lived and studied in Singapore, artificial intelligence had not yet entered clinical medicine. Hospitals concentrated on strengthening infrastructure and optimising established workflows. What stood out to me even then was the country's commitment to high clinical standards and its openness to technological innovation. These qualities made Singapore a setting where new approaches could be considered seriously and integrated pragmatically once proven effective. Seen from today's perspective, it is less surprising that the healthcare system has been able to move so quickly toward the deployment of decentralized AI in practice.

This transformation reflects qualities that I experienced first hand in Singapore's healthcare environment. Collaboration between institutions, government agencies and technology partners was already characteristic of the system. The adoption of decentralized AI builds directly on these foundations, enabling hospitals to adapt innovation to their circumstances while remaining part of a coherent national framework.

My current professional practice focuses on taxation, digital assets and emerging technologies. From this perspective, decentralized systems in finance, governance and healthcare share structural challenges. They depend on clear frameworks that preserve trust, ensure accountability and define decision rights. I have seen this dynamic in the evolution of decentralized autonomous organizations, where community participation and transparent decision making are combined with legal structures that safeguard continuity and compliance. When designed responsibly, such organizations have the capacity to complement healthcare ecosystems by enabling collaborative research funding, equitable intellectual property ownership and collective oversight of AI driven innovation. They establish a framework through which responsibility can be distributed across qualified stakeholders and governance processes can be verified transparently.

Within healthcare, decentralized organizations have the potential to function as instruments of scientific collaboration and patient engagement. Through structured voting mechanisms and traceable resource allocation, they enable environments in which researchers, clinicians and patient representatives define priorities collectively and ensure measurable progress. Such an approach strengthens integrity, inclusiveness and continuity in decision making. It positions governance as an integral part of innovation and aligns technological advancement with ethical, professional and societal objectives.

The example of Changi General Hospital together with initiatives such as VitaDAO, as well as the other case studies presented in this chapter, shows how this balance can be realised in healthcare. Whether the objective is to accelerate diagnostics, fund longevity research or coordinate multi institutional AI deployment, the same principle applies. Decentralized systems succeed when governance and accountability evolve in step with technological capability.

This chapter brings together a series of case studies that explore how decentralized validation and deployment can take shape across different contexts and maturity levels. They also illustrate that decentralisation in healthcare is not a binary condition but a spectrum. It can manifest through local data processing, distributed computation, shared governance, patient ownership of information or collective funding structures. Each model reflects different balances between autonomy and coordination, privacy and interoperability, innovation and oversight. The strength of these initiatives lies not in rejecting existing systems but in complementing them with verifiable frameworks that sustain trust and accountability as technology advances.

The views expressed are those of the author and do not necessarily reflect those of Stephanie Fuchs Consulting.

5.1 CHAPTER OVERVIEW

Chapter 5 examines decentralized validation and deployment—the critical final stage where AI models transition from development to clinical production environments. However, the boundaries between AI development stages are inherently fluid, and several initiatives featured here could equally belong in earlier chapters on data collection or model training. We position them in this chapter for two strategic reasons: first, to highlight more mature implementations that have progressed toward real-world deployment; and second, to emphasize governance models—an essential dimension that received limited attention in previous chapters despite being crucial for scaling AI systems within existing healthcare infrastructure and research programs.

The chapter presents two primary case studies and five bonus resources spanning decentralized research governance, national healthcare infrastructure, educational simulation, patient advocacy platforms, and emerging Web3 health technologies. VitaDAO demonstrates how decentralized autonomous organizations can govern AI-driven longevity research through blockchain-based decision-making, with over 10,000 community members directing \$4.1M in funding across 22+ projects. Changi General Hospital's implementation of the AimSG platform illustrates national-scale deployment where AI models for chest X-ray triage operate locally within institutional infrastructure while validation, accreditation, and oversight coordinate at the national level—exemplifying how decentralized technical architecture can coexist with centralized governance frameworks when appropriate.

Each case study follows our established structure: Organization & Challenge, The Solution, Business Implementation, and Impact & Outcomes. Beyond these mature implementations, the bonus resources deliberately span different development stages—from Tsinghua University's centralized but instructive "Agent Hospital" educational system to early-stage platforms like Lumera Health and SpectruthDAO that are still proving their models. While some may seem premature for a chapter on "deployment," they illuminate critical governance patterns emerging in decentralized healthcare AI: token-based decision-making, community-driven research prioritization, patient sovereignty over health data, and transparent accountability mechanisms. These governance innovations matter precisely because they address the most challenging aspects of bringing decentralized AI from promising prototypes to trustworthy production systems.

5.2 CASE STUDY: VITADAO – DECENTRALIZED SCIENCE FOR LONGEVITY RESEARCH (CH)

DECENTRALIZED AUTONOMOUS ORGANIZATION GOVERNING AI-DRIVEN (AND TRADITIONAL) LONGEVITY RESEARCH

Documented by: Carmen Cucul, Blockchain & AI Healthcare Consultant, Ethers HealthData Foundation

CASE STUDY AT-A-GLANCE: VITADAO (CH)

<p>MATURITY LEVEL: Production (operational since 2021)</p> <p>GEOGRAPHY: Switzerland (legal entity), global community</p> <p>REPLICABILITY: Medium (requires DAO expertise + community)</p>	
<p>THE CHALLENGE:</p> <p>Traditional longevity research funding is centralized, slow, and disconnected from patient/community priorities. AI-driven research needs transparent governance ensuring ethical oversight and stakeholder alignment.</p>	<p>THE SOLUTION:</p> <p>Decentralized Autonomous Organization (DAO) governing AI-driven longevity research through blockchain-based voting, transparent treasury management, and community participation in research prioritization.</p> <p>✅ Pillar 3: Decentralized Governance (DAO)</p>
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none">• 10,000+ community members participating in governance• \$4.1M+ deployed across 22+ research projects• Token-weighted voting on funding decisions• AI projects include drug discovery and biomarker analysis	<p>LESSONS LEARNED:</p> <p>DAOs can successfully coordinate research funding and governance, but healthcare requires hybrid models balancing community input with expert clinical/scientific oversight for safety-critical decisions.</p>

Note: While VitaDAO funds diverse longevity research spanning traditional biomedical approaches and clinical trials, this case study focuses on their AI-driven initiatives to illustrate how DAOs can govern emerging technologies in healthcare research.

1. ORGANIZATION & CHALLENGE

Organization Profile: VitaDAO is a Switzerland-based decentralized autonomous organization founded in 2021, headquartered in Neuhausen am Rheinfall. The community-owned collective has raised \$4.1M from investors including Pfizer Ventures, Vitalik Buterin, and Balaji Srinivasan, building a community of over 10,000 members with 2,000+ token holders. VitaDAO represents a pioneering model in decentralized science (DeSci), replacing traditional research funding gatekeepers with blockchain-based governance using \$VITA tokens for decision-making. While the organization funds diverse longevity research spanning traditional biomedical approaches, clinical trials, and biomarker development, VitaDAO has increasingly integrated AI agents and AI-driven research as one strategic component of its comprehensive research portfolio, currently consisting of 22+ funded projects.

The Problem: Traditional biomedical research funding face systemic barriers that limit innovation in longevity science:

- **Centralized gatekeeping:** Traditional biomedical research funding create systemic barriers limiting longevity science innovation
- **Limited Methodological Diversity:** Centralized funding bodies show conservative bias against aging research, whether traditional or AI-driven approaches
- **Fragmented research Community:** Biomedical researchers and AI/computational scientists operate in silos, with limited mechanisms for collaborative funding and knowledge sharing
- **IP Ownership Concentration:** Intellectual property for clinical research is concentrated in academic institutions or pharmaceutical companies, limiting community participation in the value creation and commercialization of breakthrough therapies

Why Decentralization?: VitaDAO democratizes longevity research funding across all methodologies while enabling community ownership of research outcomes. The DAO model provides transparent, community-driven governance where stakeholders directly influence research priorities and ensure outcomes benefit the broader community rather than only institutional shareholders.

Key Stakeholders: The ecosystem includes traditional longevity researchers, AI specialists developing computational models, biotech entrepreneurs, patient advocacy groups, and venture capital investors, all operating within Switzerland's crypto-friendly regulatory environment.

2. THE SOLUTION

VitaDAO's platform implements a sophisticated multi-stage governance framework where community members submit research proposals that undergo expert review by specialized working groups, followed by token-weighted voting to determine funding allocations. Approved projects receive funding through smart contracts with milestone-based releases, while research outputs including AI models and datasets are tokenized as IP-NFTs to enable community ownership and governance of valuable research assets. The system integrates AI agents both as research contributors (competing anonymously with human researchers for funding) and as advisors (like the Aubrey AI agent providing 24/7 scientific consultation), creating a unique ecosystem where traditional researchers, AI specialists, and community stakeholders collaborate transparently on longevity research priorities and resource allocation.

Decentralization Approach:

- ✓ Decentralized Data: Research datasets from traditional and AI projects shared through open source repositories
- ✓ Decentralized Computation: Research execution distributed across traditional laboratories and cloud-based AI development
- ✓ Decentralized Governance: All funding decisions determined through community voting, with AI agents participating as advisors

Technical Architecture - DAO Governance Process: VitaDAO implements a sophisticated multi-stage governance framework specifically designed for scientific research funding decisions:

- Proposal Submission: Community members submit research proposals through structured frameworks specifying technical objectives, funding requirements, expected outcomes, and methodology (traditional, AI-driven, or hybrid approaches). Proposals undergo initial community screening for scientific merit and alignment with longevity research goals.
- Expert Review Process: Proposals advance to technical review by specialized working groups, including the AI/ML working group for computational projects and scientific advisory board members with relevant domain expertise. This stage evaluates feasibility, innovation potential, and resource requirements.
- Multi-Stage Voting: The community participates in token-weighted voting through multiple governance cycles. \$VITA token holders vote on high-level research strategy, funding allocations, and project priorities, with voting power proportional to token holdings while implementing safeguards against governance concentration.
- Smart Contract Execution: Approved funding automatically executes through smart contracts based on predetermined milestones, ensuring accountability and transparent resource allocation. Projects must meet specified deliverables to trigger subsequent funding releases.
- IP-NFT Framework: Research outputs, including trained AI models, datasets, and traditional research discoveries, are tokenized as IP-NFTs, enabling community ownership and governance of valuable research assets. Token holders participate in decisions about IP licensing, commercialization, and revenue distribution.
- Continuous Monitoring: Agents and dashboards track performance (for instance, diagnostic accuracy), alerting to issues like biases.

AI-Integrated Research Portfolio: VitaDAO funds complementary traditional and AI-focused longevity research projects demonstrating diverse applications of artificial intelligence in aging research:

- GERO Project: This initiative leverages physics and artificial intelligence to build a proprietary platform aimed at understanding the aging process. The project uses AI models to analyze complex biological data and create a pipeline of drug candidates and therapies for age-related diseases, representing one of VitaDAO's most significant AI-driven investments.

- Scheibye-Knudsen Lab: Uses machine learning to process over 1 billion prescription records to explore how various drugs influence human lifespan, applying data-driven AI methods at unprecedented scale. This project demonstrates the power of AI to analyze massive healthcare datasets that would be impossible to process through traditional methods.
- Fang Lab: Combines AI-driven analysis and wet lab validation to identify new compounds for activating mitophagy (cellular cleanup) as potential Alzheimer's disease treatments. This hybrid approach exemplifies how AI can accelerate traditional drug discovery while maintaining experimental validation requirements.
- Humanity Project: Integrates wearable devices and digital biomarkers, applying AI to personalize aging metrics and management strategies for individuals. The project represents the application of AI to preventive longevity medicine and personalized health optimization.
- Aubrey AI Agent: VitaDAO has developed an internal AI agent trained on Aubrey de Grey's longevity research expertise, providing 24/7 scientific consultation to the community. The agent assists in study design (including documented support for the RMR2 mouse study), research proposal evaluation, and community engagement. This represents a pioneering example of AI participating directly in scientific governance and research support.
- VitaDAO Labs AI Integration: The venture studio program now enables AI agents to compete anonymously with human researchers for funding, submitting research proposals and experimental designs. This creates a unique testing ground for autonomous scientific discovery while ensuring merit-based evaluation without bias against AI-generated research.

