Etheros HealthData Foundation

Crypto Valley

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:

Etheros HealthData Foundation

Crypto Valley Association (Sustainability Working Group)

TABLE OF CONTENT

EXECUTIVE SUMMARY

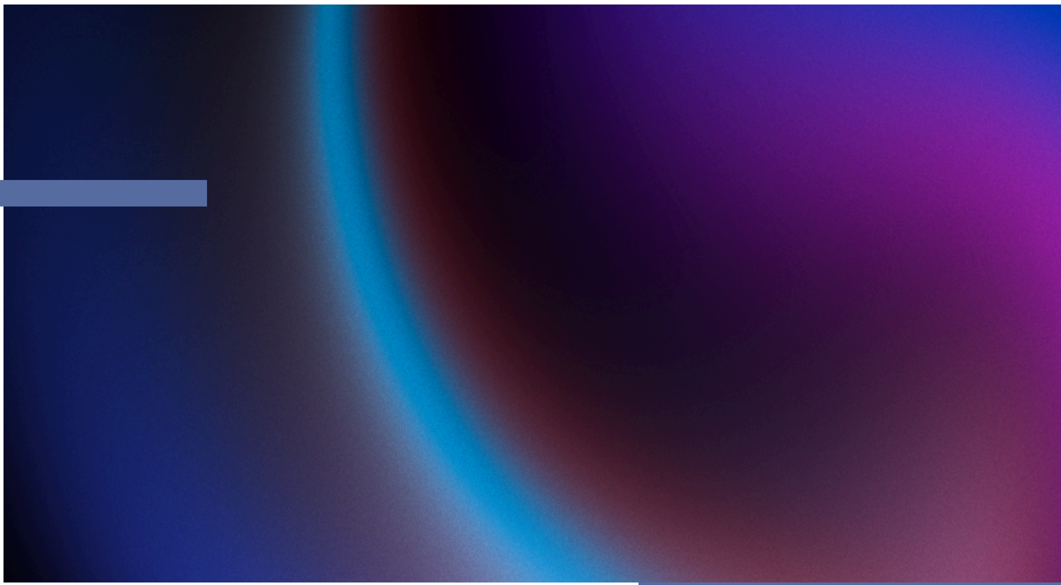

OVERVIEW

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	02	03	04
TECHNOLOGICAL MATURATION	REGULATORY MANDATES	ECONOMIC IMPERATIVES	SOCIETAL SHIFTS

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:

01

DECENTRALIZED DATA MANAGEMENT

02

DECENTRALIZED COMPUTE & CONTROL

03

DECENTRALIZED GOVERNANCE

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

Rresent 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:

Etheros 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 Etheros 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.

REQUEST THE FULL REPORT AT:
HTTPS://ETHEROSHEALTHDATA.ORG/RESEARCH
HTTPS://CRYPTOVALLEY.SWISS/