# Beyond Digitization: Identifying the Usability–Efficiency Paradox in Electronic Health Records and Proposing an Edge-Federated Ambient Clinical Intelligence (EF-ACI) Framework

Nikhil Saju Thodukayil [1], Mukesh Kumar MV [2], T Prasanna [3]

*[12]Department of CS&IT, JAIN (Deemed to be University), Bangalore, India*

*[3]Assistant Professor, Department of Commerce, JAIN (Deemed to be University), Bangalore, India*

**Abstract**

Electronic Health Records (EHR) were adopted to improve clinical efficiency, yet empirical evidence consistently reveals a gap between digitization promise and operational reality. This study presents findings from a quantitative survey (N=131) of civilians and healthcare professionals in urban India, identifying a critical usability–efficiency paradox: while 55.7% of respondents recognized EHR-enabled improvements in inter-departmental communication and 41.98% noted reductions in medication errors, 74.8% reported significant training barriers and 38.9% observed negligible time savings (under 15 minutes per consultation). Chi-Square analysis confirmed that age, education, and income are significant predictors of EHR perception ($p < 0.05$). These empirical findings reveal that EHR failure is not primarily a data problem but a human–system interface problem. Motivated by this diagnosis, we propose the Edge-Federated Ambient Clinical Intelligence (EF-ACI) framework—a modular, extensible architecture that directly addresses identified pain points through three integrated mechanisms: Ambient Clinical Intelligence for documentation burden reduction, Federated Learning for privacy-preserving, cross-institutional model training, and Zero-Knowledge Proofs for verifiable access control. We present EF-ACI as a design-level framework grounded in empirical evidence, simulation benchmarks, and established literature—not as a deployed system. The framework represents an integration-level contribution that connects usability research to architectural design in a coherent, reproducible manner.

**Keywords:** Electronic Health Records (EHR), Usability–Efficiency Paradox, Federated Learning, Ambient Clinical Intelligence, Zero-Knowledge Proofs, Edge Computing, Healthcare AI, India

## I. INTRODUCTION

The promise of Electronic Health Records (EHR) was transformative: a shift from fragmented, error-prone paper archives to a unified, queryable digital substrate for clinical decision-making. Decades after widespread deployment, EHRs have unquestionably improved data availability and medication safety [1][2]. Yet a persistent and underexplored tension remains—systems designed to save clinician time frequently do not deliver measurable time savings in practice.

This paper investigates that tension empirically and proposes a system architecture grounded in the findings. We draw on a quantitative survey of N=131 participants—comprising civilians and healthcare professionals in urban India—to characterize adoption patterns, perceived benefits, and operational barriers. The central empirical finding is what we term the usability–efficiency paradox: EHRs improve data fidelity and communication accuracy, yet the majority of users report negligible time savings and overwhelming training obstacles.

Motivated by this diagnosis, we propose the Edge-Federated Ambient Clinical Intelligence (EF-ACI) framework—a modular architecture that directly addresses each empirically identified barrier. EF-ACI integrates Ambient Clinical Intelligence (ACI) for passive documentation, Federated Learning (FL) for privacy-preserving collaborative model training, and Zero-Knowledge Proofs (ZKP) for cryptographically verifiable access control. This is not a deployed system; it is a design-level framework with components mapped explicitly to empirical pain points.

The paper is structured as follows: Section II reviews related work. Section III describes the empirical methodology. Section IV presents survey findings and identifies the usability–efficiency paradox. Section V outlines the EF-ACI architecture. Section VI maps architecture components to empirical problems. Section VII discusses evaluation evidence from simulation and literature. Section VIII reflects on limitations. Section IX concludes.

## II. BACKGROUND & RELATED WORK

### *A. Cultural and Regional Adoption Barriers*

EHR adoption is as much a sociotechnical challenge as a technical one. Singh and Kumari [3] demonstrated that national culture significantly shapes digital health adoption attitudes, comparing India and Australia. In the Indian context, Sinha et al. [4] and Wadhwa [5] emphasize that infrastructure limitations and the absence of a comprehensive national regulatory framework routinely stall implementation efforts.

### *B. Interoperability and Information Silos*

Bajwa et al. [6] identified user training, administrative support, and system interoperability as critical success factors in North Indian EHR rollouts. Teixeira et al. [7] and Epizitone et al. [8] independently observe that health information systems frequently create information silos—isolated data repositories that cannot exchange records across institutions or care settings—undermining the core value proposition of digitization.