3. BUSINESS IMPLEMENTATION

- Implementation Timeline: VitaDAO evolved from traditional biotech funding (2021) to developing the IP-NFT framework (2022), launching systematic AI research programs (2023), and scaling to 22+ projects with AI agents competing for funding (2024-2025).
- Governance Model: Community governance maintains methodology neutrality, with \$VITA token holders voting based on scientific merit rather than approach preference. The Aubrey AI agent and human advisors provide cross-disciplinary consultation, while portfolio management ensures balanced investment across traditional and AI methodologies representing approximately an increasing proportion of total funding.
- Economic Framework: Revenue streams from both traditional discoveries and AI innovations flow back to the community treasury. The IP-NFT framework accommodates intellectual property from diverse research outputs, while cross-pollination between traditional and AI research creates synergistic value.

4. IMPACT & OUTCOMES

- Quantitative Results: VitaDAO has funded 22+ longevity research projects spanning traditional biomedical research and AI-driven initiatives. The community has grown to 10,000+ members with 2,000+ token holders participating in governance. Successful outcomes include the Matrix Bio spinout from traditional research and breakthrough AI capabilities like billion-record dataset analysis acceleration from years to months.

- Qualitative Benefits: Community governance enables global participation in research direction setting across methodologies, ensuring both traditional and AI research address real-world needs. Integration of traditional and AI researchers has accelerated knowledge transfer and enabled novel hybrid approaches. Blockchain governance provides unprecedented transparency in research funding allocation, building trust in both traditional and AI-driven outcomes.
- AI-Specific Achievements: AI agents now participate in funding competition alongside traditional researchers, demonstrating autonomous research capabilities from hypothesis generation through experimental design. The Aubrey AI provides continuous community consultation while hybrid research approaches combine AI-driven analysis with traditional experimental validation.

KEY INSIGHTS SUMMARY

- Multi-Methodological DAO Viability: Successfully demonstrated that decentralized governance can effectively coordinate diverse research approaches, from traditional wet lab studies to AI-driven computational research, without favoring any single methodology
- Community-Driven Research Balance: Community governance enables more balanced research portfolios compared to traditional funding mechanisms that often bias toward established methodologies, ensuring both traditional and innovative approaches receive appropriate support
- Scalable Diverse Research Model: The integration of traditional research funding, AI development, and community governance creates a scalable model for accelerating scientific discovery that maintains methodological diversity while maximizing innovation potential

REFERENCES:

- GERO AI (2025) GERO AI Platform: VitaDAO-funded AI longevity research. Available at: <https://gero.ai/>
- Vectra Advisors (2025) 'Understanding DAOs and legal wrappers in Switzerland'. Available at: <https://vectra-advisors.com/understanding-daos-and-legal-wrappers-in-switzerland/>
- VitaDAO (2025) VitaDAO Official Website. Available at: <https://www.vitadao.com/>
- VitaDAO (2025) 'AI integration strategy presentation', YouTube. Available at: <https://www.youtube.com/watch?v=ms2dLI2-myw&t=40s>
- VitaDAO Labs (2025) Whitepaper: How the in-house Lab operates, including fellowships and roadmap. Available at: <https://drive.google.com/file/d/1zdzqoybZgxLYkiU2RTFuiMqKcHw7AoPP/view>

Disclaimer: References to the \$VITA token are descriptive only and should not be construed as financial advice. Readers should conduct independent due diligence before engaging with any platform or digital asset.

5.3 CASE STUDY: CHANGI GENERAL HOSPITAL'S DECENTRALIZED MEDICAL IMAGING WITH AI (SG)

A NATIONAL COLLABORATION THROUGH THE AIMSG PLATFORM TO IMPROVE DIAGNOSTIC EFFICIENCY IN SINGAPORE'S PUBLIC HEALTHCARE SYSTEM

Documented by: Stéphanie Fuchs, Swiss and Liechtenstein Tax Advisor, Owner of Stephanie Fuchs Consulting

CASE STUDY AT-A-GLANCE: CHANGI GENERAL HOSPITAL (SG)

<p>MATURITY LEVEL: Production (operational at national scale)</p> <p>GEOGRAPHY: Singapore (national deployment)</p> <p>REPLICABILITY: High (model applicable to other nations)</p>	
<p>THE CHALLENGE:</p> <p>Singapore needed national-scale AI deployment for medical imaging while maintaining institutional autonomy, clinical safety oversight, and interoperability across hospitals.</p>	<p>THE SOLUTION:</p> <p>Hybrid architecture: AI models deployed locally within hospital infrastructure (decentralized inference) with national validation, accreditation, and oversight coordinated centrally by Synapse and Ministry of Health.</p> <ul style="list-style-type: none">✓ Pillar 1: Decentralized Data (Local processing)✓ Pillar 2: Decentralized Computation (Institutional AI)✓ Pillar 3: Hybrid Governance (Local + National)
<p>KEY OUTCOMES:</p> <ul style="list-style-type: none">• National deployment across Singapore hospitals• Chest X-ray AI triage reducing radiologist workload• Data sovereignty maintained at institutional level• Coordinated quality monitoring without centralizing data	<p>LESSONS LEARNED:</p> <p>Successful deployment requires pragmatic balance: decentralize where it adds value (data sovereignty, local autonomy) but centralize selectively for efficiency (quality monitoring, regulatory compliance, interoperability standards).</p>

1. ORGANIZATION & CHALLENGE

Organization Profile: Changi General Hospital (CGH) is a tertiary care hospital under the SingHealth cluster, Singapore's largest public healthcare group. CGH serves a diverse patient population in eastern Singapore and is recognized for its early adoption of innovation in clinical workflows. The digital intervention described here is implemented in partnership with Synapxe (the national health technology agency), which orchestrates the creation and governance of national health IT platforms. The AimSG (AI Medical Imaging Singapore) platform is a collaborative, multi-stakeholder effort involving public health clusters, technology partners (notably Lunit -active chest X-ray deployment at CGH/Singapore-, alongside collaborations with NTT Data, DeepTek, CARPL at a platform level), and policy oversight by the Ministry of Health.

The Problem: Demand for chest X-ray imaging has increased with the complexity of patient presentations and expansion of clinical services. Manual triage created significant bottlenecks, particularly for urgent or abnormal cases, as radiologists adhered to a strict first-in, first-out (FIFO) policy. This led to delays in identifying and reporting critical cases, increasing overall turnaround time and potentially impacting clinical outcomes. The COVID-19 pandemic further strained capacity, highlighting the need for tools enabling rapid prioritization and scalable imaging management.

Why Decentralization?: Singapore's healthcare system is characterized by a network of independent public institutions operating within coordinated clusters. A decentralized solution whereby AI models are deployed, validated, and governed at the institutional level using a national framework ensures resilience, flexibility, and rapid iterative improvement. Decentralization enables institutions to adopt AI as their infrastructure and resources permit, while benefiting from common standards, accreditation, and updates provided centrally via the AimSG platform. This mitigates single-point systemic risk, increases responsiveness to local needs, and sustains regulatory compliance.

Key Stakeholders:

- Clinical teams (radiology, IT, management) at CGH.
- Synapxe (formerly IHiS), driving health digitalization and platform standards.
- Technology partners (Lunit, NTT Data, DeepTek, CARPL) for model and platform integration.
- Singapore Ministry of Health, providing oversight and investment.

2. THE SOLUTION

Platform Architecture

AimSG is a vendor-neutral, national platform designed to facilitate rapid validation, accreditation, and integration of various AI models into clinical imaging practice. It supports plug-and-play interoperability, enabling different AI solutions for radiology (chest X-ray, mammography, etc.) to be deployed across multiple institutions. The platform's architecture uses both a centralized oversight function (for accreditation, updates, security) and decentralized local deployment (for technical operations within hospital environments).

Decentralization Approach

AI model validation and application occur locally, with results monitored centrally. Hospitals can independently implement AI tools, carry out internal monitoring, and request updates or improvements via the platform. This enables site-specific adaptation while ensuring technologies meet shared national standards and data privacy expectations. Model life-cycle management (performance evaluation, versioning, and rollback) is supported at both local and platform levels.

- ✓ Pillar 1: Decentralized Data (Local processing)
- ✓ Pillar 2: Decentralized Computation (Institutional AI)
- ✓ Pillar 3: Hybrid Governance (Local + National)

Technical Architecture

At CGH, the Lunit INSIGHT CXR model is integrated on-premise into the hospital's PACS. AI triage is initiated automatically for all chest X-rays, categorizing them as 'urgent', 'abnormal', or 'normal'. Urgent/abnormal cases are flagged and moved up in the reporting queue enabling radiologists to address critical needs more rapidly. Data never leaves the hospital system without explicit cause, and technical support for both edge and hybrid cloud deployments is available for varying risk appetites and infrastructural realities.

Privacy & Security

Compliance with the national Personal Data Protection Act (PDPA) and local hospital protocols is integral. Data is processed on-site with no patient identifiers exposed externally unless subject to regulatory provisions. Each model undergoes formal technical and ethical assessment pre- and post-deployment, and centralized monitoring enables fast response to anomalies or errors. Accreditation requires ongoing demonstration of privacy robustness and technical security.

3. BUSINESS IMPLEMENTATION

- 2022: National framework and platform architecture established.
- 2023: Model validation completed, pilot rolled out at CGH and Singapore General Hospital.
- 2023-2024: Rapid scaling and optimization of clinical workflows, with iterative feedback leading to refinement.

Stakeholder training and change management activities occurred in parallel, prioritizing radiology department readiness and IT integration.

- Governance Model A dual-layer governance structure exists: Synapse coordinates overarching platform integrity, accreditation, technical support, and reporting; individual institutions maintain responsibility for local implementation, compliance, and performance monitoring. Governance committees include clinical, technical, and policy representation, with formal processes for user feedback, incidents, or model upgrades.

- Economic Framework AimSG's open ecosystem allows competitive procurement and reduces vendor lock-in. Economic efficiencies stem from centralizing some functions (validation, security, support) while preserving flexibility for local deployment. Cost-savings are realized through reduced administrative labor for manual triage, decreased unnecessary delays in acute care, and better allocation of radiology resources. Long-term economic models also anticipate increased health system capacity and improved patient outcomes as indirect benefits.
- Change Management Sustained engagement with clinicians (radiologists, ED physicians), IT departments, and administrative leads was critical. On-site training, regular workshops, iterative feedback, and visible clinical champions facilitated uptake. The platform's modularity allowed adaptation to existing workflows with minimal disruption, and post-implementation monitoring informed ongoing workflow optimization.

4. IMPACT & OUTCOMES

- Quantitative Results:
 - 97% reduction in chest X-ray triage turnaround time for inpatients at CGH.
 - 50% faster reporting of urgent cases compared to historical averages.
 - No observed increase in false negatives or clinically significant diagnostic errors post implementation over the initial reporting period.
 - Metrics are subject to ongoing evaluation and validation as system use expands.
- Qualitative Benefits:
 - Radiologists and referring clinicians report increased confidence in the reliable flagging and prioritization of abnormal cases.
 - Perceptions of improved workflow efficiency and reduced cognitive workload (triage and reporting).
 - Workflow integration was achieved with high acceptability, with broad staff buy-in attributed to systematic change management and observable improvements in clinical care dynamics.
 - Indications of wider systemic benefit, including improved patient flow, more consistent clinical response times, and an emerging culture of data-driven collaboration at the institutional level.

KEY INSIGHTS SUMMARY

The implementation of the AimSG platform and AI chest X-ray triage system at Changi General Hospital provides preliminary evidence that decentralized AI deployment within a national digital health framework may enhance operational efficiency in imaging services. Key insights include:

- Decentralized governance supports local adaptability while facilitating centralized oversight for model validation, safety, and privacy compliance.
- Open, vendor-neutral platform designs can reduce barriers to multi-vendor AI integration and scalability across institutions.

- Collaboration among clinical, technical, regulatory, and policy stakeholders is essential to enable adoption and ongoing system evolution.
- Early quantitative outcomes indicate notable reductions in triage turnaround time and accelerated reporting of urgent cases without compromising accuracy.
- Qualitative benefits observed include improved workflow prioritization, better resource allocation, and positive user engagement, all supported through effective change management.
- While these findings are encouraging, longer-term and multi-site studies are required to confirm the generalizability and sustained impact of such decentralized AI platforms in healthcare.

