Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Resolving Semantic Conflicts in E-Health Workflows Using Decentralized Autonomous Organizations and Zero-Knowledge Proofs

Aleksandr Kormiltsyn
Tallinn University of Technology
Chibuzor Udokwu
IMC University of Applied Sciences Krems
Alex Norta
Tallinn University
Vimal Dwivedi
v.dwivedi@ulster.ac.uk

University of Ulster
Jelena Ljalik
All Systems Go OÜ

Research Article

Keywords: Semantic consensus · decentralized governance · e-health systems · zero-knowledge proofs · semantic interoperability · privacy-preserving verification

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10897876/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

International Journal of Information Security manuscript No.
(will be inserted by the editor)

Resolving Semantic Conflicts in E-Health Workflows Using Decentralized Autonomous Organizations and Zero-Knowledge Proofs

Aleksandr Kormiltsyn · Chibuzor Udokwu · Alex Norta · Vimal Dwivedi · Jelena Ljalik

Received: date / Accepted: date

Abstract There is a need to address semantic conflicts that arise when heterogeneous healthcare organizations represent and interpret shared clinical data in decentralized workflows. To support verifiable decision-making without revealing sensitive information, we introduce a semantic consensus framework grounded in model-theoretic semantics that provides a formal basis for evaluating the validity of clinical and policy statements and for attesting these outcomes through privacy-preserving zero-knowledge proofs (ZKPs). Consensus mechanisms enable agreement on system state, while ontology-based approaches support semantic interoperability. ZKPs enable privacy-preserving verification of computations and protocol-level compliance. However, no existing approach jointly resolves semantic conflicts arising from different medical terminologies and data representations while preserving data privacy in decentralized governance workflows. To address this gap, we propose a semantic framework that separates semantic interpretation, cryptographic verification, and governance acceptance through a layered architecture. Semantic validity conditions are evaluated off-chain over a Decentralized Knowledge Graph (DKG) and attested using zero-knowledge Succinct Non-interactive ARguments of Knowledge (zk-SNARK) proofs without revealing underlying medical data. The governance layer comprises an automated on-chain verifier that cryptographically checks each proof and a Decentralized Autonomous Organization (DAO) that votes on updates to the shared theory when a semantic conflict cannot be resolved automatically. Agreement is reached on semantic outcomes rather than on raw data, enabling privacy-preserving coordination across heterogeneous healthcare systems. We implement the framework in an inter-organizational prescription workflow and evaluate it through circuit-level and end-to-end scenarios, circuit scalability, on-chain cost, and user-facing latency measurements.

Keywords Semantic consensus · decentralized governance · e-health systems · zero-knowledge proofs · semantic interoperability · privacy-preserving verification

A. Kormiltsyn
Department of Software Science, Tallinn University of Technology, Akadeemia tee 15A, 12816 Tallinn, Estonia
E-mail: aleksandr.kormiltson@taltech.ee

C. Udokwu
IMC University of Applied Sciences Krems, Alauntalstrasse 96, 3500 Krems, Austria
E-mail: chibuzor.udokwu@abc-research.at

A. Norta
Tallinn University, Tallinn, Estonia; Department of Informatics, University of Pretoria, South Africa, Dymaxion OÜ, Tallinn, Estonia; and Nortadesyco.xyz E-mail: alex.norta.phd@ieee.org

V. Dwivedi
Ulster University, School of Computing, Engineering & Intelligent Systems, Northern Ireland, UK
E-mail: v.dwivedi@ulster.ac.uk
Corresponding author: Vimal Dwivedi

J. Ljalik
All Systems Go OÜ, Tallinn, Estonia
E-mail: ljalik.jelena@gmail.com

1 Introduction

Healthcare data breaches have reached unprecedented levels, with serious implications for patient privacy, organizational costs, and public trust. In 2024 alone, 735 large-scale breaches affecting more than 276 million individuals were reported in the United States, representing the worst year on record for healthcare data security [3, 27]. Incidents such as the Change Healthcare ransomware attack demonstrate the scale and systemic impact of these vulnerabilities [21]. While these developments highlight the need for stronger privacy-preserving infrastructures, inter-organizational healthcare collaboration also introduces se-

mantic challenges, as different institutions may represent and interpret shared data and clinical concepts differently.

The financial consequences of healthcare data breaches are equally severe. According to IBM's 2024 Cost of a Data Breach Report, healthcare organizations face the highest average breach cost across all industries, reaching $9.77 million per incident [24]. Beyond financial losses, breaches disrupt clinical operations and patient care, with many healthcare providers reporting prolonged service interruptions following major incidents [24]. These challenges further complicate secure and trustworthy collaboration between healthcare organizations.

Furthermore, collaboration in inter-organizational decentralized e-health processes is complicated by privacy constraints, the sensitivity of health data, and legal regulations governing medical information. In addition, heterogeneous and partially trusted data sources introduce integrity, provenance, and semantic interpretation challenges that require verifiable reconciliation across institutions. Current blockchain solutions provide transparency, data integrity, and ownership guarantees, but they primarily ensure agreement on system state rather than on the semantic validity of clinical or governance decisions. As a result, existing blockchain-based approaches do not address governance-level trust and compliance in environments where different organizations may interpret the meaning and validity of shared healthcare data differently.

Although consensus mechanisms ensure agreement on transactions and system state in decentralized systems, they operate at the level of data ordering and do not address semantic interpretation [1, 40, 56, 61]. As a result, they cannot guarantee a consistent understanding of the meaning or validity of shared data among heterogeneous participants. Various studies are exploring blockchain-based e-health architectures and governance models [4, 6, 16, 26, 58]. These approaches aim to improve data security, transparency, patient control, and interoperability in healthcare systems. In parallel, standards such as HL7 FHIR are widely used to enable interoperable exchange of clinical information between healthcare systems [18]. Although such standards facilitate structural interoperability and support applications such as Clinical Decision Support Systems (CDSS), they primarily address data representation and exchange rather than semantic agreement on the validity or meaning of clinical decisions. As a result, even when shared data conform to common standards, organizations can still interpret their meaning differently, leading to unresolved semantic conflicts in decentralized healthcare workflows.

Research [35, 54] identifies zero-knowledge proof (ZKP) protocols as an important mechanism to achieve privacy-preserving verification in decentralized systems. Much of the existing work focuses on improving the efficiency and scalability of ZKP constructions [60], including approaches that reduce verification costs through techniques such as folding-based proofs [49]. In the healthcare domain, several studies explore the use of ZKP for secure data processing and privacy-preserving communication [8, 14]. These approaches demonstrate how ZKP can verify computations or compliance conditions without revealing sensitive data. However, current applications of ZKP primarily focus on validating numerical computations or protocol-level properties rather than verifying the semantic correctness of domain-specific interpretations. Consequently, the integration of ZKP-based verification with governance-level decision-making and semantic conflict resolution in decentralized healthcare settings remains largely unexplored.

Despite substantial progress in decentralized consensus, blockchain-based e-health architectures, and privacy-preserving computation, current research focuses on several critical limitations:

- *Network-layer focus without semantic interoperability:* Existing consensus mechanisms such as PoW, PoS, PBFT achieve agreement at the protocol level but do not address semantic conflicts arising from heterogeneous medical terminologies and interpretations (e.g., SNOMED-CT vs. local codes) [1, 40].
- *Transparency without sufficient privacy:* Blockchain solutions provide immutability and auditability of data, but their intrinsic transparency can expose sensitive patient information on the chain, creating challenges to comply with privacy regulations such as GDPR and HIPAA [6, 58].
- *Limitations of existing privacy-preserving techniques:* Multi-Party Computation and Homomorphic Encryption are general-purpose techniques that can protect data during semantic processing, but neither by itself yields the non-interactive, independently verifiable attestation that governance acceptance requires: MPC presumes synchronous participation of the computing parties [17], and HE returns a computed result without a compact proof that the computation followed the governance-approved rules [19].
- *Lack of governance-aware semantic consensus:* Current systems lack formal mechanisms that enable verifiable compliance with governance-approved clinical rules (e.g., treatment protocols, ontological constraints) within decentralized decision-making processes [5, 7].
- *No integration of semantic truth with cryptographic verification:* The integration of semantic conflict resolution, grounded in formal model-theoretic semantics (e.g., Tarski's theory of truth), with ZKP-based verification mechanisms remains largely unexplored [35].

These limitations make it difficult to use reliable and transparent decentralized systems to manage e-health in different organizations. This is especially true when institutions

need to resolve differences in meaning and agree on clinical decisions while still keeping data private. This paper focuses on the acceptance of semantic statements arising in cross-organizational e-health workflows (e.g., whether a prescription is valid under allergy and renal-function constraints). Such workflows often involve semantic conflicts, as different organizations may interpret clinical data differently. We model these statements using a domain-specific representation and evaluate their truth under a shared semantic model and governance-approved evaluation context based on model-theoretic semantics. The acceptance/rejection outcome provides a verifiable basis for governance-level consensus among autonomous participants.

Based on these established gaps, we derive the main research question (RQ) for the paper as follows. How can decentralized e-health workflows achieve semantic consensus on clinical statements while preserving privacy and enabling verifiable governance decisions through ZKPs? To establish a separation of concerns, we deduce the following sub-questions: RQ1: What are the formal requirements for privacy-preserving semantic consensus in decentralized e-health inter-organizational processes? RQ2: What architectural and governance mechanisms enable privacy-preserving semantic verification and acceptance of clinical decisions across heterogeneous healthcare organizations? RQ3: What cryptographic mechanisms and circuit constructions are required to encode semantic validity conditions into zero-knowledge proofs while preserving privacy, correctness, binding, and governance verifiability? The answer to the first question identifies the requirements for semantic statement evaluation, interoperability, and verifiable governance acceptance in decentralized workflows. Answering the second question specifies the architecture and governance workflow in which semantic evaluation results are anchored and verified, while ZKPs provide cryptographic evidence for private *governance policy* verification. The answer to the third question defines the technical design of the ZKP construction and verification pipeline.

The remainder of this paper is organized as follows. Section 2 analyzes related research in a similar field and introduces the e-health running case and the necessary preliminaries. Section 3 derives the requirements for semantic consensus and conflict resolution in decentralized e-health workflows. Section 4 presents the proposed architectural and governance mechanisms that address these requirements. Section 5 describes the verification pipeline based on zk-SNARK, including witness and public inputs, and its integration into the governance workflow. Section 6 implements and evaluates the proposed approach against four evaluation goals: functional correctness, circuit scalability, on-chain cost, and latency feasibility. Section 7 discusses the results, limitations, and open research directions. Finally, Section 8 concludes the paper.

2 Related Work

Research related to this work involves three main themes: decentralized consensus mechanisms and blockchain-based governance, privacy-preserving verification using ZKPs in healthcare, and semantic interoperability for heterogeneous health data. Each area contributes important building blocks to the problem addressed in this paper, however, existing research contributions have not properly integrated these topics into a coherent governance framework for healthcare processes. Section 2.1 discusses decentralized consensus and blockchain governance in healthcare. Section 2.2 reviews zero-knowledge proof systems and their application to privacy-preserving verification. Section 2.3 addresses semantic interoperability and introduces the formal model-theoretic foundations adopted in this work. A running case consolidating the identified gaps is presented in Section 2.4.

2.1 Decentralized Consensus and Blockchain Governance in Healthcare

Blockchain technology has been widely adopted to address trust, transparency, and data integrity challenges in healthcare. A survey of blockchain architectures for healthcare systems identifies security, scalability, and interoperability as persistent barriers to adoption [6]. Related work examines how blockchain can facilitate health information sharing while maintaining auditability [16], and a permissioned blockchain framework for secure EHR management demonstrates practical deployment at the institutional level [4, 29]. These works collectively establish the potential of blockchain for healthcare data management, but they operate at the level of transaction ordering and data immutability rather than on the semantic validity of clinical decisions derived from shared data.

Consensus mechanisms form the foundation of distributed agreement in blockchain systems. A comprehensive treatment of consensus in data management covers both classical commit protocols and blockchain-specific constructions from a distributed systems perspective [41]. The security implications of consensus mechanisms for permissioned blockchain systems are analyzed in [40], while hybrid consensus models that combine on-chain and off-chain components to balance performance and trust are introduced in [1]. Security vulnerabilities in blockchain consensus and corresponding mitigation strategies for open challenges are surveyed in [61]. Despite this substantial body of work, existing consensus mechanisms are designed to achieve agreement on transaction ordering and system state. They do not address the semantic layer: whether participating organizations interpret the meaning and validity of shared clinical data in the same way.

Decentralized governance through DAO structures has attracted increasing attention as a mechanism for managing shared resources without central authority. A survey of blockchain-based governance models for healthcare management systems highlights the potential for multi-stakeholder decision-making in interoperable settings [58]. The transformative role of blockchain in healthcare privacy and regulatory compliance is explored in [7]. An integrative review of DAO governance models distinguishes between purely on-chain voting and hybrid governance structures that incorporate off-chain deliberation [11]. The broader potential of DAOs to reshape healthcare management, including transparent governance and decentralized funding, is discussed in [38]. However, none of these works address how governance decisions can be grounded in formally verifiable semantic criteria rather than simple majority votes over raw or aggregated data.

2.2 Zero-Knowledge Proof Systems

ZKPs enable the verification of claims without revealing the underlying witness data. In decentralized healthcare settings, ZKPs provide a cryptographic mechanism to preserve privacy by verifying clinical constraints and semantic validity conditions derived from sensitive medical data. Several classes of ZKP systems are relevant in this context. Specifically, zk-SNARKs provide short, non-interactive proofs with efficient verification, making them suitable for on-chain governance workflows. zk-STARKs, on the other hand, avoid trusted setup assumptions and offer post-quantum security at the cost of larger proofs. Bulletproofs support efficient range and constraint proofs without trusted setup but typically incur higher verification overhead for more complex verification asks [43, 54]. These trade-offs are particularly important in governance settings, where proofs must remain compact enough for practical verification while still encoding semantically meaningful constraints. A comprehensive survey of ZKP applications in blockchain covers membership proofs, range proofs, and general-purpose arithmetic circuits [35]. Advances in ZKP constructions with particular attention to bridging theoretical efficiency and practical deployment are surveyed in [60], while folding-based proof composition as a scalable path toward incrementally verifiable computation is investigated in [49].

