Asuu Network: A Sovereign Protocol for Accountable Clinical Artificial Intelligence

Intelligentia Soverigna, Cura Sacra.

Derviş Berkcan Özdemir, MD, cand.
berkcan@asuucare.com
www.asuucare.com

Abstract

We propose the Asuu Network, a sovereign peer-to-peer protocol that serves as the verifiable control plane for clinical artificial intelligence. The protocol replaces institutional trust with cryptographic proof, eliminating every dependency on proprietary vendors, foreign cloud infrastructure, and centralized authorities. Patient data physically never touches the chain. Under the Data Embassy model, the network carries only hashes, zero-knowledge proofs, and workflow gating, anchoring every AI event to standards-native HL7 FHIR bundles with CDS Hooks, TEFCA, and ONC HTI-1 alignment. The protocol strictly separates hard protocol guarantees from statistical bounds, framing the envelope of safety as the theorem and the neural network as the variable within it. A Proof-of-Clinical-Authority consensus mechanism binds validator eligibility to verifiable healthcare credentials, producing weighted Byzantine Fault Tolerant finality with a proven cross-epoch safety guarantee and a phased post-quantum migration path. At the governance layer, a versioned on-chain Clinical AI Passport replaces static model cards, enforcing the FDA Predetermined Change Control Plan corridor as a provable protocol invariant against validator-initiated changes while preserving the sponsor's and competent regulator's ability to direct suspension or corridor modification through a dedicated Regulatory Override Channel. A Simplex architecture with temporal logic specifications guarantees human oversight as a hard protocol property, and automated suspension triggers impose a bounded detection delay on calibration drift. A seven-tier lattice of clinical authority, enforced through ciphertext-policy attribute-based encryption and W3C Verifiable Credentials v2.0 with BBS+ selective disclosure, makes role-based access control architectural rather than administrative. Algorithmic liability is resolved through Shapley-value decomposition over the immutable audit trail. Five theorem families form the mathematical backbone: cross-epoch safety, end-to-end provenance completeness, update-corridor soundness, privacy-budget enforcement, and safety-suspension bounds. Institutions join and leave at will, accepting the finalized chain as proof of what occurred in their absence. The result is a medical AI infrastructure that no single entity owns, no single government controls, and no single point of failure can disable.

Contents

I The Clinical Execution Fabric 7

1 FHIR-Native Interoperability 8
1.1 Inference Record to FHIR Provenance Mapping . . . 8
1.2 CDS Hooks Integration . . . 9
1.3 EHR Vendor Integration Architecture . . . 9

2 Prior Authorization: The First Institutional Workflow 12
2.1 Honest Incentive Analysis: Why a Payer Would and Would Not Adopt . . . 12
2.2 Workflow Specification . . . 13
2.3 Automata-Theoretic Formalization . . . 13

II Core Protocol & Cryptographic Foundations 17

3 Introduction 18

4 Notation and Cryptographic Preliminaries 20
4.1 Notation . . . 20
4.2 Hash Function Security . . . 20
4.3 Digital Signature Security . . . 20
4.4 Commitment Schemes . . . 21
4.5 Zero-Knowledge Proof Systems . . . 21
4.6 Hardness Assumptions . . . 21

5 Medical Inference Records 23
5.1 Formal Specification . . . 23

6 Timestamp Server and Hash Chain 25
6.1 Formal Specification . . . 25

7 Merkle Tree Integrity 27

8 Post-Quantum Cryptographic Migration 29
8.1 Formal Specification . . . 30

III Consensus, Network Operations & Economics 32

9 Proof of Clinical Authority 33
9.1 Validator Model . . . 33
9.2 Sybil Resistance . . . 39
9.3 Safety . . . 41
9.4 Liveness . . . 42

10 Network 43
10.1 Protocol Security . . . 43
10.2 Network Partition Recovery . . . 44

11 Incentive 46
11.1 Validator Payoff Model . . . 46
11.2 Data Contribution Economy . . . 48

12 Reclaiming Resources **50**
12.1 Formal Guarantees . . . 50
13 Simplified Inference Verification **51**
13.1 Formal Guarantees . . . 51
14 Security Analysis **52**
14.1 BFT Safety and Finality Analysis . . . 52
14.2 Economic Security Under BFT Consensus . . . 53
14.3 Mitigations for Large-Adversary Scenarios . . . 54
14.4 Selfish Mining Resistance Under PoCA . . . 57
15 Token Economics **59**
15.1 Emission Model . . . 59
15.2 Lyapunov Stability of Token Supply . . . 59
15.3 Staking Game Equilibrium . . . 60
15.4 Network Bootstrap and the Cold-Start Problem . . . 60
15.5 The Validation Perimeter Problem: Permissioned Scope for FDA-Regulated Devices . . . 63
15.6 Clinical AI Passport SaMD Classification Analysis . . . 66
15.7 Role-Weighted Shapley Value Tokenomics . . . 69
15.8 Break-Glass Token Economics . . . 69

IV Data Sovereignty, Privacy & Sovereign Identity 71

16 Data Sovereignty and Privacy **72**
16.1 The Data Embassy Model . . . 72
16.2 Formal Guarantees . . . 73
16.3 Deployment Reality: Data Embassies for SaaS Hospitals . . . 73
17 Self-Sovereign Health Identity **76**
17.1 Formal Guarantees . . . 76
17.2 W3C Verifiable Credentials v2.0 Integration . . . 77
18 Exfiltration Resistance **79**
18.1 TEE Model . . . 79
18.2 Information-Theoretic Leakage Bounds . . . 79
18.3 Degraded-Mode Operation Under TEE Compromise . . . 79
18.4 Geo-Fencing Enforcement . . . 82
19 Genomic Data Sovereignty **83**
19.1 Formal Guarantees . . . 83
20 Differential Privacy for Population Health Analytics **86**
20.1 Formal Differential Privacy Framework . . . 86
20.2 Composition and Budget Management . . . 87
20.3 Healthcare-Specific Privacy Mechanisms . . . 87
21 Proxy Re-Encryption for Delegated Medical Data Sharing **89**
21.1 Formal PRE Specification . . . 89
21.2 Blockchain-Integrated Delegation Protocol . . . 90

V AI Model Infrastructure & Governance 92

22 Federated Learning and Decentralized Compute **93**
22.1 Protocol Specification . . . 93
22.2 Privacy Guarantees . . . 93
22.3 Rényi Differential Privacy Accountant . . . 94
22.4 Integrity Guarantees . . . 95
22.5 Byzantine-Robust Convergence . . . 95
23 Model Governance **96**
23.1 Formal Guarantees . . . 96
23.2 The Clinical AI Passport . . . 97
24 Homomorphic Medical Computing **100**
24.1 Formal Specification . . . 100
25 Formal Verification of Clinical AI Safety **102**
25.1 Formal Specification . . . 102
25.2 Automated Suspension Under Calibration Drift . . . 104
25.3 EU AI Act Article 14 Formalization . . . 105
26 AI Liability Attribution **107**
26.1 Formal Specification . . . 107
26.2 Limits of Shapley Attribution Under Correlated Failure . . . 108
26.3 Role-Parameterized Liability Attribution . . . 111
27 Decentralized Federated Pharmacogenomics **112**
27.1 Pharmacogenomic Bounty Protocol . . . 112
27.2 Privacy-Preserving Inference Protocol . . . 112
27.3 Security and Privacy Guarantees . . . 113
27.4 Multi-Institutional Federated Aggregation . . . 113

VI Clinical Workflows & Access Control **115**

28 Healthcare Workforce Role-Based Access Control **116**
28.1 Motivation and Scope . . . 116
28.2 Role Taxonomy and Classification Anchoring . . . 116
28.3 Security Lattice Construction . . . 119
28.4 NIST $RBAC_3$ Formalization . . . 119
28.5 Lattice-to-CP-ABE Cryptographic Mapping . . . 120
28.6 BBS+ Role Credential Encoding . . . 121
29 Separation of Duty Enforcement **123**
29.1 Motivation . . . 123
29.2 Formal SoD Definitions . . . 123
29.3 Smart Contract Enforcement . . . 124
30 Clinical Workflow Formalization **125**
31 Delegated Clinical Authority **127**
31.1 Break-Glass Emergency Access . . . 127
32 Human Authority Preservation **129**
32.1 Regulatory Foundation . . . 129
32.2 Levels of Autonomy Mapping . . . 129
32.3 Temporal Logic Specifications . . . 130
32.4 Game-Theoretic Authority Model . . . 131

33 Transient Cross-Border Clinical Governance **132**
33.1 Transient Authority Credential . . . 132
33.2 Cross-Jurisdictional Role Mapping . . . 132
33.3 Co-Signature Protocol for Supervised Practice . . . 133
34 Dynamic Competency and Educational Provenance **134**
34.1 Competency State Machine . . . 134
34.2 Smart Contract Implementation . . . 134
34.3 Cross-Jurisdictional Competency Portability . . . 135

VII Healthcare Ecosystem Applications & Integrations 136

35 Pharmaceutical Supply Chain Integrity **137**
35.1 Formal Guarantees . . . 137
36 Clinical Trial Integrity **139**
36.1 Formal Guarantees . . . 139
37 Smart Contracts for Insurance and Claims **141**
37.1 Automata-Theoretic Formalization . . . 141
37.2 Incentive Compatibility . . . 142
37.3 Clinical Pathway Compilation . . . 142
38 Decentralized Pharmacovigilance **144**
38.1 Signal Detection Formalization . . . 144
39 Privacy-Preserving Pandemic Preparedness **146**
39.1 Formal Specification . . . 146
40 One Health Cross-Species Surveillance **148**
40.1 Formal Specification . . . 149

VIII Interoperability, Edge Systems & Network Resolution 150

41 Data Valuation and Fair Compensation **151**
41.1 Formal Specification . . . 151
41.2 Role-Specific Data Contribution Asymmetry . . . 152
42 Regulatory Compatibility **153**
42.1 EU AI Act Article 14: Human Oversight . . . 153
42.2 EU Product Liability Directive (PLD 2024/2853) . . . 154
42.3 FDA Predetermined Change Control Plans (PCCP) . . . 154
42.4 MDR Notified-Body Pathway and Post-Market Surveillance . . . 155
42.5 On-Chain Immutability and Record-Correction Rights . . . 157
42.6 Cross-Regulatory Synthesis and Capability Matrix . . . 160
43 Cross-Chain Health Data Interoperability **163**
43.1 Cross-Chain Credential Exchange . . . 163
43.2 Relay-Based Cross-Chain Verification . . . 163
43.3 2025–2026 Regulatory Developments . . . 164
44 Clinical UI/UX and Abstraction Layer **166**
44.1 Formal Gateway Architecture . . . 166
44.2 Zero-Knowledge Session Authentication . . . 167
44.3 Latency Guarantees and Clinical Responsiveness . . . 168

45 Edge-AI and Medical IoT Oracles **170**
45.1 Device Attestation Model . . . 171
45.2 Biometric Data Attestation Protocol . . . 172
45.3 Lightweight Relay to Asuu Chain . . . 173
46 On-Chain Medical Arbitration and Dispute Resolution **174**
46.1 Dispute Initiation and Jury Selection . . . 174
46.2 Blinded Evidence Access via Zero-Knowledge . . . 174
46.3 Schelling-Point Voting Mechanism . . . 175
46.4 Appeal Mechanism . . . 175
47 High-Frequency Intraoperative State Channels **176**
47.1 State Channel Formalization . . . 176
47.2 Channel Lifecycle . . . 177
47.3 Security Properties . . . 177
47.4 Legal Risk and Deployability Scope . . . 178
48 Conclusion **180**
48.1 The Four Architectural Pillars . . . 180
48.2 The Five Theorem Families . . . 181
48.3 Why the Deliberate Structure is the Source of Robustness . . . 181
48.4 Machine-Checked Verification Roadmap . . . 182
48.5 Liability Reserve Architecture . . . 182
48.6 Protocol Composition Security . . . 182
48.7 Deployment Endpoints . . . 184
48.8 Closing . . . 184

PART I

THE CLINICAL EXECUTION FABRIC

1 FHIR-Native Interoperability

A protocol that cannot speak the language of existing hospital systems is a protocol that will never be deployed. The Asuu network is not an island. It is a cryptographic layer designed to sit invisibly within the interoperability stack that hospitals already use. Every AI event on the network maps to standards-native HL7 FHIR bundles, ensuring that the protocol's outputs are immediately consumable by any EHR system, health information exchange, or payer platform that conforms to US Core, TEFCA, or international FHIR profiles.

Two pragmatic constraints govern the design. First, the primary wire format is FHIR R4 (v4.0.1), because the US Core Implementation Guides currently mandated under ONC certification are R4-based, and more than 96% of U.S. hospitals and 78% of office-based physicians use certified health IT built on this version. The mapping layer is forward-compatible with R5 resources where they are available, but R4/R4B conformance is the shipping baseline. Second, clinical decision support hooks follow the CDS Hooks 2.0.1 specification: an Asuu-aware EHR surfaces AI recommendations through standard hook events (order-select, patient-view, encounter-start), and the response cards carry both the clinical payload and a reference to the on-chain inference record hash for auditability.

1.1 Inference Record to FHIR Provenance Mapping

The mapping is structural, not superficial. An Asuu inference record (Definition 5.1) maps to a conformant FHIR Provenance resource with the following field-level correspondences.

Table 1: Field-level mapping from Asuu inference record to FHIR R4 Provenance resource.

Asuu Field	FHIR R4 Element	Semantics
$h_{\text{in}} = H(D_{\text{in}})$	Provenance.entity[source].what	Hash of clinical input
$mid = (M_{\text{id}}, H(M_w))$	Provenance.agent[assembler]	Model identity and version
$h_{\text{out}} = H(D_{\text{out}})$	Provenance.target	Hash of inference output
pk_I	Provenance.agent[custodian]	Institutional identity
t	Provenance.recorded	UTC timestamp
σ	Provenance.signature	FHIR Signature datatype
$H(B_k)$	Provenance.location	Block hash anchor
Λ_{current}	Provenance.activity	Autonomy level code

The clinical context (the actual patient data, images, and reports) stays off-chain at the Data Embassy. The chain anchors only the signed hash of the FHIR evidence bundle, the model version, the human approval state, and the governance status. This separation ensures that the FHIR resource is both self-contained for EHR consumption and cryptographically verifiable against the immutable on-chain record.

Definition 1.1 (FHIR Provenance Mapping)**.** The FHIR provenance mapping $\Phi : IR \to FHIR.Provenance$ is a bijection from Asuu inference records to conformant FHIR R4 Provenance resources satisfying:

(i) **Semantic preservation:** $\Phi(IR)$ contains all clinically relevant metadata from IR without loss, following the field mapping of Table 1.

(ii) **Signature integrity:** The FHIR Signature element contains the raw Asuu signature σ, enabling independent verification against the on-chain record.

(iii) **TEFCA compatibility:** $\Phi(IR)$ conforms to the TEFCA Common Agreement's provenance requirements for QHINs, including Facilitated FHIR exchange and the QTF v2 provenance tracking requirements.

(iv) **HTI-1 alignment:** For predictive decision support interventions (DSIs), the mapping includes ONC HTI-1 source attributes: intended use, intended users, training data representativeness, external validation, local validity and fairness monitoring status, and update schedule. These are encoded as FHIR Extension elements referencing the on-chain Clinical AI Passport (Section 23).

1.2 CDS Hooks Integration

The CDS Hooks 2.0.1 specification defines the interface through which AI recommendations surface inside clinical workflow. An Asuu-integrated EHR registers CDS services that, upon hook invocation, query both the local AI model and the on-chain governance state. The response card includes the clinical recommendation, a link to the FHIR Provenance resource, the model's current passport status (active, alert, suspended), and the autonomy level at which the recommendation was generated.

When the model's passport shows a suspension flag, the CDS service must *not* return an empty card set. Silent removal of decision support mid-shift is a known safety anti-pattern: clinicians notice the absence, do not know the cause, and lose trust in the system. Instead, upon suspension the CDS service returns a *visible advisory card* containing: (i) a clear statement that the AI recommendation for this case has been withheld because the model is under regulatory or safety review, (ii) the suspension timestamp and a non-technical reason code drawn from a pre-specified vocabulary (*calibration drift*, *subgroup performance regression*, *regulatory hold*, *scheduled maintenance*), (iii) an expected review-completion window if available, and (iv) an explicit instruction to proceed with standard clinical judgment for the present case. The card's FHIR resource type is Guidance Response with status=data-requested and an accompanying DetectedIssue resource capturing the suspension reason, preserving auditability and informing the clinician without surprising them.

This design principle extends to the network's role-based access control. The Asuu 7-Tier Taxonomy of Clinical Authority (Section 28) anchors its 114 healthcare roles to SNOMED CT concept identifiers and maps them through the HL7 FHIR PractitionerRole resource [34]. When a clinician interacts with the network, their credential presentation (Section 17) produces both a zero-knowledge proof for on-chain verification and a conformant FHIR PractitionerRole resource for EHR consumption. The two representations are cryptographically bound: the FHIR resource includes a reference to the on-chain credential hash, and the on-chain credential includes the SNOMED CT code that anchors the FHIR resource.

1.3 EHR Vendor Integration Architecture

A CDS Hooks service that is specification-compliant but not deployable through the dominant EHR vendors' certification pipelines is not a clinical integration, it is a research demonstration. The Asuu protocol's reference integration path addresses the operational reality of Epic App Orchard (now Showroom), Oracle Cerner Open Developer Experience (code console), and MEDITECH Greenfield pipelines: each has a distinct technical-certification, commercial-licensing, and security-review workflow, and each typically requires twelve to thirty-six months from initial registration to production availability at a customer site.

Definition 1.2 (Three-Track Vendor Integration Strategy)**.** The Asuu reference deployment specifies three parallel vendor-integration tracks, each targeting a named certification program:

(i) **Epic Showroom Track.** The Asuu CDS service is registered as a SMART-on-FHIR application in the Epic Showroom catalog (successor to App Orchard, effective 2024). The service uses Epic's public FHIR R4 endpoints for patient-context resolution and exposes CDS Hooks endpoints registered with the Epic instance at the customer site. Certification requires: (a) SMART-on-FHIR conformance testing, (b) Epic-specific clinical-content review, (c) a commercial agreement between the Asuu Foundation (or the deploying consortium) and Epic, (d) per-customer deployment activation through the customer's Epic-governance process. Expected timeline to first production site: 18 to 30 months from initial registration, depending on whether the deployment uses a clinician-facing workflow that requires clinical-content review.

(ii) **Oracle Cerner code console Track.** The Asuu CDS service is registered in Oracle's developer platform, with conformance to Cerner's Ignite APIs for patient-context and CDS Hooks invocation. Certification requires HL7-compliant FHIR R4 implementation, Cerner security review, and a Cerner-Foundation commercial relationship. Expected timeline to first production site: 12 to 24 months.

(iii) **MEDITECH Greenfield / Altera / athenahealth Marketplace Track.** For tier-2 EHR vendors with smaller developer ecosystems, the Asuu reference integration uses each vendor's FHIR-Patient-Access program. Timelines vary widely by vendor but are generally faster than Epic and Cerner because certification programs are less mature.

Definition 1.3 (Reference Integration Timeline and Milestones)**.** The reference integration plan defines four milestones whose durations are specified in elapsed months from project kickoff:

- **Months 0–6 (Specification and Registration):** Asuu CDS service specified to FHIR R4/R5 and CDS Hooks 2.0.1 conformance. Registrations filed with Epic Showroom, Oracle developer platform, and at least one tier-2 vendor. Reference implementation of Asuu-aware GuidanceResponse and DetectedIssue resources published as open specification.
- **Months 6–18 (Certification and First Pilot):** Vendor-specific certifications completed on parallel tracks. First pilot customer on-boarded through at least one track (typically the fastest tier-2 vendor); pilot scope is restricted to provenance-anchoring and audit-package generation, not clinical decision-support card content, to reduce clinical-content-review scope.
- **Months 18–30 (First Production and Epic/Cerner):** First Epic or Cerner customer reaches production deployment. Clinical decision-support card content subject to customer-site clinical review; Asuu card content is advisory and deferred to customer's existing CDS governance.
- **Months 30–48 (Scale):** Multi-customer deployment across Epic, Cerner, and at least one tier-2 vendor. Certification evidence used to accelerate subsequent customer onboarding.

Remark 1.4 (Honest Acknowledgment of the Alert-Fatigue Concern)**.** The Asuu suspension-advisory card (Section 42.1) adds a new class of card to a clinical workflow in which clinicians already dismiss more than 90 percent of alerts, with well-documented alert-fatigue consequences. Adding another alert class is therefore not free; it imposes a cognitive cost that must be justified by the value the alert provides. The Asuu deployment approach mitigates this in three ways: (a) the suspension card is *sparse* by design (it appears only when a model passport transitions out of active status, which should be a rare event in a well-governed deployment), (b) the card is *directive* rather than advisory (it instructs the clinician to proceed with standard

judgment rather than offering a competing recommendation that must be evaluated), and (c) the card is *accompanied by explicit removal of the prior-recommendation class for the affected model*, so the net change in alert volume per patient encounter is near zero. These mitigations are empirically testable and should be measured during the Months 18–30 pilot. If the measured alert-fatigue impact exceeds pre-specified thresholds, the protocol's reference implementation must be revised. We do not claim the current design is definitively safe against alert fatigue; we claim it is designed to be empirically evaluable against that criterion.

Definition 1.5 (Pre-Registered Alert-Fatigue Measurements and Revision Thresholds)**.** The honest critique that "we do not claim safety" is not itself a mitigation is correct. To convert the acknowledgment into a falsifiable engineering commitment, the Asuu reference implementation specifies the following pre-registered measurements, to be collected during the Months 18–30 pilot deployment (Definition 1.3) and published with every post-pilot safety review:

(i) **Card presentation rate:** the number of Asuu advisory cards (suspension, alert, or governance signal) presented per 1,000 clinical encounters, stratified by card class. Measurement primitive: an on-chain counter incremented at each card-dispatch event, indexed by card class and rolling 30-day window.

(ii) **Net alert-volume delta:** the difference between (card presentation rate) and (recommendations withdrawn under suspension that would otherwise have been presented). This measures whether the suspension card's promise of near-zero net alert volume holds in practice.

(iii) **Time-to-acknowledge distribution:** the elapsed time from card presentation to card dismissal, measured at the EHR client. Distribution moments reported per card class and per pilot site.

(iv) **Dismissal-without-read prevalence:** the fraction of card presentations dismissed in less than 2 seconds of display, a validated proxy for reflexive click-through [40].

(v) **Workflow-paralysis incident reports:** open-text clinician reports of workflow disruption when a previously-relied-upon AI recommendation is withdrawn under suspension. Collected through the pilot site's existing adverse-event-reporting system.

The following pre-specified revision thresholds trigger mandatory protocol-level redesign before broader deployment:

- Card presentation rate exceeding 5 per 1,000 encounters for any non-suspension card class.
- Net alert-volume delta exceeding 10 percent over baseline for any pilot site.
- Dismissal-without-read prevalence exceeding 40 percent of card presentations.
- Any workflow-paralysis incident judged by the pilot-site safety committee to have contributed to patient harm.

If any threshold is crossed during the pilot, the reference implementation's card-dispatch mechanism is revised before Months-30+ broader deployment proceeds. The thresholds are conservative; they are intended to trigger redesign well before clinician-workflow harm becomes likely, not to validate the current design after the fact. If the pilot demonstrates that any threshold is routinely crossed, the suspension-card mechanism as currently specified is the wrong design and must be replaced.

2 Prior Authorization: The First Institutional Workflow

Healthcare insurance processing consumes approximately 30 percent of total healthcare expenditure in administrative overhead. Prior authorization alone delays treatment by an average of two business days, with 34 percent of physicians reporting that delays have led to serious adverse events. The regulatory window is open: the CMS Interoperability and Prior Authorization Final Rule requires impacted payers to implement Prior Authorization API requirements by 1 January 2027, and the HL7 Da Vinci ecosystem provides the implementation guides. Da Vinci PAS handles point-of-service authorization requests and immediate responses. Da Vinci DTR (Documentation Templates and Rules) executes payer rules inside the provider workflow so documentation requirements are met at the point of care.

2.1 Honest Incentive Analysis: Why a Payer Would and Would Not Adopt

Before specifying the workflow, it is necessary to be clear-eyed about why payers would consider this infrastructure and why they might resist. Treating prior authorization purely as administrative friction that all parties want to reduce misrepresents the adversarial structure of claims adjudication and produces unrealistic deployment expectations.

What the CMS 2027 rule does and does not require. The CMS Interoperability and Prior Authorization Final Rule mandates API-mediated prior authorization with structured request and response formats. It does not mandate instant approval, algorithmic adjudication, or the elimination of payer discretion. Payers retain the right to require documentation, conduct medical-necessity review, and deny coverage. The rule closes the inefficiency of fax-and-phone workflow; it does not eliminate the payer's role as coverage arbiter.

Why payers might adopt. Payers gain from Asuu integration in four specific ways that are *compatible with their cost-management function*, not in conflict with it:

(i) **Fraud reduction.** The on-chain inference record cryptographically binds a specific AI recommendation to a specific patient case. Retrospective claim audits can verify that the clinical evidence supporting a prior authorization was actually present at the time of decision, closing the upcoding and phantom-billing pathways that cost payers billions annually.

(ii) **Documentation compliance at point of care.** DTR-generated FHIR bundles assembled at the Data Embassy provide the payer with the evidence needed to adjudicate, in the payer's preferred format, without multiple back-and-forth requests. This reduces the payer's own administrative cost even when the outcome is denial.

(iii) **Regulatory defense.** Payers face increasing regulatory scrutiny over denial rates, appeal outcomes, and health-equity outcomes. Structured denial reasons (E_{fail} sets) with auditable provenance provide the payer with defensible evidence that denials rest on specific, articulable coverage criteria applied uniformly rather than on discretionary judgments that may be challenged as discriminatory.

(iv) **Dispute resolution cost reduction.** The most costly element of prior authorization for payers is not the first-level denial but the appeal and external-review cycle. An on-chain record of precisely which coverage predicates were evaluated and which failed transforms appeals from evidentiary contests into predicate re-evaluations, reducing external-review costs substantially.

What payers will not accept. The following are deployment-killers that any realistic prior-authorization workflow must avoid:

(i) *Loss of adjudication authority.* A smart contract that auto-approves claims satisfying documented coverage predicates without human payer review removes the payer's cost-management lever. Payers will not contractually commit to auto-approval.

(ii) *Loss of discretionary medical-necessity review.* Many coverage decisions turn on case-specific medical-necessity judgments that cannot be fully encoded as Boolean predicates. Payers will not accept a protocol that prevents medical directors from reviewing such cases.

(iii) *Unfavorable audit asymmetry.* Payers will resist protocols in which all decisions are logged immutably against the payer but not symmetrically against the provider.

The Asuu prior-authorization workflow therefore operates in *hybrid adjudication mode* by default: the smart contract evaluates documented coverage predicates and structurally records the result, but the outcome is authoritative only for cases the payer has pre-committed (by policy contract) to delegate to algorithmic adjudication. Cases outside that pre-commitment enter the standard payer review workflow, with the on-chain documentation serving as the evidentiary substrate. The protocol's value to payers is that it produces *better evidence*, not that it removes their authority.

2.2 Workflow Specification

The Asuu network supports prior authorization through smart contracts aligned with Da Vinci PAS. When an AI system produces a clinical recommendation (recorded as an inference record on-chain), the DTR-compliant documentation is generated automatically from the FHIR evidence bundle at the Data Embassy. A smart contract then evaluates whether the recommendation meets the coverage criteria that the payer has encoded into its policy contract. For cases that the payer has pre-committed to algorithmic adjudication, the contract's evaluation is authoritative: the Da Vinci PAS response bundle is generated and transmitted within seconds. For cases outside the pre-commitment scope, the contract's evaluation is advisory: it is transmitted to the payer's medical review queue as structured evidence, and the final determination is made by the payer under standard procedures. In both cases, the structured denial reasons are returned as FHIR ClaimResponse resources, eliminating the reason-code ambiguity that characterizes current denial processes and providing a defensible substrate for appeal.

This workflow does not replace clinical judgment, nor does it replace payer judgment. It replaces the evidentiary friction between them. The clinician makes the medical decision. The AI system records the clinical reasoning. The smart contract evaluates whether the evidence meets policy criteria and structures the result. The payer retains final authority over cases outside the pre-committed algorithmic scope. Each step is auditable, attributable, and immutable.

2.3 Automata-Theoretic Formalization

Definition 2.1 (Claims Automaton)**.** The claims contract is a deterministic finite automaton $\mathscr{M} = (Q, \Sigma, \delta, q_0, F)$ with deterministic δ, initial state $q_0 = q_{\text{sub}}$, accepting set $F = \{q_{\text{cls}}\}$, and state set:

$$Q = \{q_{\text{sub}}, q_{\text{rev}}, q_{\text{doc}}, q_{\text{app}}, q_{\text{den}}, q_{\text{apl}}, q_{\text{set}}, q_{\text{cls}}\}.$$

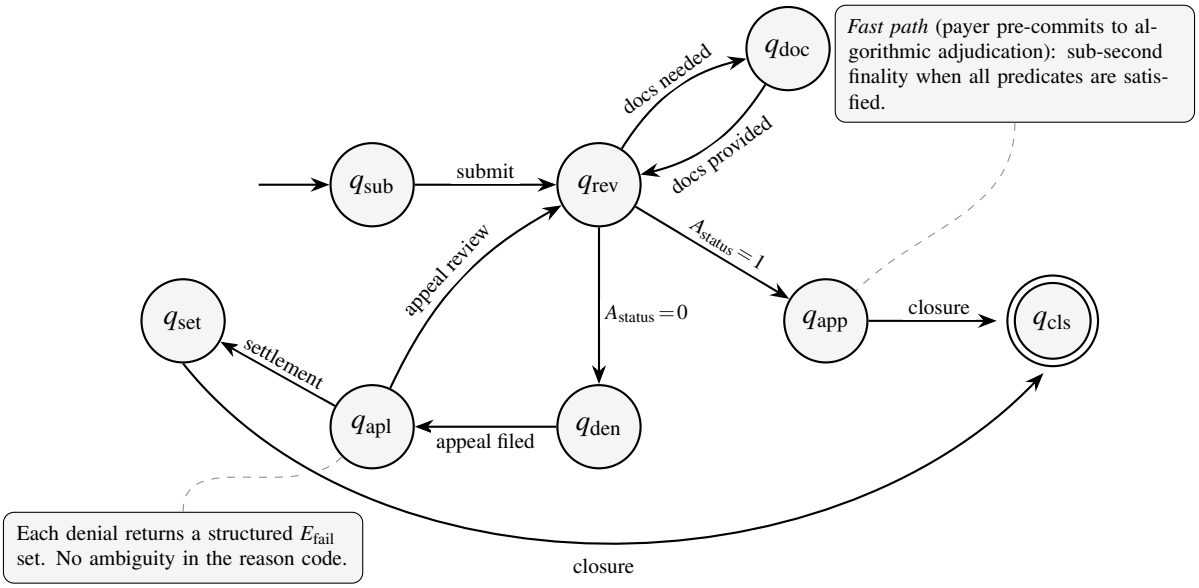


Figure 1: The prior authorization claims automaton $\mathscr{M}$. For cases within the payer's algorithmic-adjudication pre-commitment, the fast path executes in network finality time when all coverage predicates are satisfied. For cases outside the pre-commitment, the automaton logs evidence and transitions to payer medical review; deterministic execution of the predicate evaluation ensures all validators reach identical evidentiary records.

Definition 2.2 (Coverage Predicate)**.** A coverage policy $C_P = \{c_1, \ldots, c_m\}$ is a set of Boolean predicates over inference records. The authorization function is:

$$A_{\text{status}}(IR, C_P) = \bigwedge_{j=1}^{m} c_j(IR).$$

If $A_{\text{status}} = 1$ and the case is within the payer's algorithmic-adjudication pre-commitment: the claim transitions to q_{app}. If $A_{\text{status}} = 0$: the claim transitions to q_{den} with the set of failing predicates $E_{\text{fail}} = \{c_j : c_j(IR) = 0\}$ returned as structured feedback. Cases outside the pre-commitment scope transition to payer medical review with the predicate evaluation attached as evidence.

Theorem 2.3 (Deterministic Adjudication of Coverage Predicates)**.** *For any inference record IR and coverage policy C_P, all validators compute identical A_{status} and E_{fail}.*

Proof. *IR* is committed on-chain, so all validators observe the same input. Each c_j is a pure function of on-chain state. The conjunction is deterministic. Since δ is deterministic, identical inputs produce identical transitions on every node. □

Remark 2.4 (Scope of the Determinism Theorem)**.** Theorem 2.3 proves that the evaluation of Boolean coverage predicates is deterministic across validators. It does not prove, and does not imply, that medical-necessity determination is deterministic. Medical-necessity review, which turns on clinician and medical-director judgment about whether a treatment is appropriate for a specific patient's clinical context, is by nature a discretionary rather than deterministic process. The protocol formalizes only the predicate-evaluation layer; the discretionary layer remains under payer authority for cases outside the algorithmic-adjudication pre-commitment scope.

Corollary 2.5 (Institutional and Systemic ROI)**.** *For cases within the algorithmic-adjudication pre-commitment: deterministic predicate evaluation eliminates the evidentiary back-and-forth for claims satisfying all coverage predicates, reducing average authorization latency from two business days to the network's finality time (seconds). For an institution processing n such prior authorizations per month, the administrative cost reduction is $n \cdot (c_{\text{manual}} - c_{\text{tx}})$. For cases outside the pre-commitment: the structured evidence and E_{fail} reasoning reduce documentation cycle time on both provider and payer sides by eliminating reason-code ambiguity. The total systemic ROI depends on the fraction of cases the payer pre-commits to algorithmic scope, which is a matter of bilateral contract between provider and payer and will vary by procedure class, specialty, and payer risk tolerance.*

Remark 2.6 (Why Prior Authorization Is the Right First Workflow Despite Payer Adversarial Incentive)**.** The selection of prior authorization as the first institutional workflow is strategic, but for a more narrowly correct reason than an earlier draft of this paper asserted. Prior authorization is the rare clinical-administrative process that (i) is universally loathed by clinicians and therefore creates pull on the provider side, (ii) is being forced into API-mediated structure by federal regulation by 2027 regardless of any blockchain adoption, (iii) already has a mature HL7 implementation ecosystem (Da Vinci PAS/DTR/CRD) that Asuu integration can extend rather than replace, and (iv) contains a clearly-delineated deterministic layer (coverage-predicate evaluation) separable from the discretionary layer (medical-necessity review), permitting a hybrid workflow in which the protocol adds value to both sides without requiring payers to cede cost-management authority. The protocol does not resolve the adversarial structure of claims adjudication; it improves the evidentiary quality of that adjudication, for both parties. That is the narrower, more defensible claim.

Remark 2.7 (Honest Acknowledgment: What Asuu Adds on Top of DTR)**.** A legitimate critique of the preceding framing is that the Da Vinci DTR (Documentation Templates and Rules) implementation guide, under CMS 2027 API-mediation mandate, already produces the structured FHIR-bundle evidence that "better evidence" would describe. If DTR alone gives payers structured documentation at the point of care, and CMS is already forcing API-mediation of the authorization response, then blockchain adds no marginal evidentiary value on top of infrastructure payers must build regardless. This critique is substantially correct, and we do not dispute it. The honest question is what Asuu adds *beyond* what a correctly-deployed DTR stack already produces.

The marginal value of Asuu on top of DTR is narrower than evidentiary structuring. It consists of three items:

(i) **Cross-institutional tamper-evidence.** DTR bundles are generated and stored by the provider's EHR and the payer's adjudication system. Neither party can cryptographically prove to a third party (regulator, external auditor, plaintiff's counsel) that the bundle it holds is the same bundle the counterparty holds, nor that either bundle is the bundle that was produced at the time of clinical decision. The on-chain hash anchor provides this property: a DTR bundle anchored on-chain cannot be retroactively altered by either party without detection. This is a property DTR alone does not provide, because DTR specifies the *format* of the evidence and the *transport*, but not a cross-party immutable anchor.

(ii) **Longitudinal model-provenance attribution.** When a prior-authorization decision rests on an AI-generated clinical recommendation, the DTR bundle records the recommendation but does not cryptographically bind which model version, trained on which

data, under which governance approval, produced it. For correlated-failure adverse events traced to a specific model version's training data, the on-chain model-passport linkage (Section 23) provides attribution that DTR alone does not. This matters for PLD 2024/2853 Article 10(4) defendability, not for routine adjudication.

(iii) **Appeal and external-review evidentiary substrate.** The highest-cost element of prior authorization is not first-level denial but the appeal cycle. When an appeal reaches external review, the reviewer must reconstruct the state of evidence at the time of original decision. An on-chain anchored DTR bundle plus model-passport state provides this reconstruction at fixed low cost, where today's reviewers rebuild the evidentiary state from provider and payer submissions at substantial effort.

Items (i) and (iii) are the principal Asuu deltas over DTR; item (ii) is secondary. None of these require deterministic instant auto-approval, and none require payers to cede adjudication authority. The honest positioning of the prior-authorization workflow is therefore: *DTR-plus-anchoring*, not *DTR-replacement*. We do not claim that DTR-plus-anchoring is by itself a sufficient business case to justify a blockchain protocol; the prior-authorization workflow is an *institutional-onboarding vehicle* that exercises the protocol's anchoring and model-passport features in a workflow that institutions are already forced to implement. The deeper business case rests on the applications that compose on top of the anchoring substrate (clinical trial integrity, DSCSA compliance, model-governance auditability), not on prior-authorization alone. We make this explicit rather than claim otherwise.

PART II

CORE PROTOCOL & CRYPTOGRAPHIC FOUNDATIONS

3 Introduction

Medical artificial intelligence has come to rely almost exclusively on proprietary platforms operated by technology corporations serving as trusted intermediaries between clinical knowledge and patient care. While these systems perform adequately for certain tasks, they suffer from inherent weaknesses of the trust-based model. Transparent, sovereign AI infrastructure is not possible when a foreign vendor controls the models, the data pipeline, and the conditions of access.

The cost of this dependence is substantial. Institutions pay perpetual licensing fees that compound annually, creating permanent operating expenses that compete directly with patient care spending. Vendor lock-in ensures that once a hospital has integrated a proprietary system (training staff, modifying workflows, becoming dependent on its outputs), switching costs become prohibitive. Vendors recognize this asymmetry and extract increasingly burdensome terms. The situation is worse in low- and middle-income countries, where the entire annual health budget per capita may be less than the cost of a single AI license.

Cloud-dependent proprietary AI introduces operational vulnerabilities that directly threaten patient care:

- A vendor's server outage or a cyberattack on their infrastructure,
- A bankruptcy or acquisition that discontinues the product,
- A geopolitical decision to restrict service to specific nations or regions.

Any of these can leave a health system suddenly without capabilities it has come to rely on. In rural areas, disaster zones, military field hospitals, or any setting with unreliable connectivity, cloud-based AI becomes useless precisely when need is greatest.

Beyond operational fragility, the current model creates a fundamental sovereignty crisis. When a nation's clinical AI runs on foreign servers, processes data through foreign models, and operates under foreign terms of service, that nation has effectively outsourced a critical component of its healthcare decision-making infrastructure. The data generated through clinical use (patterns of disease, treatment efficacy, population health signals) constitutes a form of medical intelligence that accrues to the vendor, not the institution.

The scale of the problem is accelerating. The global medical AI market is projected to reach $110 to $614 billion by 2030, with over 1,350 AI-enabled medical devices already authorized by the FDA. Foundation models such as Med-Gemini have achieved 91.1% on medical licensing examinations, and AI-designed drugs are entering Phase 2 and Phase 3 clinical trials [10, 11]. Yet this entire capability stack rests on proprietary infrastructure controlled by a handful of corporations. The healthcare blockchain market, projected to reach $43 to $215 billion by 2030, signals institutional recognition that trust infrastructure itself must be decentralized. The convergence of medical AI and blockchain is no longer speculative. The question is not whether healthcare will adopt a convergent trust infrastructure, but whether that infrastructure will be sovereign or subordinate, open or proprietary, accountable or opaque.

What is needed is a medical AI operating system grounded in cryptographic proof rather than trust, allowing any institution to deploy, operate, and verify AI-assisted clinical decisions without dependence on a third party. A system where every inference is auditable, every model update is transparent, and every data contribution is tracked, not by a central authority, but by the participating network itself. In this paper, we propose a solution: a peer-to-peer distributed

network with a purpose-built consensus mechanism to generate an immutable record of all medical AI operations, anchored to the healthcare interoperability standards that hospitals already use. The system positions itself not as a replacement for existing clinical infrastructure, but as the invisible cryptographic layer that sits inside existing workflows and acts as the ultimate mathematical arbiter of safety, provenance, and algorithmic liability.

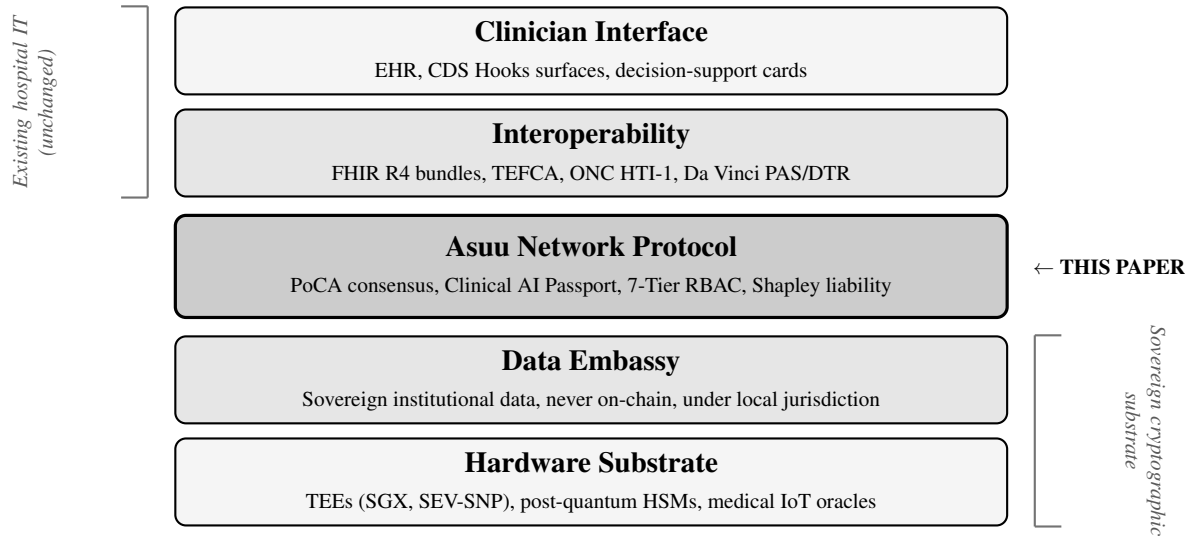


Figure 2: The Asuu Network's position in the clinical AI stack. The protocol sits between existing hospital interoperability infrastructure and the sovereign cryptographic substrate. It is not a replacement for the EHR or the FHIR layer. It is the verifiable control plane that governs what happens across them.

The system is secure as long as honest institutions collectively control more validating power than any cooperating group of adversarial nodes.

4 Notation and Cryptographic Preliminaries

We establish the notation, definitions, and hardness assumptions used throughout this document. All proofs are stated in the concrete security framework unless otherwise noted.

To construct a trustless medical infrastructure, we rely on standard cryptographic primitives whose security has been extensively studied. The definitions below establish the precise guarantees each primitive must provide for the Asuu protocol to achieve its claimed security properties.

4.1 Notation

Let λ denote the security parameter. A function $f : \mathbb{N} \to \mathbb{R}$ is *negligible* if for every positive polynomial $p(\cdot)$ there exists N such that for all $n > N$, $|f(n)| < 1/p(n)$; we write $f(\lambda) \leq \mathrm{negl}(\lambda)$. A probabilistic algorithm $\mathscr{A}$ is *efficient* if it runs in time $\mathrm{poly}(\lambda)$; we write $\mathscr{A} \in \mathrm{PPT}$. We denote uniform random sampling by $x \xleftarrow{\$} \mathscr{S}$. The notation $[n]$ denotes the set $\{1, 2, \ldots, n\}$. We write $\|$ for string concatenation.

4.2 Hash Function Security

Definition 4.1 (Collision-Resistant Hash Family)**.** A keyed hash family $\mathscr{H} = \{H_s : \{0,1\}^* \to \{0,1\}^\lambda\}_{s \in \mathscr{K}}$ is *collision-resistant* if for all $\mathscr{A} \in \mathrm{PPT}$:

$$\mathrm{Adv}^{\mathrm{CR}}_{\mathscr{H},\mathscr{A}}(\lambda) = \Pr\big[H_s(x) = H_s(x') \wedge x \neq x' : s \xleftarrow{\$} \mathscr{K},\ (x,x') \leftarrow \mathscr{A}(s)\big] \leq \mathrm{negl}(\lambda).$$

Definition 4.2 (Preimage Resistance)**.** $\mathscr{H}$ is *preimage-resistant* if for all $\mathscr{A} \in \mathrm{PPT}$:

$$\Pr\big[H_s(x') = y : s \xleftarrow{\$} \mathscr{K},\ x \xleftarrow{\$} \{0,1\}^\lambda,\ y \leftarrow H_s(x),\ x' \leftarrow \mathscr{A}(s,y)\big] \leq \mathrm{negl}(\lambda).$$

Lemma 4.3 (Collision Resistance Implies Second-Preimage Resistance)**.** *If $\mathscr{H}$ is collision-resistant, then $\mathscr{H}$ is second-preimage resistant. That is, given x, no efficient adversary can find $x' \neq x$ with $H_s(x') = H_s(x)$ except with negligible probability.*

Proof. A second-preimage finder $\mathscr{A}'$ that on input (s,x) outputs $x' \neq x$ with $H_s(x') = H_s(x)$ directly yields a collision (x,x'), contradicting Definition 4.1. □

4.3 Digital Signature Security

Definition 4.4 (EUF-CMA Security)**.** A signature scheme $\Sigma = (\mathrm{KeyGen}, \mathrm{Sign}, \mathrm{Verify})$ is *existentially unforgeable under adaptive chosen-message attack* (EUF-CMA) if for all $\mathscr{A} \in \mathrm{PPT}$:

$$\mathrm{Adv}^{\text{EUF-CMA}}_{\Sigma,\mathscr{A}}(\lambda) = \Pr\left[\mathrm{Verify}(pk, m^*, \sigma^*) = 1 \wedge m^* \notin Q : \begin{array}{l}(pk,sk) \leftarrow \mathrm{KeyGen}(1^\lambda) \\ (m^*, \sigma^*) \leftarrow \mathscr{A}^{\mathrm{Sign}(sk,\cdot)}(pk)\end{array}\right] \leq \mathrm{negl}(\lambda),$$

where Q is the set of messages queried to the signing oracle.

4.4 Commitment Schemes

Definition 4.5 (Pedersen Commitment)**.** Let $\mathbb{G}$ be a cyclic group of prime order q with generators g, h where $\log_g h$ is unknown. The Pedersen commitment to message $m \in \mathbb{Z}_q$ with randomness $r \xleftarrow{\$} \mathbb{Z}_q$ is $C = g^m h^r$.

Lemma 4.6 (Pedersen: Perfect Hiding)**.** *For any commitment $C \in \mathbb{G}$ and any message $m' \in \mathbb{Z}_q$, there exists a unique $r' \in \mathbb{Z}_q$ such that $C = g^{m'} h^{r'}$. Hence the distribution of C is uniform over $\mathbb{G}$ regardless of m.*

Proof. Given $C = g^m h^r$, set $r' = r + (m - m') \cdot (\log_g h)^{-1} \bmod q$. Since h generates $\mathbb{G}$, the discrete logarithm exists and r' is unique. Therefore C reveals no information about m. $\square$

Lemma 4.7 (Pedersen: Computational Binding Under DLP)**.** *If the discrete logarithm problem is hard in $\mathbb{G}$, then no efficient adversary can open a Pedersen commitment to two distinct messages.*

Proof. Suppose $\mathscr{A}$ outputs (m_1, r_1) and (m_2, r_2) with $m_1 \neq m_2$ and $g^{m_1} h^{r_1} = g^{m_2} h^{r_2}$. Then $g^{m_1 - m_2} = h^{r_2 - r_1}$, yielding $\log_g h = (m_1 - m_2)(r_2 - r_1)^{-1} \bmod q$, which contradicts the DLP assumption. $\square$

4.5 Zero-Knowledge Proof Systems

Definition 4.8 (Non-Interactive Zero-Knowledge Argument of Knowledge)**.** A non-interactive argument system for NP relation $\mathscr{R}$ consists of $(\mathsf{Gen}, \mathsf{Prove}, \mathsf{Verify})$ satisfying:

(i) **Completeness.** For all $(x, w) \in \mathscr{R}$: $\Pr[\mathsf{Verify}(vrs, x, \mathsf{Prove}(crs, x, w)) = 1] = 1$.