REFERENCES:

- Changi General Hospital (2024) About Changi General Hospital. Singapore: SingHealth. Available at: <https://www.cgh.com.sg/about-us>
- Lunit Inc. (2023) Lunit INSIGHT CXR: AI for Chest X-Ray Analysis. Seoul: Lunit. Available at: <https://www.lunit.io/en/products/cxr>
- Medical Imaging Singapore (AimSG) (2023) About AI Medical Imaging Platform for Singapore public healthcare (AIMSG). Singapore: Synapse. Available at: <https://www.synapse.sg/healthtech/health-ai/ai-medical-imaging-platform>
- Ministry of Health Singapore (2024) Transforming Healthcare Through Technology. Singapore: MOH Press Release, 10 October. Available at: <https://www.moh.gov.sg/newsroom/transforming-healthcare-through-technology/>
- Personal Data Protection Commission Singapore (2024) Personal Data Protection Act (PDPA). Singapore: PDPC. Available at: <https://www.pdpc.gov.sg/overview-of-pdpa/the-legislation/personal-data-protection-act>
- Synapse (2023) About Synapse: Singapore's Health Tech Agency. Singapore: Synapse. Available at: <https://www.synapse.sg/>

5.4 BONUS RESOURCE: TSINGHUA UNIVERSITY'S "AGENT HOSPITAL" (CHINA)

Documented by: LeiLei Tang, Medical Consultant, Ethers HealthData Foundation

China's Tsinghua University has unveiled one of the most ambitious healthcare AI projects to date with their "Agent Hospital" - a comprehensive virtual healthcare environment featuring 42 AI doctor agents across 21 medical departments working alongside 4 AI nursing agents. Developed by the Institute for AI Industry Research and launched in April 2025, this fully simulated hospital creates complete healthcare workflows from pre-hospital triage through consultation, examination, diagnosis, treatment, rehabilitation, and follow-up. The platform employs proprietary large language models that remain frozen during operation, with AI agents evolving through accumulated virtual medical interactions stored in "medical case bases" and "experience bases" rather than through traditional model retraining.

The results have been remarkable: AI doctors can reportedly treat 10,000 patients in just days (equivalent to two years of human physician work), and achieved 93.06% accuracy on respiratory disease questions from the MedQA dataset (U.S. Medical Licensing Examination questions), and demonstrated significant improvements in examination (88%), diagnosis (95.6%), and treatment planning (77.6%) capabilities. The system has reportedly created over 500,000 virtual patients representing diverse demographics, age groups, and medical conditions, providing medical students with risk-free training environments for rare conditions they might never encounter in traditional internships. Perhaps most significantly, the platform can simulate infectious disease emergence, spread, and containment strategies, offering valuable public health planning capabilities.

However, Agent Hospital is fundamentally a centralized AI system, not the decentralized approach this report focuses on. All processing occurs within Tsinghua's infrastructure, with primary oversight maintained through their Institute for AI Industry Research. While partnerships exist with Beijing Tsinghua Changgung Hospital for testing and deployment, there is no federated learning across multiple institutions, no distributed data governance, and no multi-stakeholder decision-making. The platform operates as a sophisticated single-institution AI development rather than the collaborative, privacy-preserving, multi-institutional approaches that define decentralized healthcare AI.

Despite this centralized architecture, Agent Hospital offers profound inspiration for future decentralized implementations. The concept of comprehensive virtual healthcare environments could be transformative when adapted to federated learning frameworks where multiple hospitals contribute to AI training while maintaining data sovereignty. Imagine Agent Hospital's simulation capabilities being distributed across a network of institutions, each contributing local expertise and patient patterns while preserving privacy through federated learning. Such an approach could combine Agent Hospital's proven educational and diagnostic capabilities with the privacy preservation, institutional autonomy, and collaborative benefits that decentralized AI enables.

Decentralized AI in Healthcare: From Theory to Practice

The regulatory context makes this particularly relevant: China's Personal Information Protection Law (PIPL) creates strong incentives for federated approaches that keep sensitive patient data within institutional boundaries. While Agent Hospital currently uses synthetic data generation to avoid these constraints, future iterations could leverage multimodal federated learning to enable real cross-hospital collaboration. This could create truly distributed virtual hospitals where AI agents learn from diverse patient populations and medical practices across institutions without compromising data sovereignty - essentially scaling Agent Hospital's remarkable simulation capabilities across China's entire healthcare ecosystem while maintaining the privacy and autonomy that decentralized approaches provide.

REFERENCES:

- TechNode (2025) 'Tsinghua University launches AI-driven hospital to train next-gen doctors', 29 April. Available at: <https://technode.com/2025/04/29/tsinghua-university-launches-ai-driven-hospital-to-train-next-gen-doctors/>
- Wang, J. et al. (2025) 'Agent Hospital: A Simulacrum of Hospital with Evolvable Medical Agents', arXiv preprint arXiv:2405.02957v3. Available at: <https://arxiv.org/pdf/2405.02957>

5.5 BONUS RESOURCE: CUREWISE – PERSONALIZED, PATIENT-CONTROLLED CANCER AI ADVOCATE (US)

Documented by: Natalia Sofia, Data & Ecosystem Lead, Ethers HealthData Foundation

Steve Brown's CureWise (<https://curewise.com/>) represents a compelling evolution in patient-facing healthcare AI, born from personal necessity and digital health expertise spanning over two decades. A "digital health OG" who previously founded Health Hero (one of the first disease management companies used by the VA) and has backgrounds in physics, filmmaking, and tech entrepreneurship, Brown developed CureWise after his own cancer diagnosis in early 2025, using AI to navigate treatment options that differed significantly from initial medical recommendations. The platform transforms how patients interact with medical information by deploying multiple specialized AI agents across oncology, pathology, clinical trials, and advocacy domains, then synthesizing insights to help patients understand their diagnosis and advocate for optimal care.

Rather than relying on a single AI perspective, CureWise navigates through medical knowledge from multiple specialized agents—oncology, pathology, clinical trials, advocacy—then synthesizes results to help patients discover new approaches while gaining complete understanding of their condition. The system allows patients to upload labs, imaging, and doctor's notes securely, receiving personalized insights tailored to their specific cancer type, stage, and treatment history. The platform employs sophisticated Retrieval-Augmented Generation (RAG) systems integrated with knowledge graphs that consolidate structured and unstructured medical data, creating a dynamic repository enabling informed healthcare decisions based on latest medical breakthroughs. Brown's personal experience exemplifies the platform's potential: by feeding his medical data into AI and asking detailed questions, he uncovered alternative tests and treatments tailored to his cancer's genetics that his doctors hadn't originally considered, ultimately leading to more effective treatment outcomes.

However, CureWise operates as a fundamentally centralized AI system, not the decentralized approach this report emphasizes. All processing occurs within CureWise's infrastructure, with Brown maintaining primary control over the platform architecture and AI agent coordination. While the multi-agent approach creates internal distribution of specialized functions, there's no federated learning across multiple institutions, no distributed patient data governance, and no multi stakeholder decision-making in platform development.

Despite this centralized architecture, CureWise offers exciting implications for future patient controlled decentralized systems. The concept of specialized AI agents advocating for individual patients could be transformative when adapted to federated frameworks where patients maintain sovereignty over their health data while accessing collective medical intelligence. Imagine CureWise's multi-agent approach distributed across networks of patient-controlled data vaults, each contributing anonymized insights to improve AI recommendations while preserving individual privacy through federated learning. Such an approach could combine CureWise's proven patient advocacy capabilities with the data sovereignty, institutional autonomy, and collaborative benefits that decentralized AI enables.

The regulatory and economic context makes this particularly compelling: healthcare systems worldwide face increasing pressure for patient-centered care and data transparency. CureWise's emphasis that "AI agents are designed to educate and inform, not to diagnose or provide medical advice" while helping patients "advocate with your doctor for the best available treatment" aligns perfectly with patient empowerment trends. Future decentralized iterations could leverage patient controlled health wallets and federated learning to enable truly personalized medical advocacy while keeping sensitive health data under individual control. This could create distributed networks of patient advocates where AI agents learn from collective medical experiences across patient communities without compromising individual privacy—essentially scaling CureWise's personalized advocacy across entire patient populations while maintaining the data sovereignty and personal control that decentralized approaches provide.

REFERENCES:

- Brown, S. (2025) 'I created my own AI medical team. It changed the way doctors treat my cancer', STAT News, 10 September. Available at: <https://www.statnews.com/2025/09/10/ai-cancer-treatment-custom-doctors-response>
- CureWise (2025) Founder Story: Steve Brown. Available at: <https://curewise.com/founder-story>
- Goldfarb, G. (2025) 'How AI helped diagnose and treat my rare cancer | Steve Brown of CureWise', Cut to the Chase Podcast, 3 September. Available at: <https://gregggoldfarb.com/podcast/episodes/how-ai-helped-diagnose-and-treat-my-rare-cancer-steve-brown-of-curewise/>
- The Health Care Blog (2025) 'Steve Brown, CureWise — AI for patients', 31 July. Available at: <https://thehealthcareblog.com/blog/2025/07/31/steve-brown-curewise-ai-for-patients/>

5.6 BONUS RESOURCE: LUMERA HEALTH – PRIVACY-FIRST AI AND TOKENIZATION (US)

Documented by: Thomas Egelhof, Chief Physician, Member of the Management Board, Merian Iselin Clinic for Orthopaedics and Surgery

Lumera Health represents an emerging approach to healthcare IT infrastructure by combining blockchain-based data management with AI-powered clinical tools. Founded in Miami in early 2025, the platform addresses persistent challenges in healthcare data fragmentation, workflow inefficiency, and restricted access to advanced AI capabilities, particularly for smaller healthcare providers. While Lumera offers broad Web3 infrastructure for healthcare, their AI-driven features demonstrate how decentralized platforms can support privacy-preserving diagnostics, workflow automation, and patient-controlled health data management.

The platform's core innovation lies in its LUR Protocol, which tokenizes health records to enable secure sharing without centralized data holders, combined with LumaNodes that distribute AI computation tasks across a decentralized network. This architecture allows healthcare providers to leverage AI tools like LumaQ (for diagnostics, report summarization, and risk flagging) and TRACE™ (for auditable decision traceability) while maintaining HIPAA and GDPR compliance. The system integrates gamified wellness features through their LOCK™ mobile app, which uses Move-to-Earn mechanics to reward users with \$LUR tokens for healthy behaviors, creating economic incentives for preventive care alongside clinical AI applications.

Lumera's governance model operates through \$LUR token staking, with community members participating in proposal submissions via Discord and Telegram, followed by expert review processes that evaluate feasibility and privacy compliance. Token holders vote on strategic priorities, resource allocation, and platform deployments through a multi-stage voting system, with decisions executed automatically via smart contracts on their custom blockchain (achieving over 2,000 transactions per second). The platform launched its \$LUR token in July 2025 and has built a community of thousands across social channels, with the token reaching a market capitalization of approximately \$34.3 million by October 2025.

The platform targets significant efficiency gains in clinical workflows—aiming for up to 67% reductions in administrative workloads and saving clinicians over 10 hours weekly through unified applications and AI-powered automation. While still in early stages with pilots in progress, Lumera's approach demonstrates the potential for tokenized healthcare systems to align incentives between patients, providers, and technology infrastructure. The decentralized data marketplace enables anonymized health data to support research while patients maintain ownership and earn rewards for participation, creating a model where privacy protection and data utility coexist through blockchain transparency and AI-driven analytics.

Looking toward future healthcare systems, Lumera's integration of privacy-first AI with tokenization frameworks illustrates how decentralized platforms might enable cross-border healthcare data exchange aligned with emerging regulations like the European Health Data Space. The combination of distributed computation through LumaNodes, auditable AI decision-making via TRACE™, and community-governed development presents a vision for healthcare IT infrastructure that balances institutional autonomy, patient sovereignty, and collaborative innovation—though realizing this vision at scale will require addressing ongoing challenges in regulatory acceptance, clinical validation, and interoperability with existing healthcare systems.

REFERENCES:

- CoinMarketCap (2025) Lumera Health Token Data. Available at: <https://coinmarketcap.com/currencies/lumera/>
- CoinTrust (2025) 'Lumera Health launches \$LUR token to power Web3 healthcare', Press Release, July. Available at: <https://www.cointrust.com/market-news/lumera-health-launches-lur-token-to-power-web3-healthcare>
- Igbozurike, C. (2025) 'Lumera: Forging a decentralized future with AI and blockchain beyond the hype', Medium, August. Available at: <https://medium.com/@chibuzorjerry045/lumera-forging-a-decentralized-future-with-ai-and-blockchain-beyond-the-hype-by-igbozurike-cb6a0fc43af6>
- Lumera Health (2025) Official Platform Overview. Miami: Lumera Health. Available at: <https://lumerahealth.ai/>

Disclaimer: This project is included as a "bonus resource" rather than a full case study due to its early development stage and associated startup risks. Inclusion does not constitute endorsement or recommendation by the authors or affiliated organizations. References to the \$LUR token are descriptive only and should not be construed as financial advice. Readers should conduct independent due diligence before engaging with any early-stage platform or digital asset.