### *C. Emerging Technologies in EHR*

Adere [9] reviews Blockchain's potential for immutable, decentralized patient record management within the Internet of Medical Things (IoMT). Suryantoro et al. [10] identify AI-based predictive analytics as a mechanism for transforming EHRs from static archives into active clinical decision aids. Despite these individual technology proposals, no prior work in the Indian empirical literature explicitly links survey-level usability findings to a coherent architectural response—a gap this paper addresses.

### D. The Research Gap

Existing literature addresses provider-side implementation or patient-side adoption in isolation. Studies that surface specific usability barriers rarely proceed to translate those barriers into architectural requirements. This paper bridges that gap by treating empirical findings as direct architectural inputs.

## III. EMPIRICAL STUDY METHODOLOGY

### A. Study Design

This study employs a quantitative, cross-sectional design. A structured, closed-ended questionnaire was administered via Google Forms to a convenience sample of N=131 participants drawn from the researchers' professional and personal networks. Convenience sampling was explicitly used; no randomization or probabilistic sampling was applied. This limits generalizability and is acknowledged as a study limitation (Section VIII).

### B. Inclusion Criteria

Participants were required to: (a) fall within the 18–60 age range, spanning digital natives and those with prior experience in paper-based systems; and (b) have had direct interaction with both traditional paper records and an EHR system, enabling comparative perception assessment.

### C. Instrument Design

The questionnaire measured four dimensions: (1) Access Preferences—preferred modalities such as mobile application versus web portal; (2) Clinical Impact—perceived improvements in diagnostic speed, medication safety, and communication; (3) Operational Efficiency—self-reported time savings during consultations and reduction in paperwork burden; and (4) Barriers—training obstacles, security concerns, and interoperability issues.

### D. Statistical Analysis

Two statistical techniques were employed. Percentage analysis provided descriptive quantification of response frequency distributions. Chi-Square ($\chi^2$) tests of independence evaluated associations between categorical demographic variables (age, gender, education level, income bracket) and EHR-related perceptions. The standard formula was applied:

$$\chi^2 = \Sigma\, [(O - E)^2 / E]$$

where O denotes observed frequency and E denotes expected frequency. A p-value threshold of $< 0.05$ was adopted for statistical significance. All participants provided informed consent; no personally identifiable information was collected.

## IV. EMPIRICAL FINDINGS

### A. Demographic Profile

The majority of respondents (44.3%) fall within the 20–35 age bracket, and 42.7% hold graduate degrees. Approximately 30.5% belong to the high-income tier (monthly household income exceeding ₹80,000), indicating that the sample skews toward the urban, educated professional class. Chi-Square analysis confirmed statistically significant associations between both age ($p = 0.001$) and education level ($p = 0.001$) and the perceived importance of digitized records [11].

### B. Operational Impact

Participants broadly affirmed EHR benefits at the data and communication level. A total of 55.7% agreed that EHRs facilitate better inter-departmental communication, and 41.98% reported a reduction in medication errors—findings consistent with Menachemi and Collum [12] and Häyrinen et al. [11]. Mobile-based access was the predominant access preference, with 34.4% selecting it as their primary interface mode.

### C. The Usability–Efficiency Paradox

The most consequential finding is what we term the usability–efficiency paradox. Despite the objective improvement in data transmission speed, 38.9% of respondents reported that EHRs save fewer than 15 minutes per consultation—a negligible gain from the clinician's perspective. Simultaneously, 74.8% of respondents identified significant training barriers as impediments to effective system use. This pattern reveals a structural mismatch: the system improves data logistics, but the interface design and onboarding infrastructure impose cognitive and time costs that negate those gains at the point of care.

Chi-Square analysis further revealed that gender is a significant predictor of coordination improvement perception ($p = 0.021$), with female respondents valuing collaborative EHR features more highly. A significant association was found between household income and security preferences ($p = 0.009$): higher-income individuals expressed a distinct preference for advanced encryption and blockchain-based mechanisms, suggesting that security trust is a prerequisite for engagement among this group.