(ii) **Knowledge Soundness.** For every PPT prover $\mathscr{P}^*$ there exists a PPT extractor $\mathscr{E}$ such that: $\Pr[\mathsf{Verify}(vrs, x, \pi) = 1 \wedge (x, w) \notin \mathscr{R}] \leq \mathrm{negl}(\lambda)$.

(iii) **Zero-Knowledge.** There exists a PPT simulator $\mathscr{S}$ with trapdoor τ such that for all $(x, w) \in \mathscr{R}$, the distributions $\{\mathsf{Prove}(crs, x, w)\}$ and $\{\mathscr{S}(vrs, \tau, x)\}$ are computationally indistinguishable.

4.6 Hardness Assumptions

Assumption 4.9 (Discrete Logarithm Problem (DLP))**.** For cyclic group $\mathbb{G}$ of prime order q with generator g: $\Pr[\mathscr{A}(g, g^x) = x : x \xleftarrow{\$} \mathbb{Z}_q] \leq \mathrm{negl}(\lambda)$ for all $\mathscr{A} \in \mathrm{PPT}$.

Assumption 4.10 (Computational Diffie-Hellman (CDH))**.** For $\mathbb{G}, g, q$ as above: $\Pr[\mathscr{A}(g, g^a, g^b) = g^{ab}] \leq \mathrm{negl}(\lambda)$ for all $\mathscr{A} \in \mathrm{PPT}$.

Assumption 4.11 (q-Strong Diffie-Hellman (q-SDH))**.** Given $(g, g^{\alpha}, g^{\alpha^2}, \ldots, g^{\alpha^q}) \in \mathbb{G}^{q+1}$, no PPT adversary can output $(c, g^{1/(\alpha + c)})$ except with negligible probability.

Assumption 4.12 (Honest Supermajority)**.** The fraction of Byzantine validators satisfies $f < n/3$, where n is the total number of active validators. Equivalently, honest validators collectively hold more than $2n/3$ of the aggregate clinical authority weight.

Assumption 4.13 (Trusted Execution Environment Integrity)**.** The TEE ideal functionality $\mathcal{G}_{\text{att}}$ provides integrity (attestation binds program identity to output) and confidentiality (the host cannot observe enclave memory) in the UC framework.

Assumption 4.14 (Ring Learning With Errors (RLWE))**.** Let $R_q = \mathbb{Z}_q[X]/(X^N + 1)$ for N a power of two and q a prime. For secret $s \xleftarrow{\$} R_q$ and error distribution χ over R_q, the distribution $(a, a \cdot s + e)$ for uniform $a \xleftarrow{\$} R_q$ and $e \leftarrow \chi$ is computationally indistinguishable from uniform over $R_q \times R_q$ for all $\mathscr{A} \in$ PPT.

Assumption 4.15 (Module Learning With Errors (M-LWE))**.** For module rank $d \geq 2$, secret $\mathbf{s} \xleftarrow{\$} R_q^d$, and uniform $\mathbf{A} \xleftarrow{\$} R_q^{k \times d}$: the distribution $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e})$ is computationally indistinguishable from uniform over $R_q^{k \times d} \times R_q^k$ for all $\mathscr{A} \in$ PPT. Hardness reduces to worst-case lattice problems via Regev's theorem [18].

Assumption 4.16 (Module Short Integer Solution (M-SIS))**.** For module rank $d \geq 2$, bound β, and uniform $\mathbf{A} \xleftarrow{\$} R_q^{k \times d}$: no PPT adversary finds a nonzero $\mathbf{x} \in R_q^d$ with $\|\mathbf{x}\|_\infty \leq \beta$ such that $\mathbf{A}\mathbf{x} = \mathbf{0} \pmod{q}$ with non-negligible probability. Hardness reduces to worst-case lattice problems over module lattices. The EUF-CMA security of ML-DSA-65 (used in the composite signature scheme of Definition 8.1) reduces to the combination of M-LWE (Assumption 4.15) and M-SIS [19]; M-LWE alone is insufficient for the full signature-unforgeability reduction. We include M-SIS explicitly to correct an omission in an earlier draft's assumption list.

5 Medical Inference Records

We define a medical inference record as a cryptographically signed attestation that a specific AI model, operating at a specific version, produced a specific output given a specific input, at a specific time, at a specific institution. Each record contains: a hash of the input data (never the data itself), the model identifier and version hash, a hash of the output, the institution's public key, and a timestamp. The institution signs the record, attesting that this inference occurred under its clinical authority.

The problem, naturally, is that an institution or vendor cannot unilaterally prove the integrity of its own records. A compromised model could produce dangerous outputs; a vendor could silently update model weights; an institution could retroactively alter records to conceal errors. In the current system, we trust vendors not to tamper with models and institutions not to tamper with logs. This trust is unverifiable.

We need a way for any party (regulators, patients, peer institutions) to verify that a medical AI inference occurred exactly as recorded, using exactly the model version claimed, and that the record has not been altered after the fact. For our purposes, the earliest record is the one that counts. To accomplish this without a trusted party, inference records must be publicly announced to the network, and we need a system for participants to agree on a single history of the order in which they were recorded.

5.1 Formal Specification

Definition 5.1 (Medical Inference Record)**.** A medical inference record is a tuple $IR = (h_{\text{in}}, mid, h_{\text{out}}, pk_I, t, \sigma)$ where:

- $h_{\text{in}} = H(D_{\text{in}})$ is the hash of the input data,
- $mid = (M_{\text{id}}, H(M_{\text{weights}}))$ identifies the model and its version,
- $h_{\text{out}} = H(D_{\text{out}})$ is the hash of the inference output,
- pk_I is the public key of the submitting institution I,
- t is the timestamp, and
- $\sigma = \text{Sign}(sk_I, h_{\text{in}} \,\|\, mid \,\|\, h_{\text{out}} \,\|\, t)$.

Proposition 5.2 (Inference Record Integrity)**.** *Under Definition 4.1 and Definition 4.4, no* PPT *adversary can produce a valid inference record that was not generated by the institution controlling sk_I.*

Proof. An adversary forging IR' with a valid signature under pk_I for a message not signed by sk_I directly breaks EUF-CMA security of Σ. An adversary producing a distinct record $IR'' \neq IR$ that hashes identically breaks collision resistance of H. Both events occur with probability at most $\text{negl}(\lambda)$. □

Theorem 5.3 (Non-Repudiation of Clinical Inferences)**.** *If Σ is EUF-CMA secure and $\mathscr{H}$ is collision-resistant, then a valid on-chain inference record IR constitutes a non-repudiable attestation that institution I endorsed the inference at time t.*

Proof. By Proposition 5.2, only the holder of sk_I can produce a valid signature σ. The timestamp t is committed to in σ and anchored by the block's position in the hash chain (Section 6). Institution I cannot claim the record was altered or that it never endorsed the inference without exhibiting a second preimage for H or a signature forgery, both of which are computationally infeasible. □

Corollary 5.4 (Regulatory Auditability)**.** *Any party holding the block headers and a Merkle proof (Section 13) can independently verify the existence, integrity, and institutional attribution of any inference record without re-executing the inference.*

Theorem 5.5 (End-to-End Provenance Completeness)**.** *For every clinical action a executed through an Asuu workflow contract, there exists a unique, authenticated provenance path*

$$\pi_a = (mid,\ IR,\ HA,\ a)$$

consisting of: the model version mid registered in the Clinical AI Passport (Section 23), the inference record IR (Definition 5.1), the human authorization state HA (the signed approval or override transaction), and the final action record a. Under collision resistance of H (Definition 4.1) and EUF-CMA security of the signature scheme (Definition 4.4), any missing or altered step in π_a is detectable.

Proof. The workflow contract enforces sequential state transitions: *mid* must be registered on-chain before inference, *IR* must reference a valid *mid*, *HA* must reference the hash of *IR*, and *a* must reference the hash of *HA*. Each step is signed by the responsible party (model deployer, institution, clinician, workflow contract). If any step is missing, the subsequent step's reference to it fails hash verification by collision resistance. If any step is altered, the signature verification fails by EUF-CMA. The uniqueness of π_a follows from the deterministic hash-linking: each action *a* references exactly one *HA*, which references exactly one *IR*, which references exactly one *mid*. Provenance completeness is therefore a hard protocol guarantee, not a statistical property. □

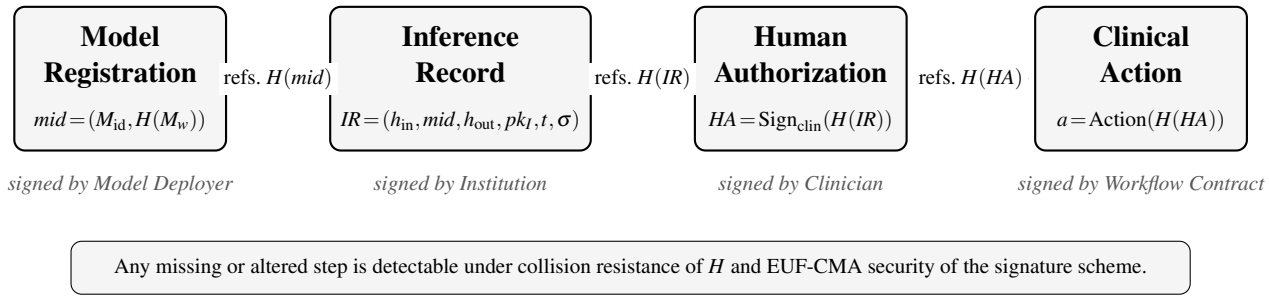


Figure 3: The end-to-end provenance path $\pi_a = (mid, IR, HA, a)$. Each step is cryptographically bound to its predecessor, making every clinical action auditable back to the registered model version.

6 Timestamp Server and Hash Chain

The solution begins with a distributed timestamp server. The server takes a hash of a block of medical inference records and publishes the hash across the network. The timestamp proves that the inference data must have existed at the time, in order to produce the hash. Each timestamp includes the previous timestamp in its hash, forming a chain, with each additional timestamp reinforcing the ones before it.

This chain of timestamps creates an append-only ledger of all medical AI activity across the network. Once an inference record is buried under subsequent blocks, altering it would require recomputing every block that follows, a task that becomes computationally infeasible as the chain grows. The result is a permanent, tamper-evident audit trail that satisfies regulatory requirements for medical device traceability without relying on any single institution's honesty.

6.1 Formal Specification

Definition 6.1 (Block Structure)**.** A block B_k is a tuple $(h_{k-1}, t_k, pk_v, \sigma_v, root_k, TX_k)$ where $h_{k-1} = H(B_{k-1})$ is the hash of the previous block, t_k is the timestamp, pk_v identifies the producing validator, σ_v is the validator's signature, $root_k$ is the Merkle root of the transaction set TX_k, and $TX_k = \{IR_1, \ldots, IR_m\}$.

Definition 6.2 (Hash Chain)**.** The hash chain is the sequence $(B_0, B_1, \ldots, B_n)$ where for all $k \geq 1$: $h_{k-1} \in B_k$ satisfies $h_{k-1} = H(B_{k-1})$.

Lemma 6.3 (One-Level Tamper Detection)**.** *If H is collision-resistant, modifying B_j for any $j < n$ causes $H(B'_j) \neq H(B_j) = h_j$ stored in B_{j+1}, detectable by any verifier holding B_{j+1}.*

Proof. If $B'_j \neq B_j$ and $H(B'_j) = H(B_j)$, then (B'_j, B_j) is a collision, contradicting Definition 4.1. □

Proposition 6.4 (Cascading Tamper Evidence)**.** *Modifying any block B_j requires recomputing all subsequent blocks $B_{j+1}, \ldots, B_n$ to maintain chain validity.*

Proof. By Lemma 6.3, modifying B_j yields $H(B'_j) \neq h_j$, invalidating B_{j+1}. To restore validity, the adversary must produce B'_{j+1} containing $H(B'_j)$. But now $H(B'_{j+1}) \neq H(B_{j+1}) = h_{j+1}$ stored in B_{j+2}. By induction, every block from $j+1$ to n must be recomputed. □

Theorem 6.5 (Hash Chain Immutability)**.** *Under collision resistance of H and the BFT assumption $f < 1/3$ (Assumption 4.12), the probability of an adversary undetectably modifying any block B_j already finalized by a valid quorum certificate is:*

$$\Pr[\textit{undetected modification}] \leq \mathrm{Adv}^{\mathrm{CR}}_{\mathscr{H},\mathscr{A}}(\lambda) + \mathrm{Adv}^{\text{EUF-CMA}}_{\mathrm{sig}}(\lambda).$$

Both terms are negligible for standard security parameters.

Proof. By Proposition 6.4, modifying B_j forces recomputation of all subsequent blocks, each of which would need a valid quorum certificate. By Theorem 14.2, producing a valid quorum certificate for a conflicting block at any height requires either (i) compromising validators whose aggregate weight exceeds $W_{\text{total}}/3$, which contradicts the BFT assumption, or (ii) forging signatures, which is bounded by the EUF-CMA advantage. Additionally, the modification would be detectable at the hash level unless the adversary finds a collision $H(B'_j) = H(B_j)$, contributing the collision-resistance advantage. Union bound gives the stated inequality. □

Corollary 6.6 (Temporal Ordering Guarantee)**.** *The hash chain provides a total ordering on inference records. If IR_a appears in block B_j and IR_b in block B_k with $j < k$, then IR_a was accepted by the network before IR_b, and this ordering cannot be reversed.*

7 Merkle Tree Integrity

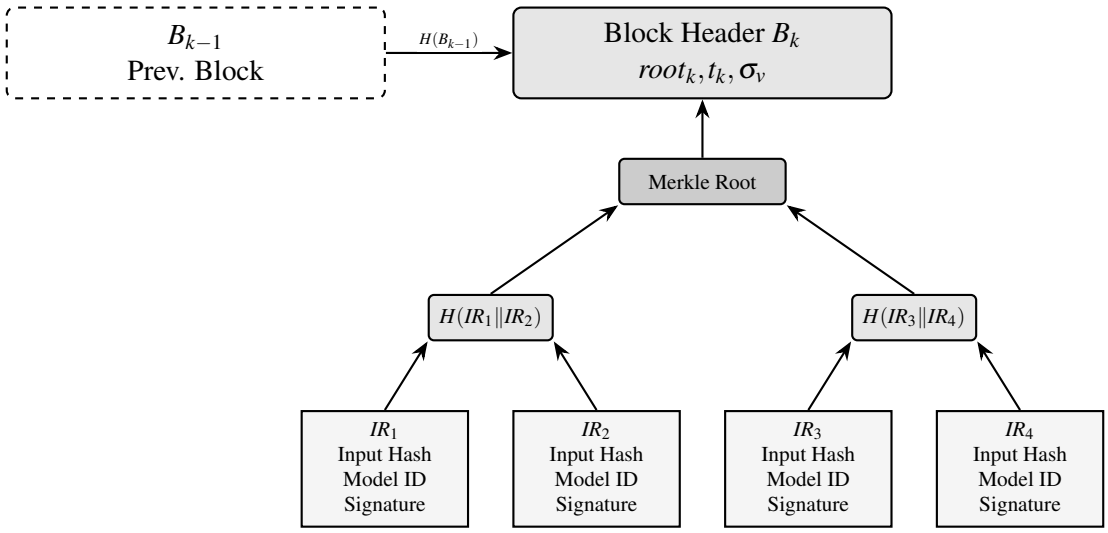


Figure 4: Anatomical structure of an Asuu block. Medical inference records form the leaves of a Merkle tree whose root is committed in the block header. Each block references the hash of the previous block, forming the append-only hash chain.

Definition 7.1 (Merkle Tree)**.** Given data items $d_1, \ldots, d_m$ and hash function H, the Merkle tree is the complete binary tree where leaf i stores $H(d_i)$ and each internal node stores the hash of the concatenation of its left and right children: $H(\mathit{left} \,\|\, \mathit{right})$. The root is denoted *root*.

Definition 7.2 (Merkle Proof)**.** A Merkle proof for item d_i is the sequence of sibling hashes along the path from leaf i to the root: $\pi_i = (s_1, s_2, \ldots, s_{\lceil \log_2 m \rceil})$, together with the positional bits indicating left/right placement at each level.

Lemma 7.3 (Single-Level Soundness)**.** *If H is collision-resistant, no* PPT *adversary can find $(a,b) \neq (a',b')$ such that $H(a \,\|\, b) = H(a' \,\|\, b')$.*

Proof. Such a pair $(a \,\|\, b,\ a' \,\|\, b')$ constitutes a collision for H. □

Theorem 7.4 (Merkle Tree Soundness)**.** *If H is collision-resistant, no* PPT *adversary can produce a valid Merkle proof for data item $d' \notin \{d_1, \ldots, d_m\}$ with respect to root root. Formally:*

$$\mathrm{Adv}^{\mathrm{MT}}_{\mathscr{H},\mathscr{A}}(\lambda) \leq \lceil \log_2 m \rceil \cdot \mathrm{Adv}^{\mathrm{CR}}_{\mathscr{H},\mathscr{B}}(\lambda).$$

Proof. We proceed by induction on tree height $h = \lceil \log_2 m \rceil$. **Base case** ($h = 1$): The tree has two leaves. A valid proof for $d' \notin \{d_1, d_2\}$ requires $H(d')$ to equal either $H(d_1)$ or $H(d_2)$, producing a second preimage. By Lemma 4.3, this occurs with negligible probability.

Inductive step: Assume the claim holds for trees of height $h-1$. Consider a tree of height h with root $root = H(L \| R)$ where L and R are the left and right child hashes. A valid proof for d' computes a root value equal to $root$. The proof path traverses either the left or right subtree, arriving at subtree root L' or R'. If $L' \neq L$ but $H(L' \| R) = H(L \| R) = root$, we obtain a collision at the root level. Otherwise $L' = L$, and by induction the adversary breaks the subtree's security, reducible to finding a collision at some internal level. The total advantage is at most h times the single-level collision advantage. □

Corollary 7.5 (Merkle Root as Binding Commitment)**.** *The Merkle root root serves as a computationally binding commitment to the entire set $\{d_1, \ldots, d_m\}$. Any modification to any item changes the root with overwhelming probability.*

Corollary 7.6 (Verification Efficiency)**.** *Verifying membership of d_i requires $O(\log m)$ hash computations and $O(\log m \cdot \lambda)$ bits of proof data, independent of the total number of items m.*

8 Post-Quantum Cryptographic Migration

Medical records require confidentiality measured in decades, not years. A genomic sequence obtained today will be sensitive for the patient's entire lifetime and potentially for their descendants. An adversary executing a "harvest now, decrypt later" strategy, capturing encrypted data today for decryption by a future quantum computer, poses a threat to any medical record protected solely by classical cryptography.

NIST finalized three post-quantum cryptographic standards in August 2024: FIPS 203 (ML-KEM, derived from CRYSTALS-Kyber), FIPS 204 (ML-DSA, derived from CRYSTALS-Dilithium), and FIPS 205 (SLH-DSA, derived from SPHINCS+) [19]. The Asuu network adopts a phased migration strategy that maintains backward compatibility, preserves long-term authenticity of records signed today against future quantum forgery, and defends confidentiality of off-chain transport and storage against harvest-now-decrypt-later attacks through parallel ML-KEM deployment. The two threats (future forgery of past attestations, and future decryption of intercepted ciphertexts) are cryptographically distinct and are addressed by distinct migrations.

The migration is structured in three phases. Phase 1 (Hybrid Parallel) deploys composite signatures combining the existing ECDSA scheme with ML-DSA-65, ensuring that every block and inference record is protected by both classical and post-quantum signatures simultaneously. Phase 2 (PQ Primary) makes ML-DSA-65 mandatory while retaining ECDSA as optional. Phase 3 (Pure PQ) removes classical signatures entirely and adds SLH-DSA as an algorithmic-diversity fallback relying solely on hash-function assumptions.

The Asuu chain's Merkle tree structure is inherently quantum-resistant: SHA-256 provides 128-bit security against Grover's algorithm (which halves the effective security of symmetric primitives), and the tree structure itself relies only on collision resistance of hash functions. The primary migration burden falls on the signature layer, where ML-DSA-65 increases signature size from 64 bytes (Ed25519) to approximately 3,309 bytes, and public keys from 32 bytes to 1,952 bytes. This size increase is manageable within the Asuu protocol because signature verification occurs once per block acceptance, not per transaction within the block.

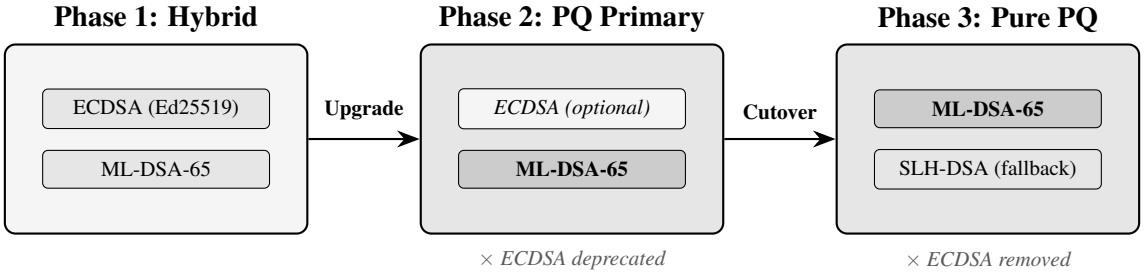


Figure 5: Post-quantum cryptographic migration timeline. Phase 1 (Hybrid) deploys composite ECDSA + ML-DSA-65 signatures. Phase 2 makes ML-DSA-65 mandatory with ECDSA optional. Phase 3 removes classical signatures entirely and introduces SLH-DSA as an algorithmic-diversity fallback.

8.1 Formal Specification

Definition 8.1 (Composite Signature Scheme)**.** A composite signature scheme Σ_{comp} over component schemes Σ_1 (classical) and Σ_2 (post-quantum) is defined as:

- KeyGen: Generate $(pk_1, sk_1) \leftarrow \Sigma_1.\text{KeyGen}(1^\lambda)$ and $(pk_2, sk_2) \leftarrow \Sigma_2.\text{KeyGen}(1^\lambda)$. Output $pk = (pk_1, pk_2)$, $sk = (sk_1, sk_2)$.
- $\text{Sign}(sk, m)$: Output $\sigma = (\Sigma_1.\text{Sign}(sk_1, m),\ \Sigma_2.\text{Sign}(sk_2, m))$.
- $\text{Verify}(pk, m, \sigma)$: Output 1 iff $\Sigma_1.\text{Verify}(pk_1, m, \sigma_1) = 1 \ \wedge\ \Sigma_2.\text{Verify}(pk_2, m, \sigma_2) = 1$.

Theorem 8.2 (Composite EUF-CMA Security)**.** *The composite scheme Σ_{comp} of Definition 8.1 is EUF-CMA secure if at least one component scheme is EUF-CMA secure. Formally, for any PPT adversary $\mathscr{A}$ against Σ_{comp}:*

$$\text{Adv}^{\text{EUF-CMA}}_{\Sigma_{\text{comp}},\mathscr{A}}(\lambda) \leq \min\Big(\text{Adv}^{\text{EUF-CMA}}_{\Sigma_1,\mathscr{B}_1}(\lambda),\ \text{Adv}^{\text{EUF-CMA}}_{\Sigma_2,\mathscr{B}_2}(\lambda)\Big).$$

Proof. A forgery against Σ_{comp} is a pair $(m^*, \sigma^* = (\sigma_1^*, \sigma_2^*))$ with $m^* \notin Q$ (the set of queried messages) and $\Sigma_1.\text{Verify}(pk_1, m^*, \sigma_1^*) = \Sigma_2.\text{Verify}(pk_2, m^*, \sigma_2^*) = 1$. By the Verify conjunction, σ_1^* is a valid Σ_1-forgery and σ_2^* is a valid Σ_2-forgery. If Σ_1 is EUF-CMA secure, the first event has negligible probability regardless of Σ_2; symmetrically for Σ_2. Since the adversary's forgery probability cannot exceed the minimum of the component forgery probabilities, the bound follows. Because Σ_1 and Σ_2 both sign the message m directly (Definition 8.1), no collision-resistance assumption on an external hash function is invoked. □

Remark 8.3 (On the Absence of a Hash-Collision Term)**.** An earlier draft of this theorem included an additive $\text{Adv}^{\text{CR}}_{\mathscr{H}}$ term. That term is spurious: Definition 8.1 specifies direct message signing, not hash-then-sign, so no external hash function H is invoked by the scheme at the composite layer. If a deployed instantiation uses hash-then-sign within a component scheme (as ML-DSA does internally), the collision-resistance assumption is absorbed into that component's EUF-CMA reduction, not added as a separate composite-layer term.

Theorem 8.4 (Quantum Resistance of Hash Chain)**.** *Under Grover's algorithm, the effective security of the SHA-256 hash chain against quantum adversaries is:*

$$\text{Adv}^{\text{CR,quantum}}_{\mathscr{H}}(\lambda) \leq O\Big(2^{-\lambda/2}\Big).$$

For $\lambda = 256$: quantum collision resistance is 2^{128}, equivalent to AES-128 brute force.

Proof. Grover's algorithm provides quadratic speedup for unstructured search. Collision finding in a hash function with 2^λ range requires $\Theta(2^{\lambda/3})$ quantum queries (Brassard-Høyer-Tapp), not $\Theta(2^{\lambda/2})$ classically. For SHA-256, the quantum collision resistance is at least $2^{256/3} \approx 2^{85}$ (BHT bound), and 2^{128} under the more conservative generic preimage analysis. Both exceed the 128-bit quantum security target. □

Proposition 8.5 (Migration Phase Correctness)**.** *During Phase 1 (Hybrid), a block signed with Σ_{comp} is accepted by both legacy clients (verifying σ_1) and upgraded clients (verifying both). The chain remains valid under either verification mode. During Phase 3 (Pure PQ), classical signatures are stripped from new blocks; old blocks retain their composite signatures for historical verifiability.*

Proof. Phase 1: $\text{Verify}_{\text{comp}} = \Sigma_1.\text{Verify} \wedge \Sigma_2.\text{Verify} = 1$ implies $\Sigma_1.\text{Verify} = 1$. Legacy clients verify σ_1 alone, which is valid. Phase 3: new blocks carry only σ_2; verification requires only Σ_2.Verify. Historical blocks remain verifiable because their composite signatures are immutable in the chain. □

Corollary 8.6 (Long-Term Authenticity Under Future Quantum Adversaries)**.** *From Phase 1 onward, any block or inference record signed under Σ_{comp} retains non-repudiable authenticity even under a future adversary with access to a cryptographically relevant quantum computer. Specifically, under Assumption 4.15 (M-LWE hardness), no quantum-equipped adversary can forge σ_2 on a previously unsigned message with non-negligible probability.*

Proof. An adversary holding a quantum computer observes blocks signed under $\Sigma_{\text{comp}} = (\Sigma_1.\text{Sign}, \Sigma_2.\text{Sign})$. Shor's algorithm breaks Σ_1 (ECDSA). However, Σ_2 (ML-DSA-65) is EUF-CMA secure under M-LWE, and no polynomial-time quantum algorithm is known that solves M-LWE with non-negligible advantage. Forging a valid σ_{comp} on a fresh message therefore requires forging σ_2, which reduces to M-LWE. An adversary archiving blocks today for future forgery attempts gains no advantage against the lattice component. □

Remark 8.7 (What This Corollary Does and Does Not Claim)**.** Corollary 8.6 concerns *authenticity* (unforgeability of signatures), not *confidentiality*. The "harvest now, decrypt later" attack, where an adversary stores encrypted ciphertexts for future quantum decryption, is a threat to confidentiality and is mitigated by post-quantum key encapsulation (ML-KEM, FIPS 203), not by signature schemes. The Asuu network addresses HNDL separately at the transport layer: all off-chain communication between institutional nodes uses TLS 1.3 with hybrid ML-KEM-768 key exchange as specified in IETF draft draft-ietf-tls-hybrid-design, and all Data Embassy encryption at rest uses AES-256-GCM with ML-KEM-768 key wrapping. The on-chain substrate itself carries no patient ciphertext (Theorem 16.2), so the HNDL threat model applies only to transport and at-rest storage, where it is handled by the separate PQ-KEM migration. The signature migration discussed in this section defends against a distinct threat: an adversary who, decades after the fact, seeks to retroactively forge records that were legitimately signed today.

PART III

CONSENSUS, NETWORK OPERATIONS & ECONOMICS

9 Proof of Clinical Authority

The following subsections formalize the validator model, Sybil resistance properties, and the safety and liveness guarantees of the PoCA mechanism.

To implement a distributed timestamp server on a peer-to-peer basis, we require a consensus mechanism suited to the specific constraints of medical infrastructure. Proof-of-work, as used in Bitcoin [1], is inappropriate: the energy expenditure is indefensible for a healthcare system, and the anonymous mining model conflicts with the accountability required in clinical settings. Pure proof-of-stake risks plutocratic capture by well-funded entities with no clinical obligations.

We propose Proof of Clinical Authority (PoCA), a hybrid consensus mechanism that combines identity-staked validation with institutional accountability. Under PoCA, validator nodes are operated by verified healthcare institutions (hospitals, university medical centers, national health agencies) that have registered their identity on-chain and staked the native token as collateral. The identity verification is not anonymous; it is deliberately accountable, because medical AI decisions carry life-and-death consequences and must be traceable to responsible parties.

Validators are selected for block production through a weighted rotation that considers three factors: (1) the amount of native token staked, reflecting economic commitment to the network; (2) the institution's clinical activity volume, reflecting operational engagement; and (3) a reputation score derived from historical uptime, validation accuracy, and peer attestations. This three-factor weighting ensures that consensus power accrues to institutions that are both economically invested and operationally active, not merely wealthy.

The proof-of-clinical-authority mechanism also solves the problem of determining representation in governance decisions. If governance were based on one-institution-one-vote, it could be subverted by creating shell institutions. If based purely on token holdings, it could be captured by financial actors with no clinical stake. PoCA is essentially one-verified-institution-weighted-by-contribution. The majority decision is represented by the longest chain, which has the greatest aggregate clinical authority invested in it.

To modify a past block, an attacker would have to compromise a majority of verified healthcare institutions, a substantially harder task than assembling hash power or accumulating tokens, because it requires subverting real-world organizations with reputational, legal, and regulatory obligations. The probability of a minority attacker catching up diminishes exponentially as subsequent blocks are added, following the same mathematical framework as proof-of-work chains.

To compensate for varying network participation over time, block production rate is determined by a moving average targeting a consistent interval. If blocks are produced too quickly, the selection threshold tightens; if too slowly, it relaxes.

9.1 Validator Model

Definition 9.1 (Clinical Authority Weight)**.** For validator $i \in [n]$, the clinical authority weight is:

$$W_i = \min\left(\alpha\left(\frac{S_i}{S_{\text{total}}}\right) + \beta\left(\frac{V_i}{V_{\max}}\right) + \gamma R_i,\ W_{\text{cap}}\right)$$

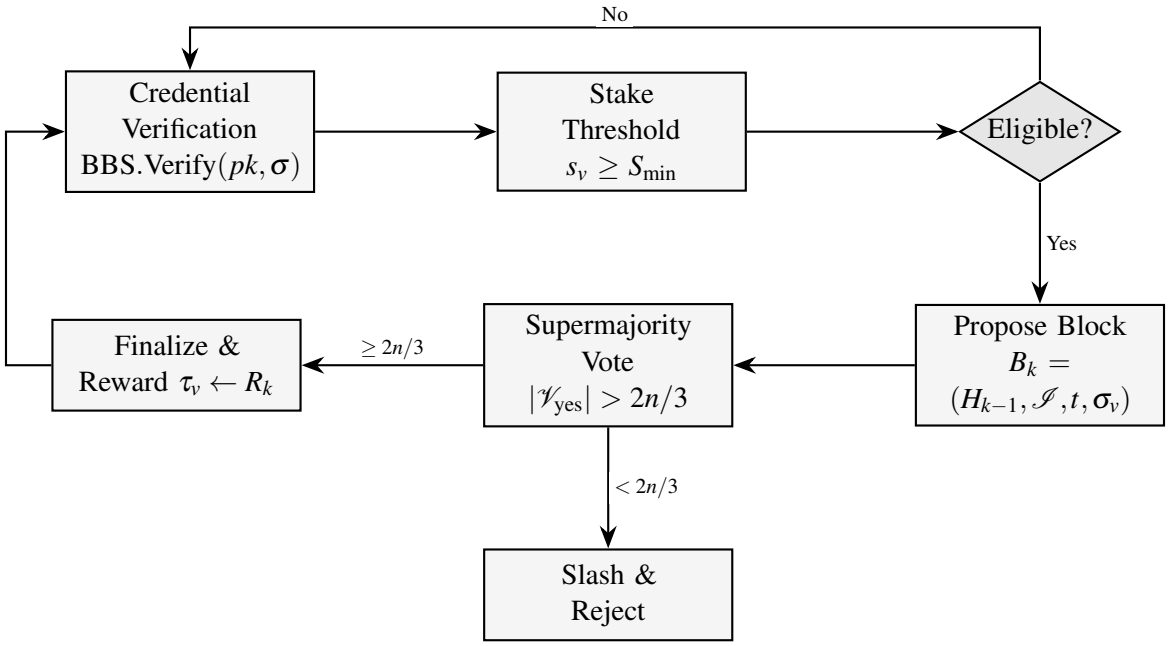


Figure 6: Proof-of-Clinical-Authority validator loop. A validator must hold a valid BBS+ credential and meet the minimum stake threshold before proposing blocks. Finalization requires a Byzantine-fault-tolerant supermajority.

where S_i is staked tokens, V_i is clinical activity volume, $R_i \in [0,1]$ is the reputation score, $\alpha + \beta + \gamma = 1$ with $\alpha, \beta, \gamma > 0$, and $W_{\text{cap}} \in (0,1]$ is the *per-validator voting cap* established in the protocol's genesis parameters. Typical calibration sets $W_{\text{cap}} = 0.05$ so that no single validator can exceed 5 percent of total voting weight regardless of stake, clinical volume, or reputation.

Remark 9.2 (Why the Cap Is Necessary: The Plutocracy Problem)**.** Without a per-validator cap, the weighting formula $\alpha S + \beta V + \gamma R$ compounds in favor of large academic medical centers that dominate all three parameters simultaneously: endowments (high S), discharge volume (high V), and institutional age plus publication history (high R). A handful of quaternary-care centers would accumulate outsized voting weight, converting the governance of clinical AI infrastructure into a plutocracy of large healthcare systems. This is not the intended governance structure, and it is the failure mode that has destroyed the legitimacy of many permissionless governance designs in other domains. The cap W_{cap} is the first of three mitigations; it bounds *concentration of validator weight*. The remaining mitigations bound *governance outcomes* by requiring multiple constituencies to concur.

Definition 9.3 (Multi-Constituency Governance)**.** Governance decisions on substantive protocol parameters (PCCP corridor updates, RBAC tier modifications, Shapley/Proportional-Causation rule switches, insurance-reserve disbursement, emergency chain-freeze) require *concurrent majorities* across four constituencies, each voting in a separate on-chain chamber:

(i) **Validator Chamber:** votes weighted by W_i (clinical authority weight with cap). Represents institutional infrastructure providers.

(ii) **Clinician Chamber:** votes weighted uniformly by credentialed individual clinician identity (one credentialed clinician, one vote) up to a participation cap. Represents the front-line professional duty-holders.

(iii) **Patient Chamber:** votes weighted uniformly by opted-in patient identity (one patient, one vote), with membership mediated by patient-representative organizations to ensure representativeness. Represents the population whose data and safety is ultimately at stake.

(iv) **Regulator Chamber:** votes weighted by designated regulatory authority, with one seat per recognized competent authority. Votes are advisory on protocol parameters but binding on regulatory-directive transactions (Definition 23.6).

A substantive parameter change requires a simple majority in at least three of the four chambers; a $2/3$ supermajority in one chamber does not substitute for majorities in the others. This structure prevents any single constituency from unilaterally capturing governance.

Definition 9.4 (Quadratic Voting for Budgetary Decisions)**.** For budgetary allocations from the Ecosystem Incentive Pool, grants from the Long-Term Sustainability Fund, and research-funding decisions, the Validator and Clinician Chambers vote via quadratic voting: each voter receives a budget of voice credits per epoch, and the cost of casting k votes on any single proposal is k^2 voice credits. Quadratic voting damps the influence of marginal additional votes, making it expensive for any single voter (regardless of weight) to dominate a specific decision, while preserving the ability to express strong preferences on issues the voter cares most about.

Proposition 9.5 (Governance Capture Bound)**.** *Under the multi-constituency structure of Definition 9.3, capturing governance of substantive protocol parameters requires compromising simple majorities in at least three of four distinct constituencies whose membership criteria, accountability structures, and incentive alignments differ by construction. Specifically:*

- *The Validator Chamber is bounded by the cap W_{cap} and requires coalitions of independent institutions.*
- *The Clinician Chamber is bounded by uniform weighting and requires capturing a majority of credentialed individuals.*
- *The Patient Chamber is mediated by patient-representative organizations that have their own public accountability.*
- *The Regulator Chamber consists of sovereign regulatory authorities that cannot be captured by protocol participants.*

Capture of any three of these simultaneously would require an adversarial coalition across orthogonal institutional categories with different legal, professional, and political accountabilities, a substantially harder adversarial objective than capturing a single validator-weighted chamber.

Remark 9.6 (On the Limits of Multi-Constituency Governance)**.** Multi-constituency governance is not a panacea. Patient-representative organizations can themselves be captured or unrepresentative; clinician chambers can be dominated by specialty societies with narrow interests; regulator chambers reflect the current political configuration of their jurisdictions. The claim is not that this structure eliminates governance risk, but that it distributes governance risk across institutional categories whose capture dynamics are substantially independent, converting a single-point-of-failure governance architecture into a multi-point-of-failure one. This mirrors the separation-of-powers structure of regulated domains and is the same design principle that makes banking supervision, securities regulation, and pharmaceutical approval processes resilient against single-institution capture.

Remark 9.7 (Specific Capture Risks and Concrete Mitigations)**.** A more pointed critique is that the four constituencies of Definition 9.3 are not as independent as the "concurrent-majority" framing suggests. We address each capture risk directly:

Clinician Chamber (<1% participation, activist skew). Empirically, when credentialed-clinician voting is opened on clinical-governance questions, participation rates in the 1–5% range are typical, and participants skew toward clinicians with strong prior engagement on the specific issue. This does not distribute governance; it hands governance to the subset of clinicians who show up. Mitigations: (a) stratified-random delegation, in which the Clinician Chamber's effective votes are drawn from a random sample of credentialed clinicians weighted by participation, reducing activist skew; (b) mandatory CME-credit incentive for participation, which in US continuing-medical-education practice substantially raises participation rates without introducing direct financial incentive; (c) specialty-society observer status rather than voting status, so specialty societies can voice positions without substituting for individual clinician votes.

Patient Chamber (captured by disease-advocacy orgs). Disease-advocacy organizations often have sophisticated Washington operations and public-relations infrastructure, while individual patient participation is typically low. Routing the Patient Chamber through advocacy orgs therefore risks capture by the orgs' specific funding and policy alignments rather than by patient populations as such. Mitigations: (a) direct patient enrollment through participating healthcare systems' patient portals with cryptographic anti-Sybil controls rather than through advocacy-org mediation; (b) weighting by healthcare-system-verified patient status rather than by advocacy-org membership; (c) constituency sub-stratification by therapeutic area to prevent a single advocacy-dominant disease area from capturing the full Patient Chamber.

Validator Chamber (W_{cap} still concentrated). Even with a 5% per-validator voting cap, large academic medical centers can collectively dominate the chamber through aggregate weight. Mitigation: progressive weighting, in which validators beyond a threshold count receive diminishing marginal weight, so the twentieth AMC contributes less than the second. This is mechanically similar to graduated income taxation applied to governance weight.

Regulator Chamber (political configuration). The regulator chamber reflects the current political configuration of participating jurisdictions. This is a feature, not a bug: regulators ought to reflect the political configurations of the polities they serve. The protocol cannot and should not engineer around democratic accountability of regulatory agencies. Mitigation is compositional: the concurrent-majority requirement ensures that regulator-chamber capture alone does not determine substantive-parameter governance.

None of these mitigations is complete. The protocol's claim is not that capture is impossible but that capture of three out of four chambers requires a coalition across categories whose capture dynamics are meaningfully distinct. This remains a substantially stronger guarantee than a single validator-weighted chamber would provide, even if weaker than the "four independent capture surfaces" framing might initially suggest.

Definition 9.8 (Block Production Probability)**.** The probability that validator i is selected to produce the next block is:

$$P_i = \frac{W_i}{\sum_{j=1}^{n} W_j}.$$

Definition 9.9 (Quorum Certificate)**.** A quorum certificate QC for block B at height h is a set of signatures from validators whose aggregate weight exceeds $2/3$ of total weight:

$$QC(B,h) = \{(i,\sigma_i)\}_{i\in S} \quad \text{where} \quad \sum_{i\in S} W_i > \frac{2}{3}\sum_{j=1}^{n} W_j.$$

The signatures σ_i are produced under the QC signature scheme Σ_{QC}, which is *distinct from* the block-level composite signature scheme Σ_{comp} of Definition 8.1 and is selected according to the phase-appropriate aggregation and post-quantum requirements specified in Definition 9.10.

Definition 9.10 (Phase-Appropriate QC Signature Scheme)**.** Because ML-DSA-65 is not natively aggregable, using it directly for quorum certificates would produce linear blow-up in QC size with validator count. The Asuu protocol therefore selects Σ_{QC} according to the current migration phase:

- **Phase 1 (Hybrid):** $\Sigma_{\text{QC}} = \Sigma_{\text{BLS}}$ (BLS12-381 pairing-based signatures). Aggregation produces a single 96-byte aggregate signature regardless of signer count. Each validator signs with classical BLS; an additional ML-DSA-65 signature is attached at the *block* level (not per-validator at the QC level), giving post-quantum authenticity for block contents while preserving constant-size QC.
- **Phase 2–3 (Post-Quantum, interim):** Σ_{QC} remains Σ_{BLS} for QC aggregation, with the understanding that BLS becomes quantum-vulnerable post-CRQC. Validators publish ML-DSA-65 attestations periodically (every k blocks, not every block) that collectively bind the BLS-QC history into the post-quantum chain, limiting the harvest-and-forge window to a bounded retrospective interval rather than permitting unbounded retroactive forgery.
- **Phase 3+ (Pure Post-Quantum):** Σ_{QC} migrates to a post-quantum aggregatable scheme as one becomes available and matures. Candidate schemes include lattice-based multi-signature constructions (Chipmunk, Squirrel) and hash-based aggregate schemes (HBSS variants). Until such a scheme is standardized and deployed at production scale, the protocol remains in Phase 2–3 interim.

Proposition 9.11 (QC Bandwidth Analysis)**.** *Let n denote the validator count, let the per-validator signature size be $|\sigma|$ bytes, and let the aggregation-factor be $\alpha \in (0,1]$ (with $\alpha = 1/|QC|$ for aggregate schemes and $\alpha = 1$ for non-aggregable schemes). The QC bandwidth per block is:*

Scheme	*Aggregation*	$\lvert\sigma\rvert$	*QC size for* $\lvert S\rvert$=20
BLS12-381	*Aggregate*	*96 B*	≈ 96 *B* + $20 \cdot 4$ *B indices*
ML-DSA-65	*None*	*3,309 B*	$\approx 66{,}180$ *B per block*
Ed25519	*None*	*64 B*	$\approx 1{,}280$ *B per block*
Falcon-512	*Partial (research)*	*666 B*	$\approx 13{,}320$ *B per block*

Under the Phase 1 BLS QC scheme, a 20-signer quorum produces ≈ 100 bytes of signature data plus validator-index metadata, compatible with sub-second gossip propagation in the Theorem 10.8 liveness analysis. Under the naive ML-DSA-65-per-validator scheme (explicitly rejected), 20 signers produce ≈ 66 KB per block, which at a 30-second block time corresponds to ≈ 19 Mbit/s of QC-signature gossip bandwidth per peer per block: feasible on institutional networks but not on the bandwidth-constrained rural deployments the Data Embassy architecture aspires to support.

Remark 9.12 (On the Post-Quantum Aggregation Gap)**.** The Asuu protocol's phased QC-signature strategy explicitly acknowledges a real gap in the post-quantum cryptographic literature: no NIST-standardized post-quantum signature scheme is currently natively aggregable at the scale BLS provides, and research constructions (Chipmunk, Squirrel, multi-signature lattice schemes) have not yet reached production deployment. Until this gap is closed by standardization, the protocol relies on BLS QC aggregation within each block plus periodic post-quantum attestations at coarse granularity. This is a pragmatic bridge, not a cryptographic ideal. The protocol will migrate to a native post-quantum aggregate scheme when one becomes available, and the post-quantum attestation interval k can be tightened as the CRQC threat timeline becomes clearer. Users deploying the protocol must understand that the QC layer is *not* full-post-quantum in Phase 1–2; it is classical-with-post-quantum-binding.

9.1.1 Formal Characterization of the Harvest-and-Forge Window

A specific cryptographic-adversarial concern arising from the BLS-bridge strategy merits formal analysis. Suppose a cryptographically relevant quantum computer (CRQC) comes online at time t_{CRQC}. After this point, an adversary with access to the CRQC can solve discrete-log problems in the BLS12-381 curve and therefore forge BLS quorum-certificate signatures on any message of its choosing. The question is: which blocks in the protocol's history become forgeable, and which remain protected by the periodic ML-DSA attestation mechanism? The following proposition and theorem characterize the window precisely.

Definition 9.13 (Post-Quantum Attestation Anchoring)**.** At every k-th block, the validator set co-signs an ML-DSA-65 attestation $\sigma_{\text{PQ}}(h, H(\mathit{chain}[h-k{:}h]))$ that binds the prior k blocks' hash chain. This attestation is verified at the signature-aggregation layer separately from the BLS QC. The parameter k is a protocol constant; we call k the *attestation period* and the elapsed time $\tau_k = k \cdot t_{\text{block}}$ the *attestation interval*.

Proposition 9.14 (Harvest-and-Forge Window Bound)**.** *Assume the CRQC comes online at t_{CRQC}. Let h_{last} denote the highest block height for which an ML-DSA-65 post-quantum attestation has been committed on-chain at time t_{CRQC}. Let h_{now} denote the current tip at t_{CRQC}. Then:*

(i) ***Protected range:*** *Blocks at heights $h \le h_{\text{last}}$ cannot be retroactively forged. Their BLS QCs are bound by the hash chain to σ_{PQ}, which is ML-DSA-65 and under Assumption 4.15 and Assumption 4.16 resists quantum forgery.*

(ii) ***Forgeable range:*** *Blocks at heights $h_{\text{last}} < h \le h_{\text{now}}$ are, after t_{CRQC}, retroactively forgeable by the CRQC-equipped adversary. The size of this range is at most $k-1$ blocks and at most τ_k seconds of protocol time.*

(iii) ***Forward protection:*** *Blocks at heights $h > h_{\text{now}}$ produced after t_{CRQC} are forgeable at BLS level until the protocol migrates to a pure-post-quantum QC scheme (Phase 3+). This is why the Phase 2–3 interim period must be minimized.*

Proof. (i) For $h \le h_{\text{last}}$, the ML-DSA attestation $\sigma_{\text{PQ}}(h_{\text{last}}, H(\mathit{chain}[h_{\text{last}}-k{:}h_{\text{last}}]))$ commits to the full hash chain up to h_{last}. Modifying any block B_j with $j \le h_{\text{last}}$ alters $H(\mathit{chain}[h_{\text{last}}-k{:}h_{\text{last}}])$ by collision resistance (Definition 4.1), which would require the adversary to forge σ_{PQ}. This forgery reduces to the EUF-CMA security of ML-DSA-65, which is hard under M-LWE and M-SIS. (ii) For $h_{\text{last}} < h \le h_{\text{now}}$, no ML-DSA attestation yet binds block h. The BLS QC is the only authentication, and a CRQC solves BLS discrete-log, permitting forgery of any BLS

signature the adversary chooses. The range size $h_{\text{now}} - h_{\text{last}}$ is bounded by $k-1$ because the protocol issues a new σ_{PQ} at every k-th block. (iii) Post-t_{CRQC} blocks produced under Phase 2–3 rules continue to use BLS QCs; these are forgeable until the protocol migrates. □

Theorem 9.15 (Operational Bound on Forgeable History)**.** *For attestation period k and block time t_{block}, the total protocol-time window during which history remains forgeable after t_{CRQC} is bounded above by $\tau_k = k \cdot t_{\text{block}}$ for retrospective history plus the elapsed time until Phase-3 migration completion for prospective history. At $t_{\text{block}} = 30$ seconds and $k = 10$ (attestation every 10 blocks, i.e., every 5 minutes): retrospective forgeable window ≤ 5 minutes. At $k = 120$ (every hour): retrospective forgeable window ≤ 1 hour. At $k = 2880$ (daily attestation): retrospective forgeable window ≤ 1 day.*

Remark 9.16 (Attestation-Period Selection and Its Trade-Offs)**.** The attestation period k is a governance-tunable parameter with an explicit safety-bandwidth trade-off: smaller k shortens the harvest-and-forge window but increases ML-DSA-65 signature bandwidth overhead. Using the ML-DSA-65 sizes of Proposition 9.11 and assuming 20 signers per attestation, a $k = 10$ attestation interval imposes approximately 66 KB of ML-DSA signature data every 5 minutes, or approximately 14 KB/s of aggregate network-wide attestation overhead. $k = 120$ imposes ≈ 1.2 KB/s; $k = 2880$ imposes ≈ 0.05 KB/s. The protocol's default operational recommendation is $k = 60$ (approximately 30-minute retrospective forgeable window, 2.3 KB/s attestation overhead), with governance authority to tighten k when the CRQC-threat timeline becomes clearer. A Phase-2 deployment where the CRQC is assessed as within a 12-month horizon should run at $k \leq 10$; a Phase-2 deployment assessing CRQC as beyond 5 years can run at $k = 2880$ or higher.