In the healthcare domain, ZKP systems have been applied primarily to protect patient identity and secure data exchange. A ZKP-based framework for privacy-preserving analytics over healthcare records stored on blockchain is proposed in [8]. Research on vehicle-to-healthcare communications demonstrates how zero-knowledge and statistical fingerprinting can protect sensitive medical data in transit [14]. A ZKP-enabled blockchain system for wearable health device data management demonstrates measurable improvements in privacy and data integrity for IoT-generated clinical evidence [28]. The MEPCA framework addresses on-chain EHR validation using cryptographic commitments and HL7 FHIR compliance checks, demonstrating that cryptographic anchoring of clinical data can support regulatory auditability [51]. A combination of ZK-Rollups with smart contracts and decentralized storage is proposed to improve scalability for cross-institutional health data exchange, decoupling proof generation from on-chain settlement latency [47]. Closest to the present work in construction are a zk-SNARK protocol pairing Groth16 with HL7 FHIR interfaces [39] and ZK-EHR, which encodes role-based access policies in a Groth16 circuit with IPFS-backed storage [2]. Both prove that a requester is authorized or that data are valid, without reconciling heterogeneous representations of the same clinical fact or governing updates to the rules against which compliance is judged.

Other privacy-preserving approaches, including Multi-Party Computation (MPC) and Homomorphic Encryption (HE), have also been applied in healthcare. A practical introduction to MPC covers its use in privacy-preserving analytics and the synchronous participation requirements that constrain its applicability [17]. Fully homomorphic encryption was established as a theoretical foundation for computation over encrypted data in [19]. The balance between patient privacy and regulatory compliance in Patient healthcare information-sharing scenarios is examined in [5]. While MPC requires synchronous participation from all parties, which is a constraint incompatible with asynchronous clinical workflows, HE introduces substantial computational overhead for general-purpose arithmetic. ZKP systems, particularly non-interactive constructions such as zk-SNARKs, avoid both limitations by generating compact, independently verifiable proofs off-chain.

These works establish that ZKP systems are technically applicable in healthcare settings. However, existing applications focus on proving properties of *data*, such as membership in a dataset, range conditions on numerical values, or identity claims, rather than on proving the semantic correctness of *clinical decisions* derived from heterogeneous ontological sources. Also, the circuit designs in these works encode numerical or structural constraints but do not capture first-order logical reasoning over domain-specific medical semantics. Table 1 summarizes the approaches used in related works.

2.3 Semantic Interoperability and Formal Foundations

Semantic interoperability remains one of the most persistent challenges in cross-institutional healthcare collaboration. Standards such as HL7 FHIR and SNOMED-CT provide frameworks for structural and terminological consis-

Table 1 Comparison with the nearest prior approaches

Approach	Semantic conflict res.	ZKP / private verif.	Authenticated evidence	Governance of rule updates	Implemented prototype	End-to-end evaluation
Semantic KG to smart contracts [57]	—	—	—	—	✓	○
MEPCA [51]	—	✓	○	—	✓	○
Wearable ZKP [28]	—	○	—	—	○	○
zk-SNARK FHIR sharing [39]	—	✓	○	—	✓	○
ZK-EHR [2]	—	✓	—	—	✓	○
This work	✓	✓	✓	✓	✓	✓

✓ = addressed; ○ = partially addressed; — = not addressed.

tency, but they do not guarantee that heterogeneous organizations will interpret shared data according to compatible semantic rules. A review of HL7 FHIR resources in digital healthcare ecosystems for chronic disease management shows that even standards-compliant systems can produce semantically inconsistent outcomes when organizations apply different local policies [18]. Research on blockchain-enabled personal health records proposes enhancements to interoperability, but addresses data representation rather than the semantic correctness of clinical decisions derived from that data [26].

Knowledge graphs and ontology-based reasoning have been proposed as tools for achieving deeper semantic integration. A framework for encoding smart contract logic using a semantic knowledge graph derived from HL7 FHIR concepts achieves structural interoperability by ensuring that contract logic reflects the vocabulary and relations of a shared domain standard [57]. This is conceptually close to our use of a Decentralized Knowledge Graph (DKG) to anchor governance-approved ontologies and theory axioms. Still, the framework does not address semantic conflicts that arise when different organizations apply heterogeneous terminologies or data representations to the same clinical facts, and it does not incorporate privacy-preserving verification mechanisms. In contrast, this paper treats semantic conflict resolution as the first concern. The framework verifies whether a clinical statement is semantically true under a shared model and governance-approved theory, and attests to this outcome cryptographically without revealing the underlying patient evidence.

Formal model-theoretic semantics, particularly Tarski's theory of truth [53], provides a rigorous foundation for evaluating the truth of statements with respect to a model and variable assignment. Model theory further develops these ideas and establishes the mathematical basis for the satisfaction relation used in this work [22]. While Tarski's framework is well established in logic and formal semantics, we did not identify prior work applying it to decentralized governance in healthcare systems (where the model is constructed dynamically from heterogeneous and partially trusted sources). Previous work on decentralized e-health systems motivates the need for architectures that verify semantic correctness without relying on centralized trust anchors [30]. Hence, we adopt Tarski's theory as the formal foundation for semantic consensus in our architecture, organized around three complementary components introduced below.

2.3.1 static semantic model (the facts)

We model shared clinical knowledge as a *semantic model*:

$$M = \langle D, I \rangle \tag{1}$$

where $D = D_P \uplus D_S \uplus D_M \uplus D_D \uplus D_R$ is a typed domain of pairwise disjoint sorts — patients, substances, medications, prescribers, and record entities for allergies and lab observations — and I is an interpretation function that assigns meaning to each relation over D.

Following the model-theoretic semantics of RDF [45] and the named-graph approach to provenance [37], clinical facts are represented as *labeled entities*: each record is a first-class entity in D whose attributes (subject, clinical content, source) are expressed through binary relations in I. For example, an allergy record a_1 stating that patient *pat1* has a documented allergy to Penicillin, recorded by hospital H, is represented as:

$a_1 \in I(isAllergy)$

$\langle a_1, pat1 \rangle \in I(hasPatient)$

$\langle a_1, Penicillin \rangle \in I(hasSubstance)$

$\langle a_1, H \rangle \in I(hasSource)$

This structure mirrors the RDF-triple representation anchored in the Decentralized Knowledge Graph (DKG), with public commitments to entity identifiers; the full attribute content resides off-chain with its respective custodians and is revealed only through zero-knowledge proofs (Section 5).

A key consequence of the labeled-fact representation is that independent observations of the same clinical condition correspond to *distinct* entities $a_1, a_2 \in I(isAllergy)$, each with its own provenance. For example, an allergy recorded by a hospital and separately confirmed by a laboratory. Both facts coexist in M, reflecting the decentralized reality that multiple independent sources may report on the same clinical condition.

2.3.2 governance-approved theory (the rules)

Clinical and governance rules are collected in a governance-approved *theory*:

$$T = \{\alpha_1, \alpha_2, \ldots, \alpha_n\} \quad (2)$$

A model M is *admissible* if it satisfies every axiom in T, written $M \models T$. The theory T is anchored in the DKG, ensuring that all participants evaluate statements against the same set of governance-approved rules.

Five classes of axioms are distinguished, each addressing a distinct form of semantic heterogeneity encountered in cross-organizational e-health workflows: *Ontological hierarchies* capture subsumption relations among clinical concepts, such as Penicillin $\sqsubseteq \beta-$lactam. *Inference rules* propagate clinical implications along hierarchical and relational structures, such as "allergy to a class implies contraindication for class members". *Terminology bridges* align equivalent predicates expressed in heterogeneous coding systems, e.g., mapping a SNOMED CT allergy predicate to a laboratory's local coding. *Numerical bridges* relate clinical values reported in different units or scales, such as converting an eGFR value computed via the Chronic Kidney Disease Epidemiology Collaboration equation (CKD-EPI) [25] and reported as an absolute value in mL/min onto the body-surface-area-indexed scale mL/min/1.73m^2 via $eGFR_{\text{idx}} = eGFR_{\text{abs}} \cdot 1.73/BSA$, where BSA is the patient's body surface area in m^2 and 1.73 m^2 is the conventional reference value, so that independent laboratory observations can be compared against a single governance-approved clinical threshold. The evaluation instantiates BSA with a population constant of 1.9 m^2; a clinical deployment would derive it from patient height and weight. *Clinical policy axioms* encode governance-approved threshold conditions on medications, e.g. Pol(Metformin, eGFR, $\geq$, 30)

2.3.3 dynamic evaluation context (the request)

In classical first-order semantics, the truth of a statement depends only on the model. In decentralized cross-organizational workflows, however, evaluating a clinical statement additionally requires reasoning about *the request itself:* who initiated it, which specific facts in M it relies upon, and when it was issued. These concerns are not properties of the world and therefore do not belong to M; nor are they rules of clinical practice belonging to T. They are evidential metadata of the current evaluation attempt, and we collect them in an *evaluation context*:

$$\Gamma = \langle cred, refs, t_{req} \rangle \quad (3)$$

Each component of Γ carries a distinct kind of request-specific information: *cred* is the *credential* presented by the initiator of the request. It is issued by a governance-approved authority and registered in the governance-approved credential registry. The credential is verified for validity without disclosing the identity or role attributes of the credential holder. *refs* is the *set of record identifiers* in M that the request explicitly relies upon. For example, the allergy records and laboratory observations invoked by a particular prescription. Referencing records explicitly, rather than leaving them implicit, makes the evidentiary basis of the request auditable and enables the patient, the governance layer, and external observers to trace which facts informed a given clinical decision. t_{req} is the *timestamp of the request*, recorded as the prescription issuance time. It is consumed downstream by operational checks such as the pharmacy's validity-window check at dispensing.

2.4 Running case

We extend a running case, originally developed in [31], to analyze semantic conflicts in inter-organizational prescription workflows in decentralized healthcare. The running case is based on the previous work and discussions with practitioners about the possible benefits of blockchain usage in healthcare processes. The scenario involves three autonomous and decentralized entities: a public hospital, a private diagnostics laboratory, and a community pharmacy. A patient authorizes data sharing among these organizations to support a coordinated prescription creation process.

Figure 1 illustrates the updated running case and the distributed interaction between the participating entities during prescription preparation. When a hospital physician initiates a digital prescription, the Health Information System (HIS) requests relevant patient information, such as laboratory results, allergy records, and current medications, from multiple external systems, including Electronic Health Records (EHR), Personal Health Records (PHR), and laboratory information systems. Although each organization maintains its own data schema and authorization policy, the workflow depends not only on interoperable data exchange but also on a semantically consistent interpretation across all participants. Figure 1 sets the context of the running case, and the architecture described in Section 4 defines the implementation that is evaluated in Section 6. Components such as prescription ownership with NFT, blockchain oracles, challenge-set libraries, and integration with the medication supply chain are outside the implementation scope and are kept for future work.

Figure 2 illustrates two representative *data-level* semantic conflicts in the prescription workflow. First, a *numerical representation conflict*: the hospital's clinical policy for Metformin requires that the patient's eGFR, computed via the CKD-EPI formula, be at least 30 mL/min/1.73m^2, while

the laboratory reports the same CKD-EPI estimate as an absolute value in mL/min, de-indexed from body surface area for dose-related use. Both values describe the same clinical quantity on different scales and cannot be directly compared against the policy threshold. Second, a *terminological conflict*: the hospital's EHR encodes the patient's allergy using a SNOMED-CT concept[1], while the laboratory annotates the same condition under a local allergen vocabulary, so that without alignment, the system cannot determine whether the two records refer to the same condition. These examples show that the syntactic interoperability of data exchange alone does not guarantee agreement on the meaning of clinical decisions.

Both conflicts above are resolved automatically when the governance-approved theory T contains the appropriate alignment axiom: a *numerical bridge* relating absolute and body-surface-area-indexed eGFR representations against the same threshold, and a *terminology bridge* aligning SNOMED CT and local allergen vocabularies. When such an axiom is missing from T, the conflict is escalated to the governance layer for review, in which the DAO either approves or rejects a new alignment axiom for inclusion in T.

Higher-level conflict types beyond data-level alignment, as well as stronger trust mechanisms beyond honest-majority DAO governance, are outside the scope of this work. This paper formalizes the *issuance step* of the workflow, with the prescription-creation decision verified against shared semantic rules. The dispensing step at the pharmacy is treated as a downstream consumer of the issuance attestation, and detailed dispensing protocols are outside the scope of this work.

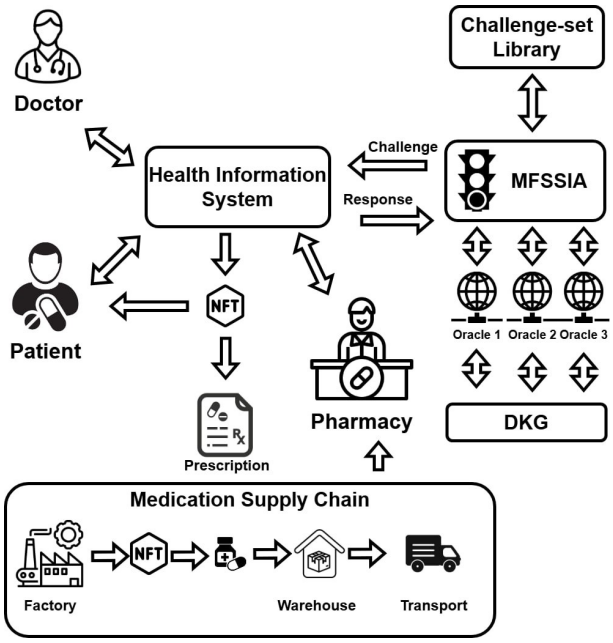


Fig. 1 Running case of inter-organizational data sharing and prescription preparation in decentralized e-health [31]

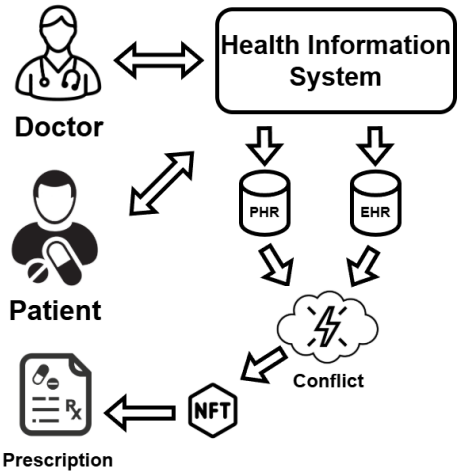


Fig. 2 Representative semantic conflicts in a decentralized digital prescription workflow [31]

3 Formal Requirements for Privacy-Aware Semantic Consensus

This section answers RQ1 by providing a formal description of semantic consensus in healthcare processes. Semantic consensus refers to governance-level agreement on whether a domain-specific statement φ is accepted under a shared semantic model and governance theory, rather than on system state. The following system requirements are derived, including: R1: Statement-level correctness, R2: trustworthiness, R3: interoperability, R4: Confidentiality, R5: replay protection, R6: minimal public disclosure, R7: robustness, R8: statement binding, and R9: auditability. These requirements are formalized in Definitions in Section 3.1 to 3.4

3.1 Conflicts about Data Representation

Representation conflicts arise when the same clinical fact is encoded under heterogeneous vocabularies, class hierarchies, formulas, or units; terminological conflicts are resolved by governance-approved alignment axioms in T, such as

$$\mathsf{Amoxicillin} \sqsubseteq \mathsf{Penicillin} \sqsubseteq \beta\text{-}\mathsf{lactam}, \tag{4}$$

which induce inference rules, such as

$$\forall x \in D_P.\ \mathrm{Allergy}(x, \mathsf{Penicillin}) \rightarrow \mathrm{Contraindication}(x, \mathsf{Amoxicillin}) \tag{5}$$

numerical bridges resolve numerical conflicts and map formula- or unit-specific measurements to a single governance-approved scale. Both classes of conflicts are resolved through

[1] https://www.snomed.org/

governance-approved alignment axioms in T. This yields **Semantic interoperability (R3)**: heterogeneous encodings of clinical concepts and numerical quantities must be reconciled via governance-approved alignment axioms in T.