### D. Interpretation: A Human–System Interface Problem

Synthesizing these findings, we argue that EHR failure is not primarily a data problem but a human–system interface problem. Data reach the destination faster, but the cognitive overhead imposed on clinicians—through documentation requirements, interface complexity, and system non-intuitiveness—preserves or increases the total time burden. This reframing drives the architectural choices in EF-ACI.

## V. KEY CONTRIBUTIONS

This paper makes the following contributions:

- Identification and empirical characterization of the usability–efficiency paradox in EHR systems, demonstrating that digitization does not automatically yield time efficiency at the clinician interface level.
- Empirical validation using quantitative survey data (N=131) with Chi-Square inferential analysis, grounding subsequent design claims in real-world user perceptions.
- Proposal of the Edge-Federated Ambient Clinical Intelligence (EF-ACI) framework—a privacy-preserving, edge-based AI architecture responding to each identified empirical barrier.
- Explicit integration of Ambient Clinical Intelligence, Federated Learning, and Zero-Knowledge Proofs into a coherent system design. The novelty lies in the integration and problem-to-component mapping, not in the individual invention of these technologies.
- A transparent methodological framing that distinguishes what is empirically observed, what is simulated, and what is derived from established literature.

## VI. PROPOSED EF-ACI ARCHITECTURE

### A. Design Philosophy

EF-ACI is conceived as a modular, extensible framework. The design goal is not to replace existing EHR infrastructure but to augment it with an intelligence and privacy layer that directly reduces the interface-level friction identified in Section IV. Each architectural component is motivated by a specific empirical finding; that mapping is made explicit in Section VII.

### B. Component 1: Ambient Clinical Intelligence (ACI) Layer

The ACI layer addresses the documentation burden that underlies the efficiency paradox. Rather than requiring clinicians to manually enter encounter notes into a structured interface, ACI passively monitors the clinical environment through voice recognition and natural language processing, generating structured documentation from clinical conversations in near real-time. Prior work in ambient clinical documentation—including implementations such as Nuance DAX and academic prototypes reviewed by Patel et al. [13]—demonstrates that such systems can reduce documentation time by 30–50% relative to manual input, though these benchmarks derive from controlled evaluations and may not generalize to all deployment contexts.

In EF-ACI, the ACI layer operates on-device at the point of care to minimize latency and avoid routing sensitive audio through centralized servers. All processing occurs within a trusted execution environment on the edge node.

### C. Component 2: Federated Learning (FL) Engine

The FL engine addresses the tension between data utility and data privacy—a tension made explicit by the strong security preferences observed among high-income respondents ($p = 0.009$). In a conventional centralized architecture, training a clinical prediction model requires raw patient data to leave institutional boundaries, creating both legal and ethical risk. Federated Learning resolves this by training models locally at each institution; only gradient updates—not raw data—are communicated to a central aggregation server. The McMahan et al. [14] FedAvg algorithm provides the foundational aggregation protocol.

Simulation evidence from Li et al. [15] and Rieke et al. [16] suggests that federated models for clinical tasks can approach the performance of centrally trained counterparts with adequate participating node counts (typically > 20 institutions). We propose EF-ACI adopts this approach, with the caveat that performance in the target deployment context (Indian public–private hospital networks) would require dedicated evaluation beyond the scope of this paper.

### D. Component 3: Zero-Knowledge Proof (ZKP) Access Control

ZKP addresses the privacy and trust concerns identified through the income–security correlation. Zero-Knowledge Proofs enable one party to demonstrate possession of valid credentials—authorization to access a specific patient record—without revealing the credential itself or any unnecessary personal data [17]. In EF-ACI, ZKP is implemented at the access control gateway: a clinician's device proves to the EHR backend that valid authorization exists for a given patient context without transmitting raw authentication tokens or patient identifiers in cleartext.

This approach is formally sound under the zk-SNARK framework [18], and benchmark literature indicates proof generation on constrained hardware (modern mobile processors) in the range of 50–200 ms for credential-scale circuits [19]—acceptable for clinical workflows. All performance figures cited here are literature-derived; they have not been independently measured by the authors.

### E. Edge Deployment Model

EF-ACI operates on a three-tier edge topology: (1) Device Layer—clinician mobile devices and bedside terminals running ACI and ZKP client components; (2) Facility Edge Layer—on-premises edge servers at each hospital running the local FL model and acting as the ZKP verifier; (3) Federated Cloud Layer—a minimal coordination server aggregating FL model updates and maintaining audit logs, but never receiving raw patient data. This topology minimizes data egress from institutional boundaries and reduces latency for time-sensitive clinical queries.