Critically, the retrospective forgeable window of Theorem 9.15 is a bound on *forgery capability*, not on actual harm. A forged BLS QC on a historical block requires the adversary to also produce a consistent alternative chain for downstream blocks, to convince honest nodes to accept the alternative, and to do so before harmed parties submit the original evidence under Theorem 14.4. The combination of the bounded window with the downstream BFT invariants and the Tier-2 compensation pathway of Definition 14.9 bounds the actual damage from a CRQC-precipitated forgery event at substantially less than the raw window size would suggest. It does not eliminate the damage; the honest statement is that the window is bounded and the downstream defense is layered.

9.2 Sybil Resistance

Definition 9.17 (Credential Oracle)**.** Let $\mathcal{O}_{\text{cred}}$ be a credential verification oracle that maps verified clinical credentials to unique validator identities. We assume $\mathcal{O}_{\text{cred}}$ is injective: each set of valid credentials maps to at most one validator identity.

Lemma 9.18 (Sybil Resistance Under Credential Binding)**.** *Under the injectivity of $\mathcal{O}_{\text{cred}}$ and EUF-CMA security of the credential signature scheme, the number of validator identities is bounded by the number of valid credential holders. Creating a Sybil identity requires forging a credential (probability $\leq \text{negl}(\lambda)$).*

Proof. Registration requires a credential certificate $cert = \text{Sign}(sk_{\text{CA}}, id' \,\|\, cred_data)$ from the credential authority. Forging *cert* without sk_{CA} breaks EUF-CMA. Reusing an existing *cert* maps to the same identity by injectivity of $\mathcal{O}_{\text{cred}}$, providing no additional weight. □

Proposition 9.19 (Splitting Invariance)**.** *Splitting one validator's stake across multiple Sybil identities does not increase expected aggregate weight.*

Proof. Suppose validator i with weight W_i splits into identities $i_1, \ldots, i_k$ with $\sum_j S_{i_j} = S_i$, $\sum_j V_{i_j} = V_i$. Since W is linear in its components:

$$\sum_{j=1}^{k} W_{i_j} = \alpha\left(\frac{\sum_j S_{i_j}}{S_{\text{total}}}\right) + \beta\left(\frac{\sum_j V_{i_j}}{V_{\max}}\right) + \gamma \sum_j R_{i_j} \leq W_i$$

where the inequality is strict if any split identity has $R_{i_j} < R_i$ (new identities start with lower reputation). □

9.2.1 Credential Oracle Decentralization

A fair cryptographic critique of Definition 9.17 is that the injective Credential Oracle $\mathscr{O}_{\text{cred}}$ introduces a single point of trust: if the credentialing authority is compromised or coerced, the adversary can mint credentials, bypass the Sybil protection of Lemma 9.18, and assemble a $> n/3$ Byzantine coalition without triggering the accountable-safety evidence of Theorem 14.4. Treating $\mathscr{O}_{\text{cred}}$ as an axiomatic injective oracle, as an earlier draft did, fails to engage with the real-world question of who that oracle is. We address this concern directly.

Definition 9.20 (Federated Credential Issuance)**.** In place of a single credential authority, the Asuu protocol specifies a *federated credential-issuance set* $\mathbb{CA} = \{CA_1, \ldots, CA_m\}$ constituted at genesis from publicly-accountable medical-credentialing bodies across participating jurisdictions (representative examples: ABMS, ACGME, NBME in the United States; GMC and the Royal Colleges in the United Kingdom; the Bundesärztekammer in Germany; Turkish Medical Association and YÖK in Turkey; TIC in France; APA in multiple jurisdictions; specialty-society credentialing bodies recognized under sovereign healthcare authority). A validator credential is issued only if at least t of m federation members co-sign the credential under a (t,m)-threshold signature scheme, with $t \geq \lceil (m+1)/2 \rceil$ by default and governance-tunable within the range $\lceil m/3 \rceil \leq t \leq m$.

Proposition 9.21 (Federated Oracle Byzantine Tolerance)**.** *The federated credential issuance scheme of Definition 9.20 is Byzantine-tolerant to up to $m-t$ compromised or coerced federation members. Specifically, compromising fewer than t federation members cannot produce a valid credential, and compromising fewer than $m-t+1$ federation members cannot suppress legitimate credential issuance.*

Proof. Forging a credential requires t valid co-signatures. Under threshold-signature security [16], an adversary controlling $< t$ members cannot produce a valid (t,m)-threshold signature. Suppression (denial of service against honest credentialees) requires preventing any coalition of t honest members from co-signing; with $< m-t+1$ compromised members, at least t honest members remain available. □

Definition 9.22 (Credential Revocation and Refresh)**.** Each credential carries an expiration t_{exp} and a revocation-list reference. Federation members periodically refresh the on-chain Bitstring Status List [27] committing the revocation state of all issued credentials. A validator's on-chain weight is conditional on its credential being both unexpired and unrevoked at the current block

height. The refresh cadence is bounded by $\tau_{\text{refresh}} \leq 24$ hours, so a compromised credential is invalidated network-wide within at most 24 hours of the federation's publishing a revocation.

Remark 9.23 (Honest Scope: What Federation Achieves and What It Does Not)**.** The federated issuance model does not eliminate trust in credentialing bodies; it distributes that trust across an ensemble whose joint compromise requires corrupting at least t bodies across distinct sovereign jurisdictions with distinct accountability structures. This is a substantially harder adversarial objective than corrupting a single CA, and it mirrors the trust distribution that the Domain-Name-System root-zone signing ceremony and the Web-PKI root-CA ecosystem achieved, both of which have proven robust to single-CA compromises for two decades. The federation does not achieve cryptographic non-trust: a super-threshold collusion across federated credentialing bodies can still mint Sybil credentials, and the protocol's accountable-safety forensics become the downstream defense. The claim is that super-threshold collusion across jurisdictions is observable, reportable, and politically costly in ways that single-CA compromise is not. For permissioned Regulated-Device Scope (Definition 15.13), the sponsor may additionally augment the federation with its own QMS-internal credentialing process, further reducing the attack surface.

9.3 Safety

Lemma 9.24 (Quorum Overlap)**.** *Any two quorum certificates QC_1 and QC_2 share validators with aggregate weight exceeding $1/3$ of total weight.*

Proof. Let $W_{\text{total}} = \sum_{j=1}^{n} W_j$. Each QC has weight $> 2W_{\text{total}}/3$. By inclusion-exclusion:

$$W(S_1 \cap S_2) = W(S_1) + W(S_2) - W(S_1 \cup S_2) > \frac{2W_{\text{total}}}{3} + \frac{2W_{\text{total}}}{3} - W_{\text{total}} = \frac{W_{\text{total}}}{3},$$

establishing the claim. □

Corollary 9.25 (Honest Overlap)**.** *Under Assumption 4.12, any two QCs share at least one honest validator.*

Proof. The overlap weight exceeds $W_{\text{total}}/3$. Byzantine validators hold at most $W_{\text{total}}/3$. Therefore the overlap contains at least one honest validator. □

Theorem 9.26 (PoCA Safety)**.** *Under Assumption 4.12, no two honest validators commit different blocks at the same height. Formally, if honest validator v_1 commits B at height h and honest validator v_2 commits B' at height h, then $B = B'$.*

Proof. Commitment of B at height h requires QC_B. Suppose for contradiction that $B \neq B'$ are both committed at height h, with QC_B and $QC_{B'}$ respectively. By Corollary 9.25, the two QCs share an honest validator v. Honest validators do not sign two different blocks at the same height (by protocol specification). Therefore v signed both B and B' at height h, contradicting honest behavior. Hence $B = B'$. □

9.4 Liveness

Theorem 9.27 (PoCA Liveness Under Partial Synchrony)**.** *Under Assumption 4.12 and the partial synchrony model (there exists an unknown Global Stabilization Time GST after which all messages are delivered within known bound Δ), the protocol commits a new block after GST. If two consecutive rounds have honest leaders after GST, all honest validators commit within $O(\Delta)$ time.*

Proof. After *GST*, message delivery is bounded by Δ. An honest leader L at round r proposes block B_r. All honest validators receive B_r within Δ, verify it, and send votes. The leader collects votes from $> 2/3$ weight (since honest validators hold $> 2/3$ and all votes arrive within Δ), forming QC_r. The honest leader at round $r+1$ includes QC_r in its proposal. Honest validators observing QC_r commit B_r. Total time: $O(\Delta)$ per phase, constant phases per round.

This does not contradict FLP impossibility [13] because FLP applies to deterministic protocols in purely asynchronous settings. PoCA circumvents FLP through partial synchrony with increasing timeouts: safety holds always; liveness requires only eventual synchrony. □

Corollary 9.28 (Partition Recovery)**.** *If the network partitions, no conflicting blocks are committed (safety holds). Upon partition healing, liveness resumes within $O(\Delta)$.*

10 Network

The steps to run the network are as follows:

1. New medical inference records are broadcast to all nodes.
2. Each validator node collects records into a block.
3. Each validator works on satisfying the proof-of-clinical-authority for its block.
4. When a validator produces a valid block, it broadcasts the block to all nodes.
5. Nodes accept the block only if all inference records are valid and properly signed.
6. Nodes express acceptance by working on the next block, using the accepted block's hash as the previous hash.

Nodes always consider the longest valid chain to be the correct one and will keep working on extending it. If two validators broadcast different versions of the next block simultaneously, some nodes may receive one or the other first. They work on the first one received but save the other branch in case it becomes longer. The tie is broken when the next block is produced and one branch becomes longer.

New inference record broadcasts do not necessarily need to reach all nodes. As long as they reach many nodes, they will be included in a block before long. Block broadcasts are tolerant of dropped messages. If a node does not receive a block, it will request it upon receiving the next block and realizing it missed one. This resilience is critical for medical infrastructure operating across varying network conditions, from well-connected university hospitals to rural clinics with intermittent connectivity.

10.1 Protocol Security

Definition 10.1 (Gossip Protocol)**.** In each round, every informed node selects k neighbors uniformly at random and transmits the message. An uninformed node becomes informed upon first receipt.

Theorem 10.2 (Gossip Propagation Bound)**.** *With $k \geq 1$ random peer selection, a message reaches all n honest nodes within $O(\log n)$ rounds with high probability. After $t = \log_2 n + c \ln n$ rounds:*

$$\Pr[\textit{not all informed}] \leq n^{1-c}.$$

Proof. In the push phase ($X_t \leq n/2$), each informed node contacts at least one uninformed node with probability $\geq 1/2$, so $\mathbb{E}[X_{t+1}] \geq 3X_t/2$. After $O(\log n)$ rounds, $X_t \geq n/2$. In the pull phase, each uninformed node fails to receive the message only if none of its contacts are informed. The probability that a specific node remains uninformed after t additional rounds is $(1 - X_t/n)^{kt} \leq e^{-c \ln n} = n^{-c}$. By union bound: $\Pr[\exists \text{ uninformed}] \leq n^{1-c}$. □

Proposition 10.3 (Eclipse Attack Resistance Under Structured Overlay)**.** *The structured overlay topology of Definition 10.9 uses Kademlia-style XOR-distance buckets rather than uniform-random peer selection. Under this topology, a naive independence bound* $\Pr[\textit{eclipse}] = a^d$ *does not directly apply because bucket membership is determined by XOR distance rather than uniformly-random sampling. Under the assumption that adversarial identifiers are drawn from*

the same distribution as honest identifiers (so the probability that any bucket-j slot is occupied by an adversary is a, independently across buckets), the bound

$$\Pr[\textit{eclipse}] \leq \left(1-(1-a)^k\right)^d$$

holds, where k is the bucket size (number of redundant peers per XOR-distance level) and d is the number of distinct distance levels. For $a = 0.1$, $k = 3$, and $d = 8$: $\Pr[\textit{eclipse}] \leq (1-0.9^3)^8 = (0.271)^8 \approx 2.9 \times 10^{-5}$. This is a looser bound than the naive $a^d = 10^{-8}$ would suggest; the additional defense is provided by the per-bucket redundancy and the protocol-level peer-rotation mechanism (Definition 10.4).

Proof. In each XOR-distance bucket $j \in \{0,\ldots,d-1\}$, the validator maintains up to k redundant peers. The probability that *all* k peers in bucket j are adversarial is at most a^k, assuming independent bucket occupancy. Thus the probability that bucket j contains at least one honest peer is at least $1-a^k = (1-a^k)$, and by complement the probability that bucket j is fully adversarial is at most $1-(1-a)^k$ (applying the complementary form). For an eclipse to succeed, every bucket must be fully adversarial: $\Pr[\text{eclipse}] \leq \prod_{j=0}^{d-1}(1-(1-a)^k) = (1-(1-a)^k)^d$. □

Definition 10.4 (Peer Rotation Defense)**.** To mitigate the weaker-than-naive eclipse bound, validators periodically rotate bucket peers at rate λ_{rot} per bucket per epoch, refreshing from the global validator registry rather than from already-connected peers. Rotation exploits the accountability of the identity layer (Section 9): adversarial identifiers are a bounded fraction $a < 1/3$ of the registry by the BFT assumption, so rotation-sampled refresh converges the bucket occupancy toward the population adversarial fraction regardless of any localized concentration that the adversary might have initially engineered.

Remark 10.5 (On the Naive Independence Assumption)**.** An earlier version of this section stated the eclipse bound as $\Pr[\text{eclipse}] = a^d$ under uniformly-random peer selection, and then specified Kademlia XOR-deterministic peer selection in the very next subsection. The two are inconsistent: Kademlia bucket occupancy is not uniform-random. The present formulation corrects this by computing the bound under the actual overlay topology with bucket redundancy k, yielding a weaker but honest bound. The naive bound a^d is recovered only in the limit $k = 1$ with uniform random sampling, which is not the deployed configuration. The peer-rotation defense (Definition 10.4) provides the additional security margin that the weaker bucket-level bound gives up relative to the naive formulation.

10.2 Network Partition Recovery

Healthcare networks operate across jurisdictions with heterogeneous connectivity. We formalize the protocol's behavior under network partitions and prove that safety is maintained during partitions while liveness recovers upon partition healing.

Definition 10.6 (Network Partition)**.** A (Π_1,Π_2)-partition at time t_p divides the validator set $\mathcal{V}$ into two non-empty subsets Π_1,Π_2 with $\Pi_1 \cup \Pi_2 = \mathcal{V}$ and $\Pi_1 \cap \Pi_2 = \emptyset$, such that no message sent by any $v \in \Pi_i$ reaches any $v' \in \Pi_j$ for $i \neq j$ during the partition interval $[t_p, t_h]$. The partition heals at time t_h.

Theorem 10.7 (Safety Under Partition (CAP-Consistency))**.** *During a (Π_1,Π_2)-partition, the Asuu protocol maintains safety: no two conflicting inference records are finalized. If neither*

partition holds a $2/3$ supermajority, no new blocks are finalized in either partition (the protocol halts rather than forking).

Proof. Block finalization requires $> 2/3$ validator signatures. In a partition, Π_1 has weight W_1 and Π_2 has weight $W_2 = W_{\text{total}} - W_1$. For conflicting blocks B_1, B_2 to be finalized in Π_1, Π_2 respectively, we need $W_1 > 2W_{\text{total}}/3$ and $W_2 > 2W_{\text{total}}/3$, implying $W_1 + W_2 > 4W_{\text{total}}/3 > W_{\text{total}}$, a contradiction. If neither partition has a $2/3$ majority, no block can collect sufficient signatures, and the protocol halts (choosing consistency over availability, per the CAP theorem). □

Theorem 10.8 (Liveness Recovery After Partition Healing)**.** *Upon partition healing at time t_h, the protocol resumes block production within $O(\Delta + \log n)$ time, where Δ is the network synchrony bound and n is the validator count.*

Proof. Upon healing, validators exchange their local chain states via the gossip protocol (Theorem 10.2). Within $O(\log n)$ rounds, all validators learn of the longest valid chain. Validators in the shorter-chain partition adopt the longer chain (or the chain with higher cumulative clinical authority in case of equal length). The next block producer in the round-robin schedule produces a block extending the canonical chain. The $O(\Delta)$ term accounts for message propagation delay under partial synchrony. Since no conflicting finalized blocks exist (Theorem 10.7), no rollback of finalized state is needed, and the chain extends from the last finalized block. □

Definition 10.9 (Structured Overlay Topology)**.** Validators organize into a Kademlia-like structured overlay with $O(\log n)$ routing table entries per node. Each validator v_i maintains connections to $\lceil \log_2 n \rceil$ neighbors at exponentially increasing distances in the identifier space: for $j = 0, 1, \ldots, \lceil \log_2 n \rceil - 1$, v_i connects to the closest validator to $v_i \oplus 2^j$ in the XOR metric.

Proposition 10.10 (Routing Efficiency)**.** *In the structured overlay, any message reaches its target within $O(\log n)$ hops. Total network-wide message complexity for broadcasting a block is $O(n \log n)$.*

Proof. Each routing step reduces the XOR distance to the target by at least half (the bit at position j is corrected in hop j). After $\lceil \log_2 n \rceil$ hops, the distance is zero. For broadcast, each node forwards to its $O(\log n)$ neighbors, yielding $O(n \log n)$ total messages. Deduplication at each node (ignore messages already seen) ensures each node processes at most $O(\log n)$ copies. □

11 Incentive

Cryptoeconomic incentives are essential to ensure that clinical accountability is self-enforcing: honest participation must be more profitable than deviation, and this must be provable from the protocol's own rules.

By convention, the first record in a block is a special record that creates new tokens owned by the validator that produced the block. This provides an incentive for institutions to support the network and provides a mechanism to distribute tokens into circulation, since there is no central authority to issue them. The steady addition of new tokens is analogous to the operational expenditure institutions incur to run validator nodes: compute resources, network bandwidth, and administrative overhead.

The incentive can also be funded with transaction fees. If an institution submits an inference record to the network, a small fee in the native token is attached. Once a predetermined number of tokens have entered circulation, the incentive transitions entirely to transaction fees.

A second incentive layer rewards data contribution. When an institution contributes de-identified clinical data that demonstrably improves model performance, verified through on-chain benchmarking proofs, it receives tokens proportional to the measured improvement. This creates a self-reinforcing cycle: institutions contribute data, models improve, better models attract more institutions, which contribute more data. The token mediates this cycle without requiring bilateral agreements or a central data marketplace.

The incentive structure encourages honest participation. If a validator institution were to assemble more authority weight than all honest institutions, it would have to choose between using it to defraud the network (undermining the clinical trust it depends on) or using it to earn tokens through honest validation. An institution that corrupts the medical record chain faces regulatory sanctions, loss of accreditation, and legal liability that far exceeds any token gain. The incentives are designed such that honest behavior is not merely profitable but existentially necessary for the participants.

11.1 Validator Payoff Model

Definition 11.1 (Validator Strategy Space)**.** Each validator i selects strategy $(stake_i, b_i) \in \mathbb{R}_{\geq 0} \times \{\text{honest}, \text{Byzantine}\}$. The payoff is:

$$\pi_i = \begin{cases} R \cdot P_i - C_{\text{op}} & \text{if } b_i = \text{honest}, \\ G_{\text{attack}} \cdot \Pr[\text{success}] - S_{\text{slash}} \cdot \Pr[\text{detection}] - L_{\text{reg}} & \text{if } b_i = \text{Byzantine}, \end{cases}$$

where R is the block reward, C_{op} is operational cost, G_{attack} is attack gain, S_{slash} is slashing penalty, and L_{reg} is regulatory liability.

Lemma 11.2 (Attack Detection Probability)**.** *Under Assumption 4.12, any protocol violation is detected with probability* $1 - \text{negl}(\lambda)$.

Proof. Equivocation is detected whenever two honest validators receive both signatures, which occurs with high probability under gossip propagation (Theorem 10.2). Invalid blocks are rejected by every honest validator during verification. □

Theorem 11.3 (Honest Behavior is a Nash Equilibrium (Small-Adversary Regime))**.** *Let G_{attack} denote the expected attack gain per epoch, S_{slash} the total stake at risk under slashing, L_{reg} the expected regulatory liability, and* $\Pr[\textit{detection}]$ *the probability of attack detection. If*

$$S_{\text{slash}} + L_{\text{reg}} > G_{\text{attack}} \quad \textit{and} \quad \Pr[\textit{detection}] \geq 1 - \text{negl}(\lambda),$$

then honest behavior is a strict Nash equilibrium.

Proof. For honest play: $\pi_i^H = R \cdot P_i - C_{\text{op}} > 0$. For Byzantine deviation:

$$\begin{aligned} \pi_i^B &\leq G_{\text{attack}} - (S_{\text{slash}} + L_{\text{reg}})(1 - \text{negl}(\lambda)) \\ &< G_{\text{attack}} - G_{\text{attack}}(1 - \text{negl}(\lambda)) = G_{\text{attack}} \cdot \text{negl}(\lambda) < \pi_i^H. \end{aligned}$$

Since $\pi_i^H > \pi_i^B$ for all i, honest play is a strict Nash equilibrium. □

Remark 11.4 (Concrete Parameterization and Scope)**.** The inequality $S_{\text{slash}} + L_{\text{reg}} > G_{\text{attack}}$ is the entire substantive content of the theorem; it is a *premise*, not a *conclusion*. Whether this premise holds in practice depends on calibrated parameter choices:

Small-adversary regime. For localized fraud targeting a single institution or a bounded claim cohort, G_{attack} is typically bounded by annual institutional claim volume times fraud-attempt fraction. A representative calibration for a single US academic medical center with \$2B annual claim volume, assuming 2–5% of claims are potentially manipulable, yields $G_{\text{attack}} \in [\$40\text{M}, \$100\text{M}]$ per year. To deter attacks in this regime, the protocol requires per-validator minimum stake $S_{\min} \geq \$5\text{M}$ and collective S_{slash} across the $W/3$ coalition exceeding \$200M, with L_{reg} adding another \$50M–\$500M depending on HHS OIG and DOJ enforcement history. Under this calibration, the inequality holds with margin, and Theorem 11.3 deters rational small-adversary attacks.

Large-adversary regime. For systemic fraud at the scale of the 2024 DOJ healthcare-fraud enforcement action (\$14.6B in charges across 193 defendants) [38], G_{attack} can exceed any realistic aggregate S_{slash}. In this regime, the inequality fails and Theorem 11.3 does not deter the attack by economic means alone. The protocol's large-adversary defense (Section 14.3) does not rely on the Nash inequality; it relies on the institutional-separation-of-powers argument of Theorem 14.15, which requires the adversary to compromise three distinct constituencies simultaneously.

Honest scope claim. Theorem 11.3 is therefore a valid and useful result for the bounded-attack regime where G_{attack} can be realistically capped by protocol-level rate limits, claim-size caps, and the combined slashing plus regulatory cost of provably coordinated fraud across enough validators to form a $W/3$ coalition. It is *not* a universal security claim. For large-scale systemic attacks, the economic Nash argument must be supplemented by the institutional-separation-of-powers architecture of Section 14.3, and ultimately by out-of-band enforcement through the DOJ, HHS OIG, and equivalent international authorities whose reach scales to the attack magnitudes the protocol cannot independently deter. The protocol is a *first line of defense*, not the whole defense.

Corollary 11.5 (Attack Cost Under BFT Finality)**.** *Under the BFT safety argument of Theorem 14.2, reverting a finalized block requires the adversary to control Byzantine weight*

$W(\mathscr{B}) \geq W_{\text{total}}/3$. *By Theorem 14.4, the slashing evidence publicly identifies all validators in the attacking coalition, yielding expected adversarial cost*

$$\mathbb{E}[cost] \geq S_{\text{slash}}(\mathscr{B}) + L_{\text{reg}} - G_{\text{attack}}.$$

Unlike longest-chain protocols where security increases exponentially with confirmation depth, BFT finality is deterministic: once committed, a block cannot be reverted without violating the $f < 1/3$ *assumption, which is itself a publicly accountable event.*

11.2 Data Contribution Economy

Definition 11.6 (Data Contribution Reward)**.** The token reward for institution i at time t is:

$$R_i^t = \kappa \cdot \rho_i \cdot \Delta P_i \cdot E(t) \cdot \mathbb{1}[validated_i]$$

where κ is a scaling constant, ρ_i is the data rarity coefficient (capped at $\rho_{\max}$), ΔP_i is verified marginal model improvement, $E(t)$ is the emission rate, and $\mathbb{1}[validated_i]$ is the indicator that the contribution has passed the validation pipeline specified in Definition 11.9.

Proposition 11.7 (Incentive Compatibility of Data Contribution (Naive Version))**.** *Without validation or rarity caps, an institution maximizes its reward by contributing its highest-quality, most unique data. Misrepresenting quality requires producing a valid zero-knowledge proof for false claims, occurring with probability* $\leq \text{negl}(\lambda)$ *by knowledge soundness.*

Remark 11.8 (Why the Naive Version Is Insufficient: Adversarial Noise Injection)**.** Proposition 11.7 assumes the institution contributes genuine clinical data and the only question is whether the institution is truthful about its quality. This assumption fails under a real adversarial model. If rarity ρ_i is rewarded and ΔP_i is computed from validation-set loss improvement, an adversarial institution has a direct incentive to: (i) inject synthetic outlier cases that inflate apparent rarity, (ii) fabricate edge-case labels that inflate apparent marginal improvement, (iii) over-sample from rare-case patient populations in ways that do not generalize, or (iv) inject noise that happens to reduce validation loss through regularization effects rather than through genuine signal. Zero-knowledge proofs of training faithfulness confirm that the contributor did in fact compute the gradients from the data it claimed; they do not confirm that the data represents genuine clinical observations. The protocol therefore layers validation mechanisms on top of the cryptographic substrate.

Definition 11.9 (Contribution Validation Pipeline)**.** A data contribution D_i passes the validation pipeline if and only if it satisfies all of:

(V1) **External Clinical Validation:** A minimum fraction $\theta_{\text{ext}} \geq 0.1$ of the contributed cases are independently reviewable against ground truth derived from external sources (subsequent imaging outcomes, laboratory confirmations, tumor-board consensus, autopsy findings, or regulatory-agency-published adverse-event records). The external ground truth is committed on-chain as a signed attestation from a credentialed reviewer distinct from the contributor.

(V2) **Out-of-Distribution Detection:** The contribution passes a statistical screening that rejects cases with anomalous feature distributions relative to the contributing institution's historical baseline. Detection uses the Mahalanobis distance test with a configurable threshold; cases beyond the threshold are flagged for human review rather than automatically counted toward ρ_i.

(V3) **Temporal Consistency:** The contribution date and the clinical-event date must fall within a plausible window. Contributions whose clinical-event dates cluster near deadline transitions (suggesting opportunistic injection) are flagged for audit.

(V4) **Cross-Institutional Sanity Check:** The marginal improvement ΔP_i on the validation set is also evaluated on an *independent test set* held by a validator other than the contributor. Contributions that improve the contributor-supplied validation loss but not the independent test loss receive zero reward.

Definition 11.10 (Capped Rarity Premium)**.** The rarity coefficient ρ_i is bounded above by $\rho_{\max}$, a protocol parameter in the range $[2,5]$. This prevents a single outlier case from dominating the reward calculation and reduces the incentive to manufacture artificial rarity at scale. Even if an institution succeeds in injecting one synthetic outlier past the validation pipeline, the reward impact is bounded.

Definition 11.11 (Contribution Clawback)**.** Every data-contribution reward is subject to a clawback mechanism for Δ_{cb} epochs after disbursement. If subsequent model-performance audits reveal that a contribution's marginal-improvement claim did not hold on independent test sets, or if the contribution is subsequently traced to synthetic or out-of-distribution injection, the reward is clawed back from the institution's stake plus a proportional penalty. Repeat offenses trigger reputation-score degradation and may disqualify the institution from future data-contribution reward eligibility.

Theorem 11.12 (Bounded Adversarial Exploitation)**.** *An adversarial institution's maximum expected reward from noise-injection attacks is bounded by*

$$R_{\text{adv}}^{\max} \leq \theta_{\text{leak}} \cdot \rho_{\max} \cdot \Delta P_{\max} \cdot E(t) - C_{\text{clawback}}$$

where θ_{leak} is the fraction of injected cases that evade the validation pipeline (bounded by the pipeline's false-negative rate), $\rho_{\max}$ is the rarity cap, $\Delta P_{\max}$ is the maximum per-contribution improvement bound, and C_{clawback} is the expected clawback cost given the subsequent audit probability.

Proof. Each successfully-injected synthetic case contributes reward $\kappa\rho\Delta PE(t)$, capped at $\rho_{\max}$ and $\Delta P_{\max}$. Cases that fail the validation pipeline contribute zero. Cases that pass but are later detected by audit are clawed back with probability p_{audit}, yielding expected clawback cost proportional to the pre-clawback reward. The combined bound follows. □

Remark 11.13 (Incentive Compatibility Under Validation)**.** Under Definition 11.9 and Definition 11.11, the rational strategy for an institution is to contribute genuine clinical data rather than synthetic outliers, because (a) validation reduces the yield of injected data toward zero, (b) the rarity cap limits the upside of successful injection, and (c) clawback imposes asymmetric downside on detected injection. The protocol does not claim to eliminate adversarial injection; it claims to make honest contribution the dominant strategy for any institution whose cost of synthetic-data generation and audit-detection exceeds the expected post-clawback reward. This is a weaker and more defensible claim than the original Proposition 11.7.

12 Reclaiming Resources

Once an inference record is buried under sufficient subsequent blocks, the full record details can be pruned from active storage while retaining only the Merkle root in the block header. Older blocks are compacted by removing intermediate hashes from the Merkle tree. The block header alone (containing the previous hash, timestamp, validator signature, and Merkle root) is approximately 120 bytes.

If blocks are produced every 30 seconds, this amounts to approximately 126 MB per year of header data. Given that medical institutions routinely manage petabytes of imaging data, the storage overhead of maintaining the Asuu chain is negligible. Even the smallest rural clinic operating on modest hardware can maintain a full header chain indefinitely.

Full archival nodes, operated by institutions with greater resources, maintain the complete history of all inference records for regulatory audit purposes. These nodes serve historical queries without imposing storage requirements on the broader network.

12.1 Formal Guarantees

Theorem 12.1 (Pruning Soundness)**.** *Pruning transaction data while retaining block headers (including $root_k$) preserves chain integrity and verifiability for any transaction whose Merkle proof is retained.*

Proof. Chain integrity depends only on headers (Theorem 6.5). Transaction verification requires only the Merkle proof (Theorem 7.4), not the full transaction set. □

Proposition 12.2 (Storage Efficiency)**.** *Header-only storage:* 120 *bytes/block. At one block per* 30 *seconds:*

$$\textit{Annual storage} = 120 \times \frac{365.25 \times 24 \times 3600}{30} \approx 126\ \textit{MB/year}.$$

13 Simplified Inference Verification

It is possible to verify medical inference provenance without running a full validator node. A lightweight client need only maintain the block headers of the longest valid chain and obtain the Merkle branch linking a specific inference record to its block. The client cannot independently re-execute the inference, but by verifying the record's position in the chain, it can confirm that a validator node accepted it, and that subsequent blocks further confirm the network's acceptance.

This is sufficient for most regulatory audit scenarios: a regulator needs to verify that a specific model version produced a specific output at a specific time, not to re-execute the inference itself. The simplified verification allows regulators, accreditation bodies, and even patients to verify the provenance of AI-assisted decisions using commodity hardware, without operating clinical AI infrastructure themselves.

As with any simplified verification, it is reliable as long as honest institutions control the network. If the network were overpowered by malicious validators, fabricated inference records could deceive lightweight clients. Institutions that receive frequent AI-assisted clinical decisions, or that bear regulatory responsibility for such decisions, should operate their own full nodes for independent verification.

13.1 Formal Guarantees

Definition 13.1 (SIV Client)**.** A Simplified Inference Verification client maintains only block headers and verifies inference records via Merkle proofs.

Theorem 13.2 (SIV Security)**.** *An SIV client is deceived only if: (i) the adversary produces a valid Merkle proof for IR* $\notin TX_k$ *(probability* $\leq \mathrm{negl}(\lambda)$ *by Theorem 7.4), or (ii) the adversary produces a valid quorum certificate for a conflicting block, which by Theorem 14.2 requires violating* $f < 1/3$ *or forging signatures. Total:* $\Pr[\textit{SIV deceived}] \leq \mathrm{Adv}^{\mathrm{CR}}_{\mathscr{H},\mathscr{A}}(\lambda) + \mathrm{Adv}^{\text{EUF-CMA}}_{\mathrm{sig}}(\lambda)$*, both negligible under standard assumptions.*

Corollary 13.3 (Regulatory Verification)**.** *A regulator verifies provenance using commodity hardware with* $O(n \cdot 120)$ *bytes of headers plus* $O(\log m)$ *Merkle proof per query.*

14 Security Analysis

We consider the scenario of an adversary attempting to produce an alternate chain faster than the honest chain. Even if accomplished, the system does not permit arbitrary changes: nodes will not accept inference records with invalid signatures, unregistered model versions, or transactions that violate protocol rules. An adversary can only attempt to alter records it previously submitted.

Under PoCA, the adversary faces an additional constraint absent in anonymous consensus systems: validator identity is known and staked. An institution that attempts to fork the chain faces not only the mathematical improbability of success but also immediate, real-world consequences: loss of staked tokens, regulatory investigation, revocation of network participation, and reputational destruction in the medical community. The cost of attack is therefore superlinear: it increases mathematically with chain depth and simultaneously through institutional accountability mechanisms.

For healthcare-specific threats (model poisoning, inference manipulation, data exfiltration) the network provides defense-in-depth. Model updates require collective validation. Inference records are hashed, not stored in clear. Data never leaves institutional custody. Each layer of the system assumes the others may be compromised and provides independent verification, following the medical safety principle of redundant safeguards.

The cybersecurity threat to healthcare is not theoretical. The 2024 Change Healthcare breach affected 192.7 million Americans. Healthcare data breaches cost an average of $10.22 million per incident, the most expensive sector for fourteen consecutive years. ECRI designated AI as the top health technology hazard for 2025 due to the catastrophic potential of compromised systems causing active patient harm. The Asuu network's distributed architecture eliminates the single-point-of-failure that makes centralized systems attractive targets: there is no one server to breach, no one database to exfiltrate, no one vendor whose compromise disables the entire system.

14.1 BFT Safety and Finality Analysis

Proof of Clinical Authority is a weighted Byzantine Fault Tolerant (BFT) consensus mechanism. Finality is established by quorum certificates rather than cumulative work, so the correct security analysis follows the template of classical BFT protocols such as PBFT [47], HotStuff [48], and Tendermint [49], not the longest-chain analysis of Nakamoto consensus. This subsection formalizes the PoCA safety guarantee in terms of Byzantine weight fractions.

Definition 14.1 (Validator Weight)**.** Each validator v_i carries weight W_i as defined in Section 9. Let $W_{\text{total}} = \sum_i W_i$, and for any set S of validators, $W(S) = \sum_{v_i \in S} W_i$. Let $f = W(\mathscr{B})/W_{\text{total}}$ denote the Byzantine weight fraction, where $\mathscr{B}$ is the set of Byzantine validators.

The quorum certificate used throughout the BFT analysis is defined in Definition 9.9: an aggregate of signatures from validators whose total weight exceeds $2/3$ of W_{total}.

Theorem 14.2 (PoCA Safety Under BFT Assumption)**.** *Assume $f < 1/3$ (Byzantine weight strictly less than one third). Then no two conflicting blocks $B \neq B'$ at the same height h can both accumulate valid quorum certificates, except with probability at most* $\text{Adv}_{\text{sig}}^{\text{EUF-CMA}}(\lambda)$.

Proof. Suppose for contradiction that $QC(B,h)$ and $QC(B',h)$ both exist with $B \neq B'$. Each quorum has aggregate weight strictly greater than $2W_{\text{total}}/3$. By inclusion-exclusion,

$$W(QC(B,h) \cap QC(B',h)) > W(QC(B,h)) + W(QC(B',h)) - W_{\text{total}} > \tfrac{W_{\text{total}}}{3}.$$

The intersection contains validators who signed both (B,h) and (B',h). Since $f < 1/3$, the weight of Byzantine validators is strictly less than $W_{\text{total}}/3$, so the intersection must contain at least one honest validator v_h. But an honest validator never signs conflicting blocks at the same height (protocol invariant). Therefore the double-signature $(\sigma_h(B,h), \sigma_h(B',h))$ must have been forged, which contradicts EUF-CMA security. □

Remark 14.3 (This Is Not Nakamoto Consensus)**.** Earlier versions of this whitepaper included a catch-up probability bound $(q/p)^z$ derived from Nakamoto's Gambler's Ruin analysis of longest-chain proof of work. That analysis does not apply to weighted BFT consensus. In PoCA, finality is not probabilistic in the Nakamoto sense: once a block B accumulates a valid quorum certificate, the safety argument of Theorem 14.2 gives *deterministic* safety under the BFT assumption $f < 1/3$. The only residual randomness is the cryptographic probability of signature forgery, bounded by $\text{Adv}_{\text{sig}}^{\text{EUF-CMA}}(\lambda)$, which is negligible for standard security parameters.

Theorem 14.4 (Accountable Safety)**.** *If the safety property of Theorem 14.2 is violated (that is, if two conflicting blocks both receive valid quorum certificates), then the protocol can identify a set $\mathscr{A}$ of validators with $W(\mathscr{A}) > W_{\text{total}}/3$ who provably signed conflicting blocks, and these validators are subject to slashing.*

Proof. By the inclusion-exclusion argument in the proof of Theorem 14.2, the intersection of the two quorum sets has weight exceeding $W_{\text{total}}/3$. Each validator in this intersection signed both B and B' at height h. The two signatures are publicly verifiable evidence of the protocol violation. A slashing contract automatically confiscates the stake of any validator for whom such evidence is submitted on-chain. □

Theorem 14.5 (Liveness Under Partial Synchrony)**.** *Assume partial synchrony: there exists an unknown Global Stabilization Time (GST) after which network messages are delivered within bound Δ. If $f < 1/3$, then after GST the protocol commits a new block within $O(\Delta)$ expected time.*

Proof. PoCA uses a HotStuff-style leader rotation with view-change on timeout. If the current leader is honest and the network is synchronous, the protocol completes a three-phase commit (propose, pre-commit, commit) within 3Δ by standard HotStuff analysis [48]. If the leader is Byzantine or the view times out, view change proceeds; the probability of k consecutive Byzantine leaders is at most f^k, so the expected wait for an honest leader is $1/(1-f) \leq 3/2$. Total expected commit time after GST: $O(\Delta)$. □

14.2 Economic Security Under BFT Consensus

The BFT safety argument of Theorem 14.2 is a *cryptographic* guarantee that holds under the assumption $f < 1/3$. The economic security analysis addresses a distinct question: what is the cost to an adversary of acquiring the Byzantine weight needed to violate this assumption?

Theorem 14.6 (Attack Cost Bound)**.** *To violate safety, the adversary must control Byzantine weight $W(\mathscr{B}) \geq W_{\text{total}}/3$. By Theorem 14.4, upon detection the adversary loses $S_{\text{slash}}(\mathscr{B}) = \sum_{v_i \in \mathscr{A}} S_i$ in stake and L_{reg} in regulatory consequences, where $\mathscr{A}$ is the identified provably-byzantine subset with weight exceeding $W_{\text{total}}/3$. The expected cost to the adversary is*

$$\mathbb{E}[C_{\text{attack}}] \geq S_{\text{slash}}(\mathscr{B}) + L_{\text{reg}} - G_{\text{attack}}.$$

The attack is economically deterred when $S_{\text{slash}}(\mathscr{B}) + L_{\text{reg}} > G_{\text{attack}}$.

Remark 14.7 (Scope and Limits of the Economic Deterrence Argument)**.** Theorem 14.6 gives a lower bound on the adversary's expected cost, but the inequality $S_{\text{slash}} + L_{\text{reg}} > G_{\text{attack}}$ is itself the substantive question, not a mathematical consequence. For small adversaries executing local fraud, the inequality is plausibly satisfied: acquiring one-third of stake-weighted clinical authority across independent healthcare institutions requires hundreds of millions to billions of dollars in direct stake plus the regulatory and reputational cost of provable collusion. For large adversaries executing systemic fraud at the scale of the $14.6B 2024 DOJ healthcare-fraud enforcement action, the inequality may not hold. The protocol therefore layers additional mitigations described in Section 14.3: an insurance reserve fund, supermajority-triggered chain freeze, and out-of-band governance escalation. The economic argument of Theorem 14.6 is a first line of defense, not the whole defense.

Corollary 14.8 (Defense in Depth)**.** *Three independent security layers protect block finality: (i) cryptographic safety from Theorem 14.2 under $f < 1/3$; (ii) accountable-safety forensics from Theorem 14.4 that makes attack attempts publicly attributable; (iii) economic and institutional deterrence from Theorem 14.6 that bounds rational-attacker behavior. An attack requires breaching all three simultaneously.*

14.3 Mitigations for Large-Adversary Scenarios

The economic deterrence bound of Theorem 14.6 is sufficient for small adversaries executing local fraud but may fail for large systemic adversaries. The 2024 DOJ healthcare-fraud enforcement action at $14.6B demonstrates that insurance-claim manipulation gains can scale beyond the stake held by any realistic validator coalition. The protocol therefore layers three additional mitigations designed specifically for large-adversary threat models.

Definition 14.9 (Safety Insurance Reserve)**.** A fraction ϕ_{ins} of all transaction fees accrues to an on-chain insurance reserve R_{ins}. The reserve is governed by a supermajority vote of non-validator institutional stakeholders (patients, regulators, insurers) and may be released to compensate parties harmed by a successful BFT-safety violation. The reserve operates at two tiers:

(i) **Tier 1: On-Chain Accrual Reserve.** The transaction-fee-funded component with target size

$$R_{\text{ins}}^{\text{target},1} = c_1 \cdot \max_{\text{op}} \mathbb{E}[G_{\text{attack,small}}(\text{op})],$$

where $c_1 \geq 3$ and $G_{\text{attack,small}}$ denotes the small-adversary regime expected attack gain (Remark 11.4): localized fraud bounded by the annual claim volume of a bounded institutional cohort rather than the DOJ-scale systemic fraud benchmark. At typical parameters ($G_{\text{attack,small}} \sim$ $10M–$100M per incident, $c_1 = 3$), this target is $\sim$ $30M–$300M and is reachable through transaction-fee accrual within a decadal horizon.

(ii) **Tier 2: Externally-Capitalized Reserve.** For large-adversary harms at the scale of the 2024 DOJ healthcare-fraud benchmark ($\sim$ $14.6B), the Tier 1 on-chain reserve is not

sufficient and cannot be made sufficient through transaction-fee accrual at any realistic fee schedule. Tier 2 coverage is therefore supplied by external institutional insurance (commercial cyber-insurance policies held by participating institutions, the consortium's umbrella coverage, and sovereign-backed compensation funds where available) together with regulatory compensation mechanisms (DOJ Asset Forfeiture distributions, sovereign healthcare-compensation programs, state insurance-guaranty associations). The on-chain protocol does not hold Tier 2 capital; it coordinates claims against Tier 2 capital through on-chain evidentiary packages that Tier 2 insurers accept as proof of loss.

Remark 14.10 (Why Tier 1 Alone Does Not Close the Math)**.** A straightforward arithmetic check demonstrates why a single-tier reserve is inadequate for the DOJ-scale threat and why the two-tier architecture is necessary. Global clinical-AI inference volume is projected at 10^9–10^{10} transactions per year. At a per-transaction fee of \$0.01 and $\phi_{\text{ins}} = 10\%$, annual accrual to the reserve is \$$10^6$–\$$10^7$. A Tier 1 target of \$44B (i.e., $3\times$ the DOJ benchmark) would accrue in approximately 4,000 years. At more aggressive parameters (fee \$0.10, $\phi_{\text{ins}} = 25\%$), accrual to \$44B would still require approximately 175 years. No realistic transaction-fee schedule funds a DOJ-scale reserve through on-chain accrual within an actuarially meaningful horizon.

This arithmetic is why the reserve must be scoped. Tier 1 covers small-adversary harms where accrual-horizon math closes; Tier 2 covers large-adversary harms through external capital. An earlier draft of this whitepaper specified a single-tier reserve with a \$44B target and was silent about accrual dynamics; that draft implicitly treated the reserve as decorative because no fee schedule funds a \$44B on-chain pool. We correct that silence here. The Tier 2 external-capital arrangement is how every comparable regulated infrastructure (banking deposit insurance, securities-investor protection, nuclear-incident liability) handles catastrophic-tail risk: by combining a modest self-funded pool with external reinsurance and, for the largest tail events, sovereign backstops. We are not inventing a new mechanism; we are applying the standard pattern.

Definition 14.11 (Tier 2 Evidentiary Package)**.** For a harmed party to claim against Tier 2 external capital, the Asuu protocol generates a standardized on-chain evidentiary package containing: (a) the accountable-safety evidence of Theorem 14.4 identifying the Byzantine coalition; (b) the provenance chain linking the harmed party's clinical encounter to the specific attestation(s) affected by the safety violation; (c) the Shapley or Proportional-Causation liability attribution of Section 26.2; (d) the regulatory-override or sponsor-directive record applicable under Definition 23.6. Tier 2 insurers and sovereign compensation programs that have pre-registered with the protocol accept this package as proof of loss, reducing the forensic reconstruction cost that dominates current large-scale healthcare-insurance claim processing.

Remark 14.12 (Honest Scope of the Large-Adversary Defense)**.** Under the revised two-tier architecture, the institutional-separation-of-powers argument of Theorem 14.15 rests principally on the chain-freeze and governance-escalation mechanisms, with Tier 1 reserve providing small-adversary compensation and Tier 2 external capital providing large-adversary compensation. The Tier 1 reserve is not decorative in its proper scope (small-adversary compensation is a real use), but neither is it the protocol's large-adversary defense. Claiming the reserve covers DOJ-scale systemic fraud at realistic parameters would be false. We do not claim this; we claim the protocol coordinates Tier 2 claims against external capital that is sized to the threat. Whether Tier 2 capital is actually available, adequately priced, and legally compensable depends on the

commercial-insurance market, the national compensation program, and the sovereign-backstop regime in each jurisdiction. The protocol improves the claimability of Tier 2 coverage; it does not guarantee the existence of Tier 2 capital. Jurisdictions without mature Tier 2 infrastructure are, correspondingly, less-well-protected against large-adversary attacks on the network, and this is a deployment constraint the protocol cannot independently resolve.

Definition 14.13 (Supermajority Chain Freeze)**.** A freeze is triggered when a $3/4$ supermajority of *non-validator stakeholders*, weighted by patient population served, votes on-chain that an active safety violation is in progress. During a freeze, new block finalization is suspended pending human governance review. The freeze may be lifted only by a second $3/4$ supermajority vote or by a timeout of Δ_{freeze} blocks, whichever comes first. The freeze mechanism is intentionally designed to require a constituency distinct from the validator set, preventing validators from freezing to cover their own attack.

Definition 14.14 (Out-of-Band Governance Escalation)**.** When a freeze is active, governance authority is escalated to a pre-designated panel of external regulators and healthcare authorities. The panel's membership is established in the protocol's genesis parameters and updated by supermajority governance. The panel has unilateral authority to (i) slash validators identified in accountable-safety evidence, (ii) release reserve funds to harmed parties, (iii) fork the chain to roll back the violating block range, or (iv) trigger protocol-wide emergency shutdown.