5.7 BONUS RESOURCE: SPECTRUTHDAO – DECENTRALIZED SCIENCE FOR PTSD TREATMENT AND TRAUMA HEALING (GLOBAL)

Documented by: Thomas Egelhof, Chief Radiologist, Member of the Management Board, Merian Iselin Clinic for Orthopaedics and Surgery

SpectruthDAO represents an emerging decentralized autonomous organization founded in early 2025 that applies blockchain and AI technologies to mental health innovation, specifically targeting PTSD treatment and trauma care. Operating without a traditional headquarters, this global community-driven initiative addresses critical gaps in mental health research where traditional approaches face barriers including fragmented data systems, privacy breaches, inefficient diagnostics that ignore biological factors like epigenetics, and limited access to innovative therapies for underserved populations. With an estimated \$232 billion annual PTSD cost in the United States alone and significant challenges in delivering effective care to veterans and trauma survivors worldwide, SpectruthDAO demonstrates how decentralized science (DeSci) frameworks can support personalized diagnostics, biomarker analysis, and community-led therapeutic innovations.

The platform's core technological approach combines blockchain-based data management with AI powered biomarker analysis through their flagship NeuroPath Navigator agent, which operates on BioProtocol and analyzes epigenetic biomarkers, physiological signals from wearables, and clinical data to identify trauma patterns supporting early intervention. SpectruthDAO's architecture tokenizes trauma biomarkers and wearable inputs for secure sharing without centralized control, distributes AI training via federated analysis of sensitive health data, and implements \$SPEC token based governance where community members vote on research priorities, funding allocation, and therapy pilot execution. The platform has developed an integrated portfolio including mobile PTSD clinics offering AI-monitored ibogaine therapy (despite its Schedule I controlled substance status in many jurisdictions), saliva-based biomarker testing for PTSD prediction using DNA methylation and microbiome analysis, and wearable integrations that detect physiological changes before symptoms manifest.

SpectruthDAO's governance operates through a multi-stage framework where members submit proposals via Discord and Telegram that undergo expert review for feasibility and privacy compliance (including GDPR), followed by token-weighted voting across strategic, allocation, and execution phases. Smart contracts execute funding releases based on milestone achievements such as biomarker validation, while AI-generated outputs including models and datasets are tokenized to enable community control over licensing, royalties, and ongoing development. The DAO launched its ambassador program on August 1, 2025, followed by the NeuroPath Navigator rollout on August 20, and has built a community of thousands across social channels including X, Discord, and Telegram, with partnerships contributing to approximately 15-20% of DeSci activity in mental health as of October 2025.

As an early-stage project, SpectruthDAO has processed initial datasets through AI pilots targeting 20-30% faster diagnostics compared to traditional methods, while the tokenized governance model aims to build trust in PTSD care through transparency and patient empowerment. The platform's approach addresses data silos in epigenetics research and enables hybrid AI-therapy deployment in underserved regions including conflict zones where traditional mental health infrastructure is limited. By combining biological prediction capabilities through biomarker AI with community-led governance that aligns economic incentives for ethical trauma data sharing, SpectruthDAO illustrates how decentralized frameworks might overcome barriers that centralized systems face in mental health innovation.

Looking toward the future of mental health care, SpectruthDAO's integration of AI-driven biomarker analysis with blockchain-based governance and tokenized incentive structures demonstrates potential pathways for scaling personalized trauma care globally while maintaining patient data sovereignty. The combination of wearable monitoring, epigenetic research, and mobile clinic deployment presents a vision for mental health infrastructure that balances rapid innovation with ethical oversight—though realizing this vision will require addressing significant regulatory challenges around ibogaine therapy access, establishing clinical validation for AI-based PTSD prediction, achieving interoperability with existing mental health systems, and navigating the complex landscape of mental health data privacy across diverse global jurisdictions.

REFERENCES:

- EIN Presswire (2025) 'Spectruth DAO launches ambassador program to empower community and accelerate mental health innovation', Press Release, August. Available at: <https://fox4kc.com/business/press-releases/ein-presswire/836110800/spectruth-dao-launches-ambassador-program-to-empower-community-and-accelerate-mental-health-innovation>
- SpectruthDAO (2025) 'PTSD treatment, wearables, and biomarkers discussion', Podcast Episode, May. Available at: <https://www.youtube.com/watch?v=rhPNTAJfefk>

Disclaimer: This project is included as a "bonus resource" rather than a full case study due to its early development stage and associated startup risks. Inclusion does not constitute endorsement or recommendation by the authors or affiliated organizations. References to the \$SPEC token are descriptive only and should not be construed as financial advice. Readers should conduct independent due diligence before engaging with any early-stage platform or digital asset.

5.8 CHAPTER INSIGHTS AND CONCLUSIONS

This chapter examined decentralized validation and deployment through implementations spanning research governance, national infrastructure, educational simulation, patient advocacy, and Web3 platforms. Several key findings guide organizations navigating the transition from experiments to production-scale systems.

Key Finding 1: Successful Deployment Requires Hybrid Architectures Balancing Decentralization with Selective Centralization

Changi General Hospital's AimSG platform demonstrates that embedding decentralized AI in existing healthcare infrastructure demands pragmatic choices. The implementation succeeds by combining local institutional autonomy (AI models deployed within hospitals, data processing executed locally) with coordinated national governance (validation and oversight managed centrally by Synapse and Ministry of Health). Organizations must identify which components genuinely benefit from decentralization (data sovereignty, local clinical decisions, institutional autonomy) versus which improve through selective centralization (quality monitoring, regulatory compliance, interoperability standards). The most successful implementations explicitly map these trade-offs rather than pursuing ideological purity.

Key Finding 2: Decentralized Governance Enables Stakeholder Participation but Requires Design for Healthcare's Unique Accountability Requirements

VitaDAO's blockchain-based research governance illustrates both transformative potential and practical complexity. The DAO successfully coordinates 10,000+ community members across funding decisions and AI deployment through transparent token-weighted voting and smart contract execution. However, healthcare's unique characteristics—regulatory requirements for clinical oversight, liability frameworks demanding identifiable decision-makers, patient safety mandates requiring rapid intervention—create tensions with pure decentralized governance. Organizations must design frameworks preserving stakeholder participation while accommodating regulatory compliance and clinical safety oversight. This often means multi-tier governance where community input shapes strategy while qualified experts retain authority over safety-critical decisions.

Key Finding 3: Certain Healthcare Domains Particularly Benefit from Decentralized Approaches Despite Early Development

The bonus resources reveal domains where decentralized approaches offer distinctive advantages. Mental health innovation (SpectruthDAO's PTSD treatment) benefits by reducing stigma through privacy-preserving data contribution and enabling survivor communities to direct research priorities. Patient advocacy (CureWise's cancer navigation) demonstrates how decentralized architectures could empower patients as active participants with collective intelligence. Chronic disease management (Lumera Health's tokenized records) illustrates how blockchain systems can align economic incentives for preventive care. These domains share characteristics: fragmented traditional systems, information asymmetries, misaligned incentives, and patient populations seeking autonomy. However, early development does not validate these approaches—regulatory acceptance, clinical validation, and operational sustainability remain unproven. Not everything should or can be decentralized.

PRACTICAL IMPLEMENTATION GUIDANCE

Translating these insights into action (as decentralized AI slowly enters mainstream operations) requires tailored approaches across the healthcare ecosystem. Healthcare delivery organizations should start with hybrid architectures and non-critical applications, establishing governance frameworks before technical deployment while investing equally in interoperability and staff training. Patient groups must engage early in governance design, demanding transparency and data sovereignty while recognizing risks in token-based systems. Technology developers should prioritize regulatory compliance from inception, design for interoperability over proprietary lock-in, and implement expert oversight rather than fully permissionless systems. Research institutions should leverage decentralized computation cost-effectively while actively shaping governance standards and maintaining clinical safety oversight. Policymakers would be advised to develop technology neutral frameworks focused on clinical outcomes, create clear approval pathways including liability guidance, and support infrastructure development while establishing firm requirements for patient safety and algorithmic accountability.

The deployment phase represents where decentralized AI transitions from potential to reality—and where the hardest challenges emerge. Technical feasibility proves insufficient; successful deployment demands regulatory acceptance, clinical validation, operational integration, and stakeholder trust. The implementations examined demonstrate both early successes and significant ongoing challenges, acknowledging that many approaches remain far from production-ready as this field rapidly evolves.

CHAPTER 6: FUTURE DIRECTIONS AND STRATEGIC RECOMMENDATIONS



Foreword by the author: Thomas Egelhof, Chief Radiologist,
Member of the Management Board, Merian Iselin Clinic for
Orthopaedics and Surgery

Picture a morning in the near future: You're starting your day when your phone chimes softly. It's not a news alert or a text, but a message from your health app: "Your sleep was slightly restless last night, and your heart rate suggests mild dehydration. Try adding an extra glass of water today, and let's schedule a quick check-in with a nutritionist." This isn't a far-off dream—it's the promise of decentralized AI in healthcare, where emerging technologies like blockchain, edge computing, and federated learning are quietly transforming how we care for ourselves.

I envision a future world where everyone carries a health companion in their pocket—an app powered by artificial intelligence that monitors vital signs through wearables, detects subtle changes in your body, and offers tailored medical advice. It flags potential issues, like irregular blood pressure or early signs of fatigue, before they escalate. More than that, it tracks your progress, adjusting recommendations based on your unique health data, from diet to recovery, all while keeping your information secure and private.

Achieving this vision requires a clear implementation roadmap. We can begin with pilot projects in diverse communities, using open-source AI tools and privacy-focused systems like decentralized data storage. Strategic recommendations include fostering collaboration between tech developers and healthcare providers to ensure seamless integration, urging policymakers to create flexible regulations that prioritize both innovation and ethics, and encouraging companies to build systems that work together, making healthcare accessible to all.

As you read on, you'll discover how we believe decentralized AI will be reshaping healthcare through secure data sharing, predictive analytics, and patient-centered care. From smarter diagnostics to streamlined insurance processes, this chapter lays out the opportunities and challenges ahead. The future of health is personal, proactive, and already taking shape. Let's explore what's next.

The views expressed are those of the author and do not necessarily reflect those of the employer, Merian Iselin - Clinic for Orthopedic and Surgery.

6.1 CHAPTER OVERVIEW

Chapter 6 synthesizes lessons from previous chapters into forward-looking strategic guidance for stakeholders navigating decentralized AI's evolution through 2030. Having examined foundational concepts, regulatory frameworks, and implementations across data collection, model development, and deployment, this chapter addresses how different stakeholders should prepare for a healthcare landscape where decentralized AI transitions from experimental pilots to operational reality.

The chapter opens with an analysis of emerging technology and regulatory trends across five interconnected domains. Technology trends examine sovereign AI imperatives driving national data sovereignty, agentic AI systems introducing autonomous decision-making, privacy-enhancing technologies maturing beyond federated learning, edge computing shifting processing closer to patients, and blockchain-based model marketplaces. Data and infrastructure trends explore critical interoperability standards, patient data wallets shifting control to individuals, synthetic data addressing training scarcity, and real-world evidence requirements driving post-deployment validation. Governance and economic trends assess DAOs as coordination mechanisms, tokenization creating incentive structures, and computation cost sustainability challenges. Sustainability constraints examine environmental impact and clinical workflow integration as persistent barriers. Regulatory landscape trends map how frameworks like the EU AI Act create clarity while liability questions for autonomous systems remain unresolved through 2030.

For each trend, the analysis explicitly identifies whether forces create tailwinds accelerating adoption or headwinds introducing resistance and complexity. This balanced assessment acknowledges that decentralized AI's trajectory depends not only on technological capability but also on regulatory evolution, economic sustainability, and organizational readiness.