3.2 Conflicts about Healthcare Policies

Policy conflicts are checked per request against patient-specific evidence and prescription-internal parameters. A prescription is acceptable only if its initiator holds a valid, governance-approved credential, the patient's clinical context supports the prescribed medication, and the workflow instance is uniquely bound to a governed statement commitment. These checks are encoded in T as clinical policy axioms of the general form

$$\mathrm{Pol}(m,\ t,\ op,\ \theta) \tag{6}$$

where $m \in D_M$ is a medication, t identifies an evaluation attribute, $op \in \{\geq, \leq, =, \neq\}$ is a comparison operator, and θ is a threshold expression. For example,

$$\mathrm{Pol}(\text{Metformin, eGFR, } \geq, 30) \tag{7}$$

requires the patient's eGFR to be at least 30. A prescription must also be bound to a unique workflow instance: uniqueness is enforced by a per-statement *nonce* in the statement commitment *stmt_hash*. The operational validity window of a prescription is consumed by downstream actors (e.g., the pharmacy at dispensing time).These checks define prescription validity, formalized as $\mathrm{RxValid}(d, p, m, n)$ in Definition 2, while source (prescriber) legitimacy is captured by CredValid in *ContextTrusted*, and the combined acceptance predicate is given in Definition 3. The proof attests to these semantic checks without revealing the patient's evidence or credential attributes. This yields three further requirements for semantic consensus: **Statement-level correctness (R1)** requires acceptance only when both context trust and semantic satisfaction hold under $M \models T$; **Context trustworthiness (R2)** requires that only trusted contexts may yield acceptance; otherwise, satisfaction is undefined; and **Replay protection (R5)** requires a downstream single-use check to prevent an already-redeemed proof from being presented a second time, distinct from the per-workflow *nonce* that binds each proof to a unique statement instance.

3.3 Privacy of Patient Evidence

Privacy is a cross-cutting constraint: while ontology and policy axioms in T are publicly anchored on the DKG, patient-specific clinical evidence remains private witness input and is never revealed on-chain. Public disclosure is limited to cryptographic commitments and the acceptance outcome, and proof verification, rather than unilateral trust, must determine whether a statement is accepted. This yields three further requirements for semantic consensus: **Witness confidentiality (R4)** requires patient-specific evidence to remain off-chain and be verified only through zero-knowledge proofs; **Minimal public disclosure (R6)** requires on-chain inputs to exclude personally identifiable clinical attributes; and **Robustness under partial collusion (R7)** requires acceptance to depend on proof verification and governance thresholds rather than on the unilateral trust of any single party.

3.4 Combined Formal Requirement

Sections 3.1, 3.2, and 3.3 identify the classes of conflicts that a decentralized e-health workflow must reconcile and the per-class requirements (R1 - R7). We now consolidate these into a single formal acceptance condition. The condition combines *context trust*, which determines whether an evaluation context is admissible, and *semantic satisfaction*, which determines whether the referenced records support the statement under the governance theory T. Their conjunction is the governance-level acceptance predicate $\mathrm{Accept}(M, \Gamma, \varphi)$.

3.4.1 Context Trust

We define $\mathrm{ContextTrusted}(\Gamma, M)$ through credential validity. $\mathrm{CredValid}(cred)$ holds iff the credential is included in the governance-approved credential registry.

Definition 1 (Context Trust): An evaluation context $\Gamma = \langle cred, refs, t_{req} \rangle$ is context-trusted with respect to M iff

$$\mathrm{CredValid}(cred). \tag{8}$$

A context failing this check is rejected before semantic evaluation. The predicate covers the credential only and asserts nothing about the records in *refs*; their trustworthiness is established separately by the governance layer at the time records are admitted to $root_M$, and Merkle membership in $root_M$ therefore subsumes the source-trust check at proof time.

3.4.2 Satisfaction Relation

The satisfaction relation $M, \Gamma \models \varphi$ is a partial relation on triples (M, Γ, φ), defined only when $\mathrm{ContextTrusted}(\Gamma, M)$ holds. Under the assumption $M \models T$, the relation is defined recursively.

Definition 2 (Satisfaction): For allergy statements,

$$\begin{aligned} &M, \Gamma \models \mathrm{Allergy}(p, s) \iff \\ &\quad \exists\, a \in D : a \in I(isAllergy) \\ &\qquad \wedge\ \langle a, p \rangle \in I(hasPatient) \\ &\qquad \wedge\ \langle a, s \rangle \in I(hasSubstance) \end{aligned} \tag{9}$$

For contraindications,

$$M, \Gamma \models \text{Contraindication}(p, m) \iff \exists\, s \in D_S : M, \Gamma \models \text{Allergy}(p, s) \wedge m \sqsubseteq s \tag{10}$$

Policy satisfaction is defined through an attribute-selection function and pointwise satisfaction of all applicable policy axioms. For a patient $p \in D_P$, a prescription dose n, and an attribute name t, the value-selection function

$$\begin{aligned}\text{AttrValue}(p,\, n,\, t) = v \iff\ & \big(\exists\, r \in D : r \in I(\mathit{isLabResult}) \\ & \wedge\ \langle r, p\rangle \in I(\mathit{hasPatient}) \\ & \wedge\ \langle r, t\rangle \in I(\mathit{hasMetric}) \\ & \wedge\ \langle r, v\rangle \in I(\mathit{hasValue}) \\ & \wedge\ \text{Latest}(r,\, p,\, t)\big) \\ & \vee\ \langle p, t, v\rangle \in I(\mathit{hasPatientAttr}) \\ & \vee\ \big(t = \mathit{dose} \wedge v = n\big)\end{aligned} \tag{11}$$

returns the value of attribute t from the most recent laboratory observation of t for p in M, a recorded patient attribute, or the prescription dose n. Here $\text{Latest}(r, p, t)$ holds if no other laboratory observation of metric t for patient p in M carries a later *measuredAt* value. The recency condition makes AttrValue a function rather than a relation: where several observations report the same metric, a prover cannot select whichever value satisfies the policy since only the latest observation is admissible as a witness. However, this condition is ensured by the prover service and not by the ZKP circuit itself. Writing $\text{PolicyAxioms}(m, T) = \{(t, op, \theta) : \text{Pol}(m, t, op, \theta) \in T\}$ for the policy axioms applicable to medication m, policy satisfaction is shown in eqn 12, and prescription validity is shown in eqn 13. When $\text{PolicyAxioms}(m,\, T) = \emptyset$, PolicySatisfied holds vacuously.

$$M, \Gamma \models \text{PolicySatisfied}(p,\, m,\, n) \iff \forall\, (t, op, \theta) \in \text{PolicyAxioms}(m, T) : \exists\, v : \text{AttrValue}(p, n, t) = v \wedge v\ op\ \theta. \tag{12}$$

$$M, \Gamma \models \text{RxValid}(d,\, p,\, m,\, n) \iff M, \Gamma \not\models \text{Contraindication}(p,\, m) \wedge M, \Gamma \models \text{PolicySatisfied}(p,\, m,\, n) \tag{13}$$

3.4.3 Governance accountability

Every acceptance decision is bound to the evaluated statement instance through

$$\mathit{stmt_hash} = H\big(\varphi^{+} \parallel \mathit{nonce}\big) \tag{14}$$

where φ^{+} is the ground statement and *nonce* is a per-workflow unique value for replay protection. The acceptance outcome is logged together with *stmt_hash* in an immutable on-chain registry. This yields **Statement binding (R8)** and **Immutable auditability (R9)**.

Definition 3 (Acceptance): A statement φ is accepted for governance purposes in the evaluation context Γ relative to model M iff

$$\text{Accept}(M, \Gamma, \varphi) \iff \text{ContextTrusted}(\Gamma, M) \wedge M, \Gamma \models \varphi. \tag{15}$$

If $\text{ContextTrusted}(\Gamma, M)$ does not hold, satisfaction is undefined and governance rejects the statement by default.

4 Architectural and Governance Mechanisms for Privacy-Preserving Semantic Consensus

This section addresses RQ2 by introducing an architectural and governance framework for semantic consensus over medical statements that are privacy-sensitive in decentralized e-health environments. The architecture operationalizes the formal specification from Section 3 through a four-layer separation in which semantic truth is evaluated off-chain, attested by zero-knowledge proofs, and accepted or rejected by on-chain governance without exposing medical data.

Section 4.1 defines the architectural scope and guiding principles, Section 4.2 describes the end-to-end workflow, Section 4.3 specifies the governance layer, and Section 4.4 discusses scope and limitations. Figure 3 summarizes the layered structure and component interactions.

4.1 Design Scope and Architectural Principles

The architecture is designed to achieve consensus over semantic statements rather than over raw medical data or numerical aggregates. It operationalizes the satisfaction relation from Section 3 through a separation of semantic interpretation, proof generation, on-chain verification, and DAO-based policy governance. This separates functional responsibilities, e. g., semantic evaluation, which is performed off-chain, from acceptance, which is determined by independent on-chain verification based on DAO-approved parameters.

The architecture targets deployment on a public, permissionless chain: network participation, transaction submission, and read access to the acceptance log and the DKG are unrestricted. Permissioning is enforced only at the application layer. The credential issuance, DKG write access, and DAO membership are governance-approved rather than at the network layer. This places the trust boundary at the credential and governance layer rather than at network admission, so that verification does not depend on membership in a consortium.

The framework is structured into four layers: domain systems, semantic interpretation, proof generation, and on-chain governance. These layers are summarized in Table 2 and illustrated in Figure 3. The on-chain governance layer comprises two functionally distinct components: an automated cryptographic verifier (active for every proof) and a DAO governance protocol (activated only on escalation when the semantic interpreter cannot resolve a conflict using the current theory T).

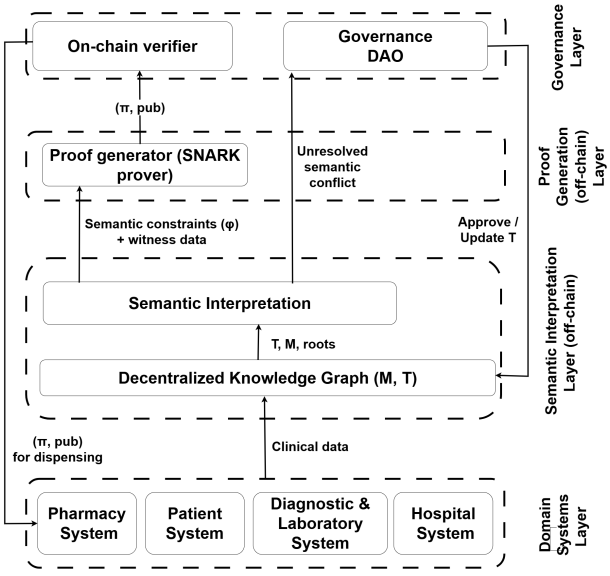


Fig. 3 Architecture of the Privacy-Preserving Semantic Consensus Framework

4.2 End-to-End Semantic Consensus Workflow

Figure 3 illustrates the end-to-end workflow for achieving semantic consensus on clinical and governance statements. The workflow begins by instantiating a domain-specific statement schema φ and an evaluation context Γ to obtain a ground statement φ^+ and a trusted request context. The semantic interpreter then checks ContextTrusted(Γ, M) and, if the context is trusted, evaluates $M, \Gamma \models \varphi^+$ under Definition 2 using theory T retrieved from the DKG; otherwise, acceptance fails. If the evaluation cannot be completed because T does not cover the case (e.g., a missing alignment axiom), the conflict is escalated to the DAO governance protocol, which votes on a corresponding update to T; the updated theory is re-anchored on the DKG and the workflow is retried. If evaluation succeeds, the acceptance condition is evaluated in a precompiled arithmetic circuit C, bound to a unique statement commitment stmt_hash, and attested by a zk-SNARK proof π; the resulting artifact (π, pub) constitutes the issuance attestation. This artifact is presented on-chain at the point of prescription dispensing.

This workflow explicitly separates semantic reasoning from governance enforcement, realizing the acceptance relation Accept(M, Γ, φ) of Definition 3 through off-chain evaluation, zk-SNARK proof generation, on-chain cryptographic verification, and DAO-based escalation for unresolved conflicts, while ensuring that acceptance decisions are cryptographically verifiable and privacy-preserving.

4.3 Governance Layer

Theory T is published on a Decentralized Knowledge Graph (DKG), where ontology schemas and policy axioms are stored as content-addressable artifacts that every participant verifies against identical commitments. Only governance-approved participants may publish or update artifacts, and conflicting axioms are handled through pre-publication governance.

The on-chain governance layer comprises two functionally distinct components. The on-chain verifier is a Solidity smart contract that performs automated cryptographic verification of each zk-SNARK proof against the committed public inputs and appends the outcome together with stmt_hash to an immutable acceptance log. The verification happens automatically, without human intervention, for every proof presented on-chain, in the current deployment, at dispensing. The *DAO governance* component is activated only when the semantic interpreter cannot resolve a conflict using the current theory T (e.g., a missing alignment axiom or a new policy proposal); DAO members vote, with a governance-approved threshold, on the corresponding update to T, which is then re-anchored on the DKG. Credential registry updates such as issuance and revocation follow the same DAO-approved publication path.

The semantic consensus guarantees established in this paper hold for any instantiation that provides (i) automated cryptographic verification of every proof, (ii) threshold-based DAO voting on escalated conflicts, and (iii) immutable logging of both verification outcomes and policy updates. The specific instantiation of voting eligibility, quorum size, and threshold values is a deployment choice that depends on the participating organizations and remains outside the scope of this protocol.

4.4 Scope and Limitations of the Architectural Model

Organizational risks such as misconfiguration, role assignment errors, and threshold manipulation remain outside the protocol. However, they can be mitigated through transparent on-chain logging and collective decision-making but are not enforced cryptographically.