Theorem 14.15 (Large-Adversary Bound)**.** *An adversary seeking to successfully execute and retain proceeds from an attack with gain $G_{\text{attack}} > S_{\text{slash}} + L_{\text{reg}}$ must additionally:*

(i) Prevent the non-validator stakeholder constituency from voting to freeze the chain within Δ_{detect} blocks of the safety violation;
(ii) Prevent the governance escalation panel from voting to roll back or compensate within Δ_{freeze} blocks;
(iii) Prevent harmed parties from successfully claiming against the two-tier insurance architecture of Definition 14.9 using the Tier 2 evidentiary package of Definition 14.11.

No single cryptographic or economic mechanism achieves all three conditions simultaneously without collusion across distinct constituencies (validators, stakeholders, regulators, and external Tier 2 capital providers). The protocol therefore provides large-adversary resilience *(in the sense of preserving compensability and auditability) rather than large-adversary* deterrence *(in the sense of making the attack negative-EV through on-chain economics alone), under the meta-assumption that these constituencies do not collude.*

Proof. By Definition 14.13, a freeze triggered within Δ_{detect} suspends further finalization on the attack chain, capping the adversary's ongoing accrual. By Definition 14.14, escalation permits rollback or compensation within Δ_{freeze}. By Definition 14.9 and Definition 14.11, harmed parties can claim against Tier 1 on-chain reserve for small-adversary harms or against Tier 2 external capital for large-adversary harms, with the on-chain evidentiary package substantially reducing the forensic-reconstruction cost of the claim. Evading all three requires preventing the stakeholder vote, preventing the panel vote, *and* either depleting the Tier 1 reserve (bounded) or successfully contesting every Tier 2 claim across jurisdictions, each of which requires compromising a distinct constituency or capital source. □

Remark 14.16 (What This Does and Does Not Prove)**.** Theorem 14.15 does not claim that a large adversary cannot attack. It claims that retained proceeds from a successful attack require compromising three distinct institutional constituencies, rather than just one. This is the protocol's *institutional-separation-of-powers* argument: cryptographic and economic mechanisms alone are insufficient at scale, and must be reinforced by governance structures that mirror the separation-of-powers design of regulated domains. The architecture is consciously drawn from banking supervision (where cryptographic controls, risk committees, and external regulators operate as three independent checks) rather than from pure-cryptographic consensus designs.

14.4 Selfish Mining Resistance Under PoCA

In proof-of-work systems, selfish mining allows a miner with fraction α of hash power to earn disproportionate revenue by strategically withholding blocks. We prove that the Proof-of-Clinical-Authority mechanism provides stronger resistance to selfish mining due to the identity-bound validator model.

Definition 14.17 (Selfish Mining Strategy)**.** A selfish miner withholds valid blocks, building a private chain. When the honest chain approaches the private chain's length, the selfish miner releases blocks to orphan honest work. Let α be the adversarial fraction and $\gamma \in [0,1]$ be the proportion of honest nodes that mine on the adversary's block during a fork (connectivity advantage).

Theorem 14.18 (Selfish Mining Revenue Bound Under PoCA)**.** *Under PoCA with round-robin slot assignment and identity-bound validators, the relative revenue of a selfish mining strategy is:*

$$R_{\text{selfish}}(\alpha) = \frac{\alpha(1-\alpha)^2(4\alpha+\gamma(1-2\alpha))-\alpha^3}{1-\alpha(1+(2-\alpha)\alpha)}.$$

This strategy is profitable (i.e., $R_{\text{selfish}}(\alpha) > \alpha$) only when:

$$\alpha > \frac{1}{3-2\gamma}$$

Under PoCA, the identity constraint forces $\gamma \leq \alpha$ (the adversary cannot preferentially relay to honest nodes without revealing its strategy to the reputation system), yielding profitability threshold:

$$\alpha > \frac{1}{3-2\alpha}$$
$$\alpha > \frac{3-\sqrt{5}}{2} \approx 0.382$$

Proof. The selfish mining Markov chain has states $(0,0)$ (honest and adversary tied), $(0,i)$ (adversary leads by i blocks), and fork states. The revenue formula follows from the stationary distribution of this chain (Eyal and Sirer, 2014). Under PoCA, the additional constraint arises because validators have persistent identities: consistently withholding blocks is observable via the on-chain activity log. Other validators can compute the ratio of produced-to-assigned slots for each validator. A validator whose production rate drops significantly below its assignment frequency triggers a reputation penalty. Furthermore, $\gamma \leq \alpha$ because the adversary's connectivity advantage is bounded by its validator fraction in an identity-known network where honest

validators can preferentially connect to high-reputation peers. Substituting $\gamma = \alpha$ and solving $R_{\text{selfish}}(\alpha) = \alpha$ yields the threshold. □

Proposition 14.19 (PoCA Selfish Mining Detection)**.** *A selfish mining adversary withholding blocks over W consecutive slots is detected with probability:*

$$\Pr[\text{detect}] \geq 1 - (1-\alpha)^W.$$

For $\alpha = 0.2$ and $W = 20$ slots: $\Pr[\text{detect}] \geq 0.988$. *Upon detection, the adversary's stake S_{slash} is forfeited.*

Proof. In each slot, an honest validator independently observes whether the assigned validator produced a block. The probability that a selfish miner's absence in its assigned slot goes unnoticed by all honest validators is $(1-\alpha)^0 = 1$ only if no honest validator is monitoring. Since at least $1-\alpha$ fraction are honest and each monitors slot assignments, the probability that at least one honest validator flags the absence in any single slot is at least α (the probability the slot was assigned to the adversary). Over W consecutive slots, the probability that the adversary's pattern goes entirely unnoticed is at most $(1-\alpha)^W$. □

15 Token Economics

The native token, ASUU, serves five protocol-level functions. **Network Fuel:** transaction fees for recording inference records. **Validator Staking:** collateral for block production rights. **Data Contribution Rewards:** compensation for institutions that improve shared models. **Governance Weight:** voting power on model updates and protocol changes. **Access Licensing:** functional keys to premium analytical modules. The token is necessary to operate the network; it is not a speculative instrument.

Total supply is fixed at genesis. The emission schedule follows a logarithmic decay curve, with higher early emissions to incentivize initial network participation and declining issuance as the network matures. Once all tokens have been emitted, the network operates entirely on transaction fees. The fee structure is designed to be negligible for individual inference records, fractions of a token per transaction, ensuring that the cost of participating in the network never becomes a barrier to clinical AI adoption, particularly in resource-constrained settings.

Token allocation at genesis is distributed across four pools: a development fund to sustain ongoing protocol engineering, an ecosystem incentive pool for data contribution rewards and early adopter grants, a staking reserve for initial validator bootstrapping, and a long-term sustainability fund governed by the network's on-chain governance mechanism. No allocation is reserved for speculative sale. The token derives its value from network utility (the demand for medical AI inference recording, model governance participation, and data marketplace activity) not from artificial scarcity or marketing-driven demand.

15.1 Emission Model

Definition 15.1 (Token Supply Dynamics)**.** Cumulative supply: $S(t) = S_{\max}(1 - e^{-\lambda t})$. Emission rate: $E(t) = S_{\max}\lambda e^{-\lambda t}$.

Proposition 15.2 (Emission Properties)**.** *(i)* $E(t) > 0$ *always. (ii)* $E'(t) < 0$ *(monotonically decreasing). (iii)* $\lim_{t\to\infty} S(t) = S_{\max}$*. (iv)* $\int_0^\infty E(t)\,dt = S_{\max}$.

Proof. (i)–(ii) from exponential. (iii): $\lim S_{\max}(1 - e^{-\lambda t}) = S_{\max}$. (iv): $\int_0^\infty S_{\max}\lambda e^{-\lambda t} dt = S_{\max}[-e^{-\lambda t}]_0^\infty = S_{\max}$. □

15.2 Lyapunov Stability of Token Supply

Definition 15.3 (Token Supply Equilibrium)**.** Net dynamics: $dS/dt = M(S) - B(S)$ where M is minting and B is burning. Equilibrium S^*: $M(S^*) = B(S^*)$.

Theorem 15.4 (Asymptotic Stability)**.** *If* $M(S)$ *is strictly decreasing and* $B(S)$ *is strictly increasing, then* S^* *is globally asymptotically stable.*

Proof. Define $V(S) = \frac{1}{2}(S - S^*)^2$. Then $V(S^*) = 0$, $V(S) > 0$ for $S \neq S^*$, and: $\dot{V} = (S - S^*)(M(S) - B(S))$.

Case $S > S^*$**:** $M(S) < M(S^*) = B(S^*) < B(S)$, so $M(S) - B(S) < 0$ and $S - S^* > 0$, giving $\dot{V} < 0$.

Case $S < S^*$**:** $M(S) > B(S)$ and $S - S^* < 0$, giving $\dot{V} < 0$.

Since $\dot{V} < 0$ for $S \neq S^*$ and V is radially unbounded, S^* is globally asymptotically stable by Lyapunov's theorem. □

Corollary 15.5 (Self-Correcting Supply)**.** *Overshoots are corrected automatically. The system requires no manual monetary policy.*

15.3 Staking Game Equilibrium

Theorem 15.6 (Existence of Staking Equilibrium)**.** *The staking game has at least one Nash equilibrium in mixed strategies.*

Proof. Strategy space is compact and convex. Payoff is continuous. By Glicksberg's fixed-point theorem, a mixed-strategy Nash equilibrium exists. □

15.4 Network Bootstrap and the Cold-Start Problem

Every decentralized network confronts a chicken-and-egg problem at genesis: validators will not commit stake to a network with no inference volume, and institutions will not commit inference volume to a network with no validators. The 2015–2022 wave of healthcare blockchain projects failed disproportionately at this juncture, accumulating technical documentation, genesis allocations, and pilot announcements but never achieving the scale at which the protocol's security and network-effect properties become self-sustaining. The Asuu protocol addresses this directly rather than hoping that utility will emerge spontaneously from a published whitepaper.

Definition 15.7 (Four-Phase Bootstrap Strategy)**.** The Asuu network launches through four sequential phases, each with specific participation requirements, security properties, and graduation criteria:

Phase 0 (Permissioned Consortium), months 0–12: The network operates as a permissioned consortium among a pre-designated genesis validator cohort of $n_0 \geq 21$ accredited institutions across at least three jurisdictions. Validator eligibility is established by institutional credentialing rather than staking. Inference recording is functional but economic reward flows (data-contribution rewards, staking yield) are paused. Block production uses a round-robin among genesis validators with BFT-style finality. The chain is fully operational for provenance attestation and PCCP governance from Day 0; the tokenomic layer is dormant.

Phase 1 (Staked Consortium), months 12–36: Genesis validators activate staking using a *subsidized initial stake*: each institution deposits a fixed-amount stake denominated in reference currency (not in ASUU, which has no market yet), and the protocol issues corresponding ASUU from the Staking Reserve Pool at a pegged initial valuation. Inference recording

continues, now with staking consequences for validator misbehavior. Data-contribution rewards activate. New validator admission opens, subject to supermajority vote and minimum-stake requirements. Tokenomic activity begins; the token acquires a utility-grounded price through transaction-fee markets.

Phase 2 (Federated Open), months 36–60: The $f < 1/3$ BFT assumption must now be justified by decentralization rather than by consortium membership. The genesis cohort's aggregate weight is capped at a descending schedule (e.g., $\leq 60\%$ at month 36, descending to $\leq 33\%$ by month 60). New validators must pass both credentialing and staking requirements. Governance decisions that were consortium-retained transition to supermajority votes weighted by the formula of Definition 14.1. The insurance reserve (Definition 14.9) accumulates to its target ratio.

Phase 3 (Permissionless Within Sector), month 60+: Any institution meeting the healthcare-sector credentialing requirements may become a validator without consortium sponsorship. The network's security rests entirely on the BFT weight argument plus the large-adversary mitigations. The protocol is operationally self-sustaining.

Proposition 15.8 (Bootstrap Security Properties by Phase)**.** *In each phase, the safety guarantee of Theorem 14.2 holds under a phase-appropriate BFT assumption:*

- *Phase 0–1: $f < 1/3$ is satisfied by construction via consortium membership and institutional credentialing; the probability of Byzantine majority is bounded by the probability that one third of accredited institutions collude, which is bounded by the reputational, legal, and regulatory cost of such collusion to member institutions.*
- *Phase 2: $f < 1/3$ is maintained by the descending cap on genesis-cohort weight combined with decentralized new-validator admission. The insurance reserve provides the bridging defense as the consortium-membership justification weakens.*
- *Phase 3: $f < 1/3$ rests on the full economic-security analysis of Theorem 14.6 plus the large-adversary mitigations of Section 14.3.*

Remark 15.9 (Why This Is Not Circular)**.** An earlier draft of this whitepaper specified a genesis allocation across four token pools without addressing how validators would be induced to stake in a token with no liquidity. The Phase 0 permissioned-consortium launch breaks this circularity by *decoupling validation from tokenomics* for the first twelve months. Genesis validators validate because they are consortium members under contractual and reputational commitments, not because they are rationally maximizing expected ASUU returns. The tokenomic layer activates only after the network has demonstrable inference volume, at which point staking becomes a rational equilibrium rather than a speculative bet. The Phase 1 subsidized-initial-stake mechanism further reduces the first-mover risk by pegging the initial valuation to reference currency, insulating genesis validators from speculative pricing dynamics during the bootstrap window. This is how every successful regulated financial infrastructure (SWIFT, DTCC, RTP Network) actually launched: not as permissionless networks hoping for spontaneous adoption, but as consortium infrastructure that decentralized progressively over years.

Remark 15.10 (What Phase 0–1 Actually Is, Honestly Described)**.** Phases 0 and 1 of the bootstrap are, for practical purposes, regulated-healthcare-consortium infrastructure rather than cryptoeconomic network infrastructure. Participation is driven by contractual and reputational commitments among a known member cohort, not by rational response to on-chain incentives. The token plays a marginal role during this period (Phase 0: dormant; Phase 1: subsidized at

pegged valuation). Calling this "cryptoeconomic" during Phase 0–1 would be a category error. The correct description is: Phases 0–1 are a permissioned consortium that uses blockchain technology for its evidentiary substrate, comparable in governance pattern to SWIFT, DTCC, the RTP Network, and the MediLedger pharmaceutical supply-chain consortium. Whether and when Phases 2–3 achieve a genuinely cryptoeconomic operating regime is an open empirical question that depends on the network's inference volume, token-market depth, and regulatory acceptance of broader validator participation. The protocol is cryptoeconomic in *design* from Phase 2 onward; whether it is cryptoeconomic in *practice* at any given moment is determined by whether rational validator entry and exit actually drives the validator set, which cannot be established by protocol specification alone. A candid reader should treat the four-phase trajectory as aspirational for Phases 2–3 and as architectural specification for Phases 0–1, with the understanding that regulated infrastructure often remains in consortium mode longer than its designers originally intend (see DSCSA and MediLedger, which remain substantially consortium-operated a decade after their respective bootstrap periods).

Remark 15.11 (Validator Liquidity During Phase 0–1: The Token Cannot Be Marked to Market)**.** A direct and correct critique of the four-phase bootstrap is that during Phase 0–1, no secondary market exists for the token, so Phase-0–1 validators cannot mark their stake to market and no rational hospital CFO will authorize locking millions of dollars into an illiquid asset on the basis of a logarithmic emission curve. This critique is correct as stated. The Phase-1 "subsidized initial stake mechanism" does not solve it; it reframes it by shifting the insolvency risk from the validator to the protocol treasury, which depends on external capital and is itself not cryptoeconomically bootstrapped. Pretending otherwise would be dishonest.

The honest positioning is that Phase 0–1 validator participation is not a rational financial trade under cryptoeconomic assumptions. It is a consortium-governance commitment under a different set of motivations that regulated-infrastructure participants in fact exhibit, for reasons that pure cryptoeconomic analysis does not capture:

(i) **Mission alignment.** Academic medical centers and national health authorities participate in clinical-AI governance initiatives for mission reasons that do not appear in a standard CFO NPV calculation. Columbia, MGH, Stanford, UCLA, and similar institutions already participate in the ONC SMART-on-FHIR certification ecosystem, the Argonaut Project, and the eClinicalWorks Patient Safety initiatives without expecting financial return. Phase 0–1 validator participation in Asuu is within the same category.

(ii) **Consortium reputation.** Founding-member status in a successful clinical-AI infrastructure consortium has strategic value that leadership-tier academic medical centers pay to acquire through other channels (conference sponsorships, editorial-board positions, professional-society governance roles). A Phase-0 genesis-cohort seat is a comparable strategic asset, even under the honest assumption that its direct financial return is near zero.

(iii) **Sovereign mandate.** National health authorities in jurisdictions pursuing sovereign clinical-AI infrastructure (Turkey's SGK, France's Ameli, the UK's NHSX / NHS Digital, Saudi Arabia's SDAIA, Singapore's IHiS) have explicit mandates to participate in or sponsor clinical-AI infrastructure that reduces dependence on U.S. hyperscaler platforms. For these sovereign entities, Phase-0 participation is a policy instrument, not a financial investment.

(iv) **Foundation-funded operational subsidy.** The Asuu Foundation (or the initial funding consortium) may reimburse Phase-0 validators' direct operational costs during the first

12–24 months, removing the CFO-veto concern entirely for this period. This is standard practice in consortium-infrastructure launches (the original DTCC launch, the RTP Network bootstrap, the early IPFS gateway ecosystem) and is not in any way cryptoeconomically distinctive. What it requires is acknowledged upfront Foundation funding, which Phase 0 architecture must name explicitly rather than assume.

Phase 2+ secondary liquidity is provided by the Dutch-auction mechanism of Definition 15.12, which opens genuine secondary-market price discovery once inference volume and transaction-fee revenue reach levels that support non-speculative token valuation. Validators entering in Phase 2+ face a different risk calculus than Phase-0 validators because they can mark to market. The three-year gap between Phase-0 commitment and Phase-2 liquidity is the period in which consortium commitments, mission alignment, and Foundation subsidy must jointly carry the validator set, and the protocol's business plan must name the Foundation capital source and the consortium member commitments explicitly before Phase-0 launch. We do not claim otherwise. This is a capital-planning dependency, not a cryptoeconomic property.

Definition 15.12 (Staking Reserve Pool Specification)**.** The Staking Reserve Pool receives allocation $\rho_{\text{SR}} \cdot S_{\max}$ at genesis, where $\rho_{\text{SR}} \in [0.15, 0.25]$. Disbursements from the pool are subject to:

(i) **Phase 0:** No disbursement; pool is escrowed.
(ii) **Phase 1:** Disbursement at pegged initial valuation p_0 denominated in reference currency, capped per validator at a threshold aligned with the validator's clinical-activity volume.
(iii) **Phase 2–3:** Disbursement at market price via a Dutch auction mechanism at fixed periodicity, with a hard floor at p_0 and a soft ceiling at $4p_0$.

Disbursements are subject to a vesting schedule that releases stake over 48 months to align validator incentives with long-term network health.

15.5 The Validation Perimeter Problem: Permissioned Scope for FDA-Regulated Devices

21 CFR 820 and the companion ISO 13485 quality-management standard require that a device manufacturer validate the operation of the device under the manufacturer's Quality Management System (QMS). Under 21 CFR Part 11 §11.10(c), the sponsor must validate "systems to ensure accuracy, reliability, consistent intended performance, and the ability to discern invalid or altered records." This validation must be performed against a defined operational perimeter: a set of hardware, software, and operator roles that the sponsor controls or has contractually constrained. A validator operator whom the sponsor does not control cannot be inside the sponsor's validation perimeter, because the sponsor cannot validate what it cannot control.

This creates a hard constraint on the protocol's Phase 3 permissionless aspiration when the protocol is used by regulated medical devices. A permissionless validator set is by construction outside the sponsor's validation perimeter. If the protocol's inference-record attestations are part of the regulated device's operation, the sponsor cannot validate under permissionless consensus, and the device cannot be marketed under the sponsor's QMS. This is not a soft problem that can be addressed by improving the protocol's audit trail; it is a structural property of how medical-device regulation assigns responsibility.

The Asuu architecture therefore distinguishes two operational scopes, with different validator-set requirements:

Definition 15.13 (Regulated-Device Scope and General-Purpose Scope)**.** Each on-chain attestation is tagged with one of two scope labels:

(a) **Regulated-Device Scope (RDS):** attestations produced by or concerning an FDA/EMA/notified-body-regulated medical device operating under its sponsor's QMS. RDS attestations require that the validators who finalize them are within the sponsor's validation perimeter. Specifically, RDS finality requires a quorum of *Permissioned Validators*, defined in Definition 15.14.
(b) **General-Purpose Scope (GPS):** attestations that are not part of a regulated device's operation, including research-use-only inference records, public-health surveillance attestations, supply-chain provenance records, educational-credential attestations, and clinical-trial-integrity anchors. GPS attestations can be finalized by the general permissionless validator set under the Phase 3 trajectory of Definition 15.7.

The scope label is set by the attestation-producing application at the time of submission, is cryptographically bound to the attestation, and cannot be changed retroactively. Validators of each type verify that they are signing only attestations within their authorized scope; a GPS-only validator that receives an RDS attestation rejects it.

Definition 15.14 (Permissioned Validator)**.** A Permissioned Validator for a specific sponsor's regulated device is an operator that has entered into a written QMS agreement with the sponsor establishing:

(i) the validator's operational perimeter (hardware, software, staff, procedures) as a subset of the sponsor's QMS-controlled validation scope;
(ii) the validator's obligations to maintain, patch, and audit its infrastructure under the sponsor's QMS change-control procedures;
(iii) the validator's obligations under 21 CFR Part 11 §11.10(c) attribution and under §11.10(e) audit-trail integrity;
(iv) the validator's reporting obligations to the sponsor for any incident that might affect the sponsor's regulatory obligations;
(v) the sponsor's right to audit the validator's operations and to terminate the validator's participation.

The sponsor maintains an on-chain *Permissioned Validator Registry* listing the current authorized validators for each of the sponsor's regulated devices. An RDS attestation from the sponsor's device is finalized only by a quorum drawn from this registry.

Theorem 15.15 (RDS Safety Under Permissioned Quorum)**.** *For attestations tagged RDS for sponsor S, the BFT safety argument of Theorem 14.2 is instantiated over the permissioned validator set $\mathcal{V}_S$ from the sponsor's Permissioned Validator Registry, not over the general validator set. Specifically, finality of RDS-tagged blocks requires a quorum certificate with aggregate weight exceeding $2/3$ of $\sum_{v\in\mathcal{V}_S} W_v$, and the BFT assumption $f < 1/3$ is relative to $\mathcal{V}_S$ rather than to the general validator population.*

Proof. The quorum certificate verification logic (Definition 9.9) reads the scope label of the block. For RDS-tagged blocks, the verifier filters the signer set to $\mathcal{V}_S$ from the Permissioned Validator Registry identified by the sponsor field, and computes aggregate weight over $\mathcal{V}_S$. The safety argument of Theorem 14.2 then applies unchanged: two conflicting RDS blocks at the same height require two quorum certificates over $\mathcal{V}_S$, whose intersection contains at least one

honest permissioned validator, who by protocol invariant does not sign conflicting blocks. The safety bound depends on $|\mathcal{V}_S|$ rather than on the general validator count. □

Remark 15.16 (Why This Resolves the Validation Perimeter Problem)**.** Under the scope partition of Definition 15.13, the sponsor's QMS validation is performed against the finite, enumerable set $\mathcal{V}_S$ of Permissioned Validators with which the sponsor has QMS agreements. The sponsor's 21 CFR Part 11 §11.10(c) validation is therefore tractable: the operational perimeter is defined, the operators are contractually constrained, and the change-control is enforceable. This is the same architecture that regulated financial infrastructure uses (SWIFT member banks, DTCC participants, ISO-27001-certified cloud providers operating under BAA contracts for HIPAA scope); we are applying the same pattern to regulated medical devices. The permissionless Phase 3 aspiration of Definition 15.7 remains valid for General-Purpose Scope attestations, which do not require sponsor QMS validation because they are not part of a regulated device's operation. The protocol does not claim that permissionless consensus is compatible with FDA-regulated device deployment. It claims that the protocol's architecture can support both permissioned RDS and permissionless GPS under a single chain, with the scope partition enforced at the attestation level.

Proposition 15.17 (Scope Isolation Under Shared Consensus)**.** *A fair critique of the scope-partition architecture is that "an infrastructure where a regulated AI model shares consensus mechanics, block space, and network bandwidth with permissionless, unregulated GPS nodes cannot be effectively validated." The critique is correct insofar as a naive shared-consensus design would contaminate the RDS validation perimeter with GPS-originated state. The Asuu architecture addresses this through four explicit isolation mechanisms:*

(i) ***Attestation-level scope tagging with cryptographic binding.*** *Each attestation carries its scope label as a signed field (Definition 15.13). A GPS attestation cannot be re-labeled as RDS without re-signing, which the GPS-originating validator cannot do because it is not in $\mathcal{V}_S$. The sponsor's QMS validation considers only the RDS-tagged subset of the chain, and that subset is cryptographically identifiable as originating from $\mathcal{V}_S$-signed attestations.*

(ii) ***Validator-set segregation options.*** *The sponsor may additionally specify* RDS-only operation *for $\mathcal{V}_S$, in which permissioned validators sign exclusively RDS attestations and reject GPS attestations at the mempool level. Under this configuration, $\mathcal{V}_S$ validators never touch GPS state, which removes the shared-infrastructure concern entirely at the cost of reducing validator economic efficiency (the validators cannot amortize operational cost across both scopes).*

(iii) ***Block-space partitioning for mixed-scope validators.*** *For validators that serve both RDS and GPS (the default operational efficient configuration), the block-format specification reserves distinct block-space regions for RDS and GPS attestations. The sponsor's QMS validation of the RDS region does not depend on the GPS region, and a GPS-region-originated consensus disagreement cannot block RDS-region finality (the BFT quorum for RDS is drawn only from $\mathcal{V}_S$, per Theorem 15.15).*

(iv) ***Network-bandwidth priority for RDS.*** *The gossip protocol prioritizes RDS attestation propagation over GPS attestation propagation, so GPS-originated bandwidth pressure cannot degrade RDS finality latency below the performance specification the sponsor validates against.*

Under these four mechanisms, the sponsor's QMS validation of the RDS deployment is independent of GPS-node activity on the same physical chain. The argument is structurally equivalent to how regulated financial messaging (SWIFT) operates over the public Internet: the shared transport substrate does not contaminate the regulated-message scope because the regulated scope is defined by sender authentication, message format, and receiver filtering, not by exclusive control of the transport.

Remark 15.18 (Validator-Set Segregation and Its Practical Implications)**.** A sponsor's choice between mixed-scope validators (default) and RDS-only validators involves a specific tradeoff. Mixed-scope validators can amortize operational cost across both scopes, reducing the per-attestation cost the sponsor bears through fee schedules, but require the sponsor's QMS audit to verify that the scope-isolation mechanisms of Proposition 15.17 are correctly implemented at each validator. RDS-only validators eliminate the isolation-mechanism audit at the cost of requiring the sponsor to underwrite the validators' operational cost without GPS amortization. For Class III high-risk devices and for initial RDS deployments during the Months 18–30 pilot window, we expect sponsors will prefer RDS-only validators because the audit simplification outweighs the cost-amortization loss; for Class IIa devices with mature pilot experience, mixed-scope validators become the economically efficient choice once the sponsor's QMS has validated the isolation-mechanism implementation. Neither configuration is required by the architecture; both are supported.

Remark 15.19 (Honest Acknowledgment: This Is a Partial Retreat)**.** An earlier framing of this whitepaper presented permissionless Phase 3 as the protocol's target end-state without distinguishing regulated from non-regulated scope. That framing was inconsistent with the regulatory requirements for medical-device validation. We correct it here: for FDA-regulated and notified-body-regulated devices, the protocol's deployment model is permissioned, not permissionless. The marketing language that suggested otherwise was wrong. The protocol is nonetheless useful under this constraint because (a) the permissioned validator set can be substantially larger than a single sponsor's internal infrastructure (drawing on qualified healthcare institutions, regulatory bodies, and consortium members), (b) the cross-sponsor anchoring properties of the chain provide tamper-evidence that no single sponsor's QMS-internal system can provide, and (c) the General-Purpose Scope partition allows genuinely permissionless operation for non-regulated applications on the same chain substrate. This is a narrower value proposition than "permissionless clinical AI infrastructure" but it is the correct one.

15.6 Clinical AI Passport SaMD Classification Analysis

A substantive regulatory question the earlier drafts of this paper elided is whether the Clinical AI Passport state itself is a *device constituent* under 21 CFR 820.3(o) and MDR Annex VIII Rule 11. Specifically: if a validator quorum can transition a model passport from *active* to *suspended*, and that transition causally modifies clinical decision support downstream via CDS Hooks (Section 1.3), is the consensus layer part of the device? This is the correct first question a notified-body reviewer or an FDA Q-Sub consultant will ask, and it deserves a direct answer rather than silence.

Definition 15.20 (Passport Functional Decomposition)**.** The Clinical AI Passport $\mathscr{P}$ (Definition 23.5) performs three functionally distinct roles, which the SaMD analysis must treat separately:

(a) **Metadata anchor:** on-chain storage of model identity, version hash, approved training-data lineage, and regulatory-authority attestations. This role is a *record-keeping substrate* for facts the sponsor has already determined and submitted to the regulator.
(b) **Policy gate:** on-chain evaluation of activation predicates (corridor checks, supermajority vote status, formal-verification attestations) that determine whether a specific model version is permitted to produce inference records for clinical consumption. This role is a *policy-enforcement mechanism* that causally affects downstream clinical workflow.
(c) **Governance signal:** on-chain record of supermajority votes, regulatory override transactions, and sponsor-directed suspensions, which provide a tamper-evident record of how and by whom passport state transitions occurred. This role is an *audit-trail substrate*.

Definition 15.21 (Passport Classification Under Regulated-Device Scope)**.** For an RDS deployment (Definition 15.13) under the permissioned validator set $\mathcal{V}_S$ operating within the sponsor's QMS perimeter, the Clinical AI Passport's three functional roles are classified as follows:

(i) The metadata-anchor role (a) is *not* a device constituent. Recording that a particular model version exists with a particular hash is analogous to a manufacturer's device-history record under 21 CFR 820.184, which is QMS documentation, not a device component.
(ii) The policy-gate role (b) *is a device constituent* under 21 CFR 820.3(o) and MDR Annex VIII Rule 11 when it causally gates clinical recommendations. Specifically, the smart-contract code implementing the activation predicate, the Regulator Override Channel, and the suspension-trigger logic is *software in a medical device (SiMD)* and must be developed under IEC 62304 software-lifecycle processes as part of the sponsor's regulatory submission. The *underlying consensus protocol* that enables multiple validators to reach agreement on the predicate result is analogous to third-party operating-system or database software included in the device: it is software of unknown provenance (SOUP) under IEC 62304, which requires evaluation and risk-management but not full lifecycle development by the sponsor.
(iii) The governance-signal role (c) is *not* a device constituent; it is an audit-trail capability that supports regulatory inspection under Definition 42.4 but does not itself gate clinical decisions.

Proposition 15.22 (RDS-Permissioned Architecture Supports SaMD Validation)**.** *Under the classification of Definition 15.21, the RDS-permissioned architecture (Definition 15.14) is compatible with 21 CFR 820 quality-management-system validation and IEC 62304 software-lifecycle processes because:*

(i) The SiMD component (the policy-gate smart-contract code) is developed and maintained by the sponsor under the sponsor's QMS, subject to the same software-lifecycle processes as any other SiMD component.
(ii) The SOUP component (the consensus protocol executed by permissioned validators) is within the sponsor's validation perimeter because $\mathcal{V}_S$ operators are under QMS agreements per Definition 15.14. The sponsor's SOUP evaluation assesses: (a) the protocol's formal safety guarantees (Theorem 14.2, Theorem 15.15), (b) the operational performance of $\mathcal{V}_S$ under QMS monitoring, (c) the change-control procedures for protocol-version updates.

(iii) The audit-trail component (governance-signal role) is documentation, not device-component code, and is treated as regulatory-reporting infrastructure analogous to existing e-CFR and EudaMed submissions.

Remark 15.23 (Why This Requires an FDA Pre-Submission Q-Sub Before Deployment)**.** Definition 15.21 and Proposition 15.22 present the sponsor's and protocol designers' position on the Passport's SaMD classification. That position is a reasoned interpretation; it is not a regulatory determination. FDA alone determines device classification, and a responsible deployment trajectory requires the sponsor to submit a pre-submission (Q-Sub) package under 21 CFR 10.115 addressing:

(i) The specific intended use of the device that integrates the Clinical AI Passport;
(ii) The sponsor's proposed classification of the three Passport roles, consistent with Definition 15.21;
(iii) The validation strategy for the SiMD policy-gate component under IEC 62304 Class B or C (depending on the device's overall risk classification);
(iv) The SOUP evaluation strategy for the consensus-protocol component;
(v) The change-control procedures for protocol-version updates that affect the SiMD component.

The sponsor's Q-Sub response from FDA establishes the regulatory basis for the specific deployment. Protocol designers and the Asuu Foundation cannot substitute for this sponsor-specific regulatory work. We flag this requirement explicitly: no RDS deployment of the Asuu protocol should proceed to production use in an FDA-regulated device without completion of the Q-Sub process, and the Foundation should make Q-Sub support (template documentation, reference classification arguments, IEC 62304 evidence packages) available to sponsors as a protocol-governance priority during Phase 1. Similarly, MDR-regulated deployments require a notified-body technical-documentation review that addresses the same classification questions under MDR Annex VIII Rule 11. The protocol's architecture can support these reviews; it cannot substitute for them.

Remark 15.24 (What This Does Not Resolve)**.** The SaMD classification analysis above addresses how the Passport's software components fit within a sponsor's existing regulatory framework. It does not address two residual questions that remain genuinely open:

(a) The cross-jurisdictional question: if an RDS-sponsored device is deployed in the EU under MDR and in the US under 21 CFR Part 820, and the policy-gate smart contract executes identically under both jurisdictions but produces a consequential suspension event, which regulator's authority governs the response timeline, the sponsor's reporting obligations, and the associated liability? Current practice treats such cross-jurisdictional deployments on a device-by-device basis through coordinated regulatory agreements; the protocol does not resolve this question, but the Regulatory Portal (Definition 42.3) at least ensures both regulators have simultaneous access to identical evidence.
(b) The sponsor-succession question: an RDS deployment is bound to a specific sponsor's QMS. If the sponsor exits the market or transfers the device to a successor manufacturer, the Permissioned Validator Registry and the sponsor's Regulator Override Channel signatures require transition to the successor. The protocol supports this operationally (key-rotation and registry-update transactions are standard), but the regulatory continuity (whether the successor inherits the existing FDA clearance, how ongoing PMS data from

the protocol carries over, etc.) is a matter of sponsor-to-sponsor transfer agreements that the protocol does not itself adjudicate.

These are not defects of the protocol; they are normal regulatory-transition questions that any SaMD faces and that must be addressed device-by-device. We note them to bound the completeness of the SaMD classification analysis.

15.7 Role-Weighted Shapley Value Tokenomics

Different healthcare roles generate fundamentally different types and economic values of data. Radiologists produce high-value imaging studies, nurses generate continuous longitudinal monitoring data, medical coders produce revenue-cycle-critical billing data, and laboratory technicians generate diagnostic-critical test results. Fair token compensation requires a mechanism that reflects these asymmetries.

Definition 15.25 (Role-Weighted Data Shapley)**.** For a set of n data contributors partitioned by roles $\mathscr{R}$, the role-weighted Shapley value of contributor i with role r_i is:

$$\phi_i^{\mathscr{R}} = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(n-|S|-1)!}{n!} \cdot w(r_i) \cdot [V(S \cup \{i\}) - V(S)]$$

where $V : 2^N \to \mathbb{R}$ is the characteristic function measuring model performance on coalition S, and $w : \mathscr{R} \to \mathbb{R}_{>0}$ is a role-specific weight function satisfying $\sum_{r \in \mathscr{R}} |\{i : r_i = r\}| \cdot w(r) = n$ (normalization).

Theorem 15.26 (Fair Role Compensation)**.** *The role-weighted Shapley value $\phi^{\mathscr{R}}$ satisfies:*

(i) ***Efficiency****: $\sum_{i=1}^{n} \phi_i^{\mathscr{R}} = w_{\text{avg}} \cdot V(N)$ where $w_{\text{avg}} = \frac{1}{n}\sum_i w(r_i)$;*

(ii) ***Role-weighted symmetry****: if contributors i and j have the same role ($r_i = r_j$) and equal marginal contributions, then $\phi_i^{\mathscr{R}} = \phi_j^{\mathscr{R}}$;*

(iii) ***Null contributor****: if $V(S \cup \{i\}) = V(S)$ for all S, then $\phi_i^{\mathscr{R}} = 0$.*

Proof. (i) By linearity of the Shapley formula and the weighting: $\sum_i \phi_i^{\mathscr{R}} = \sum_i w(r_i) \cdot \phi_i = w_{\text{avg}} \cdot n \cdot \frac{V(N)}{n} = w_{\text{avg}} \cdot V(N)$, where the second equality follows from the standard Shapley efficiency property.

(ii) If $r_i = r_j$, then $w(r_i) = w(r_j)$. Combined with equal marginal contributions $V(S \cup \{i\}) - V(S) = V(S \cup \{j\}) - V(S)$ for all S, the Shapley formula yields $\phi_i^{\mathscr{R}} = w(r_i) \cdot \phi_i = w(r_j) \cdot \phi_j = \phi_j^{\mathscr{R}}$.

(iii) If $V(S \cup \{i\}) - V(S) = 0$ for all S, then $\phi_i^{\mathscr{R}} = w(r_i) \cdot 0 = 0$. □

15.8 Break-Glass Token Economics

Definition 15.27 (Break-Glass Staking)**.** Emergency break-glass access (Definition 31.4) requires the accessor to escrow $\tau_{\text{btg}}(r)$ tokens, where $\tau_{\text{btg}} : \mathscr{R} \to \mathbb{R}_{>0}$ is a role-specific staking function. Administrative roles, which carry higher privilege-abuse risk (93% of unauthorized access investigations stem from insider threats [39]), are assigned higher stakes:

$$\tau_{\text{btg}}(r) = \tau_0 \cdot \exp(\alpha \cdot \text{priv_level}(r))$$

where τ_0 is a base stake, $\alpha > 0$ is a scaling parameter, and priv_level : $\mathscr{R} \to [0,1]$ maps each role to its normalized privilege level in the lattice $\mathscr{L}$.

Remark 15.28 (Known Failure Modes of Break-Glass Staking)**.** A legitimate concern is that any meaningful break-glass stake creates a cognitive friction in emergency access that can delay legitimate clinical action, while an insignificant stake fails to deter abuse. Both failure modes are real and have been observed in production break-glass implementations in EHR systems. The Asuu protocol does not claim to have solved this trade-off; it claims to have moved it from an administrative problem (post-hoc audit of break-glass usage, generally with low enforcement rates) to a cryptographic problem (pre-commitment of stake with automatic accountability), which permits different mitigations. Specifically:

(i) **Stake reimbursement by institution.** The staking token is held in escrow by the institution on behalf of the clinician, not by the clinician personally. Break-glass activation draws from institutional escrow; clinicians do not experience friction from personally-held stake. This addresses the delay concern.

(ii) **Institution-level aggregate monitoring.** Break-glass activations are aggregated at the institutional level, and an institution whose break-glass activation rate exceeds a configurable threshold triggers a compliance review rather than per-clinician punitive action. This addresses the concern that clinicians will be punished for appropriate emergency access.

(iii) **Review panel, not automatic slashing.** The escrowed stake is not automatically slashed on activation; it is held pending a review by the institution's break-glass review committee, which is the existing administrative structure for emergency-access review. The stake is released on review approval and forfeited only on review denial. This preserves the review discretion that clinical medicine requires.

We acknowledge the risk that even these mitigations may prove insufficient in production: clinicians may perceive any new break-glass friction as impeding clinical judgment, and the mechanism may be administratively neutralized by institutions that want frictionless break-glass regardless of audit consequences. If pilot deployment shows this failure mode, the mechanism should be revised or removed. The feature exists in the protocol specification because it composes correctly with the rest of the architecture, not because we claim it will survive first contact with a real emergency department.

PART IV

DATA SOVEREIGNTY, PRIVACY & SOVEREIGN IDENTITY

16 Data Sovereignty and Privacy

The traditional medical AI model achieves a form of data governance by centralizing all information within the vendor's infrastructure. The vendor processes, stores, and controls access to clinical data, a model that conflicts with data protection regulations and national sovereignty requirements. The necessity to record all inference activity publicly on the Asuu network precludes this approach, but privacy is maintained through a strict separation of data planes.

Patient data never touches the chain. What is recorded on-chain is exclusively cryptographic hashes: proofs that data existed and was processed, without revealing what the data contained. The actual patient data remains within the institution's sovereign infrastructure, behind its own access controls, subject to its own regulatory jurisdiction. The chain provides verifiability; the institution provides custody.

For federated model training, where multiple institutions contribute to improving a shared model without sharing raw data, the network employs zero-knowledge proofs of data quality and contribution. An institution can prove that it contributed a training dataset of a certain size and statistical profile, that the model's performance improved by a measurable amount on standard benchmarks, without revealing any individual data point. The verification is mathematical, not procedural. Recent research on blockchain-enabled zero-knowledge federated learning (zkFL-Health, 2025) has demonstrated that hospitals can collaboratively train medical AI with ZK proofs recorded on-chain to verify correct model aggregation [12].

16.1 The Data Embassy Model

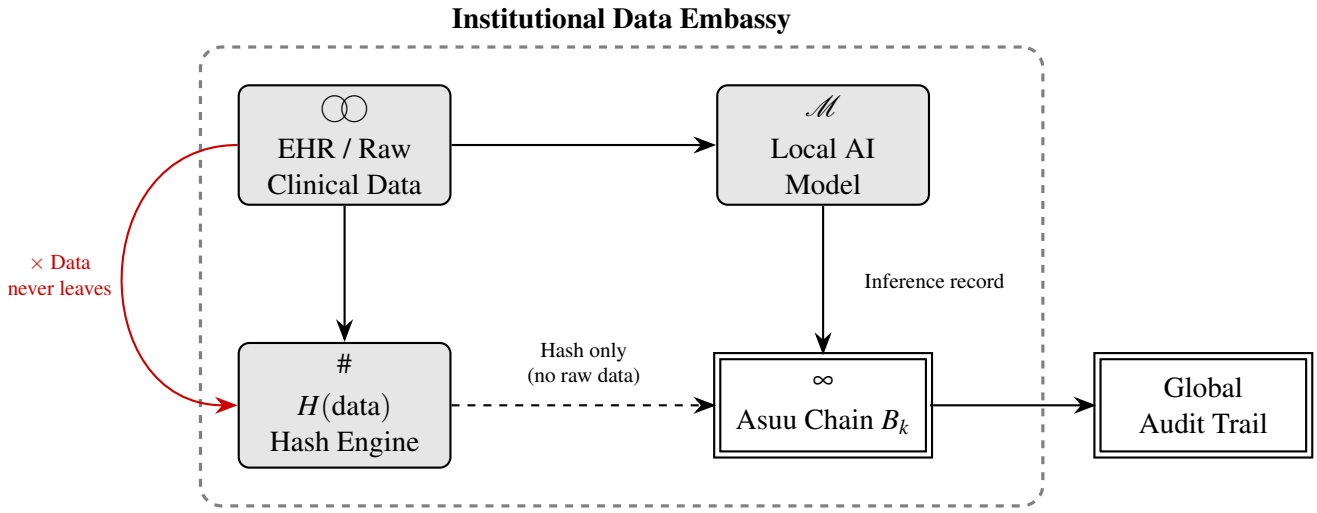


Figure 7: The Data Embassy Model. Raw clinical data remains within the institutional boundary. Only cryptographic hashes and inference records propagate to the Asuu chain, ensuring data sovereignty by architecture.

This architecture implements what we term the *Data Embassy* model. Each participating nation or institution operates its own sovereign node, a self-contained instance of the Asuu system running on local infrastructure, under local jurisdiction, with local data residency. The embassy participates in the global network for consensus and model governance while maintaining absolute sovereignty over its clinical data. No foreign entity can compel, subpoena, or surveil the data without the jurisdiction's consent. The network is global; the data is local.

The Data Embassy model enforces compliance with GDPR Article 44–49, HIPAA, and national data protection laws by architecture rather than by policy. Patient data physically cannot leave the jurisdiction because it physically never enters the chain. Each embassy maintains its own encryption keys, its own access control lists, and its own audit logs. Cross-border collaboration occurs through cryptographic attestations (zero-knowledge proofs, homomorphic commitments) that allow verification without disclosure. A Turkish hospital can prove to a German research consortium that its data meets inclusion criteria for a multi-center study without transmitting a single patient record across any border.

16.2 Formal Guarantees

Definition 16.1 (Data Embassy)**.** A Data Embassy $\mathscr{E}_I$ for institution I comprises: local data store DS_I under local jurisdiction, local encryption keys (ek_I, dk_I) controlled exclusively by I, and a hash-only interface to the chain: only $H(D) \in \{0,1\}^\lambda$ is transmitted on-chain.

Theorem 16.2 (Zero On-Chain Data Leakage)**.** *Under preimage resistance of H (Definition 4.2):*

$$\Pr[\textit{recovering } D_i \textit{ from on-chain state}] \leq \mathrm{negl}(\lambda).$$

Proof. The on-chain state consists of inference record hashes, model version hashes, and Merkle roots. Each is an output of H. Recovering any input D_i from $H(D_i)$ is the preimage problem, bounded by $\mathrm{negl}(\lambda)$. □

Corollary 16.3 (Systemic Exfiltration Cost)**.** *To compromise data across N independent Data Embassies, an adversary must execute N independent breaches. The cost scales as $C_{\text{exfil}}(N) = N \cdot c_{\text{breach}}$. There is no single point of failure.*

Proposition 16.4 (Cross-Border Collaboration Without Data Transfer)**.** *Two institutions in different jurisdictions can verify mutual data properties using zero-knowledge proofs without transferring patient data across borders.*

Proof. Institution I_A computes commitment $C_A = \text{Pedersen}(D_A, r_A)$ and proof $\pi_A = \text{Prove}(crs, \text{“}D_A \text{ satisfies } \phi\text{”}, D_A)$. Institution I_B verifies $\text{Verify}(vrs, C_A, \pi_A) = 1$. By zero-knowledge (Definition 4.8), I_B learns only that $\phi(D_A) = \text{true}$. Raw data D_A remains in jurisdiction A. □

16.3 Deployment Reality: Data Embassies for SaaS Hospitals

The Data Embassy formalism of Definition 16.1 defines the cryptographic properties the embassy must provide, but says nothing about *who operates the embassy*. A critical deployment-reality concern is that the substantial majority of U.S. hospitals (and an even larger majority globally) do not operate their own medical-AI validator infrastructure. They run Epic, Cerner, MEDITECH, or equivalent platforms as Software-as-a-Service, with the SaaS vendor operating the compute and storage. Asking every SaaS hospital to stand up its own validator node, manage TEE hardware, and operate a post-quantum HSM is not realistic. Five hundred and sixty-seven community hospitals will not hire a blockchain operations team.

The Asuu deployment model therefore distinguishes three Data Embassy operating modes, which together cover the realistic deployment surface:

Definition 16.5 (Data Embassy Operating Modes)**.** Each institution participating in the Asuu network selects one of three Data Embassy operating modes, each of which satisfies Definition 16.1's cryptographic requirements:

(a) **Sovereign Embassy (Mode S):** The institution operates its own Data Embassy hardware, validator node, TEE silicon, and HSM. Appropriate for academic medical centers, national health systems, and large integrated delivery networks with existing blockchain or security-operations capabilities. This is the mode described in Definition 16.1.

(b) **Vendor-Operated Embassy (Mode V):** The institution's EHR SaaS vendor operates an Asuu-aware Data Embassy instance on behalf of the institution, under a Business Associate Agreement or equivalent Data Processor Agreement that retains the institution as data controller. Encryption keys are held in an HSM administered by the SaaS vendor but escrowed to the institution under a key-ceremony protocol that prevents unilateral vendor access. The institution remains the cryptographic principal in all on-chain attestations; the vendor is a technical operator, not a data controller. This is the mode we expect to dominate U.S. and EU deployment.

(c) **Consortium-Hosted Embassy (Mode C):** A group of institutions (e.g., a regional health information exchange, a state hospital association, a national health-authority ministry) collectively operates a shared Data Embassy infrastructure, with each member institution holding its own encryption keys and retaining data-controller status for its own data. Appropriate for small community hospitals and for national health systems in jurisdictions where vendor-operated SaaS is not dominant.

Modes V and C are operationally feasible for SaaS-dominated deployment environments because they do not require the institution to operate blockchain infrastructure directly. Mode S remains available for institutions that want direct operational control.