The chapter then translates trend analyses into stakeholder-specific strategic implications and recommendations. Rather than prescribing universal solutions, this section frames key questions that healthcare delivery organizations, patient groups, technology developers, research institutions, and policymakers should address as they position themselves for 2025-2030. For each stakeholder group, strategic implications contextualize specific tensions, opportunities, and constraints, while key strategic questions prompt reflection on critical decisions around experimentation timing, governance design, partnership selection, capability building, and success measurement. The recommendations acknowledge uncertainty—recognizing that optimal strategies depend on organizational context, risk tolerance, and strategic objectives—while providing frameworks that stakeholders can adapt to their evolving circumstances.

6.2 EMERGING TRENDS IN DECENTRALIZED AI FOR HEALTHCARE

The authors envision that the period from 2025 to 2030 will witness accelerating transformation in decentralized healthcare AI, driven by converging technological advances, evolving regulatory frameworks, and shifting economic and social realities. This section examines emerging trends across five domains—technology, data and infrastructure, governance and economics, sustainability and constraints, and regulatory landscape—identifying which forces are likely to propel decentralized AI adoption (tailwinds) versus those that may create resistance or complexity (headwinds). These projections reflect the authors' analysis of current trajectories and should be understood as informed predictions rather than certainties in this rapidly evolving field.

A. TECHNOLOGY TRENDS

SOVEREIGN AI AND GEOPOLITICAL CONSIDERATIONS (STRONG TAILWIND)

Sovereign AI—nation-states developing indigenous AI capabilities with data remaining within national boundaries—will likely emerge as one of the most powerful forces driving decentralized healthcare AI adoption through 2030. Switzerland's Apertus project exemplifies this trajectory, having created the first sovereign large language model trained on multilingual national data. China's strategic deployment of decentralized computation infrastructure (necessitated by restricted access to advanced NVIDIA chips) demonstrates how geopolitical constraints are already accelerating alternative architectures. DeepSeek's open-source models currently enable over 100 Chinese hospitals to run state-of-the-art reasoning locally, a pattern the authors expect to proliferate globally in the future.

This trend will create strong tailwinds for decentralized approaches: national security concerns will increasingly prioritize data localization, regulatory frameworks will mandate domestic data processing, and healthcare institutions will gain leverage to maintain control rather than surrender data to foreign technology giants. However, sovereignty will also introduce headwinds through potential fragmentation—incompatible national AI ecosystems that may hinder international research collaboration and create inefficiencies through duplicated development efforts.

AGENTIC AI AND MULTI-AGENT SYSTEMS (MIXED: TAILWIND FOR INNOVATION, HEADWIND FOR REGULATION)

The authors anticipate that autonomous AI agents will meaningfully shift healthcare from passive tools to active participants in clinical workflows in the next 3-5 years. Tsinghua University's Agent Hospital demonstrates the sophistication possible through coordination across 42 specialized agents spanning 21 medical departments. CureWise's multi-agent cancer navigation and the ASI Alliance's framework for autonomous coordination across data marketplaces and computation infrastructure suggest the trajectory toward widespread agentic deployment.

For decentralized healthcare AI, agentic systems will create tailwinds by enabling sophisticated coordination without centralized control, allowing institutions to deploy specialized agents while participating in federated intelligence networks, and empowering patients through personalized AI advocates. However, significant headwinds will emerge around regulatory uncertainty—existing frameworks assume human-in-the-loop control that autonomous agents challenge—liability questions when agents make consequential decisions without direct oversight, and standard gaps for agent-to-agent communication protocols and security requirements. Additionally, architectural limitations may constrain progress: as Shojaee et al. (2025) demonstrate, current LLM architectures face reasoning complexity thresholds beyond which their accuracy collapses to zero, potentially limiting autonomous agent capabilities until next-generation architectures emerge.

PRIVACY-ENHANCING TECHNOLOGIES BEYOND FEDERATED LEARNING (STRONG TAILWIND)

Homomorphic encryption, secure multi-party computation, differential privacy, and zero-knowledge proofs will most likely transition from academic concepts to standard production implementations by 2030, in the authors' assessment. Mayo Clinic's "Data Behind Glass" architecture and European Health Data Space implementations demonstrate how secure containers and differential privacy are already enabling cross-border research collaboration while preserving individual privacy.

These technologies will create powerful tailwinds: regulatory push from GDPR, HIPAA updates, and the EU Digital Markets Act increasingly favor PETs; technical maturity will enable real-world healthcare deployments at scale; and competitive advantage will accrue to institutions demonstrating superior privacy protection. The primary headwind will remain computational overhead—some PETs significantly increase processing costs and latency, potentially limiting real time clinical applications until hardware advances (possibly including quantum computing breakthroughs by 2030) address these constraints.

EDGE AI AND ON-DEVICE HEALTHCARE MODELS (STRONG TAILWIND)

The shift from cloud to edge computation for healthcare AI will accelerate through 2030, addressing critical constraints around latency, bandwidth, and data sovereignty. We believe that wearable medical devices, point-of-care diagnostics, and clinical decision support systems will increasingly run local AI models rather than cloud-dependent inference. This trend will particularly benefit low resource settings where reliable connectivity cannot be assumed.

Edge AI will create tailwinds through reduced latency for time-critical clinical decisions, bandwidth efficiency in resource-constrained environments, privacy preservation by default through local processing, and resilience—systems continuing to function without internet connectivity. Headwinds will include limited computational power on edge devices restricting model complexity (though this may improve dramatically if quantum computing becomes viable for edge deployment by 2030), challenging model update and versioning across distributed device fleets, and fragmentation risk as device manufacturers deploy incompatible proprietary systems.

AI MODEL MARKETPLACES AND IP-NFTS (MODERATE TAILWIND, LONG-TERM)

VitaDAO and Molecule's IP-NFT framework demonstrates how blockchain-based intellectual property tokenization could enable community ownership of research outputs while preserving creator rights. The authors expect emerging decentralized marketplaces for healthcare AI models to accelerate collaborative development in the near future, incentivizing model sharing through transparent attribution and economic benefit distribution.

This trend will create tailwinds by aligning incentives for open model sharing, enabling fractional ownership of valuable AI assets, and providing transparent provenance tracking for model origins and training data. However, significant headwinds will include regulatory uncertainty around tokenized intellectual property in highly regulated healthcare contexts, potentially limited adoption due to cryptocurrency volatility and complexity, and questions about clinical validation when models trade in open marketplaces.

B. DATA & INFRASTRUCTURE TRENDS

INTEROPERABILITY STANDARDS EVOLUTION (MAJOR TAILWIND WITH IMPLEMENTATION HEADWINDS)

We believe that FHIR (Fast Healthcare Interoperability Resources) adoption will accelerate dramatically through 2030, driven by the 21st Century Cures Act in the United States and the European Health Data Space in the EU. Standardized APIs will become mandatory, while HL7 FHIR, DICOM for imaging, and IHE integration profiles will create foundations for federated AI systems.

Interoperability represents perhaps the most critical enabler: without it, decentralized AI cannot scale beyond isolated pilot projects. Regulatory mandates will create strong tailwinds, as will demonstrated cost savings from reduced data integration overhead. However, implementation headwinds might remain severe—legacy systems are expected to resist standardization, competing standards might create confusion (FHIR versions, custom extensions), and healthcare's fragmented vendor landscape could slow adoption.

PATIENT DATA WALLETS AND SELF-SOVEREIGN IDENTITY (MODERATE TAILWIND, ACCELERATING)

Personal health data vaults—such as Apple Health, Google Health Connect, and EU Health Data Space-compliant systems—have the potential to shift patients from passive data subjects to active data stewards in the next 3-5 years. Self-sovereign identity (SSI) frameworks already enable patients to control credentials and consent management across federated networks.

Patient data sovereignty is poised to create tailwinds through regulatory alignment (GDPR's data portability rights, patient access mandates), demonstrated patient demand for control over health information, and technical maturity of SSI frameworks. Future headwinds include user experience complexity—most patients will likely continue to lack technical sophistication for cryptographic key management—liability questions when patients control critical medical data, and resistance from institutions benefiting from current data asymmetries.

SYNTHETIC DATA AND DATA SCARCITY SOLUTIONS (MODERATE TAILWIND)

Generative AI is already enabling the creation of synthetic training datasets that preserve statistical properties while protecting privacy. This may address in the near future, the "peak data" problem (i.e. scarcity of quality training data, particularly for rare diseases and underrepresented populations), a theme widely present in public discourse at the time this report was written (Vetterle, 2024). Synthetic data could enable federated learning with smaller institutional datasets by augmenting limited real-world data, however a full substitute is not deemed feasible by the authors.

Tailwinds will include addressing data scarcity without privacy compromise, enabling research on rare conditions where real data remains insufficient, and regulatory acceptance growing as validation methods improve. Headwinds will involve clinical validation challenges—synthetic data must prove equivalent to real-world data for regulatory approval—quality concerns around model bias amplification, and technical limitations in capturing complex physiological interactions.

REAL-WORLD EVIDENCE VIA FEDERATED NETWORKS (STRONG TAILWIND)

FDA and EMA is likely to increasingly require post-market surveillance demonstrating AI system performance in real-world clinical settings (rather than controlled trials) through 2030. Decentralized networks might prove ideal for continuous model monitoring across diverse patient populations while respecting data privacy and institutional autonomy. Changi General Hospital's AimSG platform demonstrates federated post-deployment validation at national scale, a model the authors expect to proliferate globally during the next 3-5 years and beyond.

Real-world evidence will create tailwinds through regulatory requirements driving adoption, demonstrated value in detecting model drift and performance degradation, and institutional willingness to participate when data remains local. Minor headwinds could include standardization challenges for monitoring metrics and coordination overhead across federated networks.

C. GOVERNANCE & ECONOMIC TRENDS

DAOS FOR HEALTHCARE GOVERNANCE (MODERATE TAILWIND, HIGH UNCERTAINTY)

VitaDAO's success coordinating 10,000+ community members across research funding demonstrates blockchain-based governance viability today. SpectruthDAO's mental health innovation shows DAO frameworks extending beyond research into clinical care coordination already. The authors anticipate that in the near future, DAOs may govern significant healthcare research portfolios and potentially clinical care networks, though substantial regulatory uncertainty remains.

Future tailwinds include demonstrated community engagement and transparent decision-making, alignment with patient empowerment trends, and economic efficiency—reduced overhead versus traditional governance bureaucracies. Substantial headwinds would involve regulatory frameworks assuming identifiable legal entities, liability concerns when no central authority exists, governance attack risks (wealth concentration, coordinated manipulation), and potentially limited adoption beyond crypto-native communities.

ECONOMIC MODELS AND TOKENIZATION (MIXED: INNOVATION POTENTIAL, REGULATORY HEADWIND)

Token-based incentive structures could align stakeholder interests in ways traditional models cannot—rewarding data contribution, incentivizing model improvement, and enabling patient benefit-sharing. Lumera Health's gamified wellness tokens and SpectruthDAO's research participation rewards demonstrate emerging models that may mature towards the end of the decade.

Innovation potential creates tailwinds, but regulatory headwinds will likely dominate: healthcare will remain conservative about financial incentives potentially compromising clinical judgment, securities regulations will create legal uncertainty around utility tokens, cryptocurrency volatility will undermine economic stability, and reputational risks will deter mainstream healthcare institutions. Tokenization will likely remain peripheral for a number of years, until regulatory clarity will emerge.

D. SUSTAINABILITY & CONSTRAINTS

COMPUTATION COSTS AND FINANCIAL SUSTAINABILITY (SIGNIFICANT HEADWIND)

The Praetorian Capital analysis revealing data centers projected to spend \$40 billion in depreciation (2026) against only \$15-20 billion revenue highlights economically unsustainable trajectories. Centralized AI's concentration amplifies this problem, while decentralized computation networks (such as Akash Network, Aethir, io.net) might offer cost reduction potential of up to 85% through distributed resource utilization—a model the authors expect to gain significant traction by 2030.

This will create complex dynamics: centralized AI's unsustainability for certain use cases should become a tailwind for decentralized alternatives, demonstrated cost savings could drive institutional adoption, and utilization of existing institutional computation capacity should improve efficiency. However, coordination overhead in decentralized networks will likely introduce costs, quality assurance across heterogeneous infrastructure will prove challenging, and most institutions will likely lack expertise for federated deployment for the foreseeable future.