Performance characteristics, circuit complexity, proof size, and latency impacts are outside the architectural scope

Table 2 Architectural Layers of the Privacy-Preserving Semantic Consensus Framework

Layer	Responsibility	Description
Domain Systems Layer	Provision of medical data	Provides raw clinical data while remaining isolated from consensus and governance processes.
Semantic Interpretation Layer (off-chain)	Evaluation of semantic truth	Evaluates statements under the shared model $\boldsymbol{M}$, theory $\boldsymbol{T}$, and evaluation context Γ. Escalates unresolved conflicts to the DAO governance component.
Proof Generation Layer (off-chain)	Cryptographic proof construction	Transforms semantic evaluation results into zero-knowledge proofs (π, pub) submitted to the on-chain verifier.
On-chain Governance Layer	Automated proof verification and policy management	Comprises (i) a Solidity smart contract that automatically verifies each zk-SNARK proof and appends the outcome to an immutable acceptance log, and (ii) a DAO that votes on theory updates when the semantic interpreter cannot resolve a conflict using the current $\boldsymbol{T}$.

and are evaluated separately in Section 6. On-chain anchoring of the acceptance outcome is deferred to the first point at which the proof is presented on-chain; in the current deployment this is dispensing rather than issuance. This choice avoids an on-chain write, and its associated cost, for prescriptions that are generated but never dispensed. The trade-off is that governance-level audit visibility extends only to prescriptions that reach an on-chain interaction. For instance, a prescription that is issued, semantically evaluated, and never dispensed leaves no on-chain trace. Still, its underlying proof remains independently verifiable off-chain at any time.

5 zk-SNARK Encoding and Verification

This section addresses RQ3 by describing how the semantic acceptance condition derived in Section 3 is encoded as a zero-knowledge proof. We instantiate the construction using Groth16 over an R1CS circuit compiled with Circom. Table 3 provides the central mapping from the formal specification to the corresponding circuit constraints. The circuit verifies context trust, prescription validity, and statement binding while keeping patient-specific evidence outside the public verification interface.

5.1 Circuit Encoding and Commitments

The circuit $C(w, x) = 1$ is satisfied when the encoded decision procedure has been correctly executed under the model M and governance theory T. The valid proof ensures that the evaluation of the circuit was correct and does not imply that the prescription is accepted.

As summarized in Table 3, the construction implements four checks. First, *CredValid* proves that the prescriber's credential belongs to the governance-approved credential registry. Second, *NoContraindication* proves that the prescribed medication is not contraindicated by any committed allergy record, including contraindications derived through the governance-approved subsumption closure. Third, *PolicySatisfied* applies the governed clinical thresholds to values committed in the patient's record. Finally, the binding check associates the proof with one workflow instance through its nonce and statement commitment.

Clinical records and governance artifacts are represented by Merkle commitments. A record r is encoded as

$$\text{leaf}(r) = H(r_{\text{id}} \parallel r_{\text{type}} \parallel a_1 \parallel \cdots \parallel a_k), \tag{16}$$

where $a_1, \ldots, a_k$ are the attributes required by the corresponding semantic condition. Merkle membership is verified inside the circuit, so patient evidence can be shown to belong to the committed model without revealing other records. The component checks described in Table 3 produce the context-trust, prescription-validity, and binding flags. Their conjunction determines the public acceptance outcome.

5.2 Private Witness and Public Inputs

The private witness contains the credential, referenced clinical records, prescription content, ontology witnesses, and auxiliary values required by the circuit:

$$w = \langle w_{\text{cred}}, w_{\text{refs}}, w_{\text{stmt}}, w_{\text{ont}}, w_{\text{aux}} \rangle. \tag{17}$$

Patient identifiers, prescribed medication and dosage, clinical values, credential attributes, and Merkle authentication paths remain within this witness. The verifier receives only the circuit outcome, the statement commitment, validity metadata, and commitments to the relevant governance and record state:

Table 3 From Formal Specification to Circuit Encoding

Formal condition	Circuit encoding	Practical interpretation
Semantic model and commitments		
$M = \langle D, I \rangle$ (Eq. 1)	Per-domain Merkle trees committed in the DKG: credential registry $root_M$, *patientRecordRoot*, *labRecordRoot*, and the contraindication closure *contraindicationRoot* — with per-record leaves (Eq. 16);	Labeled records in M are hashed into leaves of the corresponding committed tree following the zk-creds pattern [48]; membership is proven without exposing remaining entries.
Context trust (Definition 1)		
CredValid(*cred*)	Merkle membership of $H(cred)$ in $root_M$	Credential inclusion in the governance-approved registry verified without exposing credential content.
Satisfaction clauses (Definition 2)		
$M, \Gamma \not\models$ Contraindication(p, m) (Eq. 10)	$no_contra = \prod_i (1 - isActive_i \cdot contraValue_i)$ via committed-closure membership	Absence of contraindication proven by membership of each allergy leaf in *patientRecordRoot* and of the triple $(substanceId_i, m, contraValue_i)$ in the governance contraindication closure (*contraindicationRoot*); the contraindication value cannot be forged.
$M, \Gamma \models$ PolicySatisfied(p, m, n) (Eq. 12)	Range-check sub-circuits over $\mathcal{T}_m$, one slot per governed policy, with the metric pinned by *labMetricIdPub*	Each policy threshold (e.g., eGFR $\geq$ 30) enforced arithmetically over *hasValue*, a private witness bound to *labRecordRoot* by Merkle membership; the metric is pinned to the one the policy was written for, and a slot is active exactly when its policy targets the prescribed medication, so neither the value's provenance nor the policy's applicability is left to the prover. The clinical value is never disclosed.
$M, \Gamma \models$ RxValid(d, p, m, n) (Eq. 13)	$rx_ok = no_contra \cdot pol_ok$	Satisfaction block aggregates the two clauses; any failure sets $rx_ok = 0$.
Acceptance and governance binding (Definition 3)		
Accept$(M, \Gamma, \varphi^{+}_{Rx})$	$outcome = ctx_ok \cdot rx_ok \cdot bind_ok$	Final acceptance requires context-trust, satisfaction, and nonce binding; published as public input.
$stmt_hash = H(\varphi^{+} \| nonce)$ (Eq. 14)	Recomputed in-circuit; constrained to equal public input	Binds proof to one prescription instance, and one workflow invocation; prevents cross-instance substitution and reuse.
$C(w, x) = 1$	$Verify_{vk}(\pi, pub) = 1$	On-chain Solidity verifier accepts using only (π, pub); no patient data disclosed to the governance layer.

$$
\begin{aligned}
pub = \langle & outcome, stmtHash, issuanceTime, validFor, \\
& validCredentialRoot, patientRecordRoot, \\
& nonce, policyDrugIds, maxDosages, \\
& contraindicationRoot, labRecordRoot, \\
& labThreshold, labRequiredOp, \\
& labAppliesToDrug, labMetricIdPub \rangle.
\end{aligned} \tag{18}
$$

The statement commitment binds the proof to the prescription and workflow instance:

$$
\begin{aligned}
stmtHash = Poseidon(& dCH, nonce, prescribedDrugIds, \\
& prescribedDosages, issuanceTime, validFor)
\end{aligned} \tag{19}
$$

where dCH is the private doctor-credential hash. Consequently, the on-chain verifier can establish that the proof concerns a specific governed statement without learning the prescription content itself.

5.3 Proof Lifecycle and Verification Guarantees

The proof (π, pub) is generated after successful semantic evaluation. It is bound to the statement through $stmtHash$, to the committed governance state through the public roots and policy parameters, and to the circuit version through the verification key vk. A change to governance parameters therefore invalidates proofs referring to the old configuration, while a structural circuit change requires a new verification key. In the current workflow, the proof is generated at prescription issuance and submitted on-chain when the prescription is dispensed. The Solidity verifier performs the Groth16 pairing check and records the accepted statement commitment. No patient-specific witness data are required by the pharmacy or the on-chain governance layer.

Under Groth16 soundness and the binding properties of the commitments, an accepted proof implies that the circuit constraints corresponding to the acceptance relation were satisfied for records committed under the approved governance state. Under the zero-knowledge property of Groth16, verification reveals no additional information about the private witness beyond the declared public inputs. These guarantees are conditional on the integrity of the circuit implementation, the trusted-setup process, and the governance-approved commitments. Their practical implications and remaining trust assumptions are discussed in Section 7.

5.4 Complexity

Circuit complexity is dominated by credential verification, Merkle membership proofs for referenced records, and range checks for governance policies. For a maximum of $N_{\max}$ reference slots,

$$n_{\text{total}} \approx n_{\text{cred}} + N_{\max} n_{\text{Merkle}} + |\text{PolicyAxioms}(m, T)| n_{\text{range}}. \tag{20}$$

The construction therefore predicts linear growth with the number of credential reference slots and applicable policy axioms. Section 6.4 evaluates this prediction empirically and examines its effect on setup cost, proof size, and on-chain verification.

6 Implementation and Evaluation

This section shows the implementation of architectural components and evaluates the artifact following the Design Science Research paradigm, structured around four evaluation goals (EG1–EG4). Section 6.1 provides an overview of the implementations necessary to realize the various architectural components in Figure 3. Section 6.2 states the evaluation design, the strategy, the evaluation goals, and the experimental setup and reproducibility conditions. Section 6.3 establishes functional correctness (EG1) through scenario-based validation at two levels: the cryptographic soundness and completeness of the proof system, and the end-to-end behavior of the full workflow. Section 6.4 quantifies the efficiency and scalability of the zero-knowledge circuit (EG2). Section 6.5 measures the on-chain cost of proof verification and of the k-of-n governance operations (EG3), and Section 6.6 assesses latency feasibility (EG4).

6.1 Components Implementation

6.1.1 Smart contract proof verifier and on-chain logging

The on-chain implementation is separated across the *Verifier.sol*, *AcceptanceGate.sol*, and *DecisionRegistry.sol* contracts.[2] The generated Groth16 verifier performs the pairing check, while *AcceptanceGate.verifyAndRecord()* provides the transaction entry point for proof-backed decisions. It verifies the proof, checks the committed governance parameters and prescription-validity metadata, and forwards the resulting *stmtHash*, outcome, and governance-state hash to the *DecisionRegistry*. The registry restricts *record()* to the acceptance-gate contract and rejects an already recorded *stmtHash*, preventing the verification path and replay check from being bypassed through a direct registry write.

6.1.2 Smart contract DAO proposal and voting

The DAO mechanism is implemented in *MinimalGovernance.sol* as a fixed-member k-of-n contract.[3] Governance changes are represented by cryptographic commitments: *propose()* registers a commitment, while *vote()* permits one vote per registered member and marks the commitment as approved once the configured threshold is reached. The same approval mechanism is used before activating a new governance vector or replacing the approved prescription validity-window set, so these configuration changes follow the same voting path as other governed policy updates.

6.1.3 zk-SNARK Circuit in Circom

The zk-SNARK implementation is defined by the parameterized *PrescriptionValidation* circuit in Circom 2.1.6 and

[2] https://github.com/medalex/ehealth-evm/tree/ac1440756ca023a75fadabfb8c711adbb5e62641/contracts

[3] https://github.com/medalex/ehealth-evm/blob/ac1440756ca023a75fadabfb8c711adbb5e62641/contracts/MinimalGovernance.sol

a NestJS-based prover service.[4] The circuit source contains reusable Boolean, comparison, Poseidon-hash, and Merkle-membership templates that are composed by the main prescription validation circuit. At runtime, *ProverService* converts the application-level prescription request and MFSSIA-supplied Merkle evidence into the fixed circuit input structure and invokes *snarkjs.groth16.fullProve()* using the compiled WebAssembly circuit and proving key. MFSSIA is a decentralized authentication service used across the project implementations [32, 42]. The resulting proof and public signals form the artifact passed to the on-chain verifier.

6.1.4 Knowledge Graph and Semantic Bridge

The knowledge-graph integration is implemented in MFSSIA through *DkgService*, which wraps the *dkg.js* client for publishing assets, retrieving them, and executing SPARQL queries against the DKG.[5] Semantic mappings are implemented separately in *NumericBridgeService* and *TerminologyBridgeService*. These services query the currently published bridge axioms, perform numerical normalization or terminology alignment, and raise a typed semantic conflict when no applicable mapping is available. For governance-added mappings, the services create a commitment for the proposed bridge, interact with the EVM governance API, and use *publishApprovedBridge()* to check DAO approval before publishing the new mapping to the DKG.

6.2 Evaluation Design and Setup

The evaluation is based on Design Science Research frameworks. The selection of evaluation methods follows [50], which provides a systematic procedure for matching an evaluation method to the goal it serves; the claims we advance are typed according to [34], which distinguishes criterion, causal, and context claims; and the validity of the resulting measurements is assessed against the five artifact validity types of [33] (technical, instrument, generalization, purpose, and context).

6.2.1 Evaluation Goals

Applying [50], we break down the evaluation into four goals. Each goal states what is to be established, the type of claim it advances in the sense of [34], and the method selected to serve it:

Functional correctness (EG1): The artifact behaves as the formal specification prescribes: the compiled circuit computes and cryptographically binds the outcome prescribed by Definition 3, whether acceptance or rejection, over the witness it is given; conditions the circuit cannot express, such as recency, are enforced by the prover service, and the deployed workflow enforces the architectural separation of Section 4. Because correctness is a property of specific behaviors rather than an aggregate, the method selected is scenario-based validation, applied at the circuit and system levels (Section 6.3).

Circuit size and scalability (EG2): The size of the compiled circuit grows linearly in the parameters Eq. 20 exposes. These include the number of reference slots $N_{\max}$ and the number of applicable policy axioms. The sweep below varies $N_{\max}$ only. The verification cost does not follow the same linear order. Establishing a parameter-cost relationship requires varying the parameter. Hence, the circuit is recompiled at each grid point, and its constraint count is recorded, with trusted-setup time and key sizes measured at the endpoints $N_{\max} = 1$ and $N_{\max} = 5$ (Section 6.4).

On-chain cost (EG3): Verification cost at the governance layer is constant in the size of the circuit, whereas the cost of settling an escalated conflict is linear in the approval threshold k, which the DAO size determines only stepwise. The first is a criterion claim, measured against a proof; the second is a causal claim, measured by deploying an independent governance contract at each of a range of member counts and driving the full resolution path from proposal to approval. Gas is deterministic given the contract bytecode and EVM rules, so both are measured deterministically rather than statistically (Section 6.5).

Latency feasibility (EG4): The interaction latency at issuance and at dispensing is measured end-to-end and judged against the response-time bounds under which interactive systems remain usable. It is decomposed to establish whether the cryptographic and on-chain components dominate it. The method selected is end-to-end timing of the two user-facing actions, decomposed into the contribution slice and the surrounding integration plumbing, and judged against the established limits of 0.1 s for instantaneous, 1 s for uninterrupted, and 10 s for retained attention [12] (Section 6.6). These are generic interaction-response guidelines rather than healthcare-specific usability thresholds. This goal is supplementary: it establishes deployability, not the efficiency of the contribution itself, which EG2 and EG3 measure in isolation.