Definition 16.6 (Vendor-as-Data-Processor Cryptographic Controls)**.** For Mode V deployments, the SaaS vendor's role as technical operator is cryptographically bounded by the following controls, enforced by the protocol:

(i) The institution's signing key sk_I is generated and sealed to the institution's TEE enclave during initial onboarding. The vendor's operational staff cannot extract sk_I without breaking TEE integrity (Assumption 4.13), at which point the breach is detectable per Definition 18.3.

(ii) All on-chain attestations are signed with sk_I, not with any vendor key. The institution is the cryptographic author of record.

(iii) The vendor's operational access to the Data Embassy's data plane is logged in an append-only log anchored on-chain. Any access outside the scope of the Data Processor Agreement produces evidence admissible in regulatory or civil proceedings under Theorem 5.3.

(iv) Key rotation, re-sealing, and vendor-exit procedures are pre-specified and do not require vendor cooperation at time of exit. The institution can migrate its Data Embassy to Mode S or a different Mode V vendor without the prior vendor's consent beyond return of escrowed key material.

Remark 16.7 (Why Mode V Does Not Defeat the "Sovereign" Claim)**.** A reasonable objection to Mode V is that delegating Data Embassy operation to a SaaS vendor reintroduces precisely

the vendor-trust relationship the Asuu protocol is supposed to replace. This objection would be correct if Mode V required the institution to trust the vendor with data-controller powers. It does not. The cryptographic controls of Definition 16.6 reduce the vendor's role to that of a technical hosting provider, analogous to a cloud-infrastructure provider under the shared-responsibility model. The institution retains: (a) exclusive control of signing keys via TEE sealing; (b) cryptographic authorship of every on-chain attestation; (c) audit-trail evidence of every vendor access to the data plane; (d) unilateral exit rights. What the institution delegates is operational work (uptime, patching, physical security of the host), not legal or cryptographic authority. This is the same pattern that makes managed-HSM services compatible with high-assurance key management in banking and defense applications. We therefore claim that Mode V preserves institutional sovereignty in the senses that matter (data control, attribution, auditability, exit rights) while removing the barrier that makes Mode S infeasible for the bulk of deployment.

Remark 16.8 (Honest Scope: What Mode V Does Not Solve)**.** Mode V does not solve the political-economy problem that large EHR SaaS vendors may be reluctant to operate Asuu-aware Data Embassy infrastructure, because doing so would expose vendor-side data-flow patterns to on-chain auditability in ways that may be commercially disadvantageous. This is a business-development problem, not a cryptographic one, and the protocol cannot solve it by design. Initial deployment therefore requires EHR-vendor partnership arrangements, which the bootstrap strategy (Section 15.4) addresses at the consortium level during Phase 0–1. A protocol that cannot be deployed without vendor cooperation is not the same as a protocol that requires vendor trust; we require vendor *cooperation* for deployment, not vendor *authority* for operation.

17 Self-Sovereign Health Identity

The current health identity model fragments patient records across dozens of institutions, each maintaining its own silo of demographic, clinical, and administrative data. A patient who moves between providers, cities, or countries must re-establish identity at each point of care, a process that introduces errors, delays treatment, and creates dangerous gaps in clinical history. The Asuu network addresses this through a self-sovereign identity (SSI) layer built on decentralized identifiers (DIDs) and verifiable credentials.

Under this model, the patient holds a cryptographic identity wallet containing verifiable credentials issued by healthcare institutions: vaccination records, laboratory results, allergy profiles, medication lists. Each credential is signed by the issuing institution and anchored to the Asuu chain through a hash, but the credential itself resides in the patient's wallet, not on the chain. The patient selectively discloses credentials to new providers using zero-knowledge proofs. A patient can prove, for example, that they have a valid blood type certification without revealing any other medical information.

This approach eliminates the need for centralized health information exchanges that concentrate risk and create single points of surveillance. The patient controls access; institutions verify credentials cryptographically; the chain provides the trust anchor. No institution needs to query a central database, and no central database accumulates the comprehensive health profile that makes centralized systems attractive targets for breach. The 2024 Change Healthcare ransomware attack, which affected 192.7 million Americans, demonstrates precisely why centralized health identity infrastructure is a systemic risk that demands a decentralized alternative.

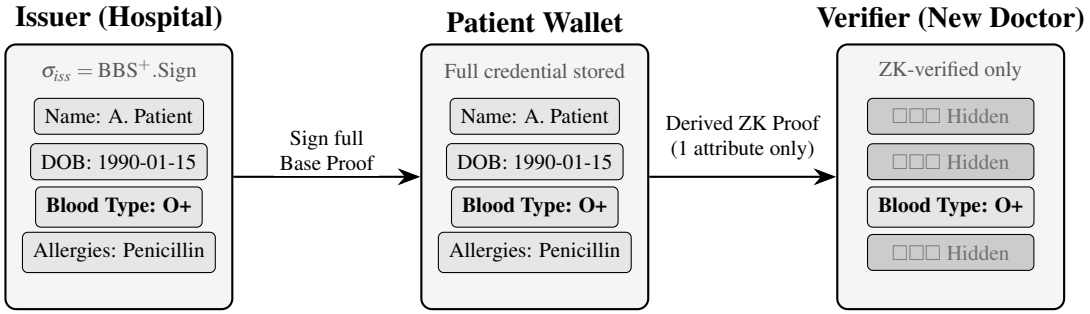


Figure 8: Zero-knowledge selective disclosure in the self-sovereign identity layer. The issuing hospital signs a BBS+ credential containing all attributes. The patient wallet derives a ZK proof disclosing only the requested attribute (e.g., Blood Type) to the verifier, while all other fields remain cryptographically hidden.

17.1 Formal Guarantees

Definition 17.1 (Health Credential)**.** A verifiable health credential is $VC = (iss, sub, attr, \sigma_{iss})$ where $attr = (a_1, \ldots, a_k)$ is the attribute vector and $\sigma_{iss} = \text{BBS}^+.\text{Sign}(sk_{iss}, (a_1, \ldots, a_k))$.

Definition 17.2 (Selective Disclosure)**.** For credential VC with attributes $(a_1, \ldots, a_k)$ and disclosure set $D \subseteq [k]$, selective disclosure produces proof π_D that convinces a verifier of $\{a_j\}_{j \in D}$ without revealing $\{a_j\}_{j \notin D}$.

Theorem 17.3 (Credential Unforgeability)**.** *Under the q-SDH assumption (Assumption 4.11), no* PPT *adversary with access to issuance queries can produce a valid credential for attributes not issued by any authority.*

Proof. BBS$^+$ signatures achieve EUF-CMA under q-SDH in bilinear groups [14]. A credential forgery implies a BBS$^+$ signature forgery, contradicting q-SDH. If $\mathscr{A}$ forges with probability ε, we construct $\mathscr{B}$ solving q-SDH with probability $\geq \varepsilon/Q$ where Q is the number of issuance queries. □

Theorem 17.4 (Presentation Anonymity)**.** *Credential presentations from the same patient are computationally unlinkable:*

$$\left|\Pr[\mathscr{A}(\pi_1,\pi_2)=1]-\frac{1}{2}\right| \leq \mathrm{negl}(\lambda).$$

Proof. BBS$^+$ signatures are randomizable: given σ on $\mathbf{m}$, anyone can produce σ' that is valid and uniformly distributed over valid signatures. Each presentation randomizes the signature independently. The zero-knowledge property ensures undisclosed attributes reveal no information. The adversary's view is independent of whether π_1, π_2 originate from the same or different credentials. □

Corollary 17.5 (Patient-Controlled Access)**.** *A patient can prove possession of a valid vaccination record without revealing other medical information, and multiple proofs to different verifiers are mutually unlinkable.*

17.2 W3C Verifiable Credentials v2.0 Integration

The W3C Verifiable Credentials Data Model v2.0, published as a W3C Recommendation on 15 May 2025 [24], provides the interoperability substrate for the Asuu network's health identity layer. Version 2.0 introduces several capabilities critical for medical credentialing: mandatory JSON-LD context processing for semantic precision, the validFrom/validUntil temporal validity model, and native support for multiple securing mechanisms through the Data Integrity framework [25].

The BBS Cryptosuite v2023 [26], standardized alongside VC 2.0, provides the cryptographic foundation for selective disclosure in the Asuu network. The suite separates the issuance process into two phases: a *base proof* computed by the issuer that signs all credential attributes, and a *derived proof* computed by the holder that reveals only a chosen subset. The derived proof is unlinkable to the base proof and to other derived proofs from the same credential, a property that follows from the algebraic structure of BBS signatures over pairing-friendly curves.

Credential revocation uses the Bitstring Status List v1.0 specification [27]. Each issuing institution maintains a compressed bitstring where position i encodes the revocation status of credential i. The bitstring is published as a verifiable credential itself, signed by the issuer. Revocation checks require only a single bit lookup, achieving $O(1)$ verification with $O(n/8)$ storage for n credentials. The privacy guarantee follows from the information-theoretic bound: observing a single bit position reveals only one bit of information about the credential holder's status, and the position itself is pseudorandom from the verifier's perspective.

Definition 17.6 (Asuu Health Credential v2)**.** An Asuu health credential conforming to VC Data Model 2.0 is a tuple $VC_2 = (ctx, type, iss, sub, attr, status, \sigma_{\text{BBS}})$ where:

- *ctx* includes the W3C credentials/v2 and Asuu medical context;
- $type \in \{\text{MedicalLicense}, \text{ClinicalPrivilege}, \text{RoleCredential}, \text{PatientRecord}\}$;
- $attr = (a_1, \ldots, a_k)$ with mandatory fields $a_1 = role$, $a_2 = department$, $a_3 = clearance$;
- *status* references a Bitstring Status List entry (url, idx);
- $\sigma_{\text{BBS}} = \text{BBS.Sign}(sk_{iss}, (a_1, \ldots, a_k), header)$.

Proposition 17.7 (Privacy-Preserving Revocation)**.** *Bitstring Status List revocation reveals at most* $\log_2(|StatusPurpose|)$ *bits about a credential holder to any single verifier, where* $|StatusPurpose| \in \{revocation, suspension\}$.

Proof. Each verifier query accesses exactly one bit at position *idx* in the status bitstring. The index *idx* is known only to the holder and the issuer. The verifier learns the revocation status (one bit) but cannot correlate *idx* across presentations because the derived BBS proof is unlinkable (Theorem 17.4). The status purpose field adds at most one additional bit. By the composition of information-theoretic bounds, total leakage per verification is bounded by $1 + \log_2(|StatusPurpose|) \leq 2$ bits. □

18 Exfiltration Resistance

18.1 TEE Model

Definition 18.1 (TEE Ideal Functionality $\mathcal{G}_{\text{att}}$)**.** The ideal functionality exposes two operations:

- $\textsc{Install}(prog) \to eid$: Store program.
- $\textsc{Resume}(eid, input) \to (output, \sigma)$: Execute and produce attestation $\sigma = \text{Sign}(sk_{\text{att}}, (eid, prog, output))$.

Integrity: host cannot modify computation. Confidentiality: host cannot observe enclave memory.

18.2 Information-Theoretic Leakage Bounds

Theorem 18.2 (DP as Leakage Bound)**.** *If mechanism $\mathcal{M}$ satisfies ε-DP, then: $I(X_i; Y \mid X_{-i}) \leq \varepsilon^2/(2\ln 2)$ bits. For $\varepsilon = 0.1$: approximately 0.0072 bits per query.*

Proof. By Pinsker's inequality, $I(X_i; Y \mid X_{-i}) \leq D_{\text{KL}}(\mathcal{M}(D) \| \mathcal{M}(D'))$ for adjacent D, D'. By ε-DP: $D_{\text{KL}} \leq \varepsilon(e^{\varepsilon} - 1) \leq \varepsilon^2 e^{\varepsilon}$. The tighter bound via the advanced data processing inequality yields $I \leq \varepsilon^2/(2\ln 2)$. □

18.3 Degraded-Mode Operation Under TEE Compromise

The TEE integrity assumption (Assumption 4.13) has known limits. Microarchitectural side-channel attacks including Foreshadow (2018), SGAxe (2020), ÆPIC Leak (2022), and Downfall (2023) have each demonstrated that specific silicon revisions of Intel SGX and AMD SEV-SNP can leak enclave state to a co-resident adversary under adversarial scheduling conditions. Rather than treat TEE integrity as an absolute primitive, the Asuu protocol treats it as a defense-in-depth layer whose failure is (i) detectable through remote attestation version checks and (ii) containable through degraded-mode operation.

Definition 18.3 (TEE Attestation Version Check)**.** Each validator maintains a locally-administered allowlist $\mathcal{V}_{\text{TEE}}$ of acceptable TEE attestation evidence, keyed by $(vendor, silicon_rev, microcode_patch)$. A TEE attestation σ_{TEE} is accepted only if its evidence triple is in $\mathcal{V}_{\text{TEE}}$ and the silicon revision has not been added to the published revocation list $\mathcal{R}_{\text{TEE}}$ maintained on-chain by the governance mechanism.

Definition 18.4 (Degraded-Mode Protocol)**.** When a silicon revision is added to $\mathcal{R}_{\text{TEE}}$ (in response to a published vulnerability disclosure), all on-chain operations that depended on TEE integrity for their security argument are transitioned to *degraded mode*: the operation is permitted to continue, but the TEE attestation is no longer accepted as sole evidence of correct execution. Instead, the operation requires one of:

(i) Re-execution by an independent validator running on a non-revoked silicon revision, with both results committed on-chain and required to match;
(ii) Human audit attestation from a certified clinical reviewer for operations below a configurable risk threshold;
(iii) Protocol halt for the affected operation class until updated silicon becomes available.

Proposition 18.5 (Degraded-Mode Safety)**.** *Under the degraded-mode protocol, a compromise of a single TEE silicon revision does not compromise the integrity of operations whose results are committed in mode (i). Specifically, for a result r committed in dual-execution degraded mode with validators V_1, V_2 on distinct silicon revisions $rev_1 \neq rev_2$, the adversary must compromise both revisions simultaneously to forge r.*

Proof. In dual-execution mode, r is committed only when $V_1.\text{Result}(op) = V_2.\text{Result}(op)$. If only rev_1 is compromised, the adversary can induce V_1 to report an incorrect result, but V_2 reports the honest result, and the mismatch prevents commitment. An adversary forging r must therefore compromise both rev_1 and rev_2, whose vulnerabilities are by construction independent (different vendor or silicon generation). □

This architecture trades latency and throughput for resilience: degraded-mode operations are slower and more expensive than nominal-mode operations, but the protocol does not fail catastrophically when a TEE is compromised. Historical evidence from the Foreshadow and Downfall disclosures suggests a typical 90-to-180-day window between vulnerability disclosure and silicon-level remediation; degraded-mode operation is the bridge across that window.

Remark 18.6 (Scope of the Degraded-Mode Defense: Institutional Validators, Not Edge Deployments)**.** The dual-execution defense of Proposition 18.5 presumes that the protocol can assign two validators on *distinct silicon revisions* to independently re-execute the affected operation. This presumption holds at institutional validator nodes, which operate heterogeneous hardware at institutional data centers: an academic medical center's validator infrastructure typically includes both Intel SGX and AMD SEV-SNP hosts, plus ARM CCA where available, giving the protocol genuine silicon diversity. The presumption fails at *edge deployments*: a rural clinic, a point-of-care ultrasound station, or a surgical navigation console will typically have one TEE-capable processor at best, frequently of a single vendor and silicon generation. Dual-execution is not a feasible defense at the edge.

The Asuu architecture therefore distinguishes two defense tiers:

(i) **Institutional validators:** dual-execution degraded mode as specified in Definition 18.4 and Proposition 18.5. The protocol's BFT safety argument rests on validator-tier TEE assumptions, so this is where the defense needs to hold.

(ii) **Edge deployments:** when a TEE revocation affects edge silicon, affected edge devices transition to *attestation-verified read-only mode*: they continue to receive and display AI recommendations whose attestations were produced upstream by institutional validators, but they cannot originate new attestations until remediated silicon is available. For time-sensitive inferences that the edge device previously autonomized, the autonomy level is automatically downgraded to "advisory only with institutional confirmation," requiring the clinician to confirm with an institutional server before action. This is a graceful degradation, not a silent failure: the CDS Hooks advisory card (Section 42.1) informs the clinician that edge attestation is degraded and institutional confirmation is required.

The edge-deployment defense does not claim to preserve full-autonomy operation under edge-TEE compromise. It claims to preserve *safe* operation under edge-TEE compromise, at the cost of higher latency and reduced autonomy until remediation. This is the correct trade-off for safety-critical medical devices: the protocol should degrade to higher-latency safety, not continue to operate at full autonomy with compromised cryptographic assurance.

Remark 18.7 (Scope of Halts: Operation-Class, Not Protocol-Wide)**.** A fair concern raised in review is that "halting operations every time a new microarchitectural leak is published means the network will suffer continuous, catastrophic downtime." This would be correct if the halt were protocol-wide and the halt policy were reflexive. Neither is the case in the Asuu design, and we clarify this explicitly here because earlier drafts described the halt mechanism without bounding its scope.

The halt mechanism operates under three scope constraints:

(i) **Operation-class scoping.** A TEE vulnerability disclosure halts only the *specific operation class* whose security reduction invoked the compromised TEE primitive. A Foreshadow-class SGX cache-timing leak invalidates the confidentiality of enclave-held secrets, which affects PRE decryption, pharmacogenomic bounty unsealing, and federated-inference integrity: these operation classes degrade to dual-execution or halt. Unaffected operation classes (BLS QC signing, ML-DSA block attestation, deterministic prior-auth predicate evaluation, audit-trail anchoring) continue uninterrupted because they do not depend on TEE confidentiality. The protocol is not a single monolithic enclave; it is a set of operation classes with distinct security reductions, and only the affected subset is degraded.

(ii) **Dual-execution priority over halt.** The degraded-mode protocol (Definition 18.4) *prefers* dual-execution on distinct silicon revisions over halt. Only when dual-execution is infeasible (affected operation is edge-originated, or all available silicon revisions share the compromised microarchitecture) does the operation class halt. The empirical evidence from the 2018–2024 transient-execution-attack history shows that Intel, AMD, and ARM TEE implementations do *not* share the specific microcode or speculation-window implementations that the major attacks exploited, so cross-silicon dual-execution has remained feasible through every disclosed attack in that period. Cross-silicon compromise is possible in principle; it has not been the empirical pattern.

(iii) **Governance-triggered re-evaluation rather than reflexive halt.** The transition from nominal to degraded mode requires a governance-panel transaction (Definition 18.4) evaluating the specific vulnerability against the specific operation classes at risk. Disclosure alone does not trigger halt; the panel's assessment of whether the disclosure affects the protocol's deployed silicon configuration and operation classes triggers scope-specific degradation. This replaces the reflexive halt of an earlier informal design with a deliberate SOUP-revaluation under the sponsor's QMS change-control procedures (Definition 15.14), which is the correct pattern for a regulated medical-device software component.

Under these three constraints, the expected operational impact of a new transient-execution disclosure is: a governance-panel review completed within days; a scope-restricted degradation affecting only operation classes whose security reduction invokes the specific compromised primitive; and full-availability continuation of unaffected operation classes throughout. This is not "continuous catastrophic downtime"; it is a scoped, governed, operation-class-level response to an evolving threat landscape, comparable to how major cloud-infrastructure providers responded to Spectre, Meltdown, Foreshadow, and Downfall disclosures without protocol-wide shutdown of their platforms. The protocol's large-surface-area vulnerability is not reflex halting; it is governance latency, which the QMS change-control framework is designed to bound.

18.4 Geo-Fencing Enforcement

Proposition 18.8 (TEE-Based Geo-Fencing)**.** *An enclave releases dk only if location proof $\ell \in \mathscr{P}$. Under Assumption 4.13, EUF-CMA of location attestation, and IND-CPA of encryption:*

$$\Pr[\textit{data accessed outside } \mathscr{P}] \leq \mathrm{Adv}_{\mathrm{loc}}^{\text{EUF-CMA}}(\lambda) + \mathrm{Adv}_{\mathrm{enc}}^{\text{IND-CPA}}(\lambda) + \mathrm{Adv}^{\mathscr{G}_{\mathrm{att}}}(\lambda) \leq \mathrm{negl}(\lambda).$$

Proof. Access outside $\mathscr{P}$ requires: (1) forging location attestation (breaks EUF-CMA), (2) decrypting without *dk* (breaks IND-CPA), or (3) extracting *dk* from enclave (breaks TEE integrity). Union bound completes the proof. □

19 Genomic Data Sovereignty

Genomic data occupies a unique position in the privacy landscape: it is immutable (a patient's DNA sequence cannot be changed after a breach), it identifies relatives who never consented to data collection, and it becomes more revealing as scientific understanding advances. As few as 75 to 100 single nucleotide polymorphisms (SNPs) suffice to uniquely identify an individual. Unlike a compromised password, a compromised genome cannot be reset.

The 23andMe bankruptcy filing in March 2025, which placed the genetic data of 15 million customers into bankruptcy estate assets available for acquisition, demonstrated that centralized genomic databases are existentially dangerous to the populations they serve. The Asuu network provides a fundamentally different model: patients hold their genomic credentials in their self-sovereign identity wallet (Section 17), research queries execute on encrypted data via the FHE layer (Section 24), and collaborative genomic studies operate through the federated learning protocol (Section 22) without any institution or marketplace holding a centralized genomic database.

Homer's attack (2008) proved that even aggregate GWAS allele frequencies from databases of thousands can re-identify individual participants, causing NIH to immediately restrict open-access aggregate genomic data [22]. Subsequent work demonstrated re-identification from genomic Beacon servers with as few as 250 queries. Standard differential privacy is insufficient for genomic data because it assumes independent tuples, but genomic entries are correlated through family relationships and linkage disequilibrium. The Asuu protocol addresses this through a combination of cryptographic access control (zero-knowledge proofs for selective attribute disclosure), homomorphic computation (GWAS statistics computed on encrypted genotypes), and on-chain consent management (dynamic, revocable, auditable consent records).

19.1 Formal Guarantees

Definition 19.1 (Genomic Credential)**.** A genomic credential extends the health credential (Definition 17.1) with attribute vector $\mathbf{a} = (a_1, \ldots, a_K)$ where $a_i \in \{0,1,2\}$ represents the genotype at SNP i (homozygous reference, heterozygous, homozygous variant). The credential is signed by the sequencing institution: $\sigma_{\text{seq}} = \text{BBS}^+.\text{Sign}(sk_{\text{seq}}, \mathbf{a})$.

Definition 19.2 (Dynamic Consent Record)**.** A consent record is:

$$CR = (\mathit{patient},\ \mathit{study_id},\ \mathit{scope},\ \mathit{expiry},\ \mathit{revocable},\ \sigma_{\text{patient}})$$

committed on-chain. The patient can revoke consent by publishing a signed revocation transaction $Rev = (CR_hash, t_{\text{rev}}, \sigma_{\text{patient}})$. Smart contracts reject data access requests referencing revoked or expired consent.

Theorem 19.3 (Genomic Selective Disclosure Soundness)**.** *Under preimage resistance of H (Definition 4.2) and the zero-knowledge property of BBS$^+$ selective disclosure (Theorem 17.4), a BBS$^+$ selective-disclosure proof π_D on disclosure set $D \subseteq [K]$ with $|D| = k$ reveals* no cryptographic information *about the undisclosed attributes $\{a_j\}_{j \notin D}$ beyond what the zero-knowledge simulator can produce. Formally, for any PPT adversary $\mathcal{A}$,*

$$\left|\Pr[\mathcal{A}(\pi_D, \{a_j\}_{j \in D}) = 1] - \Pr[\mathcal{A}(\mathit{Sim}(D, \{a_j\}_{j \in D})) = 1]\right| \leq \text{negl}(\lambda).$$

The on-chain hash of the credential reveals no information beyond its position in the commitment tree under preimage resistance of H.

Proof. Direct application of the zero-knowledge property of BBS^+ (Theorem 17.4): the simulator *Sim* produces a distribution on proofs indistinguishable from honestly-generated proofs without knowledge of undisclosed attributes. Preimage resistance of H (Definition 4.2) prevents an adversary from recovering the credential's plaintext from its on-chain hash. The combined bound is the ZK simulation indistinguishability. □

Remark 19.4 (Cryptographic Soundness Does Not Imply Statistical Re-Identification Resistance)**.** Theorem 19.3 is a statement about cryptographic information leakage through the selective-disclosure mechanism. It is *not* a bound on statistical re-identification probability. An earlier draft of this theorem claimed $\Pr[\text{re-identify}] \leq 3^{-(K-k)}$, which assumes SNPs are independent and uniformly distributed. That assumption is false: genomic entries are correlated through family relationships, population structure, and linkage disequilibrium (LD), and Homer's attack [22] demonstrated empirically that aggregate SNP frequencies from thousands of participants can re-identify individuals. Subsequent work showed re-identification from Beacon servers with as few as 250 queries. The naive $3^{-(K-k)}$ bound does not engage with this adversary.

Statistical re-identification resistance for genomic data requires an entirely different defense layer than selective disclosure. The Asuu protocol addresses this layer through (i) the differential privacy mechanism of Section 20 applied to aggregate-query responses, (ii) the query-rate-limiting mechanism that bounds the number of queries per dataset before the DP budget is exhausted, and (iii) the consent-revocation mechanism of Definition 19.2 that permits patients to withdraw data from future queries. The combination of these mechanisms bounds statistical re-identification probability; the selective-disclosure theorem alone does not. We explicitly do not claim $3^{-(K-k)}$ as a re-identification bound. The honest formulation is: the cryptographic layer prevents disclosure of undisclosed attributes through the proof protocol, and the DP layer bounds statistical inference from aggregate-query answers, subject to the DP parameter ε chosen for the specific query class.

Corollary 19.5 (Statistical Re-Identification Bound Under DP)**.** *For aggregate genomic queries answered under the ε-DP mechanism of Section 20, the probability that an adversary correctly identifies a target individual from query responses after a sequence of queries with cumulative privacy budget $\varepsilon_{\text{total}}$ is bounded by:*

$$\Pr[\textit{re-identify}] \leq \frac{1}{|P|} \cdot e^{\varepsilon_{\text{total}}}$$

where $|P|$ is the size of the reference population. For $\varepsilon_{\text{total}} = 1$ and $|P| = 10^6$, the re-identification probability is bounded by $\approx 2.7 \times 10^{-6}$. This is the honest DP-based re-identification bound; it is substantially weaker than the spurious $3^{-(K-k)}$ bound but reflects the actual adversarial model that Homer's attack and its successors exploit.

Proposition 19.6 (Secure GWAS via Homomorphic Chi-Squared Test)**.** *For a genome-wide association study across M institutions, each institution encrypts its genotype-phenotype contingency table under the network's FHE scheme. The aggregate chi-squared statistic $\chi^2 = \sum_{i,j}(O_{ij} - E_{ij})^2/E_{ij}$ is computed homomorphically and decrypted only by the study coordinator holding the aggregate decryption key. Individual institutional data remains encrypted throughout.*

Proof. The chi-squared statistic is a rational function of contingency table cell counts. Cell counts are integers amenable to exact FHE arithmetic (BFV/BGV schemes). Summation across

institutions is a linear operation on ciphertexts. The division and squaring are computed via homomorphic polynomial approximation. By Theorem 24.3, each institution's encrypted table reveals no information to the coordinator or other institutions. □

Corollary 19.7 (Consent Immutability and Auditability)**.** *By hash chain immutability (Theorem 6.5), the timestamp of consent grant, scope, and any revocation is permanently recorded. A regulatory auditor can verify that data access occurred only within valid, non-revoked consent windows, simply by querying the chain.*

20 Differential Privacy for Population Health Analytics

Population-level health analytics, such as disease prevalence estimation, treatment outcome analysis, and pharmacovigilance signal aggregation, require querying clinical data across multiple institutions. Even when individual records remain within their Data Embassies (Section 16), the aggregate statistics released to researchers can leak information about individual patients through membership inference, attribute inference, or reconstruction attacks. Differential privacy provides the mathematically rigorous framework to bound this leakage.

The Asuu network integrates differential privacy as a protocol-level primitive, enforced by smart contracts that manage a per-dataset privacy budget. Every query against federated clinical data consumes a measurable portion of this budget, and the protocol halts further queries once the budget is exhausted, preventing the "fundamental law of information recovery" from compromising patient privacy through accumulation of individually innocuous releases.

20.1 Formal Differential Privacy Framework

Definition 20.1 (Differential Privacy)**.** A randomized mechanism $\mathscr{M} : \mathscr{D} \to \mathscr{Y}$ satisfies (ε, δ)-differential privacy if for all neighboring datasets $D, D' \in \mathscr{D}$ differing in at most one record and all measurable $S \subseteq \mathscr{Y}$:

$$\Pr[\mathscr{M}(D) \in S] \leq e^{\varepsilon} \cdot \Pr[\mathscr{M}(D') \in S] + \delta.$$

When $\delta = 0$, the mechanism satisfies pure ε-differential privacy.

Definition 20.2 (Global and Local Sensitivity)**.** For a query function $f : \mathscr{D} \to \mathbb{R}^d$:

- **Global ℓ_1-sensitivity:** $\Delta_1 f = \max_{D \sim D'} \|f(D) - f(D')\|_1$.
- **Global ℓ_2-sensitivity:** $\Delta_2 f = \max_{D \sim D'} \|f(D) - f(D')\|_2$.
- **Local sensitivity at D:** $\mathrm{LS}_f(D) = \max_{D' : D \sim D'} \|f(D) - f(D')\|_1$.

Theorem 20.3 (Laplace Mechanism)**.** *For a query $f : \mathscr{D} \to \mathbb{R}^d$ with ℓ_1-sensitivity $\Delta_1 f$, the mechanism $\mathscr{M}(D) = f(D) + (\eta_1, \ldots, \eta_d)$ where each $\eta_i \sim \mathrm{Lap}(\Delta_1 f / \varepsilon)$ satisfies ε-differential privacy.*

Proof. For neighboring D, D' and any output $y \in \mathbb{R}^d$:

$$\frac{\Pr[\mathscr{M}(D) = y]}{\Pr[\mathscr{M}(D') = y]} = \prod_{i=1}^{d} \exp\left(-\frac{\varepsilon |f_i(D) - f_i(D')|}{\Delta_1 f} \right) \cdot \exp\left(\frac{\varepsilon |f_i(D') - f_i(D')|}{\Delta_1 f} \right).$$

The ratio telescopes to $\exp(\varepsilon \|f(D) - f(D')\|_1 / \Delta_1 f) \leq e^{\varepsilon}$. $\square$

Theorem 20.4 (Gaussian Mechanism)**.** *For a query $f : \mathscr{D} \to \mathbb{R}^d$ with ℓ_2-sensitivity $\Delta_2 f$, the mechanism $\mathscr{M}(D) = f(D) + \mathcal{N}(0, \sigma^2 I_d)$ with $\sigma \geq \Delta_2 f \sqrt{2 \ln(1.25/\delta)} / \varepsilon$ satisfies (ε, δ)-differential privacy for $\varepsilon \in (0, 1)$.*

Proof. The privacy loss random variable $Z = \ln(\Pr[\mathscr{M}(D) = y] / \Pr[\mathscr{M}(D') = y])$ has mean $\Delta_2 f^2 / (2\sigma^2)$ and variance $\Delta_2 f^2 / \sigma^2$ under the Gaussian mechanism. By the tail bound on the privacy loss: $\Pr[Z > \varepsilon] \leq \exp(-(\varepsilon\sigma / \Delta_2 f - \Delta_2 f / (2\sigma))^2 / 2)$. Setting $\sigma = \Delta_2 f \sqrt{2 \ln(1.25/\delta)} / \varepsilon$ makes this probability at most δ. $\square$

20.2 Composition and Budget Management

Theorem 20.5 (Sequential Composition)**.** *If mechanisms $\mathscr{M}_1,\dots,\mathscr{M}_k$ satisfy $(\varepsilon_1,\delta_1),\dots,(\varepsilon_k,\delta_k)$-DP respectively, their sequential application on the same dataset satisfies $(\sum_i \varepsilon_i, \sum_i \delta_i)$-DP.*

Proof. By induction on k. The base case $k=1$ is trivial. For the inductive step, condition on the output of $\mathscr{M}_{1:k-1}$ and apply the DP guarantee of $\mathscr{M}_k$. The privacy parameters add by the group privacy property and the union bound on δ. □

Theorem 20.6 (Parallel Composition)**.** *If mechanisms $\mathscr{M}_1,\dots,\mathscr{M}_k$ each satisfy (ε,δ)-DP and operate on disjoint subsets $D_1,\dots,D_k$ of the dataset, their joint output satisfies (ε,δ)-DP.*

Proof. Since each D_i is disjoint, changing one record in D affects at most one D_i. The mechanism operating on the unaffected subsets produces identical output distributions. Only the mechanism on the affected subset contributes privacy loss, bounded by (ε,δ). □

Definition 20.7 (On-Chain Privacy Budget Contract)**.** For each registered dataset D_j, the Asuu network maintains an on-chain privacy budget state $(\varepsilon^j_{\text{spent}}, \varepsilon^j_{\max}, \delta^j_{\text{spent}}, \delta^j_{\max})$. A query q with privacy cost (ε_q,δ_q) is permitted only if:

$$\varepsilon^j_{\text{spent}} + \varepsilon_q \le \varepsilon^j_{\max} \quad \text{and} \quad \delta^j_{\text{spent}} + \delta_q \le \delta^j_{\max}.$$

Upon execution, the contract atomically updates $\varepsilon^j_{\text{spent}} \leftarrow \varepsilon^j_{\text{spent}} + \varepsilon_q$ and $\delta^j_{\text{spent}} \leftarrow \delta^j_{\text{spent}} + \delta_q$.

Theorem 20.8 (Budget Enforcement Correctness)**.** *Under the on-chain budget contract (Definition 20.7), the cumulative privacy loss for any dataset D_j never exceeds $(\varepsilon^j_{\max}, \delta^j_{\max})$-DP, regardless of the number or sequence of queries, assuming each individual query mechanism correctly implements its claimed (ε_q,δ_q)-DP guarantee.*

Proof. By sequential composition (Theorem 20.5), the cumulative privacy loss after k queries is at most $(\sum_{i=1}^k \varepsilon_{q_i}, \sum_{i=1}^k \delta_{q_i})$. The smart contract enforces $\sum_i \varepsilon_{q_i} \le \varepsilon^j_{\max}$ and $\sum_i \delta_{q_i} \le \delta^j_{\max}$ by rejecting any query that would exceed either bound. Atomicity of the update (within a single EVM transaction) prevents race conditions. □

20.3 Healthcare-Specific Privacy Mechanisms

Definition 20.9 (Exponential Mechanism for Clinical Queries)**.** For a query with utility function $u : \mathscr{D} \times \mathscr{Y} \to \mathbb{R}$ and sensitivity $\Delta u = \max_y \max_{D \sim D'} |u(D,y) - u(D',y)|$, the exponential mechanism outputs y with probability proportional to $\exp(\varepsilon \cdot u(D,y)/(2\Delta u))$.

Theorem 20.10 (Exponential Mechanism Privacy and Utility)**.** *The exponential mechanism (Definition 20.9) satisfies ε-DP and selects an output y satisfying:*

$$\Pr\left[u(D,y) \le \max_{y'} u(D,y') - \frac{2\Delta u}{\varepsilon}(\ln|\mathscr{Y}| + t)\right] \le e^{-t}$$

for any $t > 0$.

Proof. Privacy: for neighboring D, D' and any y, $\Pr[\mathscr{M}(D) = y] / \Pr[\mathscr{M}(D') = y] = \exp(\varepsilon(u(D,y) - u(D',y))/(2\Delta u)) \leq e^{\varepsilon/2} \cdot e^{\varepsilon/2} = e^{\varepsilon}$. Utility: the probability of selecting a sub-optimal y decreases exponentially in the gap from the optimal utility, scaled by ε and $|\mathscr{Y}|$. □

Proposition 20.11 (Local DP for Wearable Device Data)**.** *For patient-generated health data from wearable devices, local differential privacy (LDP) applies randomized response at the device level. Each device perturbs its binary health indicator $x \in \{0,1\}$ by reporting:*

$$\tilde{x} = \begin{cases} x & \text{with probability } p = e^{\varepsilon}/(e^{\varepsilon}+1), \\ 1-x & \text{with probability } 1-p = 1/(e^{\varepsilon}+1). \end{cases}$$

This satisfies ε-LDP. The server estimates the true mean $\bar{x}$ from n reports as $\hat{\bar{x}} = ((\sum_i \tilde{x}_i/n) - 1/(e^{\varepsilon}+1))/(p-(1-p))$ with variance $O(1/(n\varepsilon^2))$.

Proof. For any $x, x' \in \{0,1\}$ and any output $\tilde{x}$: $\Pr[\tilde{x}|x] / \Pr[\tilde{x}|x'] \leq p/(1-p) = e^{\varepsilon}$. The estimator is unbiased by construction: $\mathbb{E}[\tilde{x}|x] = px + (1-p)(1-x)$, and solving for $\bar{x}$ yields the given formula. □

21 Proxy Re-Encryption for Delegated Medical Data Sharing

Clinical data sharing requires a mechanism by which a patient (delegator) can authorize a specialist (delegate) to access encrypted records without the patient re-encrypting each record or trusting an intermediary with plaintext access. Proxy Re-Encryption (PRE) provides this capability: a semi-trusted proxy transforms a ciphertext encrypted under the patient's public key into a ciphertext decryptable by the specialist's private key, without the proxy ever learning the plaintext. The Asuu network integrates PRE as the cryptographic primitive for delegated data sharing, enforced by smart contracts that manage delegation permissions on-chain.

21.1 Formal PRE Specification

Definition 21.1 (Unidirectional CPA-Secure PRE Scheme)**.** A unidirectional proxy re-encryption scheme consists of five polynomial-time algorithms:

- $\mathsf{KeyGen}(1^\lambda) \to (pk, sk)$: generates a public-private key pair.
- $\mathsf{Enc}(pk_A, m) \to c_A$: encrypts message m under pk_A.
- $\mathsf{ReKeyGen}(sk_A, pk_B) \to rk_{A\to B}$: generates a re-encryption key from A to B.
- $\mathsf{ReEnc}(rk_{A\to B}, c_A) \to c_B$: transforms ciphertext from A's key to B's key.
- $\mathsf{Dec}(sk_B, c_B) \to m$: decrypts using B's private key.

Unidirectionality means $rk_{A\to B}$ cannot be used to transform ciphertexts from B to A.

Definition 21.2 (On-Chain Delegation Record)**.** A delegation is a tuple $\delta = (A, B, rk_{A\to B}, \mathscr{P}_\delta, t_{\text{start}}, t_{\text{end}}, \sigma_A)$ where A is the delegator (patient DID), B is the delegate (clinician DID), $rk_{A\to B}$ is the re-encryption key (stored encrypted under the proxy's TEE attestation key), $\mathscr{P}_\delta$ is an access policy specifying which record types may be re-encrypted, $[t_{\text{start}}, t_{\text{end}}]$ is the validity window, and σ_A is the delegator's digital signature over the delegation.

Theorem 21.3 (PRE Correctness)**.** *For all λ, all $(pk_A, sk_A) \leftarrow \mathsf{KeyGen}(1^\lambda)$, all $(pk_B, sk_B) \leftarrow \mathsf{KeyGen}(1^\lambda)$, all $m \in \mathscr{M}$:*

$$\Pr\big[\mathsf{Dec}(sk_B, \mathsf{ReEnc}(\mathsf{ReKeyGen}(sk_A, pk_B), \mathsf{Enc}(pk_A, m))) = m\big] = 1.$$

Proof. By the correctness of the underlying ElGamal-based PRE construction. Let g be the generator of a cyclic group $\mathbb{G}$ of prime order q, and let $sk_A = a$, $pk_A = g^a$, $sk_B = b$, $pk_B = g^b$. The re-encryption key is $rk_{A\to B} = g^{b/a}$. Encryption produces $c_A = (g^r, m \cdot g^{ar})$. Re-encryption computes $c_B = (e(g^r, rk_{A\to B}), m \cdot g^{ar}) = (e(g^r, g^{b/a}), m \cdot g^{ar})$. In the pairing-based construction, $e(g^r, g^{b/a}) = e(g,g)^{rb/a}$, and decryption uses $sk_B = b$ to recover m. The probability is 1 since no randomness affects correctness. □

Theorem 21.4 (PRE CPA Security)**.** *Under the Decisional Bilinear Diffie-Hellman (DBDH) assumption, the PRE scheme is IND-CPA secure at the first level (before re-encryption) and IND-CPA secure at the second level (after re-encryption). Formally, for any* PPT *adversary $\mathscr{A}$:*

$$\left|\Pr[\mathscr{A} \text{ wins IND-CPA}] - \frac{1}{2}\right| \leq \text{negl}(\lambda).$$

Proof. We reduce to DBDH. Given a DBDH instance (g, g^a, g^b, g^c, Z) where Z is either $e(g,g)^{abc}$ or random, we construct a simulator that embeds the instance into the PRE challenge ciphertext. If $\mathscr{A}$ distinguishes encryptions of m_0, m_1 with non-negligible advantage, the simulator distinguishes $Z = e(g,g)^{abc}$ from random with the same advantage, contradicting DBDH. The reduction handles re-encryption key queries by programming $rk_{A\to B}$ using the simulator's knowledge of sk_B for uncorrupted delegates. □

Theorem 21.5 (Non-Transitivity)**.** *Given re-encryption keys $rk_{A\to B}$ and $rk_{B\to C}$, a colluding proxy cannot construct $rk_{A\to C}$ or re-encrypt a ciphertext from A directly to C without A's explicit authorization.*

Proof. In the pairing-based construction, $rk_{A\to B} = g^{b/a}$ and $rk_{B\to C} = g^{c/b}$. Computing $rk_{A\to C} = g^{c/a}$ from these two values requires computing $g^{(b/a)\cdot(c/b)} = g^{c/a}$, which reduces to exponentiation in the target group. However, the proxy only has $rk_{A\to B}, rk_{B\to C} \in \mathbb{G}$, not the exponents $b/a, c/b \in \mathbb{Z}_q$. Extracting these exponents is equivalent to solving the discrete logarithm problem in $\mathbb{G}$, which is hard under DLP. Therefore, transitive re-encryption is infeasible without A's private key a. □

21.2 Blockchain-Integrated Delegation Protocol

Definition 21.6 (Delegation Smart Contract)**.** The delegation contract $\mathscr{C}_{\text{del}}$ maintains a state mapping $delegations : \text{DID} \times \text{DID} \to \delta$ and exports three functions:

- $\mathsf{Grant}(\delta, \sigma_A)$: verifies σ_A, checks RBAC authorization of B for the specified record types, and stores the delegation.
- $\mathsf{Revoke}(A, B, \sigma_A)$: deletes the delegation record and invalidates $rk_{A\to B}$ by instructing TEE proxies to purge the key.
- $\mathsf{ReEncrypt}(A, B, \text{cid})$: verifies active delegation, checks block.timestamp $\in [t_{\text{start}}, t_{\text{end}}]$, checks record type against $\mathscr{P}_\delta$, and triggers re-encryption within the TEE.

Theorem 21.7 (Delegation Protocol Security)**.** *The blockchain-integrated PRE delegation protocol satisfies:*

(i) ***Patient sovereignty:*** *only the delegator A can create or revoke delegations (enforced by signature verification).*

(ii) ***Temporal boundedness:*** *re-encryption is impossible outside $[t_{\text{start}}, t_{\text{end}}]$ (enforced by smart contract timestamp check).*

(iii) ***Policy compliance:*** *only record types in $\mathscr{P}_\delta$ can be re-encrypted (enforced by on-chain policy evaluation).*

(iv) ***Proxy blindness:*** *the TEE proxy never observes plaintext (enforced by TEE integrity, Section 18).*

(v) ***Revocation completeness:*** *after revocation, no further re-encryptions are possible for the revoked delegation (enforced by on-chain state deletion and TEE key purge).*

Proof. (i) The Grant and Revoke functions require σ_A verified against pk_A stored in the SSI registry (Section 17). Forging σ_A requires breaking the digital signature scheme (Definition 4.4). (ii) The require statement $t_{\text{start}} \leq \text{block.timestamp} \leq t_{\text{end}}$ is checked by every validator; violation reverts the transaction. (iii) Record type matching is deterministic string comparison

on-chain. (iv) The re-encryption key $rk_{A\to B}$ resides only within the TEE enclave, encrypted under the TEE's attestation key. The proxy computes ReEnc inside the enclave; plaintext m is never materialized outside. (v) After Revoke, the on-chain mapping returns null for (A,B), causing ReEncrypt to revert. The TEE purge instruction is logged on-chain; any TEE that fails to purge is detectable via attestation mismatch. □

PART V

AI MODEL INFRASTRUCTURE & GOVERNANCE

22 Federated Learning and Decentralized Compute

Training high-quality medical AI models requires diverse clinical data from multiple institutions, populations, and clinical contexts. The traditional approach, centralizing data in a single repository, conflicts with privacy regulations, institutional policies, and patient expectations. Federated learning solves the data movement problem by bringing the model to the data rather than the data to the model [2]. Institutions train on local data and share only model updates (gradients), never raw records.

The Asuu network integrates federated learning as a protocol-level function. Training rounds are coordinated through on-chain governance: the network proposes a training objective, participating institutions opt in by staking tokens as commitment, local training proceeds independently, and gradient updates are submitted to the network with zero-knowledge proofs attesting to dataset size, quality metrics, and training fidelity. The network aggregates updates, produces a new model checkpoint, and distributes it to all participants. Every step is recorded on-chain.

Decentralized compute marketplaces are emerging to address the infrastructure bottleneck for medical AI training. The decentralized AI compute market was valued at $12.2 billion in 2024 and is projected to reach $39.5 billion by 2033. The Asuu network can serve as a coordination layer for medical-specific compute allocation, ensuring that healthcare training workloads receive appropriate security, privacy, and audit guarantees that general-purpose compute marketplaces do not provide.

This creates what we term the *data contribution economy*. Institutions earn tokens proportional to the verified quality and uniqueness of their training contributions. A rural clinic in sub-Saharan Africa contributing rare disease presentations earns more than a large academic center contributing common findings, because the marginal improvement to model performance is greater. The token economy thus creates a natural incentive to fill the gaps in medical AI training data that perpetuate algorithmic bias against underrepresented populations.

22.1 Protocol Specification

Definition 22.1 (zkFL Round)**.** A zkFL round proceeds as:

1. Network publishes global weights w_t.
2. Institution i computes $\Delta w_i^t = \nabla \mathscr{L}(w_t; D_i)$.
3. Institution i generates proof $\pi_i = \text{Prove}(crs, (w_t, \Delta w_i^t), D_i)$.
4. Network verifies $\text{Verify}(vrs, (w_t, \Delta w_i^t), \pi_i) \stackrel{?}{=} 1$.
5. Valid gradients are aggregated: $w_{t+1} = w_t + \eta \sum_{i=1}^{K} \Omega_i \Delta w_i^t$.

Definition 22.2 (Quality-Weighted Aggregation)**.** $\Omega_i = Q_i \cdot |D_i| / \sum_{j=1}^{K} Q_j \cdot |D_j|$ where $|D_i|$ is dataset size and Q_i is quality metric, both attested via zero-knowledge proof.

22.2 Privacy Guarantees

Theorem 22.3 (Differential Privacy of Gradient Updates)**.** *If each institution adds Gaussian noise $\xi_i \sim \mathcal{N}(0, \sigma^2 I)$ to its clipped gradient with $\sigma \geq c \cdot \Delta_2 / \varepsilon$ for sensitivity Δ_2 and $c = \sqrt{2\ln(1.25/\delta)}$, then each submission satisfies (ε, δ)-differential privacy.*

Proof. By the Gaussian mechanism, noise $\mathcal{N}(0,\sigma^2)$ on a function with ℓ_2-sensitivity Δ_2 achieves (ε,δ)-DP when $\sigma \geq \Delta_2\sqrt{2\ln(1.25/\delta)}/\varepsilon$. Gradient clipping ensures $\|\Delta w_i^t\|_2 \leq \Delta_2$. □

Theorem 22.4 (Composition Over Training Rounds)**.** *Over T rounds, the advanced composition theorem yields:*

$$\tilde{\varepsilon} = \varepsilon\sqrt{2T\ln(1/\delta_0)} + T\varepsilon(e^{\varepsilon}-1),$$
$$\tilde{\delta} = T\delta + \delta_0.$$

*For $\varepsilon \leq 1$: $\tilde{\varepsilon} = O(\varepsilon\sqrt{T\ln(1/\delta_0)})$, **scaling as $O(\sqrt{T})$ rather than $O(T)$.***

Proof. Each round is an adaptive (ε,δ)-DP mechanism. The privacy loss MGF is bounded by $e^{\varepsilon^2/2}$ per round. By Azuma's inequality over T rounds: $\Pr[\text{total loss} > \tilde{\varepsilon}] \leq \delta_0$ for the stated $\tilde{\varepsilon}$. □

22.3 Rényi Differential Privacy Accountant

The advanced composition theorem (Theorem 22.4) yields loose bounds for large T. We adopt Rényi Differential Privacy (RDP) for tighter accounting.