ENVIRONMENTAL IMPACT AND GREEN AI (EMERGING TAILWIND)

AI's energy consumption will continue to face increasing regulatory scrutiny as climate concerns intensify. Decentralized approaches may improve efficiency through distributed resource utilization, edge computing reducing data transmission overhead, and leveraging institutional infrastructure during off-peak hours. However, whether decentralized AI proves more sustainable than optimized centralized systems remains empirically unproven and will require careful measurement in the years to come.

Environmental considerations can create tailwinds through regulatory pressure for sustainable AI, institutional sustainability commitments driving technology choices, and growing public awareness of AI's carbon footprint. Headwinds however include measurement challenges—energy consumption across federated networks proves difficult to track—and uncertainty whether decentralization genuinely improves sustainability at scale. Quantum computing breakthroughs by 2030 could dramatically alter this calculus, potentially enabling far more energy-efficient AI operations or conversely requiring massive energy investments in new infrastructure. A wildcard for now.

CLINICAL WORKFLOW INTEGRATION AND WORKFORCE IMPACT (CRITICAL HEADWIND REQUIRING ATTENTION)

Physician burnout (estimated to affect ~50% of doctors and nurses) creates urgent pressure for AI assisted workflow optimization. However, poorly integrated AI worsens cognitive burden rather than reducing it. Decentralized systems will introduce additional complexity through federated coordination, heterogeneous interfaces across institutions, and training requirements for distributed systems—challenges the authors expect to persist for a while, as the learning curve stabilizes.

Workflow integration will represent a critical headwind: technologically successful systems proving clinically irrelevant if adoption fails due to poor usability. This will demand explicit design focus on clinical workflows, extensive end-user involvement in system design, and realistic deployment timelines accounting for training and adaptation periods.

E. REGULATORY LANDSCAPE

EU AI ACT AND HARMONIZED FRAMEWORKS (MODERATE TAILWIND WITH COMPLIANCE OVERHEAD)

The EU AI Act should provide regulatory clarity for high-risk healthcare AI systems, including explicit pathways for federated learning validation and recognition of decentralized approaches for privacy compliance. EHDS implementation is expected to create infrastructure for cross-border federated research while maintaining data sovereignty, however full operational deployment might take longer than expected given EU's track record of similar implementations.

Regulatory clarity however, will create tailwinds by reducing deployment uncertainty, establishing clear compliance requirements, and harmonizing approaches across EU member states. Headwinds will likely involve compliance costs, particularly burdensome for smaller organizations, and documentation requirements potentially slowing innovation velocity through the initial implementation period (2025-2027).

LIABILITY FRAMEWORKS FOR AUTONOMOUS SYSTEMS (SIGNIFICANT HEADWIND)

Existing liability frameworks assume identifiable decision-makers—physicians, device manufacturers, healthcare institutions. Decentralized autonomous agents already challenge these assumptions: when an AI agent makes a consequential clinical decision across a federated network, who will bear liability? Current legal frameworks provide no clear answers, and the authors do not anticipate resolution before 2028-2030 at earliest.

This represents one of the most significant headwinds for agentic healthcare AI through 2030: liability uncertainty will deter institutional adoption, insurance markets will lack products for distributed autonomous systems, and regulatory frameworks will require substantial evolution. Resolution likely requires 5-10 years of legal framework development, limiting widespread deployment until clarity emerges.

CROSS-BORDER DATA FLOWS AND DIGITAL SOVEREIGNTY (MIXED)

Conflicting national requirements will continue to create complexity through 2030: China's PIPL mandates data localization, GDPR restricts transfers outside the EU, and various other nations will likely impose digital sovereignty requirements, too. Federated learning can address these concerns by keeping data local while enabling collaborative model development—an approach the authors expect to become standard practice for international healthcare AI collaborations.

For decentralized AI, sovereignty requirements will likely create tailwinds—federated architectures naturally comply with data localization—but also headwinds through fragmentation risk, compliance complexity across jurisdictions, and reduced efficiency from geographic constraints on collaboration.

SYNTHESIS: NET ASSESSMENT FOR 2025-2030

The authors assess that the balance of forces will meaningfully favor decentralized healthcare AI adoption through 2030, driven primarily by sovereign AI imperatives, privacy-enhancing technology maturation, regulatory frameworks accommodating federated approaches, and real-world evidence requirements. However, significant headwinds around liability frameworks, tokenization regulatory uncertainty, workflow integration complexity, and interoperability implementation challenges will likely constrain adoption velocity through at least 2027-2028. Organizations should expect hybrid architectures combining selective decentralization with pragmatic centralization to dominate deployments in the near future, with pure decentralized systems remaining limited to specific use cases (such as research, rare diseases, mental health) where unique advantages outweigh coordination complexity.

REFERENCES:

- Abdalla, H.B., Kumar, Y., Tang, L., Marchena, J., Guzman, S., Cheraghy, M., Awlla, A. and Gheisari, M. (2025) 'The future of artificial intelligence in the face of data scarcity', *Computers, Materials & Continua*, 84(1). doi: 10.32604/cmc.2025.063551.
- China (2021) Personal Information Protection Law (PIPL). Standing Committee of the National People's Congress. Available at: https://www.gov.cn/xinwen/2021/08/20/content_5632486.htm (Chinese); English translation available at: <https://digichina.stanford.edu/work/translation-personal-information-protection-law-of-the-peoples-republic-of-china-effective-nov-1-2021/>
- ETH Zurich (2025) 'Apertus: A language model built for the public good', *ETH News*, July. Available at: <https://ethz.ch/en/news-and-events/eth-news/news/2025/07/a-language-model-built-for-the-public-good.html>
- European Commission (2024) Artificial Intelligence Act. Official Journal of the European Union. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- European Commission (2024) Regulation on the European Health Data Space (EHDS). Official Journal of the European Union. Available at: https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space_en
- European Union (2016) General Data Protection Regulation (GDPR), Regulation (EU) 2016/679. Official Journal of the European Union. Available at: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- Noor, P. (2025) 'There's a stunning financial problem with AI data centers', *Futurism*, 28 August. Available at: <https://futurism.com/data-centers-financial-bubble>
- Shojaee, P., Jain, A., Tipirneni, S. and Reddy, C.K. (2025) 'The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity', *arXiv preprint arXiv:2506.06941*. Available at: <https://arxiv.org/abs/2506.06941>
- U.S. Department of Health and Human Services (2016) 21st Century Cures Act, Public Law 114-255. Available at: <https://www.congress.gov/bill/114th-congress/house-bill/34>
- U.S. Food and Drug Administration (2024) Guidance for Industry and FDA Staff: Software as a Medical Device (SaMD): Clinical Evaluation. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/software-medical-device-samd-clinical-evaluation>

6.3 STRATEGIC IMPLICATIONS AND RECOMMENDATIONS

The emergence of decentralized AI in healthcare presents distinct challenges and opportunities for different stakeholder groups, each facing unique strategic decisions shaped by their role in the healthcare ecosystem. Rather than prescribing universal solutions, this section frames key strategic questions and considerations that each stakeholder group should address as they prepare for the 2025-2030 period, when decentralized approaches will likely transition from experimental pilots to operational deployments. The recommendations reflect lessons learned from current implementations and trend analysis, translated into forward-looking strategic frameworks that organizations can adapt to their evolving circumstances.

FOR HEALTHCARE DELIVERY ORGANIZATIONS

STRATEGIC IMPLICATIONS

Healthcare delivery organizations—hospitals, clinics, and integrated health systems—face a fundamental tension between maintaining operational control and participating in collaborative AI networks that could improve clinical outcomes. Current implementations like Changi General Hospital's AimSG platform and Mayo Clinic's federated research infrastructure demonstrate what becomes possible when institutions combine autonomy with coordinated governance. However, as decentralized AI matures through 2030, these organizations will confront intensifying pressures: legacy IT infrastructure increasingly incompatible with federated architectures, clinical workflows requiring redesign for AI-augmented decision-making, accelerating staff skill requirements, and evolving liability frameworks for systems operating across institutional boundaries. The financial calculus will shift—early adopters accepting implementation costs today position themselves advantageously, while late movers risk vendor lock-in and competitive disadvantage as decentralized approaches become standard practice.

KEY STRATEGIC QUESTIONS

As your organization prepares for a healthcare landscape where decentralized AI becomes increasingly prevalent, these questions warrant careful consideration:

Where should your organization begin experimenting with decentralized AI over the next 2-3 years to build institutional capabilities before competitive pressure forces hasty adoption?

Note: The evidence suggests starting with non-clinical applications or clinical decision support in well-defined domains, allowing time to develop governance expertise before deploying in high-stakes clinical contexts.

How can you design governance frameworks now that will scale as federated networks mature—explicitly defining which decisions remain local versus which benefit from coordination as standards emerge through 2028-2030?

What hybrid architecture positions your institution to adapt as the field evolves, avoiding premature commitment to either extreme centralization or pure decentralization while maintaining flexibility to shift as best practices crystallize?

How will you measure success beyond traditional ROI to capture emerging benefits like data sovereignty, collaborative research capabilities, and reduced algorithmic bias—and what metrics make sense when full value realization may take 5-7 years?

What partnerships position you to influence standards development rather than becoming a passive adopter of externally imposed frameworks—and which consortia and collaborative initiatives merit investment now to shape the 2028-2030 landscape?

How do you build internal expertise in federated technologies while managing existing centralized systems through what will likely be an extended 5-7 year transition period with significant hybrid complexity?

FOR PATIENT GROUPS AND END USERS

STRATEGIC IMPLICATIONS

Patient organizations and advocacy groups increasingly recognize that governance frameworks being established now will shape whether AI systems serve patient interests throughout the next decade. The shift toward patient data sovereignty—through personal health vaults, self-sovereign identity frameworks, and DAO-based research governance—creates unprecedented opportunities for meaningful participation in AI development. However, as these mechanisms proliferate through 2030, significant barriers will require proactive attention: technical complexity that excludes non expert patients, power dynamics that may co-opt patient participation for institutional legitimacy rather than genuine empowerment, and tokenization models that could create new exploitation patterns where economically vulnerable patients face pressure to commodify their health data. The regulatory landscape will continue evolving—GDPR-style patient rights spreading globally while implementation challenges persist, creating a complex patchwork requiring sophisticated navigation.

KEY STRATEGIC QUESTIONS

If you are representing a patient organization considering how to position itself for the decentralized AI future, we encourage you to reflect on the following:

How can you ensure that governance frameworks being established in 2025-2027 include authentic patient voice rather than tokenistic consultation—and what participation mechanisms remain accessible as systems grow more technically complex through 2030?

What transparency standards should you demand now from AI systems, recognizing that algorithms deployed in 2025-2026 will likely remain in use for 5-10 years and that establishing transparency precedents early proves easier than retrofitting later?

How do patient data wallets and self-sovereign identity frameworks balance individual control with collective research benefits—and what governance frameworks ensure patients who contribute data to AI development through 2030 share equitably in resulting value?

What protections prevent exploitation through tokenization schemes, and how do you distinguish genuine empowerment from extractive models as token-based healthcare systems proliferate over the next 3-5 years?

How can patient organizations build technical evaluation capacity now to meaningfully assess proposals through 2030 without creating unsustainable dependencies on external advisors?

What partnerships with academic institutions and advocacy networks provide credible guidance for navigating governance decisions that will shape healthcare AI next decade?

FOR TECHNOLOGY DEVELOPERS AND MEDTECH STARTUPS

STRATEGIC IMPLICATIONS

Technology developers face a strategic inflection point as healthcare stakeholders increasingly favor decentralized architectures, fundamentally challenging traditional software business models built on data moats and vendor lock-in. Successful implementations through 2030 will prioritize interoperability, privacy preservation, and institutional autonomy—requiring fundamentally different product architectures and business models than those optimizing for centralized control. Developers must navigate regulatory pathways that will continue evolving in the near future (EU AI Act implementation, FDA guidance updates, emerging national frameworks), design for heterogeneous deployment environments, and build sustainable business models as traditional competitive advantages erode. The talent landscape will remain challenging through 2030 as demand for federated learning, privacy-enhancing technologies, and blockchain governance expertise will likely exceed supply.

KEY STRATEGIC QUESTIONS

Technology companies positioning themselves for the decentralized AI future should consider:

How do we design products now that will remain relevant through 2030 in healthcare's hybrid reality where deployment preferences span the centralization-decentralization spectrum—and what modular architectures provide flexibility as customer requirements evolve?

What business models will prove sustainable when network effects and data moats diminish—and how do we begin transitioning toward value capture through services, expertise, and platform orchestration before market dynamics force reactive pivoting?