6.2.2 Evaluation Set up

Table 4 shows the configuration of the evaluation environment. The evaluation stack is brought up from a clean state, and all services ran as containers on a single host; the exact

[4] https://github.com/medalex/ehealth-zkp-prover/blob/0c7b5b82aac9b19c4f9607ba96cae56f4137f6f4/circuits/src/prescription_validation_poseidon_merkle.circom; https://github.com/medalex/ehealth-zkp-prover/blob/0c7b5b82aac9b19c4f9607ba96cae56f4137f6f4/src/prover/prover.service.ts

[5] https://github.com/medalex/mfssia-ehealth/blob/d3a23eecf18ca1b8e2210eba712058250bf1075b/src/providers/dkg/dkg.service.ts; https://github.com/medalex/mfssia-ehealth/tree/d3a23eecf18ca1b8e2210eba712058250bf1075b/src/modules/rx-governance

Table 4 Measurement configuration

Component	Configuration
Host	Apple M1 Pro, 10 cores, 16 GB RAM
Operating system	macOS 26.5.2
Runtime	Node.js 20.20.2
Containers	Docker 20.10.14, Compose v2.5.0
Circuit toolchain	circom 2.1.6, snarkjs 0.7.6, circomlib
Proof system	Groth16 over BN128
Contracts	solc 0.8.26, optimizer on, 200 runs
Chain	ganache, chainId 31338
Reference slots $N_{\max}$	5
Laboratory policy slots N_{LAB}	2
Prescribed items N_{PRESC}	1
Merkle depth	3
Contraindication closure depth	4
Laboratory-record depth	3
Public inputs n_{public}	23
DAO quorum	60%, $k = \lceil 0.6n \rceil$
DAO sizes swept	3–7 members
Circuit sweep (EG2)	$N_{\max} \in \{1, \ldots, 5\}$
Gas measurements	20 executions, distinct proofs
Governance settlements	11 per DAO size
End-to-end latencies	16 executions, first reported separately
Scenario validation	deterministic, one execution

commit of each service is recorded alongside the raw measurements in the evaluation repository, so that every reported figure is traceable to the stack version that produced it.

Circuit compilation and trusted setup run in an *amd64* container under emulation on this host; the resulting times are therefore conservative. These times are reported alongside the constraint counts and key sizes, which the compiler fixes independently of the execution platform. More generally, the measurements differ in how far they generalise beyond this configuration. Gas consumption is fixed by the contract bytecode and the rules of the EVM; constraint counts, and key sizes are fixed by the circuit parameters and the compiler; neither depends on the host. The proof generation time and the end-to-end latencies are characteristic of the configuration above rather than hardware-independent, and the boundaries within which they hold are stated.

6.3 Scenario-Based Evaluation

In this section, we provide a scenario-based evaluation in two steps. First, we evaluate the Zero-Knowledge Proof itself and then examine how the whole flow with the integrated ZKP works. A recorded video of the scenario demonstrations can be found in the GitHub repository[6].

6.3.1 Circuit-level validation (ZKP layer)

Twenty-four scenarios[7] exercise the context-trust, satisfaction, and binding constraints of the circuit, organized into six categories: a baseline completeness test, context-trust soundness, the two satisfaction clauses of Definition 2 taken separately, the properties enforced by the orchestration layer rather than by the circuit, and cryptographic soundness and binding.

Table 5 summarizes the results; all twenty-four passed. Every scenario generates a real Groth16 proof, except those whose expected outcome is that no witness exists. The assertion is on the circuit-computed outcome bit in most cases. The explicit pairing check is demonstrated in C12 - C15, where a proof from C1 (or C3) is verified, tampered, or rekeyed. Scenarios C16 - C18, C20 and C21 assert that witness generation or commitment fails rather than that a proof is rejected. Table 5 defines the scenario ID, description, tested circuit component, expected result and final verdict.

Additional explanations are provided for the complex scenarios. Scenario C4 demonstrates the semantic core of the construction. The patient record carries an allergy to Penicillin, while the prescription is for Amoxicillin. This is a distinct substance identifier that appears nowhere in the patient's records. The rejection therefore cannot follow from a direct comparison of identifiers. Rather, it follows from the subsumption hierarchy of Eq. 4, which the governance process materialises off-chain into a contraindication closure and commits as a Poseidon Merkle tree. The circuit proves membership of the triple $(\mathit{substanceId}, m, \mathit{contraValue})$ in that committed closure rather than consulting a prover-supplied matrix, so a prover cannot assert a fabricated contraindication value: only the leaf carrying the governance-approved value exists in the tree. C4 is thus the scenario that distinguishes semantic verification from arithmetic range checking, and it is the operational counterpart of the ontological subsumption introduced in Section 2.3.

Scenarios C16–C18 show that proving a recorded allergy blocks a prescription is insufficient if the prover can omit the allergy from the witness. The circuit therefore does not accept per-slot activation flags. It derives them from a committed record count and proves that the leaf immediately following the last active slot is the neutral leaf, with the path of that proof bound to the count itself. C16 supplies a patient whose record holds a blocking allergy and claims a count of zero. C17 declares one of three, while C18 supplies a genuine

[6] https://github.com/medalex/ehealth-evaluation/tree/main/scenarios

[7] https://github.com/medalex/ehealth-zkp-prover/blob/main/test/prover.e2e-spec.ts

Table 5 Circuit-level scenario validation results (24/24 passed)

ID	Scenario	Constraint exercised	Expected result	Verdict
Baseline (completeness)				
C1	Valid baseline (Metformin, no conflicts)	Full acceptance path	*outcome* = 1	Pass
Context trust				
C2	Credential absent from registry root	CredValid Merkle membership	*outcome* = 0	Pass
Satisfaction: contraindication				
C3	Allergy to the prescribed drug	Contraindication closure membership	*outcome* = 0	Pass
C4	Subsumption: Penicillin allergy blocks Amoxicillin	Governance-committed subsumption closure	*outcome* = 0	Pass
C5	Allergy to an unrelated drug	Active slot on a committed v=0	*outcome* = 1 (control)	Pass
C16	Claiming no allergies against a blocking record	Count commitment (zero-leaf membership)	No witness	Pass
C17	Understated count (3 allergies, claims 1)	Count commitment (zero-leaf membership)	No witness	Pass
C18	Zero-leaf path at an index other than the count	Path bound to *allergyCount*	No witness	Pass
C19	Patient with no allergies	Zero leaf at index 0	*outcome* = 1 (control)	Pass
C20	More allergies than $N_{\max}$	Capacity guard at commitment time	Loud failure	Pass
Satisfaction: policy axioms				
C6	Dose above the governed limit	Range check on Pol(m, *dose*, $\leq$, θ)	*outcome* = 0	Pass
C7	Dose within the governed limit	Range check on Pol(m, *dose*, $\leq$, θ)	*outcome* = 1 (control)	Pass
C8	eGFR below threshold (20 $<$ 30)	Range check on Pol(m, eGFR, $\geq$, 30)	*outcome* = 0	Pass
C9	eGFR at/above threshold (45 $\geq$ 30)	Range check on Pol(m, eGFR, $\geq$, 30)	*outcome* = 1 (control)	Pass
C11	Lab policy targets Metformin only	Policy scoping by medication	*outcome* = 1 (control)	Pass
C24	Blocking policy listed last	Slot activation *labIsActive* = *labApplies*	*outcome* = 0	Pass
Evidence and policy-vector construction (orchestration layer)				
C10	Most recent result selected (20 masks 45)	Recency selection by *measuredAt*	*outcome* = 0	Pass
C21	More governed policies than N_{LAB}	Capacity guard, not a silent drop	Loud failure	Pass
C22	Reordered input policies	Canonical ordering of the public vector	Identical $\vec{x}$	Pass
C23	Two-sided range on one metric	Total order under the tie-break	Identical $\vec{x}$	Pass
Cryptographic soundness and binding				
C12	Honest proof under governance *vk*	Groth16 verification	Proof accepted	Pass
C13	Tampered public signal (record root)	Public-input binding	Proof rejected	Pass
C14	Forged outcome bit (reject $\rightarrow$ accept)	Outcome-composition equality	Proof rejected	Pass
C15	Verification under a different *vk*	Verification-key binding	Proof rejected	Pass

zero leaf but at a different index than the count. All three fail during witness generation rather than producing a rejecting proof. These are distinguished from C3–C5, where the circuit computes an honest outcome. C18 in particular is the load-bearing case: without binding the path to the count, C16 and C17 would be defeated by proving the neutral leaf at any unused position.

Scenario C10 addresses evidence selection, where several laboratory observations report the same metric for the same patient. However, whichever value enters the witness determines whether the policy is satisfied. The scenario supplies two eGFR observations, an earlier value of 45, which satisfies the governed threshold, and a later value of 20. Selection by observation time is performed by the prover service when it assembles the witness. The circuit carries no

notion of time and cannot enforce it, since establishing that no later observation exists would require a non-membership argument over the committed record. What the circuit enforces is that whichever value was selected is a genuine leaf of *labRecordRoot* and is matched against the policy written for that metric. C10 is, therefore, evidence that the deployed pipeline selects correctly, not that recency is cryptographically guaranteed.

Scenarios C13 and C14 both assert that a proof is rejected, but they defeat different attacks and exercise different bindings. C13 leaves the outcome bit untouched and instead substitutes a fabricated *patientRecordRoot* in the public input vector. Rejection follows from the Groth16 pairing equations being binding over the exact vector, so a proof cannot be detached from the committed governance state under which it was generated. C14 leaves the public roots intact and flips the outcome bit of a proof produced for a contraindicated prescription, claiming acceptance where the circuit computed rejection. Rejection follows from the outcome-composition constraint, which ties the published bit to the value the circuit actually computed. Taken together, the two scenarios establish that neither the evidentiary context of a decision nor the decision itself can be altered after proof generation.

The scenarios marked *control* in Table 5 guard against a trivially sound circuit because it rejects everything. Each pairs a rejection test with a structurally adjacent case that must be admitted, differing in exactly one parameter. C5 differs from C3 only in the substance to which the recorded allergy refers, C9 differs from C8 only in the observed laboratory value, C11 differs from C8 only in the prescribed medication, confirming that a policy axiom binds only the medication it is declared for, and C19 confirms that the count commitment of C16 - C18 does not reject a patient who genuinely has no allergies.

6.3.2 System-level validation (end-to-end)

Table 6 reports the results of the system-level suite,[8] executed against the deployed stack from a clean state. Sc7 is timed with a governed validity window of 60 s, dispensing a proof immediately after issuance and again after the window had elapsed.

Scenarios Sc5 and Sc6 exercise the governance path that the circuit-level suite cannot reach. Sc5 introduces a laboratory reporting an eGFR observation under a unit that no bridge axiom in T relates to the governed scale. Sc6 introduces an allergy annotated under a local vocabulary with no alignment to the governance-approved terminology, which the semantic interpreter reports as a conflict rather than silently resolving. In both cases, a single organization then attempts to publish the missing axiom directly, and the attempt is refused. The alignment becomes part of T only after a DAO proposal reaches the governance-approved threshold, after which the updated theory is re-anchored on the DKG and the corresponding governance vector is republished on-chain. The two scenarios therefore demonstrate the escalation path of the system governance. Operationally, an unresolvable semantic conflict is detected, escalated, and settled by collective decision. These scenarios therefore substantiate **R7**, since no participant below the threshold can unilaterally alter the theory against which every subsequent proof is verified.

Table 6 System-level scenario validation results (8/8 passed)

ID	Scenario	Observed result
Issuance		
Sc1	Valid prescription	Issued
Sc2	Contraindicated prescription	Refused
Sc3	Policy violation	Refused
Access control		
Sc4	Prescriber not in registry	Denied
Semantic governance		
Sc5	Numerical bridge	Refused, then approved by quorum
Sc6	Terminology bridge	Conflict raised, then approved
Dispensing		
Sc7	Proof freshness	Stale proof refused
Sc8	Single-use dispensing	Replay refused

Scenarios Sc7 and Sc8 exercise the two downstream checks that the issuance proof does not demonstrate. Sc7 confirms the off-circuit validity-window check. With a governed window of 60 s, a proof presented immediately after issuance is dispensed, whereas the same construction presented after the window has elapsed is refused by the pharmacy against its own clock. This is because the circuit binds *issuanceTime* and *validFor* without comparing them to the current time. Sc8 confirms anti-replay at the dispensing boundary. A proof that has already been redeemed is refused on a second presentation, since the redemption log rejects a stmtHash it has already recorded. In the current implementation, this redemption log is the same as the on-chain acceptance log, populated at dispensing rather than at issuance. Anti-replay is therefore a consequence of on-chain anchoring occurring exactly once per proof, not a separately maintained structure. This is a distinct mechanism from the per-workflow nonce verified in the circuit, which prevents a proof from being reused across workflow instances but cannot by itself prevent a single valid proof from being dispensed twice.

Table 7 maps the requirements derived in Section 3 onto the scenarios that exercise them. Entries marked *by construc-*

[8] `https://github.com/medalex/ehealth-evaluation/blob/main/bench-correctness/scenarios.mjs`

Table 7 Requirement coverage across the two validation layers

Requirement	Circuit layer	System layer
R1 Correctness	C1, C3–C11, C16–C20, C24	Sc1–Sc3
R2 Context trust	C2	Sc4
R3 Interoperability	C4	Sc5, Sc6
R4 Confidentiality	by construction	—
R5 Replay protection	—	Sc8
R6 Minimal disclosure	by construction	—
R7 Partial collusion	—	Sc5, Sc6
R8 Statement binding	C13–C15	Sc8
R9 Auditability	—	—

Table 8 Circuit size and setup cost as a function of $N_{\max}$

$N_{\max}$	Constraints	Setup (s)	zkey (MB)
1	6,463	144	4.07
2	8,684	—	—
3	10,905	—	—
4	13,126	—	—
5	15,347	260	9.03

tion are established structurally rather than by a scenario: R4 follows from the zero-knowledge property of the proof system, and R6 from the composition of the public input vector (Eq. 18). The two layers of scenario evaluations are complementary rather than redundant. The circuit layer establishes that the encoded conditions evaluate correctly and that a proof cannot be detached from them. The system layer establishes that the generated proofs govern what the deployed workflow actually does.

Three requirements are not attached to any specific scenario. R4 and R6 are not properties of any single execution but of the construction itself. Witness confidentiality follows from the ZKP property and is inherited by every proof generated in Table 5. Minimal disclosure follows from the composition of the public input vector (Eq. 18). R9 is a property of the governance layer's logging discipline rather than of an individual workflow execution; its guarantee is established analytically rather than being exercised by a scenario. Hence, the suites assert the outcome of each decision rather than its subsequent appearance in the acceptance log.

6.4 ZKP Performance Evaluation

The circuit is parameterized by $N_{\max}$, and Eq. 20 predicts that its size grows linearly in that parameter. Table 8 reports the sweep: the circuit was recompiled at each grid point, and a fresh trusted setup was performed, reusing the same powers-of-tau file.