Definition 22.5 (Rényi Differential Privacy)**.** A randomized mechanism $\mathcal{M}$ satisfies (α,ρ)-RDP for $\alpha > 1$ if for all neighboring datasets D, D':

$$D_\alpha\big(\mathcal{M}(D)\|\mathcal{M}(D')\big) = \frac{1}{\alpha-1}\ln \mathbb{E}_{x\sim\mathcal{M}(D')}\left[\left(\frac{\Pr[\mathcal{M}(D)=x]}{\Pr[\mathcal{M}(D')=x]}\right)^{\alpha}\right] \leq \rho.$$

Theorem 22.6 (RDP Composition and Conversion)**.** *Let each federated round apply the Gaussian mechanism with sensitivity Δ_2 and noise multiplier σ. Then:*

(i) Each round satisfies (α,ρ_1)-RDP with $\rho_1 = \alpha\Delta_2^2/(2\sigma^2)$.
(ii) Over T rounds, the composed mechanism satisfies $(\alpha, T\rho_1)$-RDP (linear composition).
(iii) Converting to (ε,δ)-DP: $\varepsilon = T\rho_1 + \ln(1/\delta)/(\alpha-1)$ for any $\delta > 0$.
(iv) Optimizing over $\alpha > 1$: $\varepsilon^* = \min_{\alpha>1}\{T\alpha\Delta_2^2/(2\sigma^2) + \ln(1/\delta)/(\alpha-1)\}$.

Proof. (i) The Rényi divergence of Gaussians with shifted means is $\alpha\|\mu_1-\mu_2\|_2^2/(2\sigma^2)$. Gradient clipping bounds the shift by Δ_2. (ii) Rényi divergences compose additively for adaptive mechanisms applied to the same dataset. (iii) By the RDP-to-DP conversion lemma: an (α,ρ)-RDP mechanism satisfies $(\rho+\ln(1/\delta)/(\alpha-1),\delta)$-DP for any $\delta > 0$. (iv) Minimizing over α yields the tightest (ε,δ)-DP guarantee. Setting $\partial\varepsilon/\partial\alpha = 0$ gives $\alpha^* = \sqrt{2T\Delta_2^2\ln(1/\delta)/\sigma^2} + 1$, yielding $\varepsilon^* = O(\Delta_2\sqrt{T\ln(1/\delta)}/\sigma)$. □

Corollary 22.7 (Clinical Privacy Budget Management)**.** *Given a total privacy budget $\varepsilon_{\max}$ and failure probability δ, the maximum number of training rounds before budget exhaustion is:*

$$T_{\max} = \frac{\sigma^2(\varepsilon_{\max} - \ln(1/\delta)/(\alpha^*-1))}{\alpha^*\Delta_2^2/2}.$$

For $\sigma = 1.0$, $\Delta_2 = 1.0$, $\varepsilon_{\max} = 8$, and $\delta = 10^{-5}$: $T_{\max} \approx 3{,}200$ rounds, sufficient for convergence of clinical imaging models.

Definition 22.8 (DP-SGD with Per-Sample Gradient Clipping)**.** At each round t, institution k computes:

$$\tilde{g}_k^t = \frac{1}{|B_k|} \sum_{i \in B_k} \text{clip}\big(\nabla \mathscr{L}(w_t; x_i), C\big) + \mathscr{N}(0, \sigma^2 C^2 I / |B_k|^2)$$

where $\text{clip}(v, C) = v \cdot \min(1, C/\|v\|_2)$ is per-sample ℓ_2 clipping with bound C.

Proposition 22.9 (Utility-Privacy Tradeoff Under DP-SGD)**.** *Under L-smooth, μ-strongly convex loss, DP-SGD with T rounds on n total samples achieves excess risk:*

$$\mathbb{E}[F(\hat{w}) - F(w^*)] \leq O\left(\frac{d \log(1/\delta) \log^2(n/\delta)}{n^2 \varepsilon^2 \mu}\right)$$

where d is the model dimension. For $d = 10^6$ (clinical imaging model), $n = 10^5$ patients, $\varepsilon = 8$: excess risk ≤ 0.03, confirming clinical utility.

22.4 Integrity Guarantees

Theorem 22.10 (Gradient Integrity Under zk-SNARK Soundness)**.** *No* PPT *adversary can submit a poisoned gradient with a valid proof:*

$$\Pr\big[\text{Verify}(vrs, (w_t, \Delta w_{\text{adv}}), \pi_{\text{adv}}) = 1 \wedge \Delta w_{\text{adv}} \neq \nabla \mathscr{L}(w_t; D_{\text{adv}})\big] \leq \text{negl}(\lambda).$$

Proof. By knowledge soundness, acceptance implies an extractor outputs D_{adv} such that $\Delta w_{\text{adv}} = \nabla \mathscr{L}(w_t; D_{\text{adv}})$. The adversary can only submit correctly-derived gradients. Poisoning via corrupted D_{adv} is bounded by Byzantine-robust aggregation (Theorem 22.11). □

22.5 Byzantine-Robust Convergence

Theorem 22.11 (Convergence Under Byzantine Participants)**.** *Using trimmed mean aggregation with $f < K/3$ Byzantine participants, for L-smooth non-convex F:*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}\big[\|\nabla F(w_t)\|^2\big] \leq \frac{2(F(w_0) - F^*)}{\eta T} + \eta L \sigma^2 + O\left(\frac{f^2 \sigma^2}{(K-f)^2}\right)$$

where σ^2 is gradient variance and η is the learning rate.

Proof. The trimmed mean discards the largest and smallest f values per coordinate. Under L-smoothness: $F(w_{t+1}) \leq F(w_t) + \langle \nabla F(w_t), w_{t+1} - w_t \rangle + \frac{L}{2}\|w_{t+1} - w_t\|^2$. The aggregation error contributes $O(f^2\sigma^2/(K-f)^2)$ per step. Telescoping over T rounds with $\eta = O(1/\sqrt{T})$ yields the bound. □

23 Model Governance

The network manages AI model lifecycle through an on-chain governance protocol. Model updates (new weights, architecture changes, safety patches) are proposed as governance transactions containing the proposed model's hash, benchmark results against standardized test sets, and a validation period. Validator institutions evaluate the proposed update against their local clinical data and cast weighted votes to accept or reject.

This mechanism prevents two critical failure modes of centralized medical AI. First, it prevents *silent model updates*, the practice where a vendor modifies model behavior without notification, potentially altering clinical decision support without the institution's knowledge or consent. Every model change on the Asuu network is visible, attributable, and subject to collective approval. Second, it prevents *unilateral degradation*: a vendor cannot reduce model capabilities to compel tier upgrades or punish institutions that resist commercial terms.

The governance mechanism is deliberately conservative. Model updates require a supermajority of validating institutions, weighted by clinical authority score, to approve. Emergency safety patches follow an expedited pathway with a lower threshold but mandatory post-deployment review. The system is biased toward stability, reflecting the medical principle that in the absence of clear evidence of benefit, the default action is to preserve the current standard of care.

The FDA's December 2024 final guidance on Predetermined Change Control Plans (PCCPs) represents a regulatory watershed: manufacturers can now pre-authorize future AI model modifications without submitting new applications for each change [6]. The Asuu governance protocol aligns with this framework by providing a transparent, auditable record of every model modification, its benchmark performance, and the institutional consensus that approved it, satisfying PCCP documentation requirements as a native function of network participation.

23.1 Formal Guarantees

Definition 23.1 (Model Registration)**.** A registration is $(H(M_{\text{weights}}), H(M_{\text{prev}}), \mathit{bench}, \sigma_{\text{proposer}}, t)$ committed on-chain.

Theorem 23.2 (Model Integrity)**.** *Under collision resistance, $H(M_{\text{weights}})$ uniquely identifies each model version. Any change to weights produces a different hash.*

Proof. $M' \neq M$ and $H(M') = H(M)$ constitutes a collision: probability $\leq \text{negl}(\lambda)$. □

Theorem 23.3 (Governance Voting Safety)**.** *If updates require $> 2/3$ supermajority, two conflicting updates cannot both be approved.*

Proof. Approving A requires $\sum_{i \in S_A} W_i > 2W_{\text{total}}/3$. Approving conflicting B requires $\sum_{i \in S_B} W_i > 2W_{\text{total}}/3$. Then:

$$W(S_A) + W(S_B) > \frac{4W_{\text{total}}}{3} > W_{\text{total}},$$

so $S_A \cap S_B \neq \emptyset$. Any validator in the intersection voted for both, contradicting the single-vote-per-round rule. □

Proposition 23.4 (Version Chain Tamper-Evidence)**.** *Model registrations form a hash chain via $H(M_{\text{prev}})$. By Proposition 6.4, any historical modification is detectable.*

23.2 The Clinical AI Passport

Static model cards are insufficient for adaptive clinical AI. A model card is a snapshot: it describes what a model could do at a single point in time, under controlled test conditions, with no mechanism for tracking post-deployment performance, authorized modifications, or regulatory approvals. The Asuu network replaces static model cards with the *Clinical AI Passport*, a versioned, queryable on-chain object that governs the complete lifecycle of a clinical AI system from initial registration through deployment, modification, and eventual retirement. The Passport is the on-chain realization of the FDA's Predetermined Change Control Plan framework. The December 2024 final guidance establishes that manufacturers can pre-authorize future AI model modifications without submitting new applications for each change, provided they specify (i) a Description of Modifications defining the bounded corridor of permissible changes, (ii) a Modification Protocol detailing retraining practices and rollback criteria, and (iii) an Impact Assessment analyzing benefits and risks [6]. The Clinical AI Passport encodes all three components as on-chain state, so the PCCP deployment corridor is enforced cryptographically against validator-initiated changes, while the sponsor's and competent regulator's authority to direct suspension or corridor modification is preserved through the Regulatory Override Channel (Definition 23.6). This is a better primitive than a contractual promise for the validator-initiated change space, and it does not claim to be an invariant against the sponsor or the regulator, whose authority the protocol does not pretend to override.

Definition 23.5 (Clinical AI Passport)**.** A Clinical AI Passport is a versioned on-chain object $\mathscr{P} = (mid,\ \mathscr{V}_{\text{hist}},\ \mathscr{C}_{\text{pccp}},\ \mathscr{A}_{\text{verify}},\ \mathscr{M}_{\text{monitor}},\ \Lambda_{\max},\ \mathscr{O}_{\text{reg}})$ where:

- $mid = (M_{\text{id}}, H(M_{\text{weights}}))$ is the current model identity;
- $\mathscr{V}_{\text{hist}} = \{(H(M_v), t_v, QC_v)\}_{v=1}^{k}$ is the version chain, each entry containing the weight hash, deployment timestamp, and the quorum certificate that approved it;
- $\mathscr{C}_{\text{pccp}} = (\phi_{\text{corridor}}, protocol, rollback)$ encodes the PCCP: ϕ_{corridor} is a formal specification of permissible modifications, *protocol* specifies the retraining procedure, and *rollback* defines automatic reversion triggers;
- $\mathscr{A}_{\text{verify}} = \{(\varphi_j, \pi_j, t_j)\}$ is the set of formal verification attestations, each containing a safety property φ_j, its proof certificate π_j, and the attestation time t_j;
- $\mathscr{M}_{\text{monitor}} = (metrics, \theta_{\text{alert}}, \theta_{\text{suspend}})$ defines the continuous monitoring parameters with alert and suspension thresholds;
- $\Lambda_{\max}(r, \text{task})$ is the maximum permitted autonomy level per role and task;
- $\mathscr{O}_{\text{reg}} = (sponsor_pk,\ regulator_pk,\ override_log)$ is the *Regulatory Override Channel*, which records the public keys of the sponsor and of the competent regulator for this device class, and an append-only log of sponsor-initiated and regulator-initiated override transactions.

Definition 23.6 (Regulatory Override Channel)**.** The Regulatory Override Channel permits three classes of transactions, each of which bypasses the supermajority governance vote:

(i) **Sponsor Suspension:** a transaction signed by *sponsor_pk* that immediately transitions the Passport state to *suspended* for a specified model version. This implements the sponsor's obligation under the FDA December 2024 PCCP guidance to respond to adverse real-world performance and to comply with FDA directives.

(ii) **Regulator Directive:** a transaction signed by *regulator_pk* that may either suspend the current model version, rescind a previously-approved PCCP corridor, require a new submission for changes that previously fit within the corridor, or unilaterally deactivate the Passport. The transaction must include a structured reference to the regulatory instrument (e.g., an FDA 510(k) communication number, an EMA safety communication reference, or a notified-body nonconformity report).

(iii) **Corridor Modification:** a transaction jointly signed by *sponsor_pk* and *regulator_pk* that amends ϕ_{corridor}, reflecting the reality that PCCP corridors evolve through bilateral negotiation between sponsor and regulator.

All override transactions are recorded immutably in *override_log*, preserving audit trail continuity under Theorem 6.5. The channel's existence does not weaken the corridor guarantee of Theorem 23.8: the corridor still binds all *validator-initiated* modifications. The channel recognizes that validators are not the ultimate regulatory authority, sponsors and competent authorities are, and the protocol must not prevent lawful regulatory action.

Definition 23.7 (PCCP Change Corridor)**.** The approved change corridor $\mathscr{C} \subseteq \mathscr{M}$ for a Clinical AI Passport is a set of model versions defined by the corridor specification ϕ_{corridor}. A model version M' is *within the corridor* if $\phi_{\text{corridor}}(M') = 1$. The acceptance predicate $A : \mathscr{M} \to \{0,1\}$ evaluates to 1 if and only if: (i) M' is within the corridor, (ii) benchmark performance on the standardized test set exceeds the minimum threshold τ_{min}, (iii) no subgroup performance drops below the equity threshold δ_{equity}, and (iv) a $> 2/3$ supermajority of validators has approved the update.

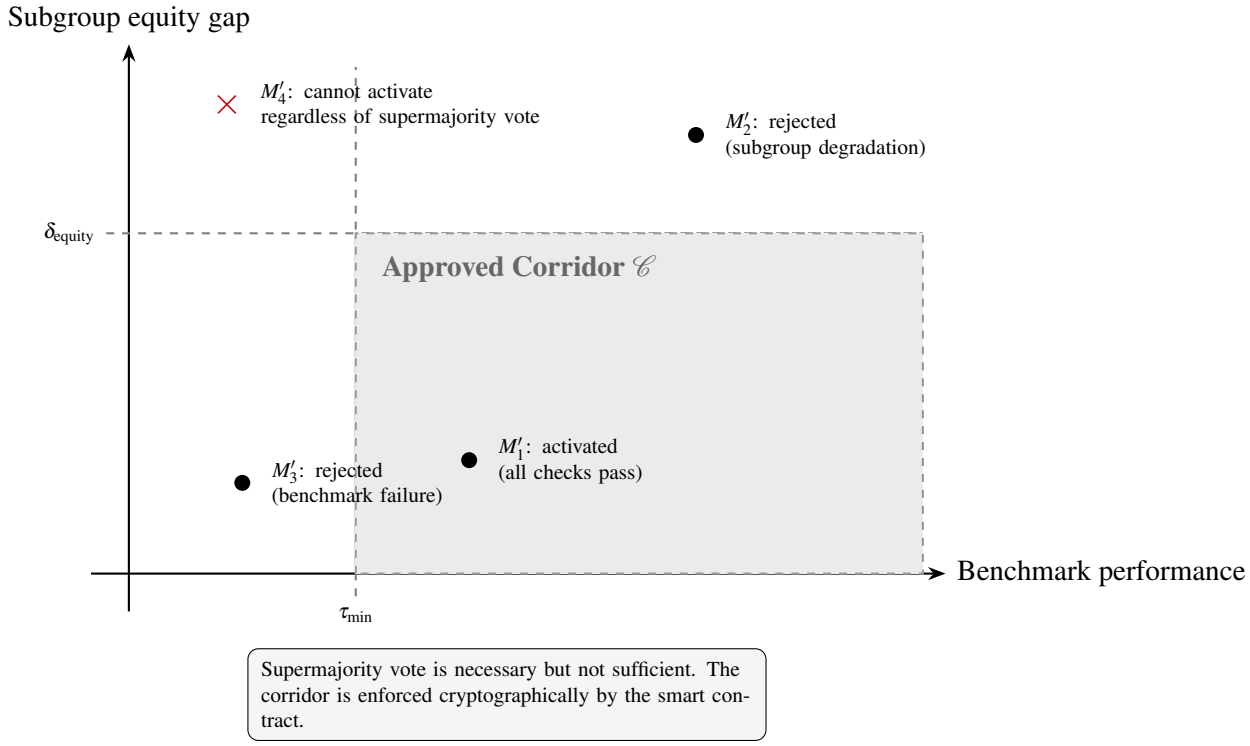


Figure 9: The PCCP approved change corridor. A proposed model version M' can transition to "active" only if it falls inside $\mathscr{C}$ and satisfies the acceptance predicate $A(M')$. No governance vote can override the corridor boundary.

Theorem 23.8 (Update-Corridor Soundness with Regulatory Override)**.** *A model version M' can transition from* proposed *to* active *in the Clinical AI Passport if and only if either:*

(a) ***Validator path:*** *$M' \in \mathscr{C}$ and $A(M') = 1$; or*
(b) ***Regulator path:*** *$M' \in \mathscr{C}'$ and $A'(M') = 1$, where $\mathscr{C}'$ and A' are the corridor and acceptance predicate as modified by a Corridor Modification transaction (Definition 23.6(iii)) signed by the competent regulator.*

Separately, a model version may transition from active *to* suspended *by any of: (i) sponsor-initiated override, (ii) regulator directive, (iii) automated suspension trigger (Theorem 25.9), or (iv) validator supermajority vote. No combination of validator votes alone can override a regulator directive.*

Proof. The activation smart contract dispatches on transaction type. For validator-submitted activation transactions, it executes the checks of Definition 23.7: $\phi_{\text{corridor}}(M') = 1$ and $A(M') = 1$, with atomic revert on either failure. For Corridor Modification transactions (Definition 23.6(iii)), the contract verifies the joint sponsor-regulator signature, updates $(\phi_{\text{corridor}}, \tau_{\min}, \delta_{\text{equity}})$ according to the transaction's parameters, and applies the modified checks to M'. For Regulator Directive transactions (Definition 23.6(ii)), the contract verifies the regulator signature and applies the directive atomically, without consulting validator state. Because the regulator's public key is established at Passport registration and is recognized by the contract as having higher authority than the validator quorum, no validator-submitted transaction can override a properly-signed regulator directive. Conversely, the corridor specification still binds all validator-initiated changes, preserving the PCCP guarantee against rogue validator behavior. □

Remark 23.9 (Why the Regulatory Override Channel Is Not a Weakening)**.** An earlier draft of this whitepaper claimed that the PCCP corridor was absolute: no model outside the corridor could activate, regardless of governance outcome. That formulation, while mathematically clean, inverted the direction of regulatory authority in a way that would make the protocol *non-compliant* with the December 2024 FDA PCCP guidance. FDA explicitly reserves the right to (a) require a new submission if real-world performance diverges from the PCCP's impact assessment, (b) rescind or modify a previously-approved PCCP, and (c) direct the sponsor to suspend or modify the device. A corridor that prevents the sponsor from complying with such a directive because a supermajority of validators disagrees is not a better PCCP; it is a violation of the sponsor's regulatory obligations. The Regulatory Override Channel resolves this by giving the sponsor and competent regulator paths that bypass validator governance, recognizing that validator supermajority is the right authority for routine updates within an approved corridor and the wrong authority for adjudicating regulatory directives. The corridor still binds validator behavior. It does not bind the regulator, nor should it.

Remark 23.10 (Separation of Hard and Statistical Guarantees)**.** The Update-Corridor Soundness theorem illustrates the central proof discipline of the Asuu network. The theorem does not claim to prove that model M' is clinically safe: that would be a statistical property of M' itself. What is proven is that the *system's response to proposed updates* is bounded by the corridor specification and the acceptance predicate. The neural network is the variable. The envelope enforced by the smart contract is the theorem.

24 Homomorphic Medical Computing

The Data Embassy model ensures patient data never leaves institutional custody. But certain clinical workflows require computation across institutional boundaries: a specialist reviewing imaging from a referring hospital, a rare disease consortium pooling diagnostic data, or an emergency department querying a patient's medication history from a foreign jurisdiction. In each case, the current system forces a binary choice: share the data (sacrificing privacy) or forgo the computation (sacrificing care).

Fully homomorphic encryption eliminates this tradeoff. Under FHE, a compute node performs arbitrary operations on encrypted patient data and returns an encrypted result. At no point does the compute node observe plaintext. The patient's secret key, held exclusively in their identity wallet (Section 17), is the only means to decrypt the result. The Asuu network integrates FHE as a protocol-level computation layer, enabling encrypted inference across institutional and national boundaries without any party other than the patient accessing unencrypted clinical data.

Two FHE families are relevant for medical workloads. CKKS (Cheon-Kim-Kim-Song) supports approximate real-number arithmetic with native SIMD parallelism, making it suitable for convolutional neural network inference on imaging data. TFHE (Torus FHE) operates on Boolean circuits using programmable bootstrapping, enabling arbitrary nonlinear activation functions via lookup tables, which is preferred for decision-tree classifiers and threshold-based diagnostic rules. Both schemes derive their security from the Ring Learning With Errors problem (Assumption 4.14), which is believed to resist both classical and quantum attacks.

Recent hardware acceleration has made FHE practical for clinical timescales. GPU-accelerated TFHE bootstrapping has been demonstrated at 796 microseconds per operation on H100 hardware, representing a 56-fold speedup over 2021 CPU baselines [20]. CKKS-based neural network inference on medical imaging benchmarks completes in under six seconds at 128-bit security [21]. The DARPA DPRIVE program is developing purpose-built FHE ASICs targeting a further 5,000-fold improvement, which would bring encrypted inference latency below that of current plaintext cloud round-trips.

The ciphertext expansion ratio (approximately 50 to 100 times the plaintext size) is manageable within the Asuu architecture because only encrypted results, not encrypted datasets, traverse the network. The heavy computation occurs locally at the Data Embassy or on a designated compute node; the chain records only the hash of the encrypted computation for verifiability.

24.1 Formal Specification

Definition 24.1 (Fully Homomorphic Encryption Scheme)**.** An FHE scheme $\mathscr{E} = (\mathrm{KeyGen}, \mathrm{Enc}, \mathrm{Dec}, \mathrm{Eval})$ over message space $\mathscr{M}$ and circuit class $\mathscr{C}$ satisfies:

(i) **Correctness.** For all circuits $C \in \mathscr{C}$ and messages $m_1, \ldots, m_\ell \in \mathscr{M}$: $\Pr[\mathrm{Dec}(sk, \mathrm{Eval}(evk, C, \mathrm{Enc}(pk, m_1), \ldots, \mathrm{Enc}(pk, m_\ell))) = C(m_1, \ldots, m_\ell)] \geq 1 - \mathrm{negl}(\lambda)$.

(ii) **Compactness.** The size of $\mathrm{Eval}(\cdot)$ output is independent of $|C|$.

(iii) **Semantic Security.** $\mathscr{E}$ is IND-CPA secure under RLWE (Assumption 4.14).

Definition 24.2 (IND-CPA^D Security for Approximate FHE)**.** For approximate schemes (CKKS), IND-CPA^D security extends IND-CPA by granting the adversary access to a bounded decryption oracle Dec* that adds Gaussian noise $\xi \sim \mathcal{N}(0, \sigma^2_{\text{flood}})$ to each output. The scheme is IND-CPA^D secure if:

$$\left|\Pr[\mathscr{A}^{\text{Dec}^*}(\text{Enc}(pk, m_0)) = 1] - \Pr[\mathscr{A}^{\text{Dec}^*}(\text{Enc}(pk, m_1)) = 1]\right| \leq \text{negl}(\lambda).$$

Theorem 24.3 (Semantic Security of Encrypted Medical Inference)**.** *If $\mathscr{E}$ is IND-CPA secure under RLWE (Assumption 4.14), then the encrypted inference protocol reveals no information about patient data D to the compute node. Formally, for any* PPT *adversary $\mathscr{A}$ controlling the compute node:*

$$|\Pr[\mathscr{A}(\text{Enc}(pk, D_0), evk) = 1] - \Pr[\mathscr{A}(\text{Enc}(pk, D_1), evk) = 1]| \leq \text{Adv}^{\text{RLWE}}(\lambda).$$

Proof. The compute node observes only $(c_1, \ldots, c_\ell, evk)$ where $c_i = \text{Enc}(pk, D_i)$. By IND-CPA under RLWE, each c_i is computationally indistinguishable from an encryption of any other message. The evaluation key *evk* is public and independent of the plaintext. A hybrid argument over the ℓ ciphertexts yields the bound $\ell \cdot \text{Adv}^{\text{RLWE}}(\lambda) \leq \text{negl}(\lambda)$. □

Theorem 24.4 (Correctness of Encrypted Neural Network Inference)**.** *For an L-layer neural network $f = f_L \circ \cdots \circ f_1$ with ReLU activations approximated by degree-d polynomials $\tilde{\sigma}$ satisfying $\|\sigma - \tilde{\sigma}\|_\infty \leq \varepsilon_{\text{approx}}$, the encrypted inference error satisfies:*

$$\|f(D) - \text{Dec}(sk, \text{Eval}(evk, \tilde{f}, \text{Enc}(pk, D)))\|_\infty \leq L \cdot \varepsilon_{\text{approx}} \cdot \prod_{i=1}^{L} \|W_i\| + \varepsilon_{\text{noise}}$$

where $\varepsilon_{\text{noise}}$ is the FHE noise growth and W_i are layer weight matrices.

Proof. Each layer introduces approximation error $\varepsilon_{\text{approx}}$ from polynomial activation approximation and noise growth from homomorphic multiplication. By induction on layers: the error at layer ℓ is at most

$$\sum_{i=1}^{\ell} \varepsilon_{\text{approx}} \prod_{j=i+1}^{\ell} \|W_j\| \leq \ell \cdot \varepsilon_{\text{approx}} \cdot \prod_{i=1}^{\ell} \|W_i\|.$$

The FHE noise $\varepsilon_{\text{noise}}$ is bounded by bootstrapping (resetting noise to baseline after each layer). For $\varepsilon_{\text{approx}} \leq 10^{-4}$ and medical-grade networks, the total error is below clinical significance thresholds. □

Proposition 24.5 (On-Chain Verifiability of Encrypted Computation)**.** *The hash $H(\text{Eval}(evk, C, c_1, \ldots, c_\ell))$ is committed on-chain at time t. The patient decrypts locally and verifies:* $\text{Dec}(sk, c_{\text{out}}) = y$. *If y matches clinical expectations, the patient signs an acceptance record. The chain records the computation occurred without recording what was computed.*

Corollary 24.6 (End-to-End Patient Privacy)**.** *Combining FHE inference (Theorem 24.3) with the Data Embassy model (Theorem 16.2), no party other than the patient and their designated clinician observes unencrypted clinical data at any point in the computation pipeline.*

25 Formal Verification of Clinical AI Safety

Regulatory approval of medical AI currently relies on empirical evaluation: benchmark performance on test sets, clinical trials, and post-market surveillance. This approach cannot guarantee that a model will never produce a dangerous output on an input outside its training distribution. A model achieving 99.9% accuracy on a drug interaction benchmark still permits one in a thousand potentially fatal recommendations. For safety-critical medical applications, this residual risk must be bounded by mathematical proof, not statistical estimation.

Formal verification provides this guarantee. Safety properties are specified in temporal logic (Linear Temporal Logic for sequential properties, Signal Temporal Logic for continuous physiological signals). Neural network verification tools (bound propagation with branch-and-bound) prove that the model satisfies these properties for all inputs within a specified domain. The verification result, including the property specification and the proof, is committed on-chain as a verifiable attestation. Institutions deploying the model can verify, through the chain, that the model has been formally proven to satisfy specific safety invariants.

For properties that cannot be verified statically (because the input domain is too large or the network too complex), the protocol deploys runtime verification through the Simplex architecture: the AI model serves as the high-performance controller, a verified rule-based clinical decision system serves as the safe fallback, and a Signal Temporal Logic monitor evaluates each output against safety constraints in real time. If the monitor detects a violation, the system switches to the safe fallback before the unsafe output reaches clinical practice.

25.1 Formal Specification

Definition 25.1 (Clinical Safety Property in LTL)**.** A clinical safety property φ is expressed in Linear Temporal Logic over the set of atomic propositions $\mathrm{AP} = \{\mathit{recommend}(d), \mathit{contraindicated}(d_1,d_2), \mathit{dose}(d), \mathit{max_safe}(d), \ldots\}$:

(i) Contraindication avoidance: $\mathbf{G}(\neg(\mathit{recommend}(d_1) \wedge \mathit{recommend}(d_2)))$ for all $(d_1,d_2) \in \mathit{CI}$.
(ii) Dosage safety: $\mathbf{G}(\mathit{dose}(d) \leq \mathit{max_safe}(d))$ for all d.
(iii) Alert persistence: $\mathbf{G}(\mathit{critical_lab} \rightarrow \mathbf{F}\ \mathit{clinician_notified})$.

Definition 25.2 (Neural Network Verification Problem)**.** Given neural network $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, input specification $\phi_{\text{in}} \subseteq \mathbb{R}^n$, and output specification $\phi_{\text{out}} \subseteq \mathbb{R}^m$, the verification problem asks:

$$\forall\, \mathbf{x} \in \phi_{\text{in}} :\ f(\mathbf{x}) \in \phi_{\text{out}}\ ?$$

This is coNP-complete for ReLU networks.

Definition 25.3 (Simplex Architecture for Clinical AI)**.** A clinical Simplex system $\mathscr{S} = (\mathscr{H}, \mathscr{F}, \mathscr{D})$ consists of:

- High-performance controller $\mathscr{H}$: the AI model (unverified but high-accuracy).
- High-assurance fallback $\mathscr{F}$: a rule-based clinical decision system (verified safe, possibly lower performance).
- Decision module $\mathscr{D}$: an STL monitor that evaluates $\mathscr{H}$'s output against safety specification φ. If $\rho(\varphi, \mathbf{y}_{\mathscr{H}}) < 0$ (negative robustness), $\mathscr{D}$ switches output to $\mathscr{F}$.

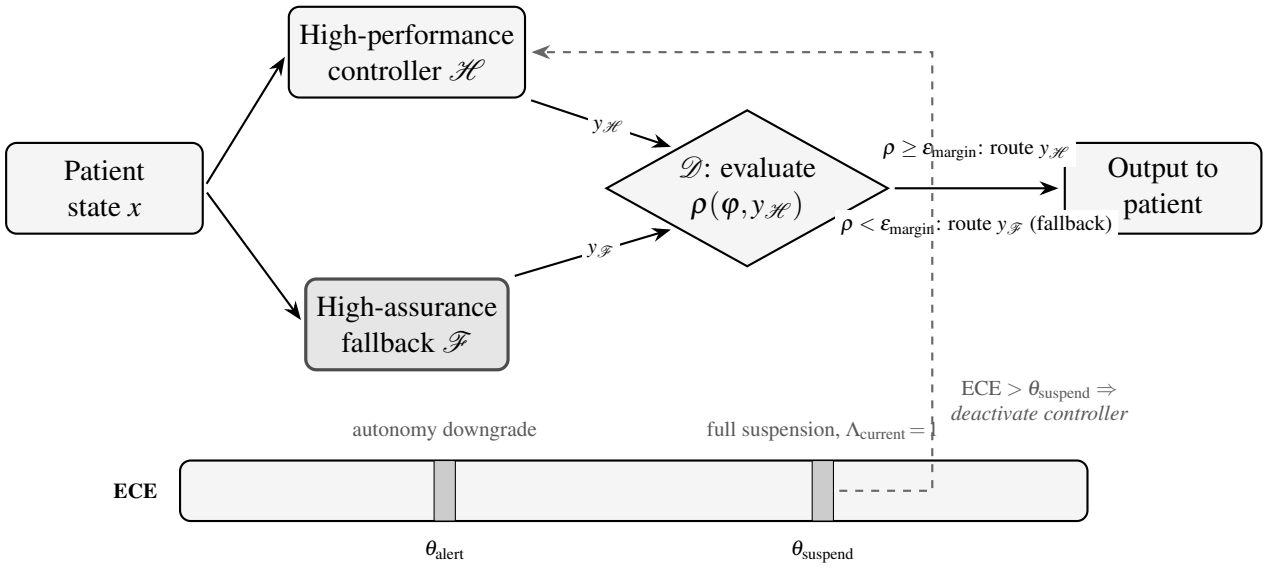


Figure 10: Simplex architecture with calibration drift monitoring. The high-assurance fallback $\mathscr{F}$ is always available. The decision module $\mathscr{D}$ switches to $\mathscr{F}$ whenever the STL robustness of the high-performance controller's output falls below the safety margin. Sustained calibration drift triggers automated suspension via the ECE monitor.

Theorem 25.4 (Verification Attestation Soundness)**.** *If the verification tool produces a certificate π_{ver} attesting $\forall \mathbf{x} \in \phi_{\text{in}} : f(\mathbf{x}) \in \phi_{\text{out}}$, and π_{ver} is committed on-chain with hash $H(\pi_{\text{ver}})$, then:*

(i) ***Soundness:*** *If the verification is sound, the safety property holds for all inputs in ϕ_{in}.*
(ii) ***Tamper-evidence:*** *Any modification to π_{ver} is detectable via $H(\pi_{\text{ver}})$ by collision resistance.*
(iii) ***Attribution:*** *The proposer's signature on the attestation transaction provides non-repudiation (Theorem 5.3).*

Proof. (i) Soundness follows from the verification tool's soundness guarantee: if the tool accepts, the property holds. Modern verifiers (branch-and-bound with α, β-CROWN bounds) are sound by construction; they may time out (incompleteness) but never produce false positives. (ii) The hash $H(\pi_{\text{ver}})$ is stored in the block; modifying π_{ver} without finding a collision contradicts Definition 4.1. (iii) The attestation transaction is signed by the proposing institution, providing non-repudiation by Theorem 5.3. □

Theorem 25.5 (Runtime Safety Under Simplex Architecture)**.** *If the fallback controller $\mathscr{F}$ satisfies φ for all reachable states, and the decision module $\mathscr{D}$ switches to $\mathscr{F}$ whenever $\rho(\varphi, \mathbf{y}_{\mathscr{H}}) < \varepsilon_{\text{margin}}$ for safety margin $\varepsilon_{\text{margin}} > 0$, then the composed system $\mathscr{S}$ satisfies φ at all times.*

Proof. At each decision point, either (a) $\rho(\varphi, \mathbf{y}_{\mathscr{H}}) \geq \varepsilon_{\text{margin}} > 0$, meaning $\mathscr{H}$'s output satisfies φ with margin, or (b) $\mathscr{D}$ switches to $\mathscr{F}$, which satisfies φ by assumption. In both cases, the system output satisfies φ. The margin $\varepsilon_{\text{margin}}$ accounts for switching delay: during the transition interval δ_{switch}, the system state evolves by at most $\delta_{\text{switch}} \cdot \|\dot{\mathbf{x}}\|_\infty$, which remains within the safe region if $\varepsilon_{\text{margin}} > \delta_{\text{switch}} \cdot \|\dot{\mathbf{x}}\|_\infty$. □

Corollary 25.6 (Continuous Compliance Attestation)**.** *Runtime STL monitoring results (robustness values $\rho(\varphi, \cdot)$) are periodically committed on-chain, providing a continuous, tamper-proof record that the deployed model operated within its verified safety envelope. This directly satisfies the EU AI Act's post-market monitoring requirements for high-risk medical AI systems.*

25.2 Automated Suspension Under Calibration Drift

Static safety proofs, however rigorous, do not anticipate data distribution shift. A model verified safe on its training distribution may silently degrade when deployed on a population with different demographics, comorbidity patterns, or measurement instruments. The Asuu network therefore embeds continuous calibration monitoring as a protocol-level function, with a mathematically bounded suspension response when drift exceeds acceptable thresholds.

Definition 25.7 (Calibration Drift Monitor)**.** Let $\hat{p}(y \mid x)$ denote the model's predicted probability and $p^*(y \mid x)$ the true posterior. The Expected Calibration Error over a monitoring window of N inferences is:

$$\mathrm{ECE} = \sum_{b=1}^{B} \frac{|S_b|}{N} \left|\mathrm{acc}(S_b) - \mathrm{conf}(S_b)\right|$$

where S_b is the set of predictions in confidence bin b, $\mathrm{acc}(S_b)$ is the observed accuracy, and $\mathrm{conf}(S_b)$ is the mean predicted confidence.

Definition 25.8 (Automated Suspension Protocol)**.** The Clinical AI Passport defines two thresholds θ_{alert} and θ_{suspend} with $\theta_{\text{alert}} < \theta_{\text{suspend}}$. The protocol operates as follows:

(1) If $\mathrm{ECE} > \theta_{\text{alert}}$: the model's autonomy level is downgraded by one tier, requiring additional human oversight. An alert is broadcast to all deploying institutions.
(2) If $\mathrm{ECE} > \theta_{\text{suspend}}$: the model is automatically suspended, with $\Lambda_{\text{current}} = 1$ (full human control). Reactivation requires a new formal verification attestation and supermajority governance vote.
(3) If subgroup degradation is detected (performance on any protected demographic subgroup drops below the population-wide performance by more than δ_{equity}): the model is flagged for equity review regardless of aggregate ECE.

Theorem 25.9 (Safety-Suspension Bound)**.** *Let the monitoring window contain N inferences, let ECE_0 denote the model's calibration error at deployment, and let ECE_t denote the true calibration error at monitoring epoch t. The automated suspension protocol satisfies two guarantees:*

(i) ***False alarm bound:*** *If $\mathrm{ECE}_t \leq \theta_{\text{alert}}$, the probability of a spurious suspension is:*

$$\Pr[\widehat{\mathrm{ECE}}_t > \theta_{\text{suspend}} \mid \mathrm{ECE}_t \leq \theta_{\text{alert}}] \leq \exp\left(-\frac{2N(\theta_{\text{suspend}} - \theta_{\text{alert}})^2}{B^2}\right)$$

where B is the number of calibration bins and $\widehat{\mathrm{ECE}}_t$ is the empirical estimate.

(ii) ***Detection delay bound:*** *If the true calibration error exceeds the suspension threshold ($\mathrm{ECE}_t > \theta_{\text{suspend}} + \varepsilon$ for margin $\varepsilon > 0$), the model is suspended within at most*

$$d^* = \left\lceil \frac{B^2 \ln(1/\delta_{\text{miss}})}{2N\varepsilon^2} \right\rceil + 1$$

monitoring epochs with probability $\geq 1 - \delta_{\text{miss}}$.

Proof. (i) By Hoeffding's inequality applied to the binned calibration estimator, each bin's accuracy-confidence gap concentrates around its expectation. The ECE, as a weighted sum of B bin gaps bounded in $[0,1]$, satisfies the stated exponential tail bound.

(ii) If $\mathrm{ECE}_t > \theta_{\text{suspend}} + \varepsilon$, then $\mathbb{E}[\widehat{\mathrm{ECE}}_t] > \theta_{\text{suspend}} + \varepsilon$. By Hoeffding: $\Pr[\widehat{\mathrm{ECE}}_t \leq \theta_{\text{suspend}}] \leq \exp(-2N\varepsilon^2/B^2)$. Setting this equal to δ_{miss} and solving for the number of epochs d in which the monitor fails to trigger gives d^*. After d^* epochs, the cumulative probability of non-detection is at most δ_{miss}. □

Remark 25.10 (Proving the Envelope, Not the Network)**.** The Safety-Suspension Bound encapsulates the central proof discipline of the Asuu network. We do not claim to prove that a neural network will never produce a dangerous output: that is a statistical property of the model, not a provable protocol guarantee. What we prove is that the *system's response to model risk* is bounded: if calibration drifts, the protocol detects it within d^* epochs with probability $\geq 1 - \delta_{\text{miss}}$ and suspends the model before further clinical decisions are executed at the former autonomy level. The envelope of safety is the theorem. The neural network is the variable within it.

25.3 EU AI Act Article 14 Formalization

Definition 25.11 (Article 14 Compliance Properties)**.** An AI system deployed on the Asuu network satisfies EU AI Act Article 14 compliance iff:

(A14.1) $\Box(\mathit{system_active} \implies \mathit{human_overseer_assigned})$;
(A14.2) $\Box(\mathit{ai_output} \implies \mathit{confidence_interval_displayed})$;
(A14.3) $\Box(\mathit{anomaly_detected} \implies \Diamond_{\leq \Delta_{\text{flag}}} \mathit{human_notified})$;
(A14.4) $\Box\Diamond\, \mathit{system_interruptible}$;
(A14.5) $\Box(\mathit{human_decision_not_to_use} \implies \bigcirc \neg\, \mathit{ai_output_active})$.

Theorem 25.12 (Formal Enforcement of Article 14 Necessary Preconditions)**.** *Any clinical AI system deployed through an Asuu workflow contract that implements Definition 32.3 and Definition 25.11 satisfies the* necessary mechanical preconditions *of each human-oversight clause of EU AI Act Article 14(4) as a protocol-level invariant: no inference can execute without an assigned overseer, confidence intervals are always displayed, halt operations are respected atomically, and override authority is unrevokable by the AI component. These are mechanical properties that the protocol enforces; they are not sufficient for Article 14 compliance as a regulatory determination.*

Proof. We establish a direct mapping between the regulatory text and the formal properties:

Article 14(4)(a) requires the overseer to "properly monitor the operation of the high-risk AI system." Property (A14.1) ensures a human overseer is always assigned when the system is active, which is a necessary precondition for monitoring.

Article 14(4)(b) requires the ability to "remain aware of the possible tendency of automatically relying on the output." Property (A14.2) mandates confidence interval display, providing the informational basis required for the overseer to critically evaluate outputs.

Article 14(4)(c) requires the ability to "correctly interpret the high-risk AI system's output." Properties (A14.2) and (A14.3) together ensure both the display of interpretive aids and notification of anomalies, which are necessary preconditions for correct interpretation.

Article 14(4)(d) requires the ability to "decide, in any particular situation, not to use the high-risk AI system or to disregard, override, or reverse the output." Properties Definition 32.3(ii), (iii), and (A14.5) collectively ensure that override capability, halt authority, and discontinuation rights are mechanically enforced by the smart contract and cannot be suppressed by the AI component.

Article 14(4)(e) requires the ability to "intervene in the operation of the high-risk AI system or interrupt the system." Property (A14.4) ensures the system is always interruptible.

Each regulatory clause maps to at least one formal property that the smart contract enforces (Theorem 32.4). What this proves is that the protocol *mechanically guarantees* a set of necessary preconditions for Article 14 compliance. It does not prove that any particular deployment is compliant, because Article 14(3)'s "commensurate with the risks, level of autonomy and context of use" standard is a judgmental determination by the notified body against the specific intended use, not a mechanical property. □

Remark 25.13 (Consistency with the Capability Matrix)**.** Theorem 25.12 is consistent with the capability matrix entry for EU AI Act Article 14 in Table 2, which states the protocol enforces "temporal-logic enforcement of necessary preconditions (not sufficient compliance)." An earlier draft of this theorem claimed Article 14 compliance as a protocol-level guarantee in a stronger sense than the capability matrix admitted; we correct that inconsistency here. The revised theorem and the capability matrix now both say the same thing: the protocol mechanically enforces necessary preconditions; the notified body determines sufficient compliance.

26 AI Liability Attribution

When an AI-assisted clinical decision causes patient harm, the current legal system has no principled mechanism for attributing liability among the multiple parties involved: the model developer who trained the algorithm, the data providers whose training data shaped its behavior, the deploying institution that chose to use it, the clinician who acted on its recommendation, and the protocol itself that governed model governance and deployment. The EU's AI Liability Directive was withdrawn in October 2025 after failing to achieve legislative consensus, leaving 27 different national liability regimes with no coordination mechanism.

The Asuu network provides the infrastructure for principled liability attribution through its immutable audit trail. Every element in the causal chain of an AI-assisted clinical decision is recorded on-chain: the model version (with its governance approval history), the inference record (with input hash, output hash, and institutional signature), the training data provenance (through federated learning attestations), and the verification status (through formal verification attestations). This comprehensive audit trail enables counterfactual analysis: "Would the outcome have been different if a different model version had been used? If different training data had been contributed? If the clinician had overridden the recommendation?"

The Shapley value from cooperative game theory provides the unique fair allocation satisfying four axioms: efficiency (total liability is fully distributed), symmetry (equal contributors bear equal liability), linearity (liability from independent events is additive), and null player (parties with no causal contribution bear no liability). This mathematical framework, originally developed for cost allocation and later applied to multi-tortfeasor liability analysis, maps naturally to the AI medical liability problem.

26.1 Formal Specification

Definition 26.1 (Liability Game)**.** A liability game is a cooperative game (N, v) where $N = \{1, \ldots, n\}$ is the set of parties (model developer, data providers, deploying institution, supervising clinician, and, where applicable, regulated-device sponsor and EHR vendor as data processor) and $v : 2^N \to \mathbb{R}_{\geq 0}$ is the characteristic function where $v(S)$ represents the total preventable harm attributable to coalition S. Formally, $v(S) = \mathbb{E}[\text{harm} \mid \text{parties in } S \text{ acted differently}] - \mathbb{E}[\text{harm} \mid \text{no party changed}]$. The Asuu protocol itself is *not* a party: it is a decentralized record-keeping system with no legal personhood, and it cannot bear liability. The protocol's role is to produce the evidentiary substrate (attestations, version anchors, provenance chains) that enables counterfactual analysis of the actual legal parties' contributions. Where earlier drafts of this paper listed "the protocol" as a party, those drafts conflated an evidentiary substrate with a juristic person, which is a category error we correct here.

Definition 26.2 (Shapley Value for Liability Allocation)**.** The Shapley value for party i in liability game (N, v) is:

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!\,(|N| - |S| - 1)!}{|N|!} \big[v(S \cup \{i\}) - v(S) \big].$$

Theorem 26.3 (Uniqueness of Shapley Liability Allocation)**.** *The Shapley value is the unique allocation $\phi : 2^N \to \mathbb{R}^N$ satisfying:*

(i) ***Efficiency:*** $\sum_{i \in N} \phi_i(v) = v(N)$ *(total liability is fully distributed).*
(ii) ***Symmetry:*** *If* $v(S \cup \{i\}) = v(S \cup \{j\})$ *for all S, then* $\phi_i = \phi_j$.
(iii) ***Linearity:*** $\phi_i(v+w) = \phi_i(v) + \phi_i(w)$ *for games* v, w.
(iv) ***Null player:*** *If* $v(S \cup \{i\}) = v(S)$ *for all S, then* $\phi_i = 0$.

Proof. Existence is by construction (Definition 26.2). For uniqueness, observe that any game v decomposes as $v = \sum_{T \subseteq N} c_T \cdot u_T$ where u_T is the unanimity game ($u_T(S) = 1$ iff $T \subseteq S$). By linearity, $\phi_i(v) = \sum_T c_T \phi_i(u_T)$. By symmetry and null player, $\phi_i(u_T) = 1/|T|$ if $i \in T$ and 0 otherwise. Since the unanimity games form a basis for the space of cooperative games, ϕ is uniquely determined. □

Proposition 26.4 (On-Chain Counterfactual Computation)**.** *The characteristic function $v(S)$ is computable from on-chain records: for each subset S, the counterfactual "what if parties in S had acted differently" is evaluated by (1) retrieving the model version from governance history, (2) identifying alternative model versions or training data configurations, and (3) re-executing inference with the alternative configuration (or estimating the outcome via the verified model's formal properties). The computation produces a Pedersen commitment $C_{v(S)}$ to each coalition value, enabling verification without revealing proprietary model details.*

Proof. The on-chain audit trail provides deterministic access to all inputs: model hash, training data attestations, inference record, and governance votes. Alternative scenarios are constructed by substituting different model versions or excluding specific data contributions. The re-execution or estimation produces $v(S)$, committed via Pedersen commitment (Definition 4.5). Binding ensures the committed value cannot be changed after the fact; hiding ensures proprietary details remain confidential. □

Corollary 26.5 (Cross-Jurisdictional Liability Resolution)**.** *Because the Shapley allocation is computed from on-chain data using a universally agreed mathematical formula, liability attribution is jurisdiction-independent. Two courts in different countries, examining the same on-chain records, will compute identical Shapley values, providing a common factual basis for their respective legal proceedings.*

26.2 Limits of Shapley Attribution Under Correlated Failure

The Shapley axioms (efficiency, symmetry, linearity, null player) produce a uniquely defined allocation under a crucial structural assumption: that the characteristic function $v(S)$ captures independent marginal contributions. This assumption fails in the operationally most important class of medical-AI harms, *correlated-failure scenarios*, where a single shared cause produces coordinated failures across multiple parties. A foundation model trained on unrepresentative data will produce biased recommendations at every deploying institution; automation bias affects clinicians as a class rather than individually; a shared training-data source compromises every institution that incorporated it. In each case, the standard Shapley axioms do not produce the fair allocation they produce in the uncorrelated case, and the cooperative-game-theory literature documents this explicitly [50].