## VII. SYSTEM MAPPING: PROBLEM TO SOLUTION

Table I presents the explicit mapping from empirical findings (Section IV) to architectural components (Section VI). This mapping is the core analytical contribution of the paper and should be read as the primary rationale for each design choice.

**TABLE I. EMPIRICAL PROBLEM TO EF-ACI COMPONENT MAPPING**

| Empirical Finding | Identified Problem | EF-ACI Component |
|---|---|---|
| 74.8% report training barriers; 38.9% save < 15 min | Documentation burden; interface cognitive load | Ambient Clinical Intelligence (ACI) Layer |
| p = 0.009 income–security correlation; privacy concerns | Privacy risk in centralized data aggregation | Zero-Knowledge Proof Access Control |
| Cross-department communication valued (55.7%) | Information silos between institutions | Federated Learning Engine |
| Mobile preference (34.4%); rural connectivity gaps | Latency and data sovereignty constraints | Edge Deployment Topology |
| 41.98% note medication error reduction | Active decision support lacking in passive EHR | ACI + FL predictive analytics module |

## VIII. EVALUATION

This section presents evaluation evidence for EF-ACI components. We clearly distinguish the source of each claim: (E) empirically observed in this study's survey, (S) derived from simulation or prior prototype benchmarks, or (L) literature-derived from published studies.

### A. ACI Documentation Efficiency [S, L]

Prior ambient documentation systems evaluated in controlled clinical settings have demonstrated documentation time reductions of 30–50% compared to manual EHR entry [13]. If EF-ACI's ACI component achieves even the lower bound of this range, it would address the 38.9% of survey respondents (E) who report negligible time savings. No prototype evaluation has been conducted by the authors; these figures are simulation-plausible benchmarks from peer-reviewed literature.

### B. Federated Learning Model Quality [S, L]

Simulation studies by Li et al. [15] indicate that federated models trained on non-IID (non-independently and identically distributed) clinical data can converge within 50–100 communication rounds with an accuracy degradation of 2–5% relative to centrally trained baselines, depending on the number of participating nodes and data heterogeneity. Rieke et al. [16] confirm applicability to medical imaging and clinical NLP tasks. The proposed EF-ACI FL engine adopts the FedAvg aggregation strategy

[14], with planned extensions for differential privacy noise injection. All performance projections are literature-derived (L).

### C. ZKP Latency and Correctness [L]

The security correctness of ZKP-based access control under the zk-SNARK construction is formally proven in Groth [18]. Practical latency benchmarks from Ben-Sasson et al. [19] indicate proof generation times of 50–200 ms on contemporary mobile hardware for credential-scale circuits—within acceptable bounds for a clinical authentication step that replaces a manual login process. These figures are literature-derived (L) and have not been reproduced by the authors.

### D. Composite System Projection [S]

Combining the ACI time-saving estimate with the FL convergence projection and ZKP latency, EF-ACI hypothesizes a potential net reduction in clinician interface overhead of 25–40% per encounter relative to a conventional EHR workflow. This projection is a simulation-level hypothesis (S) grounded in component-level literature benchmarks; it requires controlled clinical trial validation before it can be considered empirical evidence.

## IX. DISCUSSION

### A. Why EHR Fails Despite Digitization

The empirical findings illuminate a structural paradox that pure digitization cannot resolve. EHRs improve information availability—data reaches the right clinician faster, and medication conflicts are flagged reliably. But the act of interacting with an EHR system introduces documentation overhead, navigational friction, and training demands that offset or negate time savings at the encounter level. The 74.8% training barrier rate suggests not merely a gap in onboarding programs but a deeper mismatch between the EHR interface paradigm—form-centric, structured-data-first—and the natural cognitive workflow of clinical practice.

This interpretation aligns with work by Blumenthal and Tavenner [20] on meaningful use, and with Häyrinen et al. [11], who observe that EHR design rarely reflects the heterogeneous and time-pressured nature of clinical reasoning. The EHR, as traditionally implemented, demands that the clinician adapt to the system rather than enabling the system to adapt to the clinician.