Growth is linear to the constraint, at 2,221 constraints per reference slot over a fixed cost of 4,242 recovered by extrapolation to $N_{\max} = 0$; a linear fit gives $R^2 = 1.000$. The per-slot term covers one Merkle path against the patient's committed record root, together with the comparator that derives the slot's activation from the committed record count. The count commitment contributes both a fixed part, which is why introducing it raised the deployed circuit from 14,567 to 15,347 constraints, an increase of 5.4%. Setup cost and key size follow: trusted setup runs from 144 s at $N_{\max} = 1$ to 260 s at $N_{\max} = 5$, and the proving key from 4.07 to 9.03 MB. These are compile-time costs, covered once per governance-approved circuit revision. Proof size is unaffected; every instantiation produces the same 192 bytes. The evaluation results show that verification gas did not change across the same 5.4% growth.

We now instantiate the construction on the running case. The canonical proof was generated on the released stack with one governed laboratory policy in force. The policy is $\mathrm{Pol}(\mathrm{Metformin}, \mathrm{eGFR}, \geq, 30)$. It was DAO-approved and published to the DKG. The private inputs are

- *workflowId* = 424242
- *prescribedDrugIds* = [105] (Metformin)
- *prescribedDosages* = [500]

In addition there were with the physician's credential hash and its Merkle authentication path of depth 3 against the registry root.

The nonce and the statement commitment are recomputed in-circuit from these values:

$$\mathit{nonce} = \mathrm{Poseidon}(424242) = 20428 \ldots 64894, \tag{21}$$

$$\mathit{stmtHash} = \mathrm{Poseidon}(\mathit{dCH}, \mathit{nonce}, 105, 500, 1786363173, 7776000) = 17153 \ldots 67405, \tag{22}$$

where $\mathit{dCH} = 64575 \ldots 72437$. Both were verified against the public signals of the submitted proof by independent recomputation from the private inputs.

The public input vector submitted to the verifier is given in Table 9. Slot 0 carries the governed eGFR policy: threshold 30, operator code 0 (the safe condition $\geq$), scoped to drug 105, and pinned to the metric identifier id(eGFR). Slot 1 is padding. The operator field admits exactly four values, and the circuit constrains their indicators to sum to one, so a value outside $\{0, 1, 2, 3\}$ makes the witness unconstructible rather than being silently ignored.

The witness was generated by the compiled WebAssembly module *prescription.wasm*, and the proof $\pi = (\pi_a, \pi_b, \pi_c)$ was computed using *snarkjs groth16 prove*. The Groth16 verifier then evaluated the standard pairing equation

$$\begin{aligned} e(\pi_a, \pi_b) = {} & e(\mathsf{vk_alpha_1}, \mathsf{vk_beta_2}) \\ & \cdot e(\mathsf{vk_x}, \mathsf{vk_gamma_2}) \\ & \cdot e(\pi_c, \mathsf{vk_delta_2}). \end{aligned} \tag{23}$$

Table 9 Public input vector of the canonical proof instance

Idx	Signal	Value
0	*outcome*	1
1	*stmtHash*	17153 . . . 67405
2	*issuanceTime*	1786363173
3	*validFor*	7776000
4	*validCredentialRoot*	76608 . . . 12353
5	*patientRecordRoot*	11286 . . . 16792
6	*nonce*	20428 . . . 64894
7–9	*policyDrugIds*	[105, 103, 107]
10–12	*maxDosages*	[65535, 65535, 65535]
13	*contraindicationRoot*	17309 . . . 19849
14	*labRecordRoot*	30652 . . . 32557
15–16	*labThreshold*	[30, 0]
17–18	*labRequiredOp*	[0, 0]
19–20	*labAppliesToDrug*	[105, 0]
21–22	*labMetricIdPub*	[39128 . . . 75201, 0]

Slot 1 of each laboratory array is padding: its operator code of 0 is a fill value, not a second ≥ policy, since the slot's activation flag is zero. *maxDosages* carries the prover's default where no dosage-cap policy is governed in this state.

where $\mathsf{vk}_x = \mathsf{IC}_0 + \sum_{i=1}^{n_{\text{public}}} x_i \mathsf{IC}_i$ is the verification-key linear combination bound to the public input vector x. The pairing check succeeded and the proof was accepted. The proof occupies 192 bytes, corresponding to three BN128 elliptic-curve points, independently of $N_{\max}$.

6.5 On-Chain Cost: Verification and Governance

Verification and governance are measured separately: the first is a claim about a single proof, the second a claim about how settlement scales.

6.5.1 Proof verification

Figure 4 decomposes one *verifyAndRecord* transaction, averaged over 20 executions, each on a distinct freshly generated proof so that the replay guard is respected and every registry write lands in a cold slot. The components sum to the measured total; the unattributed remainder is 0.08% and is reported rather than distributed. Only calldata varies across executions, with the zero-byte composition of each proof; every execution component is constant to the gas.

The breakdown was obtained by instrumenting a mirror of the deployed contracts with stage boundaries, deployed on the same chain and reading the same verifier and governance state. Instrumentation adds 3,529 gas, measured by sending every proof through both paths rather than estimated; the components below are stated against the production total.

Two figures follow from the decomposition. The pairing check accounts for 338,596 gas and did not move by a single unit across the 20 executions, nor between the two circuit revisions reported in Section 6.4, during which the

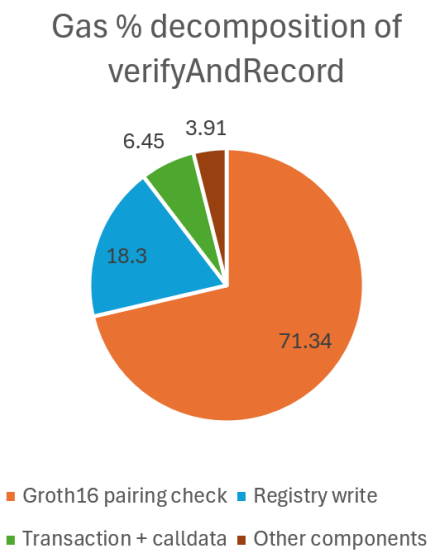


Fig. 4 Gas decomposition of one *verifyAndRecord* (mean of 20 executions)

Table 10 Governance settlement by DAO size (mean of 11 settlements per size)

n	k	**Votes** $2..k-1$	**Settlement**
3	2	—	245,848
4	3	59,909	305,759
5	3	59,909	305,759
6	4	59,909	365,668
7	5	59,909	425,579

Proposal 79,964; first vote 77,009; threshold-crossing vote 88,873. Each tier is constant across all n.

constraint count grew from 14,567 to 15,347. Verification cost is therefore constant in circuit size as a measurement, not as an inference from succinctness. What it does track is the public input vector: hashing w governance words costs $12.06\,w + 705$ gas, linear from $w = 1$ to $w = 32$.

Binding the public inputs to the committed governance state costs 5,401 gas, 1.14% of the full path. The storage read of the governance hash is excluded from that figure: each decision is stamped with that hash for the audit trail, so the read occurs whether or not the comparison is made. Charged together, the two come to 10,621 gas, or 2.24%.

6.5.2 Governance settlement

Settling an escalated conflict costs one proposal plus k votes, where the approval threshold $k = \lceil 0.6\,n \rceil$ follows from the DAO size n. Table 10 reports the sweep over $n \in \{3, \ldots, 7\}$, with 11 settlements per size on distinct proposals.

Votes are not interchangeable. The first costs 77,009 gas, intermediate votes cost 59,909. Incrementing it, the vote that crosses the threshold costs 88,873, flipping the approval flag and emitting the event. Each tier is constant across all n.

Settlement is therefore linear in the threshold,

$$gas = 59{,}910 \cdot k + 126{,}028, \qquad R^2 = 1.000, \tag{24}$$

where the slope is the intermediate tier and the intercept collects the two fixed surcharges. Fitting against membership size instead gives $R^2 = 0.94$: n is the wrong regressor. Because k advances stepwise, $n = 4$ and $n = 5$ settle for identical gas, adding a fifth member costs nothing per settlement. What membership does affect is deployment, at 23,110 gas per member, paid once. The first proposal on a freshly deployed instance costs 97,060 rather than the steady-state 79,964, as the proposal array grows from zero; the two are reported separately because their mean describes neither.

6.6 Latency Feasibility

Table 11 reports end-to-end latency for the two user-facing flows executed 16 times against the canonical Metformin instance of Section 6.4; the first execution is reported separately from the steady-state statistics of the remaining 15, following the same convention already used for the first governance proposal.

Consistent with the deferred-anchoring choice of Section 4.4, issuance produces (π, pub) without an on-chain call; on-chain verification and recording occur when the pharmacy presents (π, pub) at dispensing. Issuance therefore decomposes into the semantic interpretation and DKG queries that precede proof generation, together with witness and proof generation themselves. Dispensing decomposes into the on-chain *verifyAndRecord* call together with the freshness and replay checks.

Issuance of prescription completes with a steady-state median of 11,288.1 ms (min 11,119.1, max 11,528.9, IQR 100.5), exceeding the 10 s retained-attention bound. Figure 5 decomposes this total: witness and proof generation together are 1,848.3 ms (16.4%), a minor share as anticipated. The numeric-bridge normalization against the DKG accounts for 9,276.6 ms (82.2%), confirmed independently at ≈9.3 s over three standalone runs (IQR 24.6 ms). The bottleneck is therefore the semantic-interoperability layer rather than the cryptographic one.

Dispensing of prescription completes with a steady-state median of 1,551.6 ms (min 1,511.9, max 1,821.2, IQR 54.7), within the 10 s bound but above the 1 s uninterrupted-flow threshold. On-chain verification and recording account for 1,476.7 ms (95.2%) of this total. Here the cryptographic and on-chain components are the dominant share, since they are the mechanism by which the deferred anchoring is realized. The Groth16 pairing check itself is 385.0 ms; the remainder is transaction submission and receipt polling against the local development chain. Isolating verification from the log write gives a steady-state median of 463.7 ms (min 418.6, max 623.0, IQR 31.0), within the 1 s bound. The freshness and replay checks together cost 0.5 ms.

Table 11 End-to-end latency by user-facing flow ($n = 16$ executions per flow)

Metric (ms)	Issuance	Dispensing
First run	12,072.6	1,600.1
Steady median	11,288.1	1,551.6
Min	11,119.1	1,511.9
Max	11,528.9	1,821.2
IQR	100.5	54.7

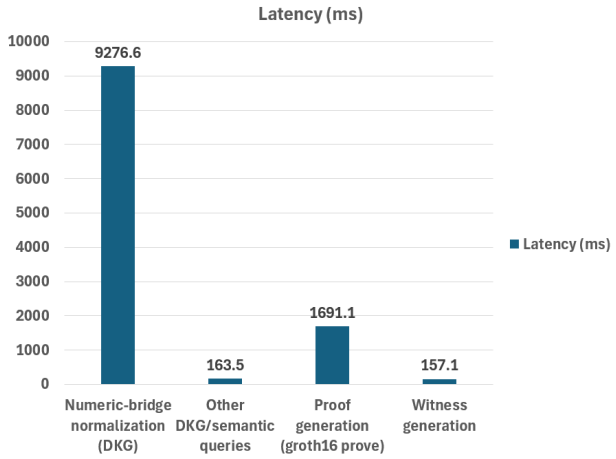


Fig. 5 Latency decomposition of the issuance flow (steady-state median)

First-run overhead is present but not uniform across flows: issuance is 784.5 ms slower on the first execution (12,072.6 ms vs. the steady-state median) than dispensing, which is comparatively unaffected (48.5 ms on the full record path, 122.8 ms on verification in isolation).

7 Discussions

7.1 Security, Privacy, and Trust Analysis

Threat analysis for privacy-preserving collaborative systems extends beyond identifying individual attackers. Privacy-oriented frameworks consider information disclosure, linkability, identifiability, and detectability as distinct privacy threats [15], whereas secure multiparty systems additionally emphasize confidentiality, integrity, and availability under mutually distrustful participants [23]. Decentralized trust-management research further highlights the importance of credentials, policy enforcement, and the distribution of trust across organizational boundaries [9]. Accountability and resistance to colluding or faulty participants are also important where independent organizations collaborate without a single trusted authority [13, 59]. We therefore assess the proposed framework against six properties that recur across

these bodies of work: confidentiality, unlinkability, integrity and authenticity, accountability, resistance to collusion, and availability.

7.1.1 Confidentiality and minimal disclosure

This is one of the strongest properties of the design implemented in this paper. Patient allergies, laboratory measurements, physician credential content, and other clinical evidence remain private witness values, while the verifier receives only the proof and declared public inputs. Under the zero-knowledge property of Groth16 [20], neither the pharmacy nor an observer of the verification transaction needs access to these underlying records. This directly limits disclosure across organizational boundaries and is preferable to a workflow in which the pharmacy receives the patient's complete clinical context. The guarantee is nevertheless limited to information hidden by the circuit interface. Metadata deliberately exposed as public inputs, including record roots, statement commitments, validity information, and the acceptance outcome, remain observable.

7.1.2 Unlinkability and metadata privacy

The design implemented in this paper performs less strongly against the linkability threats emphasized in [15]. Although the patient's identity and clinical values are not public, the `patientRecordRoot` and `labRecordRoot` variables are public proof inputs. If the same record root is used for several prescriptions, an observer can determine that the transactions rely on the same committed record state even without learning the patient's identity. The statement hash, issuance metadata, and transaction timing provide additional correlation points.

7.1.3 Integrity, authenticity, and semantic correctness

The framework performs strongly in binding a decision to the evidence and governance state that were actually evaluated. Merkle membership constrains clinical evidence to committed roots, the statement hash and nonce bind the proof to a workflow instance, and the on-chain acceptance gate verifies both the Groth16 proof and the DAO-approved governance-vector commitment before a decision is recorded. These mechanisms substantially reduce the ability of a dishonest prover to alter evidence or independently select more favourable governance parameters. However, the guarantee is not fully end-to-end. Credential signatures and source attestations are checked when records enter the governed registries rather than inside the proof, and selection of the most recent laboratory observation remains an orchestration-layer operation. Consequently, a valid proof establishes correctness relative to the committed inputs supplied to the circuit; it does not independently establish that no newer clinical record exists. The current off-chain handling of patient-root admission represents a similar trust boundary. These are the principal remaining integrity assumptions of the implementation.

7.1.4 Accountability and auditability

Accountability is particularly important in healthcare, where audit mechanisms complement preventive privacy controls [10] and, more generally, enable the examination of inappropriate information use or decisions after execution [59]. The implementation performs well on this property. A proof-backed decision is written through a single acceptance gate to an append-only registry, together with the statement commitment and the hash of the governance state in force at acceptance. Repeated statement hashes are rejected, and administratively entered decisions are distinguishable from proof-verified decisions. A limitation is that anchoring is deferred until dispensing; a prescription that is issued but never presented on-chain therefore has no governance-level audit record. Auditability is thus strong for completed on-chain interactions but not complete for every issuance event.