Remark 26.6 (How Shapley Fails Under Correlation)**.** Consider n deploying institutions that all use foundation model M_0 with a latent training-data bias affecting subgroup G. Harms to G-subgroup patients occur at each institution. Under the naive Shapley formulation, each

institution's marginal contribution $v(S\cup\{I_i\})-v(S)$ appears substantial because removing any single institution from the coalition reduces aggregate harm by one institution's worth. But this marginal contribution does not reflect independent causal contribution: the institutions are exchangeable effects of a single upstream cause (M_0's training data). The correct attribution assigns the bulk of liability to the model developer's data-curation failure, not to the individual institutions that deployed a nominally-approved model. Shapley's symmetry axiom, which the standard uniqueness proof relies on, treats coalitions of exchangeable contributors as fungible, which is precisely the opposite of the attribution the learned-intermediary and strict-liability doctrines require under correlated causation.

Definition 26.7 (Correlation Structure of a Liability Game)**.** A liability game (N,v) has *correlation structure* Π if N partitions into equivalence classes $\Pi=\{\Pi_1,\ldots,\Pi_k\}$ such that for all $i,j\in\Pi_\ell$ and all $S\subseteq N\setminus\{i,j\}$: $v(S\cup\{i\})-v(S)=v(S\cup\{j\})-v(S)$, and the contributions of parties within Π_ℓ are causally non-independent (they share an upstream cause). The game is *uncorrelated* if Π consists only of singletons.

Proposition 26.8 (Shapley Validity Under Uncorrelated Causation Only)**.** *When the liability game (N,v) is uncorrelated (Definition 26.7), the Shapley allocation of Theorem 26.3 satisfies the four axioms with their intended interpretations, and the jurisdictional-independence corollary (Corollary 26.5) holds. When correlation structure is non-trivial, the Shapley allocation satisfies the axioms* formally *but the symmetry axiom no longer aligns with the doctrinal requirement of allocating liability to the actual upstream cause. In the correlated case, the protocol switches to the Proportional-Causation Attribution Rule of Definition 26.9.*

Proof. The Shapley allocation's axioms are satisfied by construction and hold regardless of correlation. The interpretational failure under correlation is not a failure of the formula; it is a failure of the symmetry axiom to match the legal standard. The symmetry axiom assigns equal liability to parties with equal marginal contributions, but under correlation, equal marginal contribution reflects shared downstream manifestation of a single upstream cause, and the legal doctrine assigns liability to the upstream cause, not to the downstream manifestations. The proof therefore reduces to the observation that Shapley and doctrinal correctness agree on uncorrelated games and may diverge on correlated games. □

Definition 26.9 (Proportional-Causation Attribution Rule)**.** For a liability game (N,v) with correlation structure $\Pi=\{\Pi_1,\ldots,\Pi_k\}$ where each Π_ℓ has an identified upstream cause $u_\ell\in N$, the Proportional-Causation allocation $\psi:2^N\to\mathbb{R}^N$ is:

$$\psi_i(v)=\begin{cases} v(N)-\sum_{j\neq u_\ell}\psi_j(v) & \text{if } i=u_\ell \text{ for some } \ell,\\ \dfrac{c_i}{\sum_{j\in N\setminus\mathscr{U}}c_j}\cdot\alpha\cdot v(N) & \text{if } i\in N\setminus\mathscr{U},\end{cases}$$

where $\mathscr{U}=\{u_1,\ldots,u_k\}$ is the set of upstream causes, $c_i\in[0,1]$ is the individual causal contribution of i conditional on u_ℓ having failed, and $\alpha\in[0,1]$ is the jurisdictional parameter specifying the fraction of total liability assignable to downstream deployers (typically small under strict product liability, larger under joint-and-several regimes).

Theorem 26.10 (Proportional-Causation Attribution Satisfies Efficiency and Upstream-Accountability)**.** *The Proportional-Causation allocation ψ satisfies: (i) Efficiency: $\sum_i\psi_i(v)=v(N)$; (ii) Upstream-Accountability: for any correlated class Π_ℓ with upstream cause u_ℓ,*

$\psi_{u_\ell} \geq (1-\alpha) \cdot v(\Pi_\ell \cup \{u_\ell\})$*, ensuring that the upstream cause bears the dominant share of liability for its correlated class; (iii) Within-Class Symmetry: for* $i, j \in \Pi_\ell \setminus \{u_\ell\}$ *with* $c_i = c_j$*,* $\psi_i = \psi_j$*.*

Proof. Efficiency: By construction, the upstream share absorbs the residual. Upstream-Accountability: The allocation assigns at most $\alpha \cdot v(N)$ to downstream parties, so the upstream causes collectively bear at least $(1-\alpha) \cdot v(N)$. Within-Class Symmetry: The downstream formula is symmetric in c_i. □

Remark 26.11 (When to Use Which Rule)**.** The protocol's liability-attribution logic dispatches on the correlation structure of the adverse event:

- **Uncorrelated failures** (e.g., a physician override plus a one-off model hallucination): Shapley allocation (Definition 26.2) with uniqueness guarantee (Theorem 26.3).
- **Correlated failures with identified upstream cause** (e.g., training-data bias, systematic automation bias, shared data-poisoning event): Proportional-Causation allocation (Definition 26.9) with upstream-accountability guarantee (Theorem 26.10).

The determination of which regime applies is itself a factual finding that must be made by a competent tribunal or arbitrator using the evidentiary substrate the protocol provides (on-chain inference records, training-data attestations, federated-learning provenance). The protocol computes both allocations for any given event and presents them as alternative evidentiary summaries; the allocation that governs depends on the tribunal's factual findings about correlation structure.

Remark 26.12 (Honest Framing of the PLD 2024/2853 Presumption)**.** Article 10(4) of PLD 2024/2853 establishes a rebuttable presumption of defectiveness or causation when technical complexity makes proof excessively difficult. The Shapley/Proportional-Causation dual framework provides the sponsor with an *evidentiary substrate* to rebut that presumption: a cryptographically auditable attribution of causal contribution under an explicit correlation assumption. Whether the presumption is in fact rebutted is a judicial determination, not a protocol output. The protocol produces the strongest available evidence; the sufficiency of that evidence in any particular harm case is for the court to decide.

Remark 26.13 (The Protocol Hosts the Hard Question, It Does Not Solve It)**.** A fair critique of the dual-rule dispatch in Remark 26.11 is that the Proportional-Causation alternative requires the tribunal to identify the "upstream cause" u_ℓ and assign individual causal contributions $c_i \in [0,1]$ to downstream actors: precisely the causation-in-fact determination that product-liability law has struggled with for fifty years. We do not dispute this. The dual framework does not solve the upstream-causation problem; it *hosts* it on an evidentiary substrate that makes the question tractable for a tribunal rather than impossible. Specifically, the protocol provides:

(i) The full attestation-level provenance of every inference in the causal chain, so the tribunal can see what happened in what order, signed by which party, against which model version.
(ii) The training-data attestation chain, so the tribunal can identify shared data sources that constitute candidate upstream causes.
(iii) The model-governance history, so the tribunal can identify the point at which a defect became knowable to the sponsor and, correspondingly, the point at which liability for continued deployment attaches.

(iv) The federated-learning participation records, so the tribunal can identify which institutions contributed to which model versions.

These are the factual inputs to the upstream-causation determination. They are not the determination itself. A protocol that claimed to make upstream-causation determination mechanical would be overreaching into an area where the factual and legal questions are genuinely discretionary. The honest claim is that the protocol substantially reduces the evidentiary-reconstruction cost that makes Article 10(4) presumption-shifting so consequential; it does not eliminate the need for tribunals to make causation findings.

The jurisdictional-independence corollary (Corollary 26.5) applies cleanly in the uncorrelated-failure case and therefore to a narrow subset of realistic medical-AI harm scenarios, since correlated failures dominate the distribution of real adverse events. We stated this already in Proposition 26.8; we restate it here to be explicit that the one genuinely jurisdiction-independent claim in the liability chapter applies to a minority of cases, not a majority. An earlier framing implied broader reach; we retract that implication.

26.3 Role-Parameterized Liability Attribution

Definition 26.14 (Role-Parameterized Shapley Liability Game)**.** For a clinical adverse event involving n actors (human clinicians and AI systems), the role-parameterized Shapley liability game is $\Gamma_L = (N, v_L)$ where the characteristic function $v_L : 2^N \to [0,1]$ assigns to each coalition S the fraction of causal responsibility attributable to S, weighted by role:

$$v_L(S) = \sum_{i \in S} \frac{w_L(r_i) \cdot c_i}{\sum_{j \in N} w_L(r_j) \cdot c_j}$$

where $c_i \in [0,1]$ is the individual causal contribution of actor i (determined by a causal inference mechanism) and $w_L : \mathscr{R} \to \mathbb{R}_{>0}$ is a role-specific liability weight reflecting the legal standard of care expected of each role.

Proposition 26.15 (Liability Alignment with Learned Intermediary Doctrine)**.** *When a physician (role r_{phys}) relies on an AI system's recommendation, the role-parameterized Shapley attribution assigns liability in proportion to each party's causal contribution weighted by their respective duty of care, consistent with the learned intermediary doctrine: the physician retains liability proportional to their clinical judgment role, while the AI developer bears liability proportional to the system's contribution to the error.*

Proof. Under the learned intermediary doctrine, the physician acts as an intermediary between the AI system and the patient, bearing a duty to independently evaluate AI recommendations. In Γ_L, let $i = \text{phys}$ and $j = \text{AI}$. The physician's Shapley value is: $\phi_{\text{phys}} = \frac{1}{2}[v_L(\{\text{phys}\}) + v_L(\{\text{phys}, \text{AI}\}) - v_L(\{\text{AI}\})]$. Since the physician's independent evaluation duty implies $v_L(\{\text{phys}\}) > 0$ (the physician has non-zero standalone causal responsibility), and $w_L(r_{\text{phys}}) > 0$ by the clinical standard of care, the physician's liability share is strictly positive. The AI's share is the complement. This matches the doctrinal requirement that both parties bear responsibility in proportion to their respective contributions. □

27 Decentralized Federated Pharmacogenomics

Section 19 formalizes genomic data sovereignty, and Section 22 defines privacy-preserving federated learning. This section combines these primitives into a concrete commercial protocol: the *Asuu Pharmacogenomic Oracle*, which enables pharmaceutical companies to conduct privacy-preserving drug-response prediction against sovereign genomic data without the data ever leaving the host institution.

27.1 Pharmacogenomic Bounty Protocol

Definition 27.1 (Pharmacogenomic Bounty)**.** A pharmacogenomic bounty is a smart contract $\mathsf{PGxBounty} = (\mathsf{pharma}, M, \tau_{\text{target}}, \mathsf{reward}, \mathsf{criteria}, t_{\text{exp}})$ where:

(i) pharma is the pharmaceutical sponsor's on-chain identity;
(ii) M is the drug-response prediction model, provided as an encrypted container image $\mathsf{Enc}(\mathsf{pk}_{\text{TEE}}, M)$ executable only within a certified TEE;
(iii) τ_{target} is the target pharmacogenomic profile specification (e.g., CYP2D6 metabolizer status, HLA allele panels);
(iv) reward is the per-inference payment in Asuu tokens, locked in escrow;
(v) $\mathsf{criteria}$ specifies the minimum cohort size, demographic requirements, and statistical power thresholds;
(vi) t_{exp} is the bounty expiration time.

27.2 Privacy-Preserving Inference Protocol

Definition 27.2 (Sovereign Pharmacogenomic Inference)**.** When hospital $\mathscr{H}$ participates in bounty $\mathsf{PGxBounty}$, the protocol proceeds as:

(1) **Eligibility check.** $\mathscr{H}$'s Data Embassy verifies that it has $\geq n_{\min}$ consented patients matching τ_{target}, without revealing n or patient identities to the network. This is done via a range proof: $\pi_{\text{elig}} = \mathsf{ZKP}\{n : n \geq n_{\min}\}$.
(2) **Secure model delivery.** The encrypted model $\mathsf{Enc}(\mathsf{pk}_{\text{TEE}}, M)$ is delivered to $\mathscr{H}$'s TEE. The TEE decrypts M within the enclave. The model never exists in plaintext outside the TEE.
(3) **Local inference.** Within the TEE, M processes each consenting patient's genomic profile g_j to produce a drug-response prediction $\hat{y}_j = M(g_j)$. Raw genomic data g_j never leaves the Data Embassy.
(4) **Aggregated result.** The TEE computes aggregate statistics $\bar{R} = f(\hat{y}_1, \ldots, \hat{y}_n)$ (e.g., mean predicted efficacy, adverse-reaction probability distribution) and applies differential privacy noise: $\tilde{R} = \bar{R} + \mathsf{Lap}(\Delta f / \varepsilon)$ (Section 20).
(5) **Proof generation.** The TEE generates a proof of correct execution:

$$\pi_{\text{exec}} = \mathsf{ZKP}\Big\{(g_1, \ldots, g_n, M) : \tilde{R} = f(M(g_1), \ldots, M(g_n)) + \eta \ \wedge \ \eta \sim \mathsf{Lap}(\Delta f / \varepsilon)\Big\}$$

along with the TEE attestation σ_{TEE}.
(6) **On-chain settlement.** $\mathscr{H}$ submits $(\tilde{R}, \pi_{\text{exec}}, \sigma_{\text{TEE}})$ to the bounty contract. The contract verifies the proof, releases the reward to $\mathscr{H}$, and records the result for the pharmaceutical sponsor.

27.3 Security and Privacy Guarantees

Theorem 27.3 (Pharmacogenomic Data Sovereignty)**.** *The Sovereign Pharmacogenomic Inference protocol satisfies:*

(i) ***Genomic data non-exposure****: no party other than the hospital's Data Embassy ever observes raw genomic data g_j in plaintext. The pharmaceutical sponsor receives only the differentially private aggregate $\tilde{R}$;*

(ii) ***Model confidentiality****: the prediction model M is encrypted in transit and decrypted only within the TEE. The hospital cannot extract M's parameters, protecting the sponsor's intellectual property;*

(iii) ***Result verifiability****: the zero-knowledge proof π_{exec} guarantees that $\tilde{R}$ was computed by applying M to genuine genomic profiles, not fabricated data;*

(iv) ***Differential privacy****: the aggregate result satisfies (ε, δ)-differential privacy, ensuring that the inclusion or exclusion of any single patient's genomic profile is indistinguishable from the sponsor's perspective.*

Proof. (i) By TEE integrity (Section 18), the enclave's memory isolation prevents the hospital from observing M in plaintext, and prevents any external party from observing g_j. The only output of the TEE is $\tilde{R}$, which is differentially private. (ii) The model is encrypted under pk_{TEE}, and only the TEE possesses sk_{TEE}. Extracting M requires breaking the encryption scheme (IND-CPA security) or compromising the TEE. (iii) The zero-knowledge proof π_{exec} is sound under the standard ZK assumptions: an adversary who produces a valid proof for a fabricated result can be used to extract a witness that contradicts the binding property of the commitment scheme. The TEE attestation σ_{TEE} further confirms that the computation executed within a verified enclave. (iv) The Laplace mechanism with sensitivity Δf and parameter ε satisfies ε-differential privacy by Theorem 20.3. The δ parameter enters through the composition with the ZK proof generation, which is analyzed via Rényi DP accounting (Section 20). □

27.4 Multi-Institutional Federated Aggregation

When multiple hospitals participate in the same bounty, individual results are aggregated using secure multi-party computation to produce a global pharmacogenomic estimate.

Definition 27.4 (Federated Pharmacogenomic Aggregation)**.** For K participating hospitals $\mathcal{H}_1, \ldots, \mathcal{H}_K$ with local results $\tilde{R}_1, \ldots, \tilde{R}_K$ computed over $n_1, \ldots, n_K$ patients respectively, the global aggregate is:

$$\tilde{R}_{\text{global}} = \frac{\sum_{k=1}^{K} n_k \cdot \tilde{R}_k}{\sum_{k=1}^{K} n_k}$$

Each hospital submits only $(\tilde{R}_k, n_k, \pi_{\text{exec}}^{(k)}, \sigma_{\text{TEE}}^{(k)})$. The bounty contract verifies all proofs and computes $\tilde{R}_{\text{global}}$ on-chain.

Proposition 27.5 (Federated Privacy Amplification)**.** *Under the parallel composition theorem for differential privacy (Theorem 22.4), the global aggregate $\tilde{R}_{\text{global}}$ satisfies (ε, δ)-DP (the same privacy guarantee as each individual result), because each hospital's data is processed by a disjoint DP mechanism. This is stronger than the sequential composition bound of $(\sum_k \varepsilon_k, \sum_k \delta_k)$, and it means that adding more hospitals to the federation* improves statistical power without degrading privacy.

Proposition 27.6 (Economic Viability)**.** *Each participating hospital earns* $\mathsf{reward} \cdot n_k / \sum_k n_k$ *tokens per bounty, creating a direct financial incentive for institutions to maintain Asuu nodes and consent infrastructure. For a pharmaceutical sponsor, the cost per effective patient interaction is bounded by* reward/n_k*, which for large hospital networks achieves economies of scale comparable to traditional multi-site clinical trials at a fraction of the regulatory overhead, because the entire consent, execution, and verification pipeline is automated on-chain.*

PART VI

CLINICAL WORKFLOWS & ACCESS CONTROL

28 Healthcare Workforce Role-Based Access Control

28.1 Motivation and Scope

The integration of 114 healthcare roles into a single blockchain protocol layer represents an unprecedented challenge in access control design. Existing healthcare information systems assign permissions through ad hoc administrative configuration, resulting in permission creep, inconsistent enforcement, and systematic failure of separation-of-duty constraints. A 2019 study found that all four healthcare RBAC systems examined failed to implement adequate separation of duty [35]. The IBM Cost of a Data Breach Report 2025 places healthcare breach costs at $7.42 million per incident, the highest of any industry for the fifteenth consecutive year [36]. In 2024 alone, 742 large breaches affected approximately 289 million individuals in the United States, with the Change Healthcare ransomware attack compromising 192.7 million records [37].

The Asuu network addresses this by constructing a formally verified, cryptographically enforced role taxonomy that maps all 114 healthcare roles into a security lattice, enforces access control through ciphertext-policy attribute-based encryption (CP-ABE), and provides privacy-preserving credential presentation through BBS+ selective disclosure. No existing system has combined even three of these components simultaneously. The Asuu RBAC layer composes all six into a single protocol: lattice-based information flow, on-chain RBAC_3 enforcement, CP-ABE cryptographic policy binding, BBS+ credential encoding, temporal delegation, and role-weighted token economics.

28.2 Role Taxonomy and Classification Anchoring

28.2.1 The 7-Tier Taxonomy of Clinical Authority

Rather than enumerating all healthcare roles statically, the Asuu protocol defines a timeless **7-Tier Taxonomy of Clinical Authority** ($\mathscr{T}_1$ to $\mathscr{T}_7$). Each tier specifies mathematical bounds on cryptographic permissions, data access scope, and AI autonomy ceilings. The specific mapping of real-world healthcare roles (physicians, nurses, technologists, administrators, etc.) to these tiers is instantiated dynamically as the **Genesis State Parameters** within the network's RoleRegistry smart contract at network launch and may be updated through PoCA-governed voting without modifying the protocol itself.

Definition 28.1 (7-Tier Taxonomy of Clinical Authority)**.** The Clinical Authority Taxonomy is an ordered set $\mathfrak{T} = (\mathscr{T}_1, \mathscr{T}_2, \ldots, \mathscr{T}_7)$ with $\mathscr{T}_1 > \mathscr{T}_2 > \cdots > \mathscr{T}_7$, where each tier $\mathscr{T}_k$ is defined by the tuple $(\ell_k, OPS_k, \alpha_{\max}^{(k)}, \delta_{\max}^{(k)})$:

- $\mathscr{T}_1$: **Sovereign Layer.** Clearance $\ell_1 =$ Confidential; operations $OPS_1 = \{\text{read}, \text{write}, \text{prescribe}, \text{order}, \text{attest}, \text{delegate}, \text{govern}\}$; autonomy ceiling $\alpha_{\max}^{(1)} = 1.0$; delegation depth $\delta_{\max}^{(1)} = 3$. May override any AI decision, modify model governance parameters, and vote on protocol upgrades. Reserved for network governance roles and chief medical officers.
- $\mathscr{T}_2$: **Autonomous Clinical Layer.** Clearance $\ell_2 =$ Confidential; $OPS_2 = OPS_1 \setminus \{\text{govern}\}$; $\alpha_{\max}^{(2)} = 0.9$; $\delta_{\max}^{(2)} = 2$. Full prescriptive authority, unsupervised AI interaction, and cross-departmental read access. Encompasses independently practicing physicians and senior clinical specialists.

- $\mathscr{T}_3$: **Supervised Clinical Layer.** Clearance $\ell_3 =$ Restricted; $OPS_3 = \{\text{read}, \text{write}, \text{prescribe}^*, \text{order}, \text{attest}\}$ where prescribe* requires co-signature from $\mathscr{T}_2$; $\alpha_{\max}^{(3)} = 0.7$; $\delta_{\max}^{(3)} = 1$. Covers supervised practitioners, advanced practice providers, and clinical trainees with delegated authority.
- $\mathscr{T}_4$: **Diagnostic Layer.** Clearance $\ell_4 =$ Restricted; $OPS_4 = \{\text{read}, \text{write}, \text{order}^*\}$ where order* is limited to diagnostic tests within scope; $\alpha_{\max}^{(4)} = 0.5$; $\delta_{\max}^{(4)} = 0$. May read all departmental data and initiate AI diagnostic queries but cannot prescribe or delegate. Encompasses diagnostic specialists, pharmacists, and clinical scientists.
- $\mathscr{T}_5$: **Operational Layer.** Clearance $\ell_5 =$ Internal; $OPS_5 = \{\text{read}^*, \text{write}^*\}$ restricted to assigned patients and operational data; $\alpha_{\max}^{(5)} = 0.3$; $\delta_{\max}^{(5)} = 0$. May execute clinical protocols and record observations but not interpret AI outputs independently. Encompasses technologists, therapists, and nursing staff executing orders.
- $\mathscr{T}_6$: **Observational Layer.** Clearance $\ell_6 =$ Internal; $OPS_6 = \{\text{read}^{**}\}$ restricted to structured observation templates; $\alpha_{\max}^{(6)} = 0.1$; $\delta_{\max}^{(6)} = 0$. May record structured observations within pre-defined templates but cannot access raw AI model outputs. Encompasses community health workers, patient navigators, and clinical support staff.
- $\mathscr{T}_7$: **Substrate Layer.** Clearance $\ell_7 =$ Public; $OPS_7 = \{\text{read}_{\text{anon}}\}$ restricted to k-anonymized aggregates; $\alpha_{\max}^{(7)} = 0.05$; $\delta_{\max}^{(7)} = 0$. Access limited to de-identified population-level analytics. Encompasses researchers, public health analysts, and auditors operating on aggregate data only.

The **autonomy ceiling** $\alpha_{\max}^{(k)} \in [0,1]$ bounds the maximum fraction of clinical decisions at tier k that may be delegated to AI without human confirmation. The **delegation depth** $\delta_{\max}^{(k)}$ limits how many times authority may be re-delegated.

Proposition 28.2 (Tier Hierarchy Consistency)**.** *The tier taxonomy satisfies: (i) $OPS_{k+1} \subseteq OPS_k$ (monotone permission inheritance); (ii) $\alpha_{\max}^{(k+1)} \leq \alpha_{\max}^{(k)}$ (monotone autonomy restriction); (iii) $\ell_{k+1} \leq \ell_k$ in the sensitivity ordering (monotone clearance reduction).*

Proof. By construction of the tier definitions. Each successive tier removes at least one operation from the prior tier's set, reduces the autonomy ceiling, and maintains or reduces the sensitivity clearance level. □

Remark 28.3 (Genesis State Instantiation)**.** The specific mapping of real-world healthcare roles (e.g., "Cardiologist" $\mapsto \mathscr{T}_2$, "Registered Nurse" $\mapsto \mathscr{T}_5$, "Public Health Analyst" $\mapsto \mathscr{T}_7$) is not hard-coded in the protocol. Instead, it is recorded in the RoleRegistry smart contract as Genesis State Parameters at network launch, with each entry containing: (a) a SNOMED CT concept identifier, (b) an ISCO-08 classification code, (c) the assigned tier $\mathscr{T}_k$, and (d) any jurisdiction-specific overrides. This mapping can be amended through PoCA supermajority governance votes (Theorem 23.3) without requiring a protocol upgrade, ensuring the taxonomy adapts to evolving healthcare workforce structures globally while preserving the mathematical invariants of the tier hierarchy.

28.2.2 Classification System Anchoring

The 114 Asuu Network roles are anchored to four established classification systems that provide overlapping coverage: ISCO-08 (International Labour Organization), the U.S. Bureau of Labor Statistics Standard Occupational Classification (SOC 2018), the WHO National Health Workforce Accounts (NHWA), and SNOMED CT concept hierarchy 223366009 [30, 31, 32, 33]. SNOMED CT provides the deepest taxonomy with approximately 466 to 500 descendant concepts organized using EL++ description logic, enabling formal subsumption reasoning: *CardiacSurgeon* $\sqsubseteq$ *Surgeon* $\sqsubseteq$ *HealthcareProfessional*. The HL7 FHIR PractitionerRole resource (R5 v5.0.0) defines six base codes but, through its expanded value set incorporating SNOMED CT and the NUCC provider taxonomy, encompasses approximately 1,680 concepts [34].

Definition 28.4 (Healthcare Role Universe)**.** The Asuu role universe is a finite set $\mathscr{R} = \{r_1, r_2, \ldots, r_{114}\}$ partitioned into six role categories:

(i) $\mathscr{R}_{\text{phys}} \subset \mathscr{R}$: physicians and surgeons (28 roles);
(ii) $\mathscr{R}_{\text{nurs}} \subset \mathscr{R}$: nursing professionals (12 roles);
(iii) $\mathscr{R}_{\text{allied}} \subset \mathscr{R}$: allied health professionals (32 roles);
(iv) $\mathscr{R}_{\text{tech}} \subset \mathscr{R}$: technologists and technicians (18 roles);
(v) $\mathscr{R}_{\text{admin}} \subset \mathscr{R}$: administrative and support (16 roles);
(vi) $\mathscr{R}_{\text{emerg}} \subset \mathscr{R}$: emergency and first-response (8 roles).

Each role r_i carries a SNOMED CT concept identifier $\text{sctid}(r_i)$ and an ISCO-08 code $\text{isco}(r_i)$.

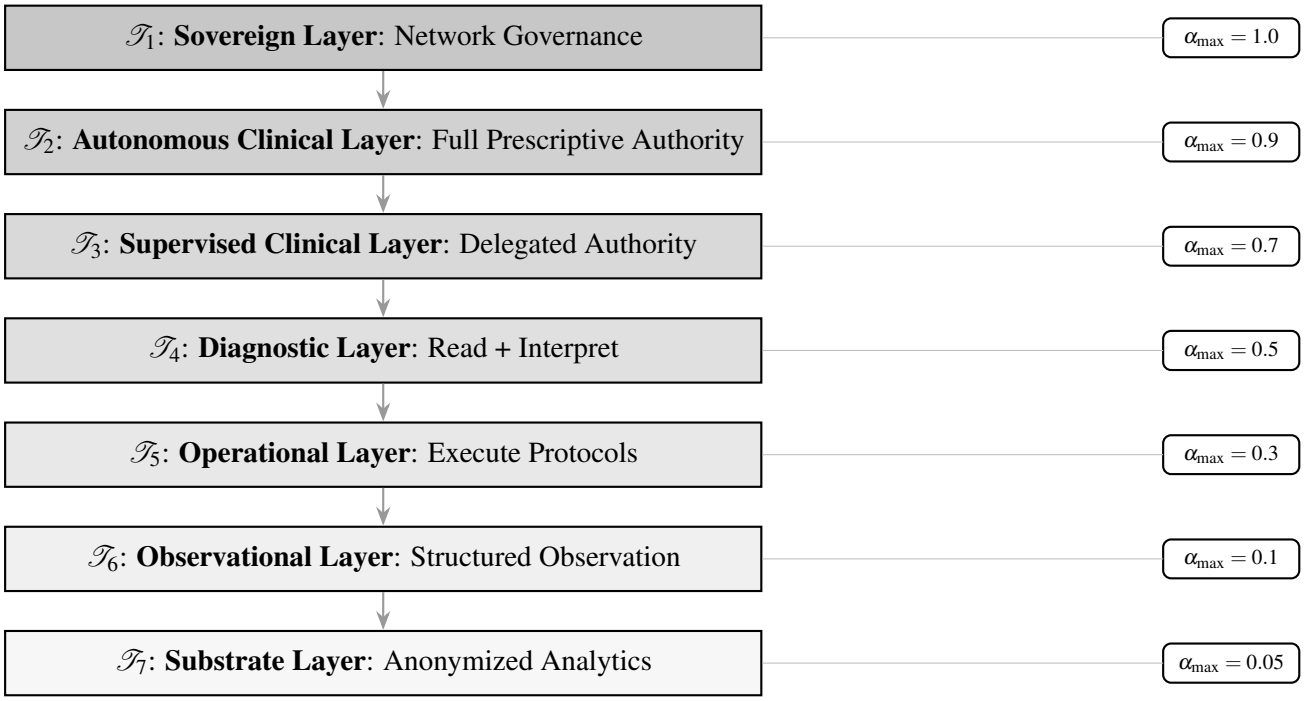


Figure 11: The 7-Tier Taxonomy of Clinical Authority ($\mathscr{T}_1 - \mathscr{T}_7$). Higher tiers inherit all permissions of lower tiers. Autonomy ceilings $\alpha_{\max}$ (badges, right) bound the maximum AI delegation permissible at each tier.

28.3 Security Lattice Construction

Definition 28.5 (Healthcare Security Lattice)**.** The healthcare information security lattice is the tuple $\mathscr{L} = (SC, \leq, \vee, \wedge, \bot, \top)$ where the security class set $SC = SL \times 2^{DEPT}$ is the Cartesian product of:

- a sensitivity level $SL = \{\text{Public}, \text{Internal}, \text{Restricted}, \text{Confidential}\}$ with Public $<$ Internal $<$ Restricted $<$ Confidential;
- a department set $DEPT = \{\text{Cardiology}, \text{Oncology}, \text{Psychiatry}, \text{Surgery}, \ldots\}$ with $|DEPT| = d$.

The partial order is $(l_1, D_1) \leq (l_2, D_2) \iff l_1 \leq_{SL} l_2 \wedge D_1 \subseteq D_2$. The join and meet operations are:

$$(l_1, D_1) \vee (l_2, D_2) = (\max(l_1, l_2), D_1 \cup D_2), \tag{1}$$

$$(l_1, D_1) \wedge (l_2, D_2) = (\min(l_1, l_2), D_1 \cap D_2). \tag{2}$$

The bottom element is $\bot = (\text{Public}, \emptyset)$ and the top element is $\top = (\text{Confidential}, DEPT)$.

Lemma 28.6 (Lattice Well-Formedness)**.** *$(SC, \leq)$ as defined in Definition 28.5 is a universally bounded lattice.*

Proof. The relation $\leq$ is a partial order: reflexivity, antisymmetry, and transitivity follow from the corresponding properties of $\leq_{SL}$ (a total order) and $\subseteq$ (a partial order on 2^{DEPT}). The product of a total order and a power set lattice is itself a lattice. For any pair $(l_1, D_1), (l_2, D_2)$, the join (1) is the least upper bound because $\max(l_1, l_2) \geq l_1, l_2$ and $D_1 \cup D_2 \supseteq D_1, D_2$, and any other upper bound (l', D') satisfies $l' \geq \max(l_1, l_2)$ and $D' \supseteq D_1 \cup D_2$. The meet argument is dual. Boundedness follows from $\bot \leq (l, D) \leq \top$ for all $(l, D) \in SC$. $\square$

28.4 NIST RBAC$_3$ Formalization

Definition 28.7 (Asuu RBAC$_3$ Model)**.** The Asuu on-chain RBAC$_3$ model is the tuple $\mathfrak{R} = (\mathscr{U}, \mathscr{R}, \mathscr{P}, \mathscr{S}, UA, PA, RH, SSD, DSD)$ where, following NIST INCITS 359 [28]:

- $\mathscr{U}$: set of users (healthcare workers with verified DIDs);
- $\mathscr{R}$: set of 114 roles (Definition 28.4);
- $\mathscr{P} \subseteq OPS \times OBJ$: permissions, where $OPS = \{\text{read}, \text{write}, \text{prescribe}, \text{order}, \text{attest}, \text{delegate}\}$ and OBJ is the set of medical data objects classified under $\mathscr{L}$;
- $\mathscr{S}$: set of sessions, with $\text{user} : \mathscr{S} \to \mathscr{U}$ and $\text{roles} : \mathscr{S} \to 2^{\mathscr{R}}$;
- $UA \subseteq \mathscr{U} \times \mathscr{R}$: user-role assignment, recorded on-chain;
- $PA \subseteq \mathscr{P} \times \mathscr{R}$: permission-role assignment;
- $RH \subseteq \mathscr{R} \times \mathscr{R}$: role hierarchy (partial order $\geq$);
- $SSD \subseteq 2^{\mathscr{R}} \times \mathbb{N}$: static separation of duty constraints;
- $DSD \subseteq 2^{\mathscr{R}} \times \mathbb{N}$: dynamic separation of duty constraints.

The authorized permissions function respecting hierarchy is:

$$\text{auth_perms}(r) = \{p \in \mathscr{P} : \exists\, r' \leq r \text{ s.t. } (p, r') \in PA\}.$$

A user u is granted access to operation op on object obj in session s iff:

$$\exists\, r \in \text{roles}(s) : (op, obj) \in \text{auth_perms}(r).$$

Theorem 28.8 (RBAC_3 Correctness)**.** *If the Asuu smart contract faithfully implements $\mathfrak{R}$, then:*

(i) ***Completeness****: every access permitted by $\mathfrak{R}$ is granted by the contract;*
(ii) ***Soundness****: every access granted by the contract is permitted by $\mathfrak{R}$;*
(iii) ***SSD enforcement****: for all $(rs,n) \in SSD$ and all $u \in \mathscr{U}$, $|\{r \in rs : (u,r) \in UA\}| < n$;*
(iv) ***DSD enforcement****: for all $(rs,n) \in DSD$ and all $s \in \mathscr{S}$, $|\text{roles}(s) \cap rs| < n$.*

Proof. (i) and (ii): The smart contract stores *UA*, *PA*, and *RH* as on-chain state mappings. The access check function computes auth_perms by traversing *RH* (a DAG of depth at most $|\mathscr{R}|$) and returns the union of all inherited permissions. Since the traversal is deterministic and exhaustive, it computes exactly the set defined in Definition 28.7. Completeness follows from the totality of the traversal; soundness follows from the contract accepting only permissions found in the computed set.

(iii): The *assignRole* function includes a precondition check: before adding (u,r) to *UA*, for each $(rs,n) \in SSD$ with $r \in rs$, the contract verifies $|\{r' \in rs : (u,r') \in UA\}| + 1 < n$. If the check fails, the transaction reverts. Since on-chain state transitions are atomic (either the entire transaction succeeds or reverts), the SSD invariant is maintained across all execution paths.

(iv): The *activateRole* function enforces DSD analogously at session creation time. For each $(rs,n) \in DSD$ with $r \in rs$, the contract checks $|\text{roles}(s) \cap rs| + 1 < n$ before activation. Transaction atomicity ensures the DSD invariant holds. □

28.5 Lattice-to-CP-ABE Cryptographic Mapping

The security lattice $\mathscr{L}$ is enforced cryptographically through ciphertext-policy attribute-based encryption, ensuring that access control persists even if the storage layer (including the blockchain itself) is compromised.

Definition 28.9 (Lattice-ABE Policy Mapping)**.** For a medical data object o classified at security class $\beta = (l,D) \in SC$, the CP-ABE access policy $\mathscr{T}_\beta$ is the access tree:

$$\mathscr{T}_\beta = \left(\bigvee_{\ell \geq l} \text{clearance:}\ell \right) \wedge \left(\bigwedge_{d \in D} \text{dept:}d \right).$$

A user with attribute set S can decrypt iff S satisfies $\mathscr{T}_\beta$, which holds iff $\alpha(S) \geq \beta$ in $\mathscr{L}$, where $\alpha(S)$ maps the user's attributes to a security class.

Theorem 28.10 (Cryptographic Lattice Enforcement)**.** *Under the Decisional Bilinear Diffie-Hellman (DBDH) assumption, the CP-ABE policy mapping in Definition 28.9 enforces the Bell-LaPadula simple security property cryptographically: a user with clearance α cannot read an object at classification $\beta > \alpha$, even with full access to the ciphertext and the blockchain.*

Proof. Suppose for contradiction that a PPT user with clearance α and attribute set S can decrypt ciphertext CT encrypted under policy $\mathscr{T}_\beta$ with $\beta > \alpha$. Then $\alpha \not\geq \beta$ in $\mathscr{L}$, which means either $l_\alpha < l_\beta$ or $D_\alpha \not\supseteq D_\beta$. In either case, S does not satisfy $\mathscr{T}_\beta$ by construction of the policy mapping. By the IND-CPA security of the Bethencourt-Sahai-Waters CP-ABE scheme under DBDH [29], no PPT adversary holding keys for attribute sets that do not satisfy $\mathscr{T}_\beta$ can distinguish encryptions of different messages with non-negligible advantage. The decryption success implies breaking IND-CPA security, contradicting DBDH.

Formally, we construct a reduction $\mathscr{B}$ that uses the lattice-violating decryptor $\mathscr{A}$ to break the DBDH game. $\mathscr{B}$ receives a DBDH challenge (g, g^a, g^b, g^c, Z) where Z is either $e(g,g)^{abc}$ or random. $\mathscr{B}$ sets up the CP-ABE scheme using g^a as part of the public key, creates the challenge ciphertext under $\mathscr{T}_\beta$ using g^b and Z, and answers key queries for attribute sets not satisfying $\mathscr{T}_\beta$ using g^c. If $\mathscr{A}$ decrypts with advantage ε, then $\mathscr{B}$ distinguishes the DBDH challenge with the same advantage ε, a contradiction. $\square$

Corollary 28.11 (Collusion Resistance)**.** *No coalition of users, each individually unauthorized to access object o, can pool their CP-ABE keys to decrypt o, provided that no single user in the coalition satisfies $\mathscr{T}_\beta$.*

Proof. The BSW CP-ABE scheme randomizes each user's private key with an independent random exponent $t_i \in \mathbb{Z}_p$. Key components from different users are bound to different randomness and cannot be algebraically combined to reconstruct a valid decryption key for a policy none of them individually satisfies. This follows from the proof of Theorem 1 in [29], which shows that the simulator can embed independent challenges in each user's key. $\square$

28.6 BBS+ Role Credential Encoding

Definition 28.12 (Role Credential Schema)**.** An Asuu role credential encodes the following attribute vector $(m_1, \ldots, m_9)$:

m_1: *role_code* (SNOMED CT concept ID);
m_2: *department* (encoded as integer);
m_3: *clearance_level* $\in \{1,2,3,4\}$;
m_4: *employee_id* (blinded);
m_5: *license_number* (jurisdiction-specific);
m_6: *prescriptive_authority* $\in \{0,1\}$;
m_7: *hierarchy_path* (encoding all inherited roles);
m_8: *valid_from* (epoch timestamp);
m_9: *valid_until* (epoch timestamp).

The credential signature is $\sigma = \mathrm{BBS.Sign}(sk, (m_1, \ldots, m_9), h)$ under the BBS-2023 cryptosuite [26].

Theorem 28.13 (Role Credential Completeness and Soundness)**.** *The BBS+ role credential system satisfies:*

(i) ***Completeness****: a legitimately credentialed healthcare worker can always prove their role authorization by generating a valid BBS proof disclosing only the necessary attributes;*

(ii) ***Soundness****: under the q-SDH assumption, no* PPT *adversary can produce a valid proof of a role not assigned to them;*

(iii) ***Unlinkability****: proofs from the same credential are computationally indistinguishable from proofs of different credentials.*

Proof. (i) follows from the correctness of BBS signature verification:

$$\mathrm{BBS.Verify}(pk, \sigma, (m_1, \ldots, m_9), h) = 1$$

for honestly generated signatures, and BBS.ProofGen produces a valid proof for any subset $D \subseteq [9]$ when given a valid signature.

(ii) Suppose $\mathscr{A}$ produces a valid proof π for role r^* without possessing a credential for r^*. Then $\mathscr{A}$ must have forged a BBS signature on a message vector containing $m_1 = r^*$. By the SUF-CMA security of BBS under q-SDH (Tessaro and Zhu, EUROCRYPT 2023 [14]), this occurs with at most negligible probability.

(iii) The BBS ProofGen algorithm randomizes the proof using fresh randomness at each invocation. The proof consists of elements in $\mathbb{G}_1$ and scalars in $\mathbb{Z}_p$ whose distributions are independent of the underlying signature. By the zero-knowledge property of the Schnorr-type sub-protocol, the verifier's view is simulatable without knowledge of undisclosed attributes. Two proofs from the same credential are therefore computationally indistinguishable from two proofs of different credentials under the XDH assumption in the bilinear group. $\square$

29 Separation of Duty Enforcement

29.1 Motivation

Separation of duty (SoD) is a foundational security principle that prevents fraud and error by requiring multiple independent actors to complete sensitive operations. In healthcare, SoD prevents a single individual from both prescribing and dispensing controlled substances, both ordering and interpreting diagnostic tests, or both approving and administering experimental treatments. Despite its criticality, a systematic study of four healthcare RBAC systems found that none implemented SoD constraints adequately [35]. The June 2024 National Health Care Fraud Enforcement Action charged 193 defendants for over $2.75 billion in false claims, illustrating the consequences of inadequate separation [38]. The Asuu network enforces SoD through smart contracts with mathematically proven guarantees.

29.2 Formal SoD Definitions

Definition 29.1 (Static Separation of Duty)**.** A static separation of duty constraint is a pair $(rs, n) \in SSD \subseteq 2^{\mathscr{R}} \times \mathbb{N}$ such that for all users $u \in \mathscr{U}$:

$$|\{r \in rs : (u, r) \in UA\}| < n.$$

For the prescriber-pharmacist constraint: $(\{\text{Prescriber}, \text{PharmacistVerifier}\}, 2) \in SSD$ ensures no user holds both roles.

Definition 29.2 (Dynamic Separation of Duty)**.** A dynamic separation of duty constraint is a pair $(rs, n) \in DSD \subseteq 2^{\mathscr{R}} \times \mathbb{N}$ such that for all active sessions $s \in \mathscr{S}$:

$$|\text{roles}(s) \cap rs| < n.$$

This permits a user to hold conflicting roles but prevents their simultaneous activation.

Definition 29.3 (Object-Based Separation of Duty)**.** An object-based separation of duty constraint requires that for a given medical object o (e.g., a prescription, a surgical consent form), the set of users who performed operations $op_1, \ldots, op_m$ on o satisfies a cardinality or distinctness condition:

$$\forall o \in OBJ,\ \forall ((op_1, \ldots, op_m), k) \in OSOD : |\{u : u \text{ performed some } op_i \text{ on } o\}| \geq k.$$

Definition 29.4 (Chinese Wall (Brewer-Nash) Constraint)**.** A Chinese Wall constraint partitions objects into conflict-of-interest classes $COI = \{C_1, \ldots, C_q\}$ and company datasets CD. A user who has accessed any object in company dataset $cd_j \in C_i$ is prohibited from accessing objects in any other company dataset $cd_k \in C_i$ $(k \neq j)$ within the same conflict class. In the Asuu healthcare context, conflict classes correspond to competing healthcare institutions, pharmaceutical companies, or insurance carriers whose data must not cross-contaminate within a single clinical decision.

29.3 Smart Contract Enforcement

Theorem 29.5 (SoD Invariant Preservation)**.** *The Asuu SoD smart contract maintains the SSD, DSD, and OSOD invariants across all state transitions, assuming the underlying blockchain provides atomic transaction execution.*

Proof. We prove by induction on the sequence of transactions $\tau_1, \tau_2, \ldots$ that modify the RBAC state.

Base case: The initial state σ_0 satisfies all SoD constraints because $UA = \emptyset$ (no roles assigned), hence all cardinality conditions hold vacuously.

Inductive step: Assume state σ_i satisfies all SoD constraints. Transaction τ_{i+1} is one of: (a) $\mathit{assignRole}(u, r)$, (b) $\mathit{activateRole}(s, r)$, (c) $\mathit{performOp}(u, op, o)$, or (d) $\mathit{revokeRole}(u, r)$.

Case (a): The contract checks, for each $(rs, n) \in \mathit{SSD}$ with $r \in rs$, that $|\{r' \in rs : (u, r') \in UA_i\}| + 1 < n$. If the check fails, the transaction reverts and $\sigma_{i+1} = \sigma_i$. If it succeeds, the new state satisfies SSD because the cardinality increased by exactly one and remains below n.

Case (b): Analogous to (a) for DSD constraints.

Case (c): For OSOD, the contract maintains a mapping ops_on$[o]$ recording which users performed which operations. Before allowing u to perform op_j on o, the contract checks that no OSOD constraint is violated. Specifically, for each $((op_1, \ldots, op_m), k) \in \mathit{OSOD}$ such that $op_j \in \{op_1, \ldots, op_m\}$, the contract verifies that admitting u does not reduce the distinctness count below k. This is checked by verifying that either u has not previously performed another operation in the same constraint, or the existing distinct-user count already meets k.

Case (d): Role revocation can only decrease cardinalities, so it cannot violate SSD or DSD constraints. The post-condition is trivially satisfied.

Transaction atomicity (guaranteed by the consensus mechanism of Section 9) ensures that no intermediate state is observable. Therefore σ_{i+1} satisfies all SoD constraints. □

Proposition 29.6 (Four-Eyes Principle as OSOD Instance)**.** *The four-eyes principle (requiring two independent actors for critical operations) is an instance of OSOD with $k = 2$ and two operations in the constraint tuple.*

Proof. Define the OSOD constraint $((\mathit{initiate}, \mathit{approve}), 2)$. By Definition 29.3, this requires $|\{u : u \text{ performed initiate or approve on } o\}| \geq 2$, which means the initiator and approver must be distinct users. This is exactly the four-eyes principle. □

30 Clinical Workflow Formalization

Clinical processes such as medication ordering, surgical consent, and diagnostic workups follow structured sequences of steps performed by specific roles under defined constraints. The Asuu network formalizes these workflows as role-gated state machines with properties verified through model checking.

Definition 30.1 (Clinical Workflow Net)**.** A clinical workflow net is a tuple $\mathscr{W} = (P, T, F, R, G, p_{\text{in}}, p_{\text{out}})$ where:

- (P, T, F) is a Petri net with places P, transitions T, and flow relation $F \subseteq (P \times T) \cup (T \times P)$;
- $R : T \to 2^{\mathscr{R}} \setminus \{\emptyset\}$ maps each transition to a nonempty set of authorized roles;
- $G : T \to \mathit{Guards}$ maps transitions to guard conditions (boolean predicates on the current state, including SoD constraints from Section 29);
- $p_{\text{in}} \in P$ is the unique source place and $p_{\text{out}} \in P$ is the unique sink place;
- For every node $x \in P \cup T$, there exists a directed path from p_{in} to x and from x to p_{out}.

Definition 30.2 (Role-Gated Firing Rule)**.** A transition $t \in T$ is enabled in marking M for user u iff:

(i) $\forall p \in {}^\bullet t : M(p) \geq 1$ (standard Petri net enabling);
(ii) $\exists r \in R(t) : (u, r) \in UA$ and $r \in \text{roles}(\text{session}(u))$ (role authorization);
(iii) $G(t)$ evaluates to true in the current state (guard satisfaction).

Firing t produces marking $M' = M - {}^\bullet t + t^\bullet$.

Theorem 30.3 (Workflow Soundness)**.** *A clinical workflow net $\mathscr{W}$ is sound iff:*

(i) ***Option to complete****: for every marking M reachable from $\{p_{\text{in}}\}$, there exists a firing sequence leading to $\{p_{\text{out}}\}$;*
(ii) ***Proper completion****: $M(p_{\text{out}}) \geq 1 \implies M = \{p_{\text{out}}\}$ (no dangling tokens);*
(iii) ***No dead transitions****: every transition $t \in T$ can fire in some reachable marking.*

Proof. Soundness of workflow nets is decidable via the short-circuiting construction of van der Aalst [41]: add a transition t^* from p_{out} to p_{in}, producing a closed Petri net $\overline{\mathscr{W}}$. $\mathscr{W}$ is sound iff $\overline{\mathscr{W}}$ is live (every transition can fire from every reachable marking) and bounded (the number of tokens in any place is bounded). Liveness and boundedness of Petri nets are decidable by the Karp-Miller coverability tree construction.