How do we build regulatory compliance into development roadmaps from the start, anticipating regulatory maturation rather than treating compliance as a future problem that constrains already-developed products?

How do we attract and retain scarce talent in decentralized technologies through a 5-7 year period when competing against well-funded centralized AI companies?

What partnerships accelerate development while reducing duplicated effort—and how do we structure intellectual property frameworks that enable collaboration without constraining business model evolution through 2030?

FOR RESEARCH INSTITUTIONS AND CONSORTIA

STRATEGIC IMPLICATIONS

Research institutions occupy a unique position—simultaneously developing technologies, stewarding data, training future practitioners, and bridging academic research with clinical implementation. Successful initiatives like MELLODDY, VitaDAO, and federated clinical trial networks demonstrate collaborative models that will likely become standard practice in the near future. However, academic incentive structures often conflict with long-term collaborative commitments that decentralized approaches require, while resource constraints around infrastructure investment, technical expertise, and legal capacity persist. As decentralized approaches mature through 2030, research institutions that establish collaborative frameworks and technical capabilities early will significantly influence how the field evolves.

KEY STRATEGIC QUESTIONS

Research institutions preparing for the next phase of healthcare AI development might ask:

What governance frameworks enable multi-institutional collaboration without sacrificing academic freedom—and how do we establish agreement structures now that remain viable as projects scale and commercial opportunities emerge through 2030?

How do we participate in standards development actively enough to influence protocols reflecting academic values around openness and equity as standards crystallize between 2026-2028?

What educational programs prepare students for decentralized paradigms that will define their careers through 2030 and beyond, rather than teaching primarily centralized approaches that may prove obsolete?

How do we balance open publication with data protection and institutional interests in federated projects—and what frameworks established now scale gracefully as our institution's collaborative research portfolio expands through 2030?

What institutional review board processes address decentralized research considerations, and how do we develop ethical frameworks now that remain robust as technology capabilities and deployment contexts evolve?

What funding models support sustained participation when most grant cycles emphasize novel findings over the infrastructure maintenance and standards development that decentralized approaches require?

FOR POLICYMAKERS AND REGULATORS

STRATEGIC IMPLICATIONS

Policymakers and regulators face perhaps the most complex challenge—creating frameworks that encourage innovation while protecting patients as decentralized AI transitions from experimental to operational deployment between 2025 and 2030. Current regulatory disparities are stark: the EU has developed frameworks explicitly accommodating decentralized approaches while many jurisdictions lack any AI-specific healthcare regulations. The decisions made now will profoundly shape the field through 2030 and beyond—premature regulation risks stifling beneficial innovation before approaches prove themselves, while delayed regulation allows potential harms to scale and creates difficult-to-correct path dependencies. The liability question grows more urgent as autonomous systems proliferate, yet resolution likely requires 5-10 years of legal framework development.

KEY STRATEGIC QUESTIONS

As policymakers and regulators prepare for decentralized AI's maturation, we encourage you to consider these forward-looking questions:

How do you develop technology-neutral frameworks now that remain relevant by 2030 as specific technologies evolve, focusing regulation on measurable outcomes and validation processes rather than mandating architectural approaches that may prove obsolescent?

What approval pathways should you establish in 2025-2027 to accommodate decentralized systems where traditional assumptions no longer apply—and how do pilot programs inform scaled frameworks by 2028-2030?

What liability frameworks can you begin developing now to address distributed autonomous systems, recognizing that comprehensive frameworks likely require 5-10 years but that partial guidance would reduce current uncertainty?

How do you facilitate international regulatory coordination as federated networks increasingly operate across borders—and what mechanisms established now enable cooperation as systems scale through 2030?

How do you create regulatory sandboxes that allow experimentation without prematurely validating technologies requiring further proof, and what graduated pathways transition successful pilots to scaled deployment by 2028-2030?

What monitoring frameworks provide visibility into decentralized system performance without intrusive oversight that undermines privacy and autonomy benefits?

SYNTHESIS ACROSS STAKEHOLDER GROUPS

Several themes emerge across stakeholder categories that suggest opportunities for productive collaboration as decentralized AI matures through 2030. All groups will benefit from interoperability standards development—yet standards require sustained investment that no single stakeholder can justify individually, suggesting the need for multi-stakeholder consortia to be established immediately in order to shape the 2028-2030 landscape. Liability frameworks for autonomous decentralized systems create uncertainty for every stakeholder, indicating that regulatory clarity should be a collective advocacy priority with the recognition that comprehensive frameworks likely won't emerge until 2028-2030. The tension between individual institutional interests and collective benefits from federation recurs across contexts, suggesting that successful decentralized AI requires explicit governance frameworks aligned with incentives—and that frameworks established now might prove difficult to revise later. Finally, the gap between technical possibility and organizational readiness appears consistently, suggesting that adoption through 2030 depends as much on change management, training, and cultural transformation as on technology development itself.

6.4 CHAPTER INSIGHTS AND CONCLUSIONS

This chapter examined the future trajectory of decentralized AI in healthcare through 2030, analyzing emerging trends and translating them into strategic guidance for diverse stakeholders. Several critical insights emerge that frame both the opportunities and constraints shaping the field's evolution over the next 3-5 years.

Key Finding 1: Hybrid Architectures Will Dominate Through 2030, Not Pure Decentralization

The trend analysis reveals that successful implementations through 2030 will combine selective decentralization with pragmatic centralization rather than pursuing ideological purity. Sovereign AI imperatives, privacy-enhancing technologies, and real-world evidence requirements create powerful tailwinds for decentralized approaches—yet equally significant headwinds around liability frameworks, workflow integration complexity, and interoperability challenges constrain wholesale adoption. Organizations claiming "fully decentralized" systems often obscure centralized components essential for regulatory compliance, quality assurance, or economic viability. The evidence suggests that explicitly mapping which functions genuinely benefit from decentralization (data sovereignty, local clinical decision-making, institutional autonomy) versus which improve through selective centralization (quality monitoring, regulatory compliance, interoperability standards) will prove more valuable than defaulting to either extreme. Changi General Hospital's AimSG platform exemplifies this pragmatism—local AI deployment with centralized national validation—demonstrating how hybrid architectures balance innovation with accountability.

Key Finding 2: Two Foundational Enablers Will Determine Whether Decentralized AI Scales: Interoperability Standards and Liability Frameworks

Two foundational challenges prevent decentralized AI from scaling beyond experimental deployments: interoperability standards and liability frameworks. These prove deeply interconnected—without interoperability, federated systems cannot communicate; without liability clarity, organizations cannot deploy autonomous agents.

Despite FHIR adoption mandates, legacy systems resist standardization while competing versions create fragmentation. Standards established in 2025-2027 will shape healthcare AI through 2030—once solidified, revision becomes exponentially harder. Stakeholders must engage now in standards consortia and regulatory sandboxes to influence rather than inherit governing frameworks. Similarly, when AI agents make decisions across federated networks with no controlling entity, existing legal frameworks provide no guidance. Organizations defer deployment fearing uninsurable risks while patients face unclear recourse. This uncertainty will persist through 2028-2030 as legal frameworks require 5-10 years to evolve.

KEY FINDING 3: ORGANIZATIONAL READINESS GAPS EXCEED TECHNICAL LIMITATIONS AS THE PRIMARY CONSTRAINT

The gap between technical capability and organizational capacity to implement decentralized AI emerged consistently across case studies and trend analysis. Federated learning, privacy-enhancing technologies, and blockchain governance frameworks prove technically viable today—yet most healthcare organizations lack the expertise to deploy them, governance frameworks to manage them, and change management capacity to integrate them into clinical workflows. Physician burnout affecting a large proportion of clinicians creates urgent pressure for AI assistance, yet poorly integrated systems worsen cognitive burden rather than relieving it. This suggests that investment in training, workflow redesign, and organizational change management may prove more valuable than additional technical development for many organizations. The implication challenges conventional assumptions: success through 2030 depends less on developing more sophisticated decentralized AI technologies than on building institutional capabilities to implement existing approaches effectively.

PRACTICAL IMPLICATIONS FOR STRATEGIC PLANNING

These findings translate into several practical considerations for organizations navigating the 2025–2030 period. Organizations should start with small pilots in non-clinical domains or well-defined clinical contexts to build capabilities without existential risk, while engaging immediately in standards consortia and policy development to influence rather than inherit frameworks. Interoperability investments should be treated as essential infrastructure rather than optional overhead, recognizing that without standards, decentralized AI cannot scale. Governance frameworks defining decision authorities, consent management, and benefit-sharing should be designed during low-stakes pilot phases rather than amid operational pressure. Organizations must plan for extended hybrid operation managing centralized and decentralized systems simultaneously through 2030, measuring success through capability building and strategic positioning rather than only near-term operational metrics. Finally, collaborative advocacy through professional societies and industry consortia might prove essential for accelerating liability framework clarity that no single organization can achieve alone.

CLOSING REMARKS

This report has examined decentralized AI in healthcare as it transitions from theoretical promise to practical implementation, documenting the convergent forces driving adoption, the architectural foundations enabling collaboration, and the real-world deployments demonstrating viability across data collection, model development, and validation. At the moment of this report's publication (November 2025), the evidence reveals a field at an inflection point: technological capabilities have matured, regulatory frameworks increasingly accommodate federated approaches, and economic pressures make centralized AI's sustainability questionable. Yet significant barriers around liability frameworks, interoperability standards, and organizational readiness will shape whether decentralized AI fulfills its potential or remains confined to narrow applications.

We, the authors, believe that the trajectory through 2030 will follow an evolutionary rather than revolutionary path. Rather than wholesale replacement of centralized systems, healthcare will witness gradual integration of decentralized components where they provide decisive advantages—data sovereignty addressing regulatory requirements, federated learning enabling multi-institutional collaboration without compromising privacy, edge computing supporting real-time clinical decisions, and DAO governance coordinating research consortia. Centralized elements will persist where they prove superior for quality assurance, regulatory compliance, or operational efficiency. This hybrid reality requires organizations to develop sophisticated judgment about which functions genuinely benefit from decentralization versus which improve through selective centralization—a capability built through experimentation rather than theoretical analysis.

The outcome depends not on technological determinism but on strategic decisions that diverse stakeholders make individually and collectively over the next 3-5 years. Healthcare delivery organizations that experiment strategically while building governance capabilities position themselves to shape emerging standards rather than inheriting externally imposed frameworks. Patient organizations that engage early in governance design can ensure authentic voice rather than tokenistic consultation as decentralized systems proliferate. Technology developers that prioritize interoperability and regulatory compliance from inception will thrive as vendor lock-in strategies lose effectiveness. Research institutions that establish collaborative frameworks now will influence how academic values around openness and equity manifest in operational systems. Policymakers that create adaptive regulatory frameworks balancing innovation with patient protection will enable beneficial deployments while constraining potential harms.

Decentralized AI offers healthcare a path toward addressing longstanding challenges—enabling collaborative research while preserving privacy, empowering patients while maintaining safety, democratizing access while ensuring quality, and respecting institutional sovereignty while capturing collective intelligence. Whether this potential materializes or whether coordination complexity, regulatory barriers, and organizational inertia constrain progress to incremental improvements remains uncertain. What proves certain is that the decisions made in 2025-2027—which standards to adopt, which governance frameworks to establish, which partnerships to form, which capabilities to build—will shape healthcare AI's evolution through 2030 and beyond. The field's future is not predetermined but rather will emerge from the choices that stakeholders across the healthcare ecosystem make as they navigate the opportunities and constraints this report has examined.