7.1.5 Resistance to collusion and governance capture

The k-of-n DAO removes unilateral control over changes to the shared theory, policy commitments, and other governed parameters. This is consistent with decentralized trust-management principles in which authorization depends on explicit credentials and policies rather than trust in one organization [9]. A single participant, or a coalition below the approval threshold, cannot activate a governance update. This protection should not, however, be interpreted as Byzantine fault tolerance in the sense of protocols such as PBFT [13]. Once an adversarial coalition controls the required threshold, it can approve an incorrect semantic bridge or governance state through the legitimate protocol. The implementation therefore provides resistance to unilateral and sub-threshold manipulation, but not to governance capture at or above the voting threshold. Compromise of sufficient member keys has the same consequence.

7.1.6 Availability and robustness

Availability is the weakest of the properties considered here. Secure multiparty systems commonly treat high availability together with confidentiality and integrity as a core requirement [23], while Byzantine-resilient protocols explicitly reason about continued service in the presence of faulty participants [13]. The present framework instead concentrates on semantic correctness, privacy, and governance verifiability. Retry mechanisms mitigate transient DKG failures, but the

protocol provides no replicated availability guarantee for the semantic interpreter, DKG access, prover, or on-chain verification path. Moreover, a governance update requiring k votes cannot complete if too few eligible members are available. Denial-of-service and operational key failure are also outside the current threat model. Thus, the framework can fail closed when required components or governance participants are unavailable, which protects correctness but may interrupt collaboration.

7.2 Privacy-Preserving Collaboration Beyond Healthcare

7.2.1 Transferable Collaboration Use Cases

Although the approach described in the current paper was instantiated for healthcare, the combination of governed semantic knowledge, authenticated input commitments, zero-knowledge verification, and collective governance is not inherently domain-specific. The architecture is most directly transferable to inter-organizational collaboration in which three conditions hold: the participating organizations can agree on a shared semantic model and governance rules, the evidence used for a decision can be authenticated before it enters the shared knowledge layer, and one authorized prover can legitimately assemble the private witness. Examples include cross-organizational qualification or certification checks, regulatory eligibility decisions, and consortium-based compliance assessment. In such settings, a knowledge graph can represent heterogeneous concepts and their agreed mappings, a DAO can govern changes to these mappings or policies, and a ZKP can attest that an entity satisfies the resulting rule without revealing the underlying evidence. Related credential systems similarly show that complex assertions can be proven over authenticated credentials without disclosing their complete contents [48].

7.2.2 Alternative ZKP-enabled Privacy Architectures

The boundary of the current design becomes important when no single participant should possess all private inputs. Consider a financial-market index jointly computed from confidential values supplied by several panel members. Sending the raw contributions to a conventional ZKP prover would reveal them to that prover, which could create confidentiality, commercial, and collusion risks. A better design will use cryptographically masked inputs through secure MPC, allowing the aggregate to be computed without revealing individual contributions. A proof can subsequently attest to the correctness of the computation. Prior work shows both that inter-organizational analytics can be executed with maliciously secure MPC [46] and that mutually distrusting parties can jointly construct zk-SNARKs without first revealing their distributed secrets to one prover [44].

A different adaptation is required for supply-chain proofs such as verifying that a product contains a required proportion of certified or recycled material [55]. The important trust problem here is not about hiding the product material compositions. It is also necessary to verify who originally asserted each data input as goods move through multiple independent custodians in a supply chain network. Allowing one downstream actor (prover) to create and authenticate all input records would weaken this chain of responsibility. Instead, each supplier can issue a signed Verifiable Credential (VC) for the material or transformation for which it is responsible [52], while provenance relations between suppliers, transformations, and resulting products can be represented using a provenance model [36]. The final prover then generates a ZKP over these issuer-authenticated claims, for example proving that the accumulated material composition satisfies a threshold without revealing exact supplier-specific quantities.

8 Conclusion

Heterogeneous healthcare organizations increasingly exchange clinical information across organizational boundaries; yet, structural interoperability does not ensure that the participating systems interpret the information or the validity of a clinical decision consistently. Existing blockchain mechanisms primarily establish agreement on transactions and system state, while privacy-preserving techniques such as zero-knowledge proofs do not by themselves resolve semantic disagreements or determine how unresolved conflicts should be governed. Against this background, this work asked how decentralized e-health workflows can achieve semantic consensus on clinical statements while preserving privacy and enabling verifiable governance decisions. This main question was decomposed into three research questions concerning: (RQ1) the formal requirements for privacy-preserving semantic consensus, (RQ2) the architectural and governance mechanisms required to realize such consensus across heterogeneous organizations, and (RQ3) the cryptographic mechanisms needed to encode and verify semantic validity conditions without disclosing the underlying clinical evidence.

For RQ1, we used model-theoretic semantics to distinguish the shared semantic model M, the governance-approved theory T, and the request-specific evaluation context Γ. From this foundation, the work derived requirements covering statement correctness, context trustworthiness, semantic interoperability, confidentiality, replay protection, minimal disclosure, resistance to unilateral manipulation, statement binding, and auditability. Their combination yields an explicit acceptance relation in which a clinical statement is accepted only when the request context is trusted and the statement is satisfied under the shared model

and governance theory. For RQ2, these requirements were operationalized through a layered architecture separating domain systems, semantic interpretation, proof generation, and on-chain governance. Semantic conflicts that can be resolved using the current theory are handled off-chain through DKG-anchored terminology and numerical bridges, whereas unresolved conflicts are escalated to a k-of-n DAO before an update to the theory can become effective. For RQ3, the acceptance conditions were compiled into a Groth16 zk-SNARK circuit implemented in Circom. Merkle commitments bind credentials, patient records, laboratory evidence, and governance artifacts to the proof, while patient-specific evidence remains private. A statement commitment and nonce bind each proof to a specific workflow instance, and an on-chain verifier checks the proof and governance state before recording the decision.

The implementation was evaluated against four goals. For functional correctness (EG1), all 24 circuit-level scenarios and all 8 end-to-end system scenarios produced the expected results, covering valid and invalid prescriptions, context trust, policy and contraindication checks, proof manipulation, semantic-conflict escalation, freshness, and replay protection. For circuit scalability (EG2), increasing the reference capacity from $N_{\max} = 1$ to $N_{\max} = 5$ increased the circuit from 6,463 to 15,347 constraints, corresponding to 2,221 additional constraints per reference slot with $R^2 = 1.000$. Trusted-setup time and proving-key size increased accordingly, while the Groth16 proof remained constant at 192 bytes. For on-chain cost (EG3), the complete `verifyAndRecord` transaction consumed 474,557 gas, including 338,596 gas for the Groth16 pairing check. Verification cost remained unchanged when the circuit grew from 14,567 to 15,347 constraints, whereas DAO settlement cost increased linearly with the approval threshold k. For interaction latency (EG4), steady-state median latency was 11.29 s for prescription issuance and 1.55 s for dispensing. Importantly, 82.2% of issuance latency resulted from DKG-based numerical normalization, while witness and proof generation together accounted for 16.4%. Isolated verification had a median latency of 463.7 ms. Thus, the principal performance limitation of the current implementation lies in semantic-data retrieval and normalization rather than in zero-knowledge proof generation.

Some limitations still exist in this work. First, the framework was instantiated and evaluated using a single inter-organizational drug-prescription workflow. Although the formal model and architectural separation are intended to be reusable, the present evaluation does not establish that the same requirements, circuit structures, or governance mechanisms generalize without modification to other healthcare workflows. Second, the threat assessment identifies security properties that remain only partially addressed. While the design provides strong confidentiality of clinical witness data, public record commitments and transaction metadata can permit correlation between interactions and therefore do not provide full unlinkability. Evidence freshness and parts of patient-root admission are also enforced outside the circuit, leaving residual trust in the orchestration layer. The DAO prevents unilateral and sub-threshold changes but cannot protect against governance capture when an adversarial coalition controls the approval threshold. Finally, availability and denial-of-service resistance are not provided by the protocol; failure or unavailability of the semantic interpreter, DKG, prover, governance participants, or associated key-management infrastructure can interrupt the workflow even when confidentiality and integrity remain preserved. Although EG4 evaluates technical interaction latency, actual user experience assessment is still missing. No clinicians, pharmacists, healthcare administrators, or other end users participated in the present evaluation.

Declarations

Funding No funding was received to assist with the preparation of this manuscript.

Competing interests The authors have no relevant financial or non-financial interests to disclose.

Ethics approval This study did not involve real patient data or human participants; all clinical values used in the evaluation scenarios (e.g., laboratory observations, prescription records) are synthetic and were generated solely for demonstration purposes. Formal ethics approval was therefore not required.

Consent to participate Not applicable.

Consent for publication Not applicable.

Data availability The synthetic evaluation data and scenario scripts generated and analysed during the current study are archived in Zenodo under DOI `https://doi.org/10.5281/zenodo.22237226` and are also available at `https://github.com/medalex/ehealth-evaluation`. No real patient data were used; all clinical values are synthetic and were generated solely for demonstration purposes.

Code availability The circuit and prover implementation is archived in Zenodo under DOI `https://doi.org/10.5281/zenodo.22237251` (`https://github.com/medalex/ehealth-zkp-prover`), the smart-contract and governance implementation under DOI `https://doi.org/10.5281/zenodo.22237247` (`https://github.com/medalex/ehealth-evm`), and the underlying MFSSIA authentication infrastructure under DOI `https://doi.org/10.5281/zenodo.22237239` (`https://github.com/medalex/mfssia-ehealth`).

Author contributions A.K. conceived the study, developed the formal model-theoretic framework and the semantic model, implemented the proof-of-concept, and wrote the first draft

of the manuscript. C.U. designed and implemented the zero-knowledge proof component, including the zk-SNARK circuit. A.K. and C.U. jointly developed the smart-contract implementation. A.N. contributed to the design of the decentralized governance mechanism. V.D. contributed to the governance framework and critically revised the manuscript for important intellectual content. J.L. designed the evaluation methodology, selected the testing approach, and executed the evaluation runs. All authors commented on previous versions of the manuscript, read and approved the final version, and agree to be accountable for all aspects of the work.

References

1. Al-awamy AA, Al-shaibany N, Sikora A, Welte D (2025) Hybrid consensus mechanisms in blockchain: A comprehensive review. International Journal of Intelligent Systems 2025(1):5821997, DOI 10.1155/int/5821997, article ID 5821997
2. Albalwy F (2026) Zero-knowledge-based policy enforcement for privacy-preserving cross-institutional health data sharing on blockchain. Systems 14(4):385, DOI 10.3390/systems14040385
3. Alder S (2026) Healthcare data breach statistics – updated for 2026. https://www.hipaajournal.com/healthcare-data-breach-statistics/, hIPAA Journal, posted June 18, 2026, statistics last updated August 4, 2026, accessed August 25, 2026
4. Alruwaill MN, Mohanty SP, Kougianos E (2025) hchain 4.0: A permissioned blockchain framework for secure, privacy-preserving, and scalable ehr management. In: 2025 1st International Conference on Secure IoT, Assured and Trusted Computing (SATC), IEEE, pp 1–6
5. Amin MA, Tummala H, Shah R, Ray I (2024) Balancing patient privacy and health data security: The role of compliance in protected health information (PHI) sharing. arXiv preprint arXiv:240702766 Available at https://arxiv.org/abs/2407.02766
6. Andrew J, Isravel DP, Sagayam KM, Bhushan B, Sei Y, Eunice J (2023) Blockchain for healthcare systems: Architecture, security challenges, trends and future directions. Journal of Network and Computer Applications 215:103633, DOI 10.1016/j.jnca.2023.103633
7. Angurala M, Singh S, Rajpoot NK, Rattan S (2025) The transformative role of blockchain in healthcare privacy and compliance. In: Securing Healthcare: Leveraging Blockchain for Data Integrity in Healthcare Services, Bentham Science Publishers, chap 10, pp 211–231, DOI 10.2174/9789815238433125010014
8. Babu SB, Jothi K (2024) A secure framework for privacy-preserving analytics in healthcare records using zero-knowledge proofs and blockchain in multi-tenant cloud environments. IEEE Access
9. Blaze M, Feigenbaum J, Lacy J (1996) Decentralized trust management. In: Proceedings of the 1996 IEEE Symposium on Security and Privacy, IEEE, pp 164–173, DOI 10.1109/SECPRI.1996.502679
10. Blocki J, Christin N, Datta A, Sinha A (2011) Audit mechanisms for privacy protection in healthcare environments. In: 2nd USENIX Workshop on Health Security and Privacy (HealthSec 11), USENIX Association, San Francisco, CA
11. Cabello C, Mikalef P (2024) Exploring decentralized autonomous organization (DAO) governance: An integrative literature review. In: Proceedings of the 18th Mediterranean Conference on Information Systems (MCIS 2024), URL https://aisel.aisnet.org/mcis2024/27/
12. Card SK, Robertson GG, Mackinlay JD (1991) The information visualizer, an information workspace. In: Proceedings of the SIGCHI Conference on Human factors in computing systems, pp 181–186
13. Castro M, Liskov B (1999) Practical byzantine fault tolerance. In: Proceedings of the Third Symposium on Operating Systems Design and Implementation (OSDI '99), USENIX Association, New Orleans, LA, pp 173–186
14. Chaudhry JA, Saleem K, Alazab M, Zeeshan HMA, Al-Muhtadi J, Rodrigues JJPC (2021) Data security through zero-knowledge proof and statistical fingerprinting in vehicle-to-healthcare everything (v2hx) communications. IEEE Transactions on Intelligent Transportation Systems 22(6):3869–3879, DOI 10.1109/TITS.2021.3066487
15. Deng M, Wuyts K, Scandariato R, Preneel B, Joosen W (2011) A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements. Requirements Engineering 16(1):3–32, DOI 10.1007/s00766-010-0115-7
16. Elvas LB, Serrão C, Ferreira JC (2023) Sharing health information using a blockchain. Healthcare 11(2):170
17. Evans D, Kolesnikov V, Rosulek M (2018) A pragmatic introduction to secure multi-party computation. Foundations and Trends in Privacy and Security 2(2-3):70–246, DOI 10.1561/3300000019
18. Gazzarata R, Almeida J, Lindsköld L, Cangioli G, Gaeta E, Fico G, Chronaki CE (2024) Hl7 fast healthcare interoperability resources (hl7 fhir) in digital healthcare ecosystems for chronic disease management: Scoping review. International journal of medical informatics 189:105507
19. Gentry C (2009) A fully homomorphic encryption scheme. PhD dissertation, Stanford University, available at https://crypto.stanford.edu/craig/craig-thesis.pdf