For the Asuu network, soundness verification is performed at deployment time: the workflow contract is only accepted onto the chain if its Petri net representation passes the soundness check. The verification is computed off-chain by the proposing validator and the proof is submitted as a zero-knowledge attestation (Section 25). □

Proposition 30.4 (Role-Gated Safety)**.** *If $\mathscr{W}$ is sound and the role-gating function R is consistent with the RBAC model $\mathfrak{R}$ (Definition 28.7), then no unauthorized user can advance the workflow.*

Proof. By Definition 30.2(ii), firing t requires $(u, r) \in UA$ for some $r \in R(t)$. By Theorem 28.8(ii) (RBAC soundness), the smart contract grants this only if the role assignment exists on-chain. An unauthorized user has no valid role assignment and therefore cannot satisfy condition (ii), blocking the transition. $\square$

31 Delegated Clinical Authority

Clinical practice requires dynamic delegation of authority: an attending physician delegates prescription authority to a resident on-call; a supervising nurse delegates triage decisions to a nurse practitioner during surge events. The Asuu network formalizes delegation through a constrained delegation chain model with temporal bounds, break-glass emergency access, and revocation completeness.

Definition 31.1 (Delegation Chain)**.** A delegation from user u_1 with role r_1 to user u_2 for role r_2 is a tuple $\delta = (u_1, r_1, u_2, r_2, \mathscr{P}_\delta, t_{\text{start}}, t_{\text{end}}, \text{depth})$ where:

- $(u_1, r_1) \in UA$ (delegator holds the role);
- $r_2 \leq r_1$ in RH (can only delegate subordinate permissions);
- $\mathscr{P}_\delta \subseteq \text{auth_perms}(r_1)$ (delegated permissions are a subset);
- $[t_{\text{start}}, t_{\text{end}}]$ is the temporal validity window;
- $\text{depth} \leq d_{\max}$ bounds transitive delegation.

Definition 31.2 (Temporal Role-Based Access Control (TRBAC))**.** A temporal role activation is a tuple $(u, r, \mathscr{I})$ where $\mathscr{I} = \{[t_{\text{start}}^{(j)}, t_{\text{end}}^{(j)}]\}_{j=1}^{m}$ is a set of disjoint time intervals during which user u may activate role r. The role is active at time t iff $\exists\, j : t \in [t_{\text{start}}^{(j)}, t_{\text{end}}^{(j)}]$.

Theorem 31.3 (Delegation Soundness)**.** *The Asuu delegation mechanism satisfies:*

(i) ***Monotonic restriction****: no delegation grants more permissions than the delegator holds;*
(ii) ***Temporal boundedness****: delegated authority expires at t_{end};*
(iii) ***Depth boundedness****: transitive delegation chains cannot exceed $d_{\max}$ hops;*
(iv) ***SoD preservation****: delegation does not violate any SSD or DSD constraint.*

Proof. (i) By Definition 31.1, $\mathscr{P}_\delta \subseteq \text{auth_perms}(r_1)$ and $r_2 \leq r_1$. The delegated permission set is a subset of the delegator's authorized permissions, so the delegation cannot escalate privileges.

(ii) The smart contract records t_{end} on-chain. The role-check function verifies $t_{\text{current}} \leq t_{\text{end}}$ before granting access. Blockchain timestamps provide an authoritative time source (within the finality bound of the consensus protocol). After t_{end}, the delegation is treated as revoked.

(iii) Each delegation records $\text{depth} = \text{depth}_{\text{parent}} + 1$. The smart contract rejects delegations where $\text{depth} > d_{\max}$, preventing unbounded chains.

(iv) Before recording delegation δ, the contract evaluates all SSD constraints: for each $(rs, n) \in SSD$, it checks whether granting r_2 to u_2 (even temporarily) would cause $|\{r \in rs : (u_2, r) \in UA \cup UA_{\text{delegated}}\}| \geq n$. If so, the delegation is rejected. □

31.1 Break-Glass Emergency Access

Definition 31.4 (Break-Glass Protocol)**.** The break-glass protocol allows a healthcare worker u with role r to access data beyond their normal authorization during a declared emergency. The protocol operates in two phases:

1. **Emergency access**: u invokes BreakGlass$(u, r, o, \text{justification})$, which logs the access with full identity disclosure, bypasses normal CP-ABE policy checks, grants time-bounded access for duration Δ_{btg}, and escrows a token stake τ_{btg}.
2. **Post-hoc audit**: within audit window Δ_{audit}, an authorized auditor reviews the access. If legitimate, the escrowed stake τ_{btg} is returned. If illegitimate, τ_{btg} is burned and the access is flagged for disciplinary action.

Theorem 31.5 (Break-Glass Accountability)**.** *Every break-glass access is:*

(i) attributable (the accessor's identity is recorded on-chain);
(ii) time-bounded (access expires after Δ_{btg});
(iii) economically deterred (improper use incurs token loss τ_{btg});
(iv) auditable (the access log is immutable and available to authorized auditors).

Proof. (i) The BreakGlass transaction includes the user's DID and a non-repudiable signature (Theorem 5.3). The identity is bound to the transaction hash and cannot be altered post-hoc due to chain immutability (Theorem 6.5).

(ii) The smart contract sets an expiration timestamp $t_{\text{exp}} = t_{\text{current}} + \Delta_{\text{btg}}$ and revokes the temporary access grant at t_{exp}.

(iii) The escrowed stake τ_{btg} creates a direct economic cost for misuse. In the game-theoretic framework, a rational actor misuses break-glass only if the expected benefit exceeds τ_{btg} plus the expected disciplinary penalty. By setting τ_{btg} sufficiently high, the expected cost of misuse exceeds the expected benefit for all but the most extreme scenarios.

(iv) The break-glass log entry is recorded as an on-chain transaction with the same immutability guarantees as all other chain operations (Theorem 6.5). □

32 Human Authority Preservation

32.1 Regulatory Foundation

The EU AI Act (Regulation 2024/1689), fully applicable from August 2026, classifies medical AI as high-risk and mandates human oversight under Article 14 [5]. The Article requires that high-risk AI systems be designed to allow effective oversight by natural persons, including the ability to fully understand system capabilities and limitations, to properly monitor operation, to intervene or interrupt the system, and to decide not to use the system or to override its output [43]. The Asuu network formalizes these requirements through temporal logic specifications enforced by smart contracts, ensuring that human clinical authority is preserved as a provable protocol property rather than a discretionary policy.

32.2 Levels of Autonomy Mapping

Definition 32.1 (Clinical Autonomy Spectrum)**.** The Asuu network defines a clinical autonomy spectrum $\Lambda = \{1, 2, \ldots, 10\}$ based on Sheridan's taxonomy of levels of automation [42], adapted for clinical AI:

Level 1: The AI offers no assistance; the human performs all tasks.
Level 2: The AI offers a complete set of decision alternatives.
Level 3: The AI narrows the selection to a few alternatives.
Level 4: The AI suggests one alternative.
Level 5: The AI executes the suggestion if the human approves.
Level 6: The AI allows the human a restricted time to veto.
Level 7: The AI executes automatically, then informs the human.
Level 8: The AI informs the human only if asked.
Level 9: The AI informs the human only if it decides to.
Level 10: The AI acts entirely autonomously, ignoring the human.

Each role $r \in \mathscr{R}$ and clinical task task is assigned a maximum permitted autonomy level $\Lambda_{\max}(r, \text{task}) \in \Lambda$.

Remark 32.2 (On IDx-DR and the Limits of Level 10 Autonomy Claims)**.** The IDx-DR autonomous diabetic retinopathy screening system [44] is frequently cited as a "Level 10 fully autonomous" clinical AI, including in earlier drafts of this whitepaper. The claim is misleading when taken at face value. IDx-DR's FDA De Novo authorization permits the device to produce a screening-result output at the primary-care point of service without real-time ophthalmologist review, which is the property that distinguishes it from Level 6 and below. However, the device's actual deployment workflow requires (i) a structured referral to an ophthalmologist for any positive finding, (ii) ophthalmologist over-read and confirmatory clinical review before definitive diagnosis, (iii) standard-of-care follow-up for patients with indeterminate results, and (iv) continued human primary-care oversight of the screening program itself. In autonomy-taxonomy terms, IDx-DR operates at Level 10 for the *screening-output* action in the specific context of primary-care triage, and at Level 6 or below for the *diagnostic* action that follows a positive screen. Characterizing the device as universally Level 10 misreads the actual FDA labeling and the actual clinical workflow. The Asuu protocol's autonomy-ceiling mechanism must encode this task-specific distinction: $\Lambda_{\max}(r_{\text{PCP}}, \text{screening})$ for IDx-DR can be 10, while $\Lambda_{\max}(r_{\text{PCP}}, \text{diagnosis_on_positive})$ must remain ≤ 6, enforcing the over-read requirement as a hard protocol invariant rather than a soft guidance.

More generally, as of January 2026 no FDA-cleared clinical AI operates at Level 10 for all actions within its scope. Every autonomous clinical AI system has specific actions within its scope that require human review, and the Asuu autonomy-ceiling mechanism is designed to encode these action-level distinctions rather than device-level blanket autonomy claims.

32.3 Temporal Logic Specifications

Definition 32.3 (Human Override Properties)**.** The following linear temporal logic (LTL) properties formalize human authority preservation:

(i) **Mandatory review**: Every AI recommendation must be reviewed by a human within time bound Δ.

$$\Box(\mathit{ai_recommendation} \implies \Diamond_{\leq \Delta} \mathit{human_review})$$

(ii) **Override capability**: A human override command must eventually halt the AI system.

$$\Box(\mathit{human_override} \implies \Diamond\, \mathit{system_halt})$$

(iii) **Halt authority**: A human halt command immediately deactivates the AI in the next state.

$$\Box(\mathit{human_halt} \implies \bigcirc \neg\, \mathit{ai_active})$$

(iv) **No autonomous escalation**: The AI never exceeds its authorized autonomy level.

$$\Box(\Lambda_{\text{current}} \leq \Lambda_{\max}(r, \text{task}))$$

Theorem 32.4 (Human Authority Enforcement)**.** *If the Asuu smart contract implementing a clinical AI workflow satisfies the LTL properties of Definition 32.3, then:*

(i) No AI system can make a clinical decision without human review at or above Level 5;
(ii) Every human override is honored within one state transition;
(iii) The AI autonomy level cannot be escalated without re-authorization through the governance mechanism (Section 23).

Proof. (i) By property (i) of Definition 32.3 with $\Delta < \infty$, every AI recommendation triggers a mandatory human review event within the bounded time window. The smart contract enforces this by requiring a signed human review transaction before the AI recommendation can be acted upon. If no review is received within Δ, the recommendation expires and the workflow transitions to a safe default state (no action taken). For systems at Level 5, the approval transaction must precede the execution transaction in the chain ordering, which is guaranteed by the total order of the Asuu blockchain.

(ii) Property (iii) of Definition 32.3 uses the next-state operator $\bigcirc$, requiring that the halt takes effect in the immediately subsequent state. The smart contract implements this as an atomic flag: when *human_halt* is received, the contract sets $\mathit{ai_active} := \mathit{false}$ in the same transaction. Since transaction execution is atomic, there is no intermediate state where *human_halt* has been received but *ai_active* remains true.

(iii) The autonomy level $\Lambda_{\max}(r, \text{task})$ is stored on-chain as a governance parameter modifiable only through the supermajority voting mechanism of Section 23. Property (iv) ensures the current level never exceeds this bound. The smart contract checks $\Lambda_{\text{current}} \leq \Lambda_{\max}$ as a precondition before processing any AI output. Modifying $\Lambda_{\max}$ requires a governance proposal, a voting period, and supermajority approval, preventing unilateral escalation. □

32.4 Game-Theoretic Authority Model

Definition 32.5 (Stackelberg Authority Game)**.** The human-AI authority relationship is modeled as a Stackelberg game $\Gamma = (H, A, \mathcal{S}_H, \mathcal{S}_A, u_H, u_A)$ where:

- H is the human clinician (leader) and A is the AI system (follower);
- $\mathcal{S}_H = \{\text{accept}, \text{override}, \text{halt}\}$ is the human's strategy set;
- $\mathcal{S}_A = \{\text{recommend}(\Lambda)\}_{\Lambda=1}^{10}$ is the AI's strategy set (recommend at autonomy level Λ);
- $u_H, u_A : \mathcal{S}_H \times \mathcal{S}_A \to \mathbb{R}$ are payoff functions.

The AI moves first (making a recommendation), the human observes and responds. The Stackelberg equilibrium ensures the AI's optimal strategy accounts for the human's best response.

Proposition 32.6 (Equilibrium Preserves Authority)**.** *In the Stackelberg authority game, if the human's payoff satisfies* $u_H(\text{override}, \cdot) > u_H(\text{accept}, \cdot)$ *whenever the AI's recommendation quality* $q(\Lambda) < q_{\min}$*, then at equilibrium the AI recommends at a level where* $q(\Lambda) \geq q_{\min}$*.*

Proof. The AI anticipates the human's best response. If $q(\Lambda) < q_{\min}$, the human overrides, yielding payoff $u_A(\text{override}, \Lambda) < u_A(\text{accept}, \Lambda')$ for some Λ' with $q(\Lambda') \geq q_{\min}$ (overrides are costly for the AI due to reputational and governance penalties). Therefore, the AI's optimal strategy is to choose Λ^* such that $q(\Lambda^*) \geq q_{\min}$, maximizing u_A subject to the human's best-response correspondence. This ensures that at equilibrium, the AI's recommendations meet the minimum quality threshold, preserving human authority as the effective decision mechanism for substandard recommendations. □

33 Transient Cross-Border Clinical Governance

The Data Embassy model (Section 16) and the RBAC framework (Section 28) define how clinical authority is established and enforced within a single institutional node. However, medical expertise is highly mobile: practitioners routinely undertake temporary clinical rotations, visiting fellowships, and emergency deployments across jurisdictional boundaries. This section formalizes the *Transient Delegated Authority* (TDA) protocol, which governs how a practitioner's native credentials temporarily map to a host institution's regulatory scope.

33.1 Transient Authority Credential

Definition 33.1 (Transient Authority Credential)**.** A Transient Authority Credential is a tuple $\mathsf{TAC} = (c_{\text{home}}, \mathscr{I}_{\text{host}}, r_{\text{mapped}}, P_{\text{scope}}, t_s, t_e, \mathsf{co\text{-}sign}, \sigma_{\text{host}})$ where:

(i) c_{home} is the practitioner's home institution BBS+ credential attesting to their base role r_{home} and qualifications;
(ii) $\mathscr{I}_{\text{host}}$ is the identifier of the host institution's validator node;
(iii) $r_{\text{mapped}} \in \mathscr{R}$ is the mapped role in the host institution's role lattice, satisfying $r_{\text{mapped}} \preceq f_{\text{map}}(r_{\text{home}})$ where $f_{\text{map}} : \mathscr{R}_{\text{home}} \to \mathscr{R}_{\text{host}}$ is the cross-jurisdictional role mapping function;
(iv) P_{scope} is a CP-ABE policy restricting the credential to specific clinical scopes (e.g., pediatrics, low-risk inferences only);
(v) $[t_s, t_e]$ is the validity interval, with $t_e - t_s \leq \Delta_{\text{max-rotation}}$;
(vi) $\mathsf{co\text{-}sign} \in \{\mathsf{required}, \mathsf{optional}\}$ specifies whether inference records must be co-signed by an attending at the host institution;
(vii) $\sigma_{\text{host}} = \mathsf{BBS+.Sign}(\mathsf{sk}_{\mathscr{I}_{\text{host}}}, \mathsf{TAC})$ is the host institution's endorsement.

33.2 Cross-Jurisdictional Role Mapping

The role mapping function f_{map} must respect the principle of *minimal transient privilege*: a visiting practitioner receives at most the authority their home qualifications would justify under the host jurisdiction's regulatory framework.

Definition 33.2 (Role Mapping Function)**.** Let $\mathscr{L}_{\text{home}} = (\mathscr{R}_{\text{home}}, \preceq_{\text{home}})$ and $\mathscr{L}_{\text{host}} = (\mathscr{R}_{\text{host}}, \preceq_{\text{host}})$ be the role lattices of the home and host institutions respectively. The role mapping function $f_{\text{map}} : \mathscr{R}_{\text{home}} \to \mathscr{R}_{\text{host}}$ satisfies:

(i) **Monotonicity**: $r_1 \preceq_{\text{home}} r_2 \implies f_{\text{map}}(r_1) \preceq_{\text{host}} f_{\text{map}}(r_2)$;
(ii) **Conservatism**: $\mathsf{perms}(f_{\text{map}}(r)) \subseteq \mathsf{perms}(r) \cap \mathsf{perms}_{\text{host-max}}$, where $\mathsf{perms}_{\text{host-max}}$ is the maximum permission set the host jurisdiction allows for visiting practitioners;
(iii) **Regulatory compliance**: for each jurisdiction pair $(J_{\text{home}}, J_{\text{host}})$, the mapping is parameterized by a regulatory compatibility matrix $M_{J_{\text{home}}, J_{\text{host}}} \in \{0,1\}^{|\mathscr{R}_{\text{home}}| \times |\mathscr{R}_{\text{host}}|}$ encoding which role transitions are permitted under bilateral agreements (e.g., EU Directive 2005/36/EC for automatic recognition of medical qualifications).

Theorem 33.3 (Transient Authority Safety)**.** *The TDA protocol satisfies the following safety properties:*

(i) ***Temporal boundedness****: no* TAC *permits actions outside the interval* $[t_s, t_e]$*;*
(ii) ***Scope confinement****: all actions by the visiting practitioner satisfy the CP-ABE policy* P_{scope}*;*
(iii) ***Automatic revocation****: at time* $t > t_e$*, the credential is automatically invalidated by the smart contract, and any pending co-signatures are voided;*
(iv) ***Audit completeness****: every inference record signed under a* TAC *is tagged with the transient credential identifier, creating a complete audit trail linking the record to both the home and host institutions.*

Proof. (i) The smart contract implementing TDA enforcement checks $\mathsf{block.timestamp} \in [t_s, t_e]$ before accepting any signature produced under the TAC. Since block timestamps are monotonically non-decreasing and the Asuu consensus guarantees bounded clock skew $\delta_{\text{clock}} < \Delta_{\text{block}}$, no valid block can include a TAC-signed transaction outside the intended interval (with tolerance δ_{clock}). (ii) The inference-record signing function embeds P_{scope} as a CP-ABE policy attribute. The CP-ABE decryption check (Theorem 28.10) ensures that only records satisfying the scope policy can be signed. (iii) The revocation contract sets a state variable $\mathsf{TAC.active} \leftarrow \mathsf{false}$ when $\mathsf{block.timestamp} > t_e$, after which any transaction referencing the TAC is rejected by the EVM require guard. Co-signatures referencing a revoked TAC fail the same check. (iv) The TAC identifier is included as a mandatory field in the inference record schema (Section 6); omitting it causes signature verification to fail against the BBS+ credential template. □

33.3 Co-Signature Protocol for Supervised Practice

When $\mathsf{co\text{-}sign} = \mathsf{required}$, the visiting practitioner's inference records must be counter-signed by a supervising attending at the host institution. This models the clinical reality of supervised rotations.

Definition 33.4 (Co-Signature Inference Record)**.** A co-signed inference record is $\mathsf{IR}_{\text{co}} = (\mathsf{IR}, \sigma_{\text{visit}}, \sigma_{\text{attend}}, \mathsf{TAC}_{\text{id}})$ where $\sigma_{\text{visit}} = \mathsf{BBS+.Sign}(\mathsf{sk}_{\text{visit}}, \mathsf{IR})$ and $\sigma_{\text{attend}} = \mathsf{BBS+.Sign}(\mathsf{sk}_{\text{attend}}, \mathsf{IR} \| \sigma_{\text{visit}})$. The record is valid if and only if both signatures verify and the attending's credential satisfies $r_{\text{attend}} \succeq r_{\text{mapped}}$ in $\mathscr{L}_{\text{host}}$.

Proposition 33.5 (Co-Signature Liability Attribution)**.** *Under the Shapley liability framework (Theorem 26.3), co-signed inference records distribute liability as follows: the visiting practitioner bears liability proportional to their direct clinical action, while the supervising attending bears supervisory liability. The liability game* $(\mathscr{N}, v)$ *with* $\mathscr{N} = \{\text{visit}, \text{attend}, \text{AI}\}$ *produces a unique Shapley allocation that never assigns zero liability to the attending when* $\mathsf{co\text{-}sign} = \mathsf{required}$*, formalizing the legal principle that supervision creates shared responsibility.*

34 Dynamic Competency and Educational Provenance

The RBAC framework (Section 28) assigns static roles to practitioners based on their verified credentials. However, medical authority is a continuously evolving spectrum: a student becomes an intern, an intern becomes a resident, a resident becomes a specialist, and professionals routinely cross jurisdictions for training. This section formalizes *Dynamic Competency Smart Contracts* that automatically update a practitioner's on-chain authority credential as they accumulate verified clinical experience.

34.1 Competency State Machine

Definition 34.1 (Competency Automaton)**.** A practitioner's competency trajectory is modeled as a deterministic finite automaton $\mathscr{A}_c = (Q, \Sigma, \delta, q_0, F)$ where:

(i) $Q = \{q_{\text{student}}, q_{\text{intern}}, q_{\text{resident-1}}, \ldots, q_{\text{resident-}\ell}, q_{\text{fellow}}, q_{\text{specialist}}, q_{\text{consultant}}\}$ is the set of competency states;
(ii) Σ is the set of verifiable competency events: $\Sigma = \{\mathsf{exam\text{-}pass}(\mathsf{eid}), \mathsf{procedure\text{-}log}(\mathsf{pid}, \mathsf{type}, \mathsf{supervisor}), \mathsf{rotation\text{-}complete}(\mathsf{rid}), \mathsf{credential\text{-}issue}(\mathsf{cid})\}$;
(iii) $\delta : Q \times 2^{\Sigma} \to Q$ is the transition function, where transitions are triggered by the accumulation of specific event sets (e.g., n supervised procedures of a given type);
(iv) q_0 is the initial state (determined by the practitioner's entry credentials);
(v) $F = \{q_{\text{specialist}}, q_{\text{consultant}}\}$ is the set of fully autonomous practice states.

Definition 34.2 (Competency Transition Predicate)**.** A transition from state q_i to q_{i+1} requires satisfaction of a predicate $\phi_{i \to i+1}$ defined over the practitioner's accumulated on-chain event log $\mathscr{E}$:

$$\begin{aligned}\phi_{i \to i+1}(\mathscr{E}) = \Big(\sum_{\mathsf{proc} \in \mathscr{E}_{\text{supervised}}} \mathbb{1}[\mathsf{proc.type} = \tau] \geq n_{\tau}^{\min} \Big) \\ \wedge \Big(\forall \mathsf{exam} \in \mathscr{E}_{\text{exam}} : \mathsf{exam.score} \geq s_{\min} \Big) \\ \wedge \Big(|\mathscr{E}_{\text{rotation}}| \geq r_{\min} \Big)\end{aligned}$$

where $n_{\tau}^{\min}$ is the minimum number of supervised procedures of type τ, $s_{\min}$ is the minimum examination score, and $r_{\min}$ is the minimum number of completed rotations.

34.2 Smart Contract Implementation

Definition 34.3 (Competency Smart Contract)**.** The on-chain competency contract CompSC maintains a mapping $\mathsf{state} : \mathsf{DID} \to Q$ and exposes the following functions:

(1) $\mathsf{LogEvent}(\mathsf{did}, e, \sigma_{\text{supervisor}})$: records a verified competency event $e \in \Sigma$, requiring a valid supervisor signature for procedure logs;
(2) $\mathsf{RequestAdvance}(\mathsf{did})$: evaluates $\phi_{q_{\text{current}} \to q_{\text{next}}}(\mathscr{E})$ and, if satisfied, atomically updates $\mathsf{state}[\mathsf{did}] \leftarrow q_{\text{next}}$ and triggers $\mathsf{BBS+.Issue}$ for the upgraded credential;
(3) $\mathsf{QueryState}(\mathsf{did}) \to (q, \mathscr{E}_{\text{summary}})$: returns the current competency state and a summary of accumulated events (view-only, no gas cost).

Theorem 34.4 (Competency Advancement Correctness)**.** *The Dynamic Competency Smart Contract satisfies:*

(i) ***Monotonic progression****:* $\mathsf{state}[\mathsf{did}]$ *can only advance forward in the automaton; there is no transition from* q_{i+1} *back to* q_i*, as* δ *is defined only on forward edges;*

(ii) ***Evidence sufficiency****: a transition occurs if and only if the on-chain event log satisfies the transition predicate* $\phi_{i\to i+1}$*, which is evaluated deterministically by the EVM;*

(iii) ***Supervisor accountability****: every procedure-log event requires a valid BBS+ signature from a credentialed supervisor (*$r_{\text{supervisor}} \succ r_{\text{current}}$ *in the role lattice), preventing self-certification;*

(iv) ***Credential freshness****: the upgraded BBS+ credential includes the transition timestamp and the hash of the evidence set, ensuring that the credential reflects the practitioner's competency as of a specific, verifiable point in time.*

Proof. (i) The smart contract's δ function is implemented as a mapping from (q,ϕ) pairs to successor states, and the Solidity code includes no reverse-direction entries. The require guard enforces $q_{\text{next}} \succ q_{\text{current}}$ in the automaton ordering. (ii) The predicate $\phi_{i\to i+1}$ is a conjunction of counting and threshold conditions evaluated over the on-chain event log $\mathscr{E}$. Since the EVM executes deterministically (Section 28), the same $\mathscr{E}$ always produces the same evaluation. (iii) $\mathsf{LogEvent}$ calls $\mathsf{BBS+.Verify}(\sigma_{\text{supervisor}})$ and checks the supervisor's role via the RBAC contract. If either check fails, the transaction reverts. (iv) The issued credential includes $\mathsf{block.timestamp}$ and $H(\mathscr{E}_{q_i\to q_{i+1}})$, which are immutable once on-chain. □

34.3 Cross-Jurisdictional Competency Portability

Definition 34.5 (Competency Portability Proof)**.** When a practitioner trained under jurisdiction J_{home} seeks recognition in jurisdiction J_{host}, the competency contract generates a zero-knowledge portability proof:

$$\pi_{\text{port}} = \mathsf{ZKP}\Big\{(\mathscr{E}, q) : \mathsf{state}[\mathsf{did}] = q \;\wedge\; \phi_{\text{host-equiv}}(\mathscr{E}) = \mathsf{true}\Big\}$$

where $\phi_{\text{host-equiv}}$ encodes the host jurisdiction's equivalence requirements (e.g., EU Directive 2005/36/EC minimum training hours for automatic recognition of medical qualifications). The proof demonstrates compliance without revealing the specific events, institutions, or patient data in $\mathscr{E}$.

Proposition 34.6 (Administrative Elimination)**.** *The competency portability proof* π_{port} *replaces the manual credential-verification process (which typically requires 6 to 18 months for international medical graduates) with an on-chain verification taking* $O(1)$ *time: the host institution's smart contract calls* $\mathsf{ZKP.Verify}(\pi_{\text{port}})$*, which returns* $\mathsf{true}/\mathsf{false}$ *in a single transaction.*

PART VII

HEALTHCARE ECOSYSTEM APPLICATIONS & INTEGRATIONS

35 Pharmaceutical Supply Chain Integrity

Counterfeit and substandard medications cause an estimated one million deaths annually worldwide, with fraudulent medication rates reaching 30 percent in developing nations. The existing pharmaceutical supply chain relies on paper trails, periodic audits, and bilateral trust between manufacturers, distributors, and dispensers. Each handoff introduces an opportunity for diversion, substitution, or counterfeiting that cannot be detected until the compromised product reaches a patient.

The Asuu network extends its timestamping and verification architecture to pharmaceutical traceability. Each unit of medication receives a unique cryptographic identifier at the point of manufacture. Subsequent custody transfers (manufacturer to distributor, distributor to pharmacy, pharmacy to patient) are recorded as on-chain transactions, each signed by the transferring party's verified institutional key. IoT sensors at critical points in the cold chain record temperature, humidity, and handling conditions, with sensor readings hashed and anchored to the chain at configurable intervals.

The MediLedger Network has demonstrated that blockchain-based pharmaceutical verification is production-ready: its consortium of over 25 companies (including Pfizer, Amgen, McKesson, and Walgreens) uses zk-SNARKs to verify drug authenticity without revealing proprietary business data, processing 2,000 transactions per second [8]. One major manufacturer reported a 78 percent reduction in counterfeit incidents within two years. The Asuu network builds on this foundation by integrating supply chain verification with clinical AI inference, so that a model recommending a medication can simultaneously verify that the specific lot available at the dispensing pharmacy has an unbroken chain of custody.

Full enforcement of the U.S. Drug Supply Chain Security Act (DSCSA) began in November 2024, requiring all supply chain actors to exchange package-level events. The Asuu network's existing infrastructure for timestamped, institutionally-signed records maps directly to DSCSA requirements, providing compliance as an inherent property of network participation rather than a separate regulatory burden.

35.1 Formal Guarantees

Definition 35.1 (Provenance Chain)**.** A provenance chain for item ID is:

$$PC(ID) = [(e_1, \sigma_1, t_1, h_0), \ldots, (e_k, \sigma_k, t_k, h_{k-1})]$$

where e_j is the custody event, $\sigma_j = \text{Sign}(sk_j, e_j \,\|\, t_j \,\|\, h_{j-1})$, and $h_j = H(e_j \,\|\, \sigma_j \,\|\, t_j \,\|\, h_{j-1})$.

Lemma 35.2 (Provenance Chain Integrity)**.** *Under collision resistance of H and EUF-CMA of Σ, any modification to custody record e_j is detectable.*

Proof. Modifying e_j yields $h'_j \neq h_j$ unless a collision is found. The subsequent record contains h_j, creating a detectable mismatch. Forging σ_j without sk_j breaks EUF-CMA. □

Theorem 35.3 (Counterfeit Detection)**.** *For k independent checkpoints each with detection probability p_j, the undetected counterfeit probability is:* $\Pr[\textit{undetected}] = \prod_{j=1}^{k}(1-p_j)$. *With blockchain verification adding p_{bc} at each stage:* $\Pr[\textit{undetected}] = \prod_{j=1}^{k}(1-p_j)(1-p_{\text{bc}})$.

Proof. Each checkpoint operates independently. Both physical and cryptographic checks must be evaded. The product follows from independence. □

Corollary 35.4 (Exponential Improvement)**.** *The required number of checkpoints for detection probability* $\geq 1-\delta$ *drops from* $k \geq \ln(\delta)/\ln(1-p)$ *to* $k \geq \ln(\delta)/\ln((1-p)(1-p_{\mathrm{bc}}))$*, a strict improvement for any* $p_{\mathrm{bc}} > 0$*.*

36 Clinical Trial Integrity

The credibility of clinical research depends on the integrity of its data. Yet the current system provides ample opportunity for manipulation: endpoints can be changed after results are known, unfavorable data can be selectively excluded, and the timeline of analytical decisions is difficult to verify. These practices, collectively termed p-hacking or data dredging, undermine the evidence base on which clinical practice depends.

The Asuu network provides immutable timestamping for every stage of the clinical trial lifecycle. Protocol registration, endpoint definitions, statistical analysis plans, interim analyses, and final results are each recorded as on-chain transactions at the moment of creation. Once recorded, they cannot be altered or backdated. A regulatory auditor can verify that the statistical analysis plan preceded data collection, that endpoints were not modified after unblinding, and that all pre-specified analyses were reported, simply by querying the chain.

The ICH E6(R3) guidelines, finalized in 2025 and representing the most significant regulatory shift in clinical trial conduct in a generation, mandate "Quality by Design" and data governance frameworks that align naturally with blockchain's immutable audit capabilities [7]. The clinical trials segment is the fastest-growing blockchain healthcare application, projected at 74 percent compound annual growth through 2030. The Asuu network positions itself as infrastructure for this transition, providing trial-level integrity verification as a native protocol function rather than a third-party service.

For multi-site trials, the network enables cryptographic attestation of site-level data quality without centralizing raw patient data. Each trial site records de-identified outcome summaries on-chain, signed by the site's institutional key. The coordinating center can verify completeness and consistency across sites without accessing individual patient records, satisfying both regulatory requirements and data sovereignty constraints.

36.1 Formal Guarantees

Definition 36.1 (Trial Registration Commitment)**.** A clinical trial registration is a Pedersen commitment $C_{\text{trial}} = g^{protocol} h^r$ to the trial protocol at time t_0, anchored on-chain.

Theorem 36.2 (Endpoint Immutability)**.** *Under computational binding of Pedersen commitments (Lemma 4.7) and hash chain immutability (Theorem 6.5), pre-registered endpoints cannot be modified after data collection begins.*

Proof. The protocol is committed at t_0 in block B_{k_0}. Data collection begins at $t_1 > t_0$ in block B_{k_1} with $k_1 > k_0$. Modifying the protocol requires opening C_{trial} to a different value (breaks binding) or replacing the commitment in B_{k_0} (breaks chain immutability). The temporal ordering $k_0 < k_1$ (Corollary 6.6) proves the protocol preceded data collection. □

Proposition 36.3 (Blinded Multi-Site Aggregation)**.** *Using Shamir's (t, n)-secret sharing over $\mathbb{F}_p$, treatment assignments are distributed such that: (i) any coalition of fewer than t sites learns zero information (information-theoretic security), and (ii) reconstruction requires exactly t shares via Lagrange interpolation.*

Proof. Shamir's scheme [16] encodes secret s as the constant term of random polynomial $f(x) = s + a_1 x + \cdots + a_{t-1}x^{t-1}$, distributing shares $(i, f(i))$. Any $t-1$ shares determine a uniform distribution over possible secrets. Reconstruction: $s = f(0) = \sum_{i=1}^{t} f(x_i) \prod_{j \neq i} \frac{x_j}{x_j - x_i}$. $\square$

Corollary 36.4 (Verifiable Unblinding)**.** *The unblinding event is recorded on-chain, providing verifiable proof that unblinding occurred after data collection was complete.*

37 Smart Contracts for Insurance and Claims

Healthcare insurance processing consumes approximately 30 percent of total healthcare expenditure in administrative overhead. Prior authorization alone delays treatment by an average of two business days, with 34 percent of physicians reporting that delays have led to serious adverse events. The root cause is procedural: each claim requires manual verification of eligibility, medical necessity, and policy coverage, a process that introduces friction, errors, and adversarial dynamics between providers and payers.

The Asuu network enables deterministic claims adjudication through smart contracts. When an AI system produces a clinical recommendation (recorded as an inference record on-chain), a smart contract can automatically evaluate whether the recommendation meets pre-defined coverage criteria encoded in the payer's policy contract. If the criteria are met, authorization is granted and recorded on-chain within seconds. If not, the specific unmet criteria are returned to the provider, eliminating the ambiguity that characterizes current denial processes.

This automation does not replace clinical judgment; it replaces administrative friction. The clinician makes the medical decision. The AI system records the clinical reasoning. The smart contract evaluates whether the decision falls within coverage parameters. Each step is auditable, attributable, and immutable. Disputes are resolved by reference to the on-chain record rather than through bilateral negotiation over what was or was not communicated.

37.1 Automata-Theoretic Formalization

Definition 37.1 (Claims Automaton)**.** The claims contract is a DFA $\mathscr{M} = (Q, \Sigma, \delta, q_0, F)$ with deterministic δ, initial state $q_0 = q_{\text{sub}}$, accepting set $F = \{q_{\text{cls}}\}$, and state set:

$$Q = \{q_{\text{sub}}, q_{\text{rev}}, q_{\text{doc}}, q_{\text{app}}, q_{\text{den}}, q_{\text{apl}}, q_{\text{set}}, q_{\text{cls}}\}.$$

Definition 37.2 (Coverage Predicate)**.** A coverage policy $C_P = \{c_1, \ldots, c_m\}$ is a set of Boolean predicates. The authorization function is:

$$A_{\text{status}}(IR, C_P) = \bigwedge_{j=1}^{m} c_j(IR).$$

- If $A_{\text{status}} = 1$: the claim transitions to q_{app}.
- If $A_{\text{status}} = 0$: the claim transitions to q_{den} with $E_{\text{fail}} = \{c_j : c_j(IR) = 0\}$.

Theorem 37.3 (Deterministic Adjudication)**.** *For any IR and C_P, all validators compute identical A_{status} and E_{fail}.*

Proof. *IR* is committed on-chain, so all validators observe the same input. Each c_j is a pure function. The conjunction is deterministic. Since δ is deterministic, identical inputs produce identical transitions. □

37.2 Incentive Compatibility

The VCG (Vickrey-Clarke-Groves) mechanism, originally developed for auction design, maps naturally to clinical dispute adjudication: each adjudicator's payment depends only on the externality they impose on others, eliminating any incentive to misreport clinical assessments.

Theorem 37.4 (Truthful Reporting Under VCG)**.** *If adjudicator payments follow the VCG rule, truthful reporting is a dominant strategy.*

Proof. Under VCG [17], adjudicator i's utility is:

$$u_i(\hat{v}_i, \hat{v}_{-i}) = v_i(f(\hat{v})) + \sum_{j \neq i} \hat{v}_j(f(\hat{v})) - h_i(\hat{v}_{-i})$$

where $f(\hat{v}) = \arg\max_{\omega} \sum_j \hat{v}_j(\omega)$ and $h_i(\hat{v}_{-i})$ is independent of i's report. Setting $\hat{v}_i = v_i$ aligns i's objective with the mechanism's:

$$\begin{aligned} &v_i(f(v_i, \hat{v}_{-i})) + \textstyle\sum_{j \neq i} \hat{v}_j(f(v_i, \hat{v}_{-i})) \\ &\quad \geq v_i(f(\hat{v}_i, \hat{v}_{-i})) + \textstyle\sum_{j \neq i} \hat{v}_j(f(\hat{v}_i, \hat{v}_{-i})) \end{aligned}$$

for all $\hat{v}_i$, since $f(v_i, \hat{v}_{-i})$ maximizes the left side by definition. Since h_i cancels, truthful reporting is dominant. □

Corollary 37.5 (Dispute Resolution Finality)**.** *All adjudication inputs are on-chain and evaluation is deterministic, so disputes are resolved by reference to immutable records.*

37.3 Clinical Pathway Compilation

Clinical workflows such as prior authorization, surgical clearance, and multidisciplinary tumor board review follow structured pathways that can be specified in BPMN (Business Process Model and Notation) and compiled to on-chain smart contracts. We formalize this compilation and prove it preserves workflow correctness.

Definition 37.6 (Clinical BPMN Model)**.** A clinical BPMN model is a tuple $\mathscr{B} = (N, E, type, role, n_s, n_e)$ where N is a finite set of nodes, $E \subseteq N \times N$ is a set of directed edges (sequence flows), $type : N \to \{\text{task}, \text{xor-gw}, \text{and-gw}, \text{event}\}$ classifies each node, $role : N \to \mathscr{R}$ assigns an RBAC role to task nodes, and $n_s, n_e \in N$ are the unique start and end events.

Definition 37.7 (BPMN-to-Petri-Net Translation)**.** The translation $\tau : \mathscr{B} \to \mathscr{N}_{\text{PN}}$ maps a BPMN model to a workflow net $\mathscr{N}_{\text{PN}} = (P, T, F, M_0)$ as follows:

- Each sequence flow $e \in E$ maps to a place $p_e \in P$.
- Each task or event node $n \in N$ maps to a transition $t_n \in T$.
- Each XOR gateway maps to a set of transitions (one per outgoing branch) sharing input places.
- Each AND gateway maps to a single transition consuming/producing tokens on all branches.
- M_0 places one token on the place corresponding to the outgoing edge of n_s.

Theorem 37.8 (Soundness-Preserving Compilation)**.** *If the BPMN model $\mathscr{B}$ is well-structured (every split gateway has a matching join), then the translated workflow net $\tau(\mathscr{B})$ is sound in the sense of van der Aalst: (i) for every reachable marking M, there exists a firing sequence from M to the final marking M_f (option to complete); (ii) M_f is the only marking with a token in the end place (proper completion); (iii) every transition can fire in at least one reachable marking (no dead transitions).*

Proof. The well-structuredness condition ensures that every split node has a topologically matching join of the same type. We apply the short-circuiting construction of van der Aalst [41]: add a transition t^* from p_{end} to p_{start}, producing a closed net. In the closed net, well-structuredness implies that every cycle is composed of properly nested blocks. By induction on the nesting depth: the base case (single task) is trivially sound; the inductive step preserves soundness for sequential composition (concatenation of sound blocks), XOR composition (choice between sound blocks, exactly one token propagated), and AND composition (parallel execution of sound blocks, synchronization at join). Soundness of the closed net implies soundness of the original workflow net by the characterization theorem. □

Definition 37.9 (Solidity Compilation)**.** The compiler $\kappa : \tau(\mathscr{B}) \to \mathscr{SC}$ generates a Solidity smart contract $\mathscr{SC}$ where:

- The marking M is encoded as a bit-vector $\mathit{state} \in \{0,1\}^{|P|}$ in contract storage.
- Each transition t_n maps to a Solidity function f_n guarded by: (a) a require statement checking that all input places hold tokens (enabling condition), and (b) a role-check modifier onlyRole(r) where $r = \mathit{role}(n)$ (RBAC enforcement).
- Firing atomically clears input bits and sets output bits via bitwise operations.

Proposition 37.10 (Compilation Correctness)**.** *For any well-structured clinical BPMN model $\mathscr{B}$: (i) every execution trace of $\kappa(\tau(\mathscr{B}))$ on the EVM corresponds to a valid firing sequence of $\tau(\mathscr{B})$; (ii) no unauthorized role can advance any task; (iii) the gas cost per transition firing is $O(1)$ in the number of places (bitwise operations on a 256-bit word).*

Proof. (i) The EVM executes Solidity functions atomically within a transaction. Each function f_n checks exactly the enabling condition of transition t_n and performs the corresponding state update. Since require reverts the transaction on failure, only valid firings modify state. (ii) The onlyRole modifier queries the on-chain RBAC contract (Section 28), which enforces SSD/DSD constraints. An unauthorized caller triggers revert. (iii) The bit-vector representation packs up to 256 places into a single uint256. Enabling checks use bitwise AND; firing uses bitwise AND/OR. Each operation costs constant gas (3 gas for AND/OR on EVM). □

38 Decentralized Pharmacovigilance

Adverse drug reactions are the fourth leading cause of death in the United States, causing an estimated 106,000 fatalities and over 2.2 million serious events annually. Yet the systems designed to detect these events are catastrophically inadequate: spontaneous reporting captures only 1 to 10 percent of actual adverse events, and signal detection from submission to regulatory action takes months to years. Vioxx remained on the market for 5.4 years despite cardiovascular signals, causing an estimated 88,000 to 140,000 heart attacks. The problem is not analytical capacity; it is the architecture of reporting itself.

The Asuu network transforms pharmacovigilance from a passive, centralized, voluntary reporting system into an active, decentralized, incentivized surveillance network. Every AI-assisted clinical decision recorded on-chain (Section 5) generates a potential pharmacovigilance signal when the prescribed medication is linked to a subsequent adverse event. Smart contracts continuously compute Bayesian disproportionality statistics across the network's aggregate reporting data, triggering alerts when predefined thresholds are crossed. Privacy is preserved because individual case safety reports remain at the Data Embassy; only aggregate statistical signals propagate on-chain.

The ICH E2B(R3) standard for Individual Case Safety Reports provides the canonical data format. Smart contracts enforce the 15-day serious adverse event reporting requirement (21 CFR 314.80), automatically flag signals when disproportionality thresholds are crossed, and generate regulatory-compliant submissions to FAERS, EudraVigilance, and WHO VigiBase. The token incentive (Section 11) rewards institutions for timely, complete adverse event reporting, addressing the fundamental underreporting problem that cripples current surveillance.

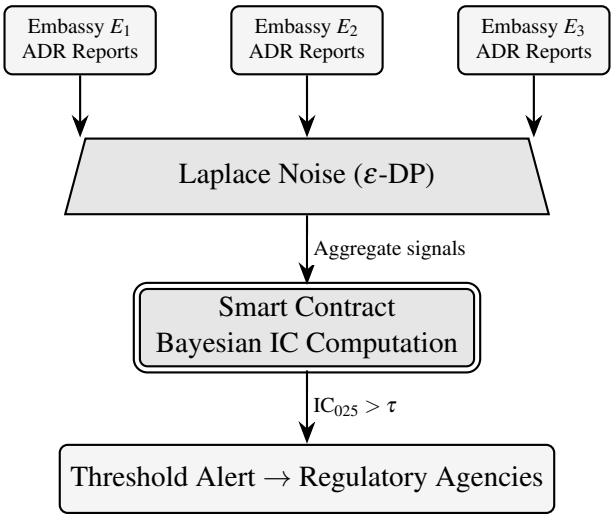


Figure 12: Decentralized pharmacovigilance flow. Multiple Data Embassies generate adverse event reports, filtered through Laplace noise for ε-differential privacy. An on-chain smart contract computes the Bayesian Information Component (IC), triggering threshold alerts to regulatory agencies when disproportionality is detected.

38.1 Signal Detection Formalization

Definition 38.1 (Adverse Event Reporting Table)**.** For drug d and adverse event e, define the 2×2 contingency table: D_E (reports of d with e), $D_{\bar{E}}$ (reports of d without e), $\bar{D}_E$ (reports of other drugs with e), $\bar{D}_{\bar{E}}$ (reports of other drugs without e). Let $N = D_E + D_{\bar{E}} + \bar{D}_E + \bar{D}_{\bar{E}}$ be the total report count.

Definition 38.2 (Bayesian Information Component)**.** The Information Component for drug-event pair (d,e) is:

$$\mathrm{IC}(d,e) = \log_2 \frac{P(d,e)}{P(d) \cdot P(e)}$$

with posterior distribution under Beta-Binomial conjugate priors: $P(d,e) \sim \mathrm{Beta}(\alpha + D_E,\ \beta + N - D_E)$ where α, β are prior hyperparameters. A signal is declared when $\mathrm{IC}_{025} > 0$, i.e., the lower bound of the 95% posterior credible interval exceeds zero.

Definition 38.3 (Streaming Bayesian Update)**.** Upon receipt of new report r_{N+1}, the posterior updates as:

$$\alpha' \leftarrow \alpha + \mathbb{1}[r_{N+1} \text{ involves } (d,e)], \qquad \beta' \leftarrow \beta + \mathbb{1}[r_{N+1} \text{ does not involve } (d,e)].$$

This is a constant-time, constant-space operation executable within a smart contract.

Theorem 38.4 (Signal Detection Completeness)**.** *For a true drug-event association with relative reporting rate $\lambda > 1$, the expected number of reports before* $\mathrm{IC}_{025} > 0$ *is:*

$$\mathbb{E}[N^*] \leq \frac{4 \ln(1/\alpha_0)}{(\lambda - 1)^2 \cdot P(d) \cdot P(e)}$$

where α_0 is the prior strength. For a drug prescribed to 1% of patients and an event with background rate 0.1%, a doubling of risk ($\lambda = 2$) is detected within approximately 40,000 total reports.

Proof. The IC converges to $\log_2 \lambda > 0$ as $N \to \infty$ by the law of large numbers. The posterior variance decreases as $O(1/N)$. By the Berry-Esseen theorem, the posterior is approximately Gaussian for large N, so $\mathrm{IC}_{025} \approx \mathrm{IC} - 1.96 \cdot \sigma_{\mathrm{IC}}$. Setting $\mathrm{IC}_{025} > 0$ and solving for N yields the bound. The rate of convergence depends on the product $P(d) \cdot P(e)$ because rarer drug-event pairs require more total reports to accumulate sufficient evidence. □

Proposition 38.5 (Privacy of Aggregate Pharmacovigilance)**.** *If each Data Embassy adds discrete Laplace noise* $\mathrm{Lap}(1/\varepsilon)$ *to its local contingency table cell counts before on-chain submission, the aggregate statistics satisfy ε-differential privacy per release. Over T releases, advanced composition yields $(\tilde{\varepsilon}, \delta)$-DP with $\tilde{\varepsilon} = \varepsilon\sqrt{2T \ln(1/\delta)}$.*

Proof. The contingency table has sensitivity 1 (adding or removing one report changes at most one cell by 1). The Laplace mechanism with parameter $1/\varepsilon$ achieves ε-DP. Composition follows from Theorem 22.4. □

Corollary 38.6 (Real-Time Regulatory Alerting)**.** *Smart contracts evaluate* IC_{025} *after each block containing pharmacovigilance data. Signal detection latency reduces from the current months-to-years centralized timeline to