REFERENCES

CHAPTER 1: SETTING THE STAGE

- Boniol, M., Kunjumen, T., Nair, T.S., Siyam, A., Campbell, J. and Diallo, K. (2022). The global health workforce stock and distribution in 2020 and 2030: a threat to equity and 'universal' health coverage? *BMJ Global Health*, 7(6), e009316. <https://doi.org/10.1136/bmjgh-2022-009316>
- Chateau Health & Wellness (2025). Addressing Healthcare Worker Burnout in 2025. Available at: <https://www.chateaufrecovery.com/addressing-healthcare-worker-burnout-2025>
- Crypto Valley Association (2023). Next Step: Sustainable AI - RESTART Framework. Available at: https://d1c2gz5q23tkk0.cloudfront.net/assets/uploads/3594882/asset/Sustainability_WG_AI_White_paper_v5.0.pdf?1690344196
- Pardo Martin, C., Lamb, J., Dahab, A., Jones, J. and Bhasker, S. (2025). Generative AI in healthcare: Current trends and future outlook. McKinsey & Company, 26 March. Available at: <https://www.mckinsey.com/industries/healthcare/our-insights/generative-ai-in-healthcare-current-trends-and-future-outlook>

CHAPTER 2: FOUNDATIONAL FRAMEWORK FOR DECENTRALIZATION IN HEALTHCARE

- Teo, J.T.H., Lim, J.L.W., Lee, J.Y., et al. (2024). Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture. *Cell Reports Medicine*, 5, 101419. <https://doi.org/10.1016/j.xcrm.2024.101419>
- Additional resource on federated learning architectures. Available at: <https://www.mdpi.com/2079-9292/13/24/4926>

CHAPTER 3: DECENTRALIZED DATA COLLECTION AND PROCESSING

MELLODDY CONSORTIUM

- Innovative Health Initiative (2025). MELLODDY Project Factsheet. Available at: <https://www.ih.europa.eu/projects-results/project-factsheets/melloddy>
- Kubermatic (2024). MELLODDY: Turning pharma competition into 'coopetition'. Available at: <https://www.kubermatic.com/blog/melloddy-turning-pharma-competition-into-coopetition/>
- Oldenhof, M., et al. (2024). Industry-scale orchestrated federated learning for drug discovery. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(13), 15576-15584. <https://doi.org/10.1609/aaai.v37i13.26847>
- Owkin (2024). MELLODDY project hits its year one goal: First secure platform for federated learning in drug discovery. Available at: <https://www.owkin.com/newsfeed/melloddy-project-meets-its-year-one-objective-deployment-of-the-worlds-first-secure-platform-for-multi-task-federated-learning-in-drug-discovery>

MAYO CLINIC PLATFORM

- Mayo Clinic Platform (2023). Mayo Clinic Platform Keeps Data Behind Glass® [White Paper]. Available at: https://www.marco.health/content/files/2023/05/Platform_DBG-White Paper.pdf
- Mayo Clinic Platform (2025). Discover Platform. Available at: <https://www.mayoclinicplatform.org/discover/>
- NCBI (2025). Mayo-Google partnership: Sharing health data. Available at: <https://www.ncbi.nlm.nih.gov/books/NBK594445/>

SEOVE - CADEIA DE CUIDADOS

- Brazil (2018). Lei Geral de Proteção de Dados Pessoais (LGPD), Lei nº 13.709. Available at: https://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/l13709.htm
- Interview with Dr. Andreia de Bem Machado, Postdoctoral Researcher, Federal University of Santa Catarina, SEOVE's Treasury Board Member and Financial Delegated Responsible. Conducted for Decentralized AI in Healthcare Report.
- Interviews with Andréia Carla Tonin, SEOVE Technical Manager and clinical staff. Conducted for Decentralized AI in Healthcare Report.
- SEOVE (2025). Sociedade Espírita Obreiros da Vida Eterna. Available at: <https://seove.org.br/>

ACOER CLINICAL TRIALS PLATFORM

- Acoer (2024). Federated learning and cryptographic data mesh: Technical explanation of privacy-preserving AI implementation. Acoer Blog, July. Available at: https://www.acoer.com/news/blog/federated-learning_
- Acoer (2025). Acoer Official Website: Company overview, products, and partnership information. Available at: <https://www.acoer.com>
- Acoer (2025). Cancer Trials Explorer: Clinical trials data aggregation and visualization platform. Available at: <https://cancertrials.acoer.com/>
- Acoer (2025). Health Data Explorer Platform: Medicolegal death investigation and public health reporting. Available at: <https://www.healthdataexplorer.io/>
- Hedera (2021). Acoer's RightsHash™ builds on Hedera to pioneer decentralized management and protection of user's rights. Hedera Blog, July. Available at: <https://hedera.com/blog/acoers-rightshash-builds-on-hedera>

PROJECT MONAI (BONUS RESOURCE)

- MONAI Consortium (2025). MONAI Deploy Documentation. Available at: <https://docs.monai.io/projects/monai-deploy/en/stable/>
- NVIDIA (2025a). MONAI: Medical Open Network for AI. Available at: <https://monai.io/>
- NVIDIA (2025b). NVIDIA FLARE Documentation. Available at: <https://nvflare.readthedocs.io/en/main/>
- Patel, S. et al. (2022). The federated tumor segmentation (FeTS) tool. Frontiers in Neuroinformatics. Available at: <https://www.frontiersin.org/journals/neuroinformatics>

KNOWS PLATFORM (BONUS RESOURCE)

- Interview with Steven Yuwen, Founder of KnowS. Interviewed by LeiLei Tang for Decentralized AI in Healthcare Report.
- MedKnowS (2025). MedKnowS Official Website. Available at: <https://www.medknows.com/>

CHAPTER 4: DECENTRALIZED MODEL DEVELOPMENT AND TRAINING

DEEPSEEK R1/VL

- Gradient Flow (2025). DeepSeek in Action: Practical AI Applications Transforming Chinese Healthcare, 7 March. Available at: <https://gradientflow.com/deepseek-in-action-practical-ai-applications-transforming-chinese-healthcare/>
- Healthcare IT News (2025). Chinese health players begin integrating DeepSeek. Available at: <https://www.healthcareitnews.com/news/asia/chinese-health-players-begin-integrating-deepseek>
- PandaYoo Analysis (2025). DeepSeek Diagnosis: Is AI Finally Revolutionizing Healthcare in China? Available at: <https://pandayoo.com/post/deepseek-diagnosis-is-ai-finally-revolutionizing-healthcare-in-china/>
- Tordjman, M., Liu, Z., Yuce, M., Forghani, R., Tang, A., Kadoury, S. and Duron, T. (2025). Comparative benchmarking of the DeepSeek large language model on medical tasks and clinical reasoning. Nature Medicine. <https://doi.org/10.1038/s41591-025-03726-3>

AKASH NETWORK AND AETHIR

- Aethir (2025). Aethir Official Website: Enterprise cloud infrastructure documentation. Available at: <https://www.aethir.com>
- Akash Network (2025). Akash Network Official Website: Decentralized cloud computing marketplace. Available at: <https://akash.network/>
- Akash Network (2025). Akash Network and Solve.Care bring true patient data ownership to the healthcare industry. Akash Network Blog. Available at: <https://akash.network/blog/akash-network-and-solve-care-bring-true-patient-data-ownership-to-the-healthcare-industry/>

BITTENSOR

- Bittensor (2025). Bittensor Official Website: Decentralized machine learning network. Available at: <https://bittensor.com/>
- Bittensor (2025). Bittensor Documentation: Technical architecture and subnet protocols. Available at: <https://docs.bittensor.com/>

ASI ALLIANCE

- ASI Alliance (2025). ASI Alliance Official Website. Available at: <https://asi.xyz/>
- ASI Alliance (2025). Healthcare breakthrough announcement, 7 January. Twitter/X. Available at: https://x.com/ASI_Alliance/status/1915439728658198904
- CoinBureau (2025). ASI Alliance review: Comprehensive alliance overview. Available at: <https://coinbureau.com/review/asi-alliance-review/>

- Fetch.ai (2025). ASI Alliance revolutionizes AI earning with ASI Train: Unlock the power of DeSci models. Available at: <https://fetch.ai/blog/asi-alliance-revolutionizes-ai-earning-with-asi-train-unlock-the-power-of-de-sci-models>
- Fetch.ai (2025). Press & Media Kit: Overview of autonomous agents in healthcare systems. Available at: <https://fetch.ai/press-media>

CHAPTER 5: DECENTRALIZED VALIDATION AND DEPLOYMENT

VITADAO

- Gero.ai (2025). Gero AI: Longevity research and biomarker development. Available at: <https://gero.ai/>
- Molecule (2025). IP-NFT Framework: Intellectual property tokenization for research. Available at: <https://www.molecule.xyz/>
- Vectra Advisors (2025). Understanding DAOs and legal wrappers in Switzerland. Available at: <https://vectra-advisors.com/understanding-daos-and-legal-wrappers-in-switzerland/>
- VitaDAO (2025). VitaDAO Official Website. Available at: <https://www.vitadao.com/>
- VitaDAO (2025). AI integration strategy presentation. YouTube. Available at: <https://www.youtube.com/watch?v=ms2dLI2-myw&t=40s>
- VitaDAO Labs (2025). Whitepaper: How the in-house Lab operates, including fellowships and roadmap. Available at: <https://drive.google.com/file/d/1zdzqoybZgxLYkiU2RTFuiMqKcHw7AoPP/view>

CHANGI GENERAL HOSPITAL AIMSG PLATFORM

- Changi General Hospital (2024). About Changi General Hospital. Singapore: SingHealth. Available at: <https://www.cgh.com.sg/about-us>
- Lunit Inc. (2023). Lunit INSIGHT CXR: AI for Chest X-Ray Analysis. Seoul: Lunit. Available at: <https://www.lunit.io/en/products/cxr>
- Medical Imaging Singapore (AimSG) (2023). About AI Medical Imaging Platform for Singapore public healthcare (AIMSG). Singapore: Synapse. Available at: <https://www.synapse.sg/healthtech/health-ai/ai-medical-imaging-platform>
- Ministry of Health Singapore (2024). Transforming Healthcare Through Technology. Singapore: MOH Press Release, 10 October. Available at: <https://www.moh.gov.sg/newsroom/transforming-healthcare-through-technology/>
- Personal Data Protection Commission Singapore (2024). Personal Data Protection Act (PDPA). Singapore: PDPC. Available at: <https://www.pdpc.gov.sg/overview-of-pdpa/the-legislation/personal-data-protection-act>
- Synapse (2023). About Synapse: Singapore's Health Tech Agency. Singapore: Synapse. Available at: <https://www.synapse.sg/>

CHAPTER 6: FUTURE DIRECTIONS AND STRATEGIC RECOMMENDATIONS

- Abdalla, H.B., Kumar, Y., Tang, L., Marchena, J., Guzman, S., Cheraghy, M., Awlla, A. and Gheisari, M. (2025). The future of artificial intelligence in the face of data scarcity. *Computers, Materials & Continua*, 84(1). <https://doi.org/10.32604/cmc.2025.063551>
- China (2021). Personal Information Protection Law (PIPL). Standing Committee of the National People's Congress. Available at: https://www.gov.cn/xinwen/202108/20/content_5632486.htm (Chinese); English translation available at: <https://digichina.stanford.edu/work/translation-personal-information-protection-law-of-the-peoples-republic-of-china-effective-nov-1-2021/>
- ETH Zurich (2025). Apertus: A language model built for the public good. *ETH News*, July. Available at: <https://ethz.ch/en/news-and-events/eth-news/news/2025/07/a-language-model-built-for-the-public-good.html>
- European Commission (2024). Artificial Intelligence Act. Official Journal of the European Union. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- European Commission (2024). Regulation on the European Health Data Space (EHDS). Official Journal of the European Union. Available at: https://health.ec.europa.eu/ehealth/digital-health-and-care/european-health-data-space_en
- European Union (2016). General Data Protection Regulation (GDPR), Regulation (EU) 2016/679. Official Journal of the European Union. Available at: <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
- Noor, P. (2025). There's a stunning financial problem with AI data centers. *Futurism*, 28 August. Available at: <https://futurism.com/data-centers-financial-bubble>
- Shojaee, P., Jain, A., Tipirneni, S. and Reddy, C.K. (2025). The illusion of thinking: Understanding the strengths and limitations of reasoning models via the lens of problem complexity. arXiv preprint arXiv:2506.06941. Available at: <https://arxiv.org/abs/2506.06941>
- U.S. Department of Health and Human Services (2016). 21st Century Cures Act, Public Law 114-255. Available at: <https://www.congress.gov/bill/114th-congress/house-bill/34>
- U.S. Food and Drug Administration (2024). Guidance for Industry and FDA Staff: Software as a Medical Device (SaMD): Clinical Evaluation. Available at: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/software-medical-device-samd-clinical-evaluation>

THANK YOU!

THE AUTHORS

- Antonio Pesqueira
- Carmen Cucul
- LeiLei Tang

- Natalia Sofia
- Stephanie Fuchs
- Thomas Egelhof

CONTACT US



MAIL

info@etheroshealthdata.org
info@cryptovalley.swiss



WEBSITE

<https://etheroshealthdata.org>
<https://cryptovalley.swiss>