20. Groth J (2016) On the size of pairing-based non-interactive arguments. In: Advances in Cryptology – EUROCRYPT 2016, Springer, Lecture Notes in Computer Science, vol 9666, pp 305–326, DOI 10.1007/978-3-662-49896-5_11
21. Healthcare IT News (2024) Top 15 largest u.s. healthcare provider data breaches of 2024. URL https://www.healthcareitnews.com/news/top-15-largest-us-healthcare-provider-data-breaches-2024, accessed: 2025-01-30
22. Hodges W (1997) A Shorter Model Theory. Cambridge University Press
23. Howard H, Alder F, Ashton E, Chamayou A, Clebsch S, Costa M, Delignat-Lavaud A, Fournet C, Jeffery A, Kerner M, Kounelis F, Kuppe MA, Maffre J, Russinovich M, Wintersteiger CM (2023) Confidential consortium framework: Secure multiparty applications with confidentiality, integrity, and high availability. Proceedings of the VLDB Endowment 17(2):225–240, DOI 10.14778/3626292.3626304
24. IBM Security, Ponemon Institute (2024) Cost of a data breach report 2024. Tech. rep., IBM Corporation, URL https://www.ibm.com/reports/data-breach, 19th annual report
25. Inker LA, Eneanya ND, Coresh J, Tighiouart H, Wang D, Sang Y, Crews DC, Doria A, Estrella MM, Froissart M, Grams ME, Greene T, Grubb A, Gudnason V, Gutiérrez OM, Kalil R, Karger AB, Mauer M, Navis G, Nelson RG, Poggio ED, Rodby R, Rossing P, Rule AD, Selvin E, Seegmiller JC, Shlipak MG, Torres VE, Yang W, Ballew SH, Couture SJ, Powe NR, Levey AS (2021) New creatinine- and cystatin c–based equations to estimate gfr without race. New England Journal of Medicine 385(19):1737–1749, DOI 10.1056/NEJMoa2102953, URL https://www.nejm.org/doi/full/10.1056/NEJMoa2102953, https://www.nejm.org/doi/pdf/10.1056/NEJMoa2102953
26. Jha V, Maurya A, Baghel V, Vathkar S (2025) Blockchain empowered personal health records: Enhancing security, privacy, and interoperability in healthcare management. Privacy, and Interoperability in Healthcare Management (April 14, 2025)
27. Journal H (2024) The biggest healthcare data breaches of 2024. URL https://www.hipaajournal.com/biggest-healthcare-data-breaches-2024/, accessed: 2025-01-30
28. Khan AA, Laghari AA, Almansour H, Kumar T, Jaghdam IH, Hajjej F, Mohamed MA (2025) Leveraging blockchain with zero knowledge proofs in wearable health technologies for personalized healthcare. Scientific Reports 15(1):45672
29. Kormiltsyn A, Udokwu C, Dwivedi V, Norta A, Nisar S (2023) Privacy-conflict resolution for integrating personal and electronic health records in blockchain-based systems. Blockchain in Healthcare Today 6:10–30953
30. Kormiltsyn A, Norta A, Nisar S, Dwivedi V (2024) Preventing data-security breaches and patient-safety risks in cross-blockchain e-healthcare systems. In: Chbeir R, Benslimane D, Zervakis M, Manolopoulos Y, Ngyuen NT, Tekli J (eds) Management of Digital EcoSystems, Springer Nature Switzerland, Cham, pp 41–54
31. Kormiltsyn A, Avanzo S, Dwivedi V, Norta A, Schifanella C (2025) Requirements for decentralized consensus in e-health workflows: Towards trustworthy interorganizational coordination. In: Proceedings of the 2025 International Conference on Information Technology for Social Good, pp 204–212
32. Kormiltsyn A, Iqbal M, Dwivedi V, Norta A, Matulevičius R (2026) MFSSIA: Ontology of blockchain-based multi-factor challenge set self-sovereign identity authentication. Telematics and Informatics Reports 22:100322, DOI 10.1016/j.teler.2026.100322
33. Kroop S (2025) Artifact validity in design science research (dsr): A comparative analysis of three influential frameworks. In: Local Solutions for Global Challenges: 20th International Conference on Design Science Research in Information Systems and Technology, DESRIST 2025, Montego Bay, Jamaica, June 2–4, 2025, Proceedings, Part I, Springer-Verlag, Berlin, Heidelberg, pp 199–215, DOI 10.1007/978-3-031-93976-1_13, URL https://doi.org/10.1007/978-3-031-93976-1_13
34. Larsen KR, Lukyanenko R, Mueller RM, Storey VC, Parsons J, VanderMeer D, Hovorka DS (2025) Validity in design science. MIS quarterly 49(4):1267–1294
35. Lavin R, Liu X, Mohanty H, Norman L, Zaarour G, Krishnamachari B (2024) A survey on the applications of zero-knowledge proofs. arXiv preprint arXiv:240800243 Available at https://arxiv.org/abs/2408.00243
36. Lebo T, Sahoo S, McGuinness D (2013) PROV-O: The PROV ontology. W3c recommendation, World Wide Web Consortium, URL https://www.w3.org/TR/prov-o/
37. Massari A (2023) Representing provenance and track changes of cultural heritage metadata in RDF: a survey of existing approaches. arXiv preprint arXiv:230508477 DOI 10.48550/arXiv.2305.08477, URL https://arxiv.org/abs/2305.08477
38. Mateus S, Sarkar S (2023) Can decentralized autonomous organizations (daos) revolutionize healthcare. California Management Review 65(1)
39. Myeong GE, Ram KS (2025) Blockchain based zero knowledge proof protocol for privacy preserving healthcare data sharing. Journal of Technology Informatics and Engineering 4(1):171–189, DOI 10.51903/jtie.v4i1.296

40. Nasir NM, Hassan S, Zaini KM (2024) Securing permissioned blockchain-based systems: An analysis on the significance of consensus mechanisms. IEEE Access 12:138211–138238, DOI 10.1109/ACCESS.2024.3468902, published online October 2024
41. Nawab F, Sadoghi M (2023) Consensus in data management: From distributed commit to blockchain. Found Trends Databases 12(4):221–364, DOI 10.1561/1900000075
42. Norta A, Kormiltsyn A, Udokwu C, Dwivedi V, Aroh S, Nikolajev I (2022) A blockchain implementation for configurable multi-factor challenge-set self-sovereign identity authentication. In: 2022 IEEE International Conference on Blockchain (Blockchain), IEEE, pp 455–461
43. Oude Roelink B, El-Hajj M, Sarmah D (2024) Systematic review: Comparing zk-SNARK, zk-STARK, and bulletproof protocols for privacy-preserving authentication. Security and Privacy 7(1):e401, DOI 10.1002/spy2.401
44. Ozdemir A, Boneh D (2022) Experimenting with collaborative zk-SNARKs: Zero-Knowledge proofs for distributed secrets. In: 31st USENIX Security Symposium (USENIX Security 22), USENIX Association, Boston, MA, pp 4291–4308
45. Patel-Schneider PF, Hayes PJ (2026) RDF 1.2 Semantics. W3C Candidate Recommendation Snapshot, https://www.w3.org/TR/rdf12-semantics/
46. Poddar R, Kalra S, Yanai A, Deng R, Popa RA, Hellerstein JM (2021) Senate: A Maliciously-Secure MPC platform for collaborative analytics. In: 30th USENIX Security Symposium (USENIX Security 21), USENIX Association, pp 2129–2146, URL https://www.usenix.org/conference/usenixsecurity21/presentation/poddar
47. Raghav A, Tripathi AM, Wani NA, Ahmad N, Javed Y, Ladan M (2026) Secure, scalable, and interoperable healthcare data exchange using layer-2 zk-rollups, smart contracts, and ipfs. Scientific Reports
48. Rosenberg M, White J, Garman C, Miers I (2023) zkcreds: Flexible anonymous credentials from zkSNARKs and existing identity infrastructure. In: 2023 IEEE Symposium on Security and Privacy (SP), IEEE, pp 790–808, DOI 10.1109/SP46215.2023.10179430
49. Sakwa CO, Anyembe AO, Li F (2026) A survey of folding-based zero-knowledge proofs. Information Sciences 724:122698, DOI https://doi.org/10.1016/j.ins.2025.122698, URL https://www.sciencedirect.com/science/article/pii/S002002552500831X
50. Schäffer T, Leyh C (2026) Holistic evaluations in design science research: A framework for systematic selection of evaluation method and effective artifact evaluation. In: Ziemba EW, Chmielarz W, Wątróbski J (eds) Information Technology for Management: Intelligent Alignment of IT with Business and Society, Springer Nature Switzerland, Cham, pp 155–175
51. da Silva Vanin FN, Melo B, da Rosa Righi R, da Costa CA, Antunes RS (2025) Enhancing health data integrity through on-chain verification and cryptographic proofs. Blockchain: Research and Applications p 100373
52. Sporny M, Longley D, Chadwick D, Herman I (2025) Verifiable credentials data model v2.0. W3c recommendation, World Wide Web Consortium (W3C), URL https://www.w3.org/TR/vc-data-model-2.0/
53. Tarski A (1944) The semantic conception of truth: and the foundations of semantics. Philosophy and phenomenological research 4(3):341–376
54. Udokwu C (2025) Zero knowledge proof solutions to linkability problems in blockchain-based collaboration systems. Mathematics 13(15):2387
55. Udokwu C, Craß S (2026) Zero-knowledge proof extensions for digital product passports in sustainability claims reporting and verifications. Electronics 15(4):745
56. Udokwu C, Kormiltsyn A, Thangalimodzi K, Norta A (2018) The state of the art for blockchain-enabled smart-contract applications in the organization. In: 2018 Ivannikov Ispras Open Conference (ISPRAS), IEEE, pp 137–144
57. Van Woensel W, Seneviratne O (2025) Semantic interoperability on blockchain by generating smart contracts based on knowledge graphs. Blockchain: Research and Applications p 100320
58. Villarreal ERD, García-Alonso J, Moguel E, Alegría JAH (2023) Blockchain for healthcare management systems: A survey on interoperability and security. IEEE Access 11:5629–5652, DOI 10.1109/ACCESS.2023.3236370
59. Weitzner DJ, Abelson H, Berners-Lee T, Feigenbaum J, Hendler JA, Sussman GJ (2008) Information accountability. Communications of the ACM 51(6):82–87, DOI 10.1145/1349026.1349043
60. Xie T (2023) Advances in zero-knowledge proofs: Bridging the gap between theory and practices. PhD thesis, UC Berkeley
61. Yakubu MM, Hassan FB, Danyaro KU, Junejo AZ, Siraj M, Yahaya S, Adamu S, Abdulsalam K, et al (2024) A systematic literature review on blockchain consensus mechanisms' security: Applications and open challenges. Computer Systems Science & Engineering 48(6)

Notations reference

Table 12 Consolidated Notation Reference

Symbol	Reference	Description
Semantic Model and Domain		
$M = \langle D, I \rangle$	Eq. 1	Shared semantic model: domain D and interpretation function I
$T = \{\alpha_1, \ldots, \alpha_n\}$	Eq. 2	Governance-approved theory: structural axioms, alignment bridges, clinical policy axioms
$M \models T$	§ 2.3	Model M satisfies all axioms in theory T (admissibility precondition)
$M, \Gamma \models \varphi$	Def. 2	Extended Tarski satisfaction (partial relation: defined only when ContextTrusted(Γ, M))
φ; φ^+	§ 3, 4.2	Statement schema; ground instantiation with concrete identifiers
$p \in D_P$; $s \in D_S$; $m \in D_M$; $d \in D_D$; $r \in D_R$	Eq. 1	Typed entities in D: patient; substance/drug class; medication; prescriber; clinical record entity
Labeled Facts in M (RDF-style entity-attribute structure)		
$I(isAllergy)$, $I(isLabResult)$	§ 2.3	Type-membership relations identifying labeled record entities in M
$I(hasPatient)$, $I(hasSubstance)$, $I(hasMetric)$, $I(hasValue)$	§ 2.3	Attribute relations linking record entities to clinical content
$I(hasPatientAttr)$	Eq. 11	Ternary relation linking patient to attribute name and value (age, weight, pregnancy status)
$I(hasSource)$	§ 2.3	Provenance attribute attached to each labeled record
Evaluation Context and Trust Predicates		
$\Gamma = \langle cred, refs, t_{req} \rangle$	Eq. 3	Evaluation context: credential, references to records in M, request timestamp
CredValid$(cred)$	Def. 1	Credential is included in the governance-approved credential registry committed in $root_M$
ContextTrusted(Γ, M)	Def. 1	Context-trust predicate; holds iff CredValid$(cred)$
Accept(M, Γ, φ)	Def. 3	Governance acceptance: ContextTrusted$(\Gamma, M) \wedge M, \Gamma \models \varphi$
Clinical Predicates		
Allergy(p, s)	Def. 2	Labeled allergy record links patient p to substance s
Contraindication(p, m)	Def. 2	Medication m contraindicated for p (via allergy + subsumption)
PolicySatisfied(p, m, n)	Def. 2	All policy axioms for m satisfied by p's evidence (lab results, patient attributes) and prescribed dose n
RxValid(d, p, m, n)	Def. 2	Prescription validity: ¬Contraindication ∧ PolicySatisfied
Pol(m, t, op, θ)	Eq. 6	Clinical policy axiom in T: medication m subject to t op θ
PolicyAxioms(m, T)	Def. 2	Set of policy triples (t, op, θ) in T applicable to m
BSA	§ 2.3	Patient body surface area in m^2; divisor of the numerical bridge normalizing absolute eGFR onto the 1.73 m^2 reference scale
$\sqsubseteq$	Eq. 4	Subsumption relation in the ontology (governance-approved hierarchy in T)
Cryptographic and Circuit Notation		
π; vk	§ 5	Zero-knowledge proof; verification key
$w = \langle w_{\text{cred}}, w_{\text{refs}}, w_{\text{stmt}}, w_{\text{ont}}, w_{\text{aux}} \rangle$	Eq. 17	Private witness: credential, references, statement, ontology, auxiliary
$pub = \langle outcome, stmtHash, issuanceTime, validFor, validCredentialRoot, patientRecordRoot, nonce, policyDrugIds, maxDosages, contraindicationRoot, labRecordRoot, labThreshold, labRequiredOp, labAppliesToDrug, labMetricIdPub \rangle$	Eq. 18	Public input vector for on-chain verification ($n_{\text{public}} = 23$)
$stmtHash = \text{Poseidon}(dCH, nonce, prescribedDrugIds, prescribedDosages, issuanceTime, validFor)$	Eq. 19	In-circuit commitment binding prescription content, nonce, and validity window
$outcome \in \{0, 1\}$	§ 5	Accept (1) or reject (0) decision produced by the circuit
$C(w, x)$	§ 5	Arithmetic circuit: outputs 1 iff semantic constraints hold
id$(\cdot)$; $H(\cdot)$	§ 5	Identifier encoding; hash function