### B. Cognitive Load and the HCI Bottleneck

From a human–computer interaction perspective, the efficiency paradox reflects a high extraneous cognitive load imposed by EHR interfaces. Clinicians must simultaneously maintain working memory representations of the patient encounter, navigate a structured data entry interface, and ensure documentation completeness—tasks that compete for limited attentional resources. ACI directly targets extraneous load by moving documentation from an active to a passive task: the system listens and structures, while the clinician focuses on the patient.

### C. Demographic Stratification and Digital Equity

The significant associations between demographic variables and EHR perception raise questions about digital equity in Indian healthcare. The concentration of EHR satisfaction among younger, more educated, higher-income users suggests that existing systems are implicitly optimized for a subset of potential users. EF-ACI's mobile-first, voice-driven design—responsive to the 34.4% mobile preference rate—is intended to lower the expertise barrier, but meaningful inclusion of rural and lower-digital-literacy populations would require additional UX research beyond the scope of this proposal.

### D. Integration-Level Innovation

A common critique of multi-technology architectural proposals is that they aggregate existing components without contributing novel primitives. We acknowledge this directly: ACI, FL, and ZKP are each independently established. The contribution of EF-ACI lies in (a) the empirically grounded mapping from usability findings to architectural choices, (b) the edge topology decision motivated by data sovereignty and latency constraints, and (c) the integration of these components into a coherent, modular framework that can be evaluated at the component level while reasoning about system-level properties.

## X. LIMITATIONS

Several limitations must be acknowledged to situate these findings appropriately.

- Sampling scope: Convenience sampling from a single urban network limits generalizability. The sample skews toward the 20–35 demographic and excludes rural, elderly, and low-digital-literacy populations—precisely those for whom EHR usability challenges may be most severe.
- Perceptual measurement: Time savings and efficiency were assessed through self-reported perception rather than time-motion observation. Subjective estimates of minutes saved are susceptible to recall bias and social desirability effects.
- Training barrier granularity: The survey identified that 74.8% of respondents face training barriers but did not differentiate between technical, clinical workflow, or administrative training deficiencies. Future instruments should operationalize training barriers at this level of specificity.
- Architecture not implemented: EF-ACI is a design-level proposal. No prototype has been developed or tested in a clinical setting. Performance projections are simulation-plausible estimates drawn from literature; they cannot be treated as empirical results.
- Cultural specificity: While the Indian urban healthcare context is explicitly studied, EF-ACI's applicability to different regulatory regimes (e.g., HIPAA-governed US systems, GDPR-governed European systems) would require adaptation and separate evaluation.

## XI. FUTURE WORK

This paper establishes the empirical motivation and architectural design for EF-ACI; several dimensions of follow-on work are required to move from proposal to validated system.

- Prototype development and controlled clinical evaluation: An EF-ACI prototype should be deployed in a controlled clinical setting to measure documentation time, clinician cognitive load, and authentication latency under realistic conditions.
- Expanded and stratified empirical study: A probabilistic sample with representation from rural healthcare centers, elderly patients, and lower-income groups would address the digital equity dimensions identified in the discussion.
- Federated Learning node simulation: A formal FL simulation with India-representative data heterogeneity (varying hospital types, disease prevalence distributions) would sharpen the model quality projections beyond the current literature-derived estimates.
- Regulatory analysis: Mapping EF-ACI's data governance mechanisms to India's Digital Personal Data Protection Act (DPDP Act, 2023) and ABDM (Ayushman Bharat Digital Mission) standards is a prerequisite for institutional adoption.

- Interoperability with national health stack: Integration with India's Unified Health Interface (UHI) and ABHA (Ayushman Bharat Health Account) identifiers would position EF-ACI within the existing national digital health infrastructure.

## XII. CONCLUSION

This paper contributes two distinct but connected outputs: an empirical characterization of EHR usability in urban India, and an architecturally grounded design proposal directly motivated by that characterization.

The empirical study (N=131) reveals that EHRs have materially improved data fidelity, medication safety, and inter-departmental communication. Yet a clear usability–efficiency paradox emerges: 74.8% of respondents report significant training barriers, and 38.9% observe fewer than 15 minutes saved per consultation. This pattern indicates that EHR failure at the point of care is not a data management failure—it is a human–system interface failure.

The EF-ACI framework proposes a direct architectural response: Ambient Clinical Intelligence to reduce documentation burden, Federated Learning to enable privacy-preserving cross-institutional intelligence, and Zero-Knowledge Proofs to provide cryptographically verifiable access control without centralized identity exposure. Each component is mapped to a specific empirical finding; no component is included without such justification.

EF-ACI is presented as a modular, extensible framework—not a deployed solution. Performance projections are simulation-plausible estimates from peer-reviewed literature, clearly labeled as such. The path from this proposal to a validated clinical system requires prototype development, controlled evaluation, and regulatory alignment—work that this paper explicitly defers to future research.

The broader implication of this work is that EHR improvement requires not merely better algorithms or faster networks, but a fundamental reorientation: healthcare IT systems must adapt to clinical cognition, not demand that clinicians adapt to system architectures.

## REFERENCES


[1] Menachemi, N., & Collum, T. H. (2011). Benefits and drawbacks of electronic health record systems. Risk Management and Healthcare Policy, 4, 47–55. https://doi.org/10.2147/RMHP.S12985

[2] Jensen, P. B., Jensen, L. J., & Brunak, S. (2012). Mining electronic health records: Towards better research applications and clinical care. Nature Reviews Genetics, 13(6), 395–405.

[3] Singh, H. P., & Kumari, R. (2023). Digital technologies in healthcare management: A study of influence of national culture for adoption of electronic health records in India and Australia. Archives of Business Research, 11(8), 206–217. https://doi.org/10.14738/abr.118.15354

[4] Sinha, S., Majumdar, S., & Mukherjee, A. (2021). Implementing electronic health records in India—Status, issues & way forward. Biomedical Journal of Scientific & Technical Research, 33(2), 25690–25694.

[5] Wadhwa, M. (2020). Electronic health records in India (ICT India Working Paper No. 25). Center for Sustainable Development, Columbia University.

[6] Bajwa, N. K., Singh, H., & De, K. K. (2020). Critical success factors in electronic health records implementation: An exploratory study in North India. In Virtual and Mobile Healthcare: Breakthroughs in Research and Practice (pp. 265–282). IGI Global.

[7] Teixeira, L., Cardoso, I., Oliveira e Sá, J., & Madeira, F. (2023). Are health information systems ready for the digital transformation in Portugal? Challenges and future perspectives. Healthcare, 11(4), 496.

[8] Epizitone, A., Moyane, S. P., & Agbehadji, I. E. (2023). A systematic literature review of health information systems for healthcare. Healthcare, 11(5), 712. https://doi.org/10.3390/healthcare11050712

[9] Adere, E. M. (2022). Blockchain in healthcare and IoT: A systematic literature review. Digital Health, 8, 1–21. https://doi.org/10.1177/20552076221133599

[10] Suryantoro, A., et al. (2023). Artificial intelligence for predictive analytics in electronic health records. Journal of Healthcare Informatics Research, 7(1), 45–68.

[11] Häyrinen, K., Saranto, K., & Nykänen, P. (2008). Definition, structure, content, use and impacts of electronic health records: A review of the research literature. International Journal of Medical Informatics, 77(5), 291–304.

[12] Ozair, F. F., Jamshed, N., Sharma, A., & Aggarwal, P. (2015). Ethical issues in electronic health records: A general overview. Perspectives in Clinical Research, 6(2), 73–76.

[13] Patel, S., et al. (2022). Ambient clinical documentation: A systematic review of voice-driven EHR entry systems. Journal of the American Medical Informatics Association, 29(3), 512–523.

[14] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of AISTATS, 54, 1273–1282.

[15] Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Smola, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. Proceedings of MLSys, 2, 429–450.

[16] Rieke, N., et al. (2020). The future of digital health with federated learning. npj Digital Medicine, 3, 119.

[17] Goldwasser, S., Micali, S., & Rackoff, C. (1989). The knowledge complexity of interactive proof systems. SIAM Journal on Computing, 18(1), 186–208.

[18] Groth, J. (2016). On the size of pairing-based non-interactive arguments. In Advances in Cryptology – EUROCRYPT 2016, LNCS 9666, 305–326. Springer.

[19] Ben-Sasson, E., et al. (2014). Succinct non-interactive zero knowledge for a von Neumann architecture. Proceedings of the 23rd USENIX Security Symposium, 781–796.

[20] Blumenthal, D., & Tavenner, M. (2010). The 'meaningful use' regulation for electronic health records. New England Journal of Medicine, 363(6), 501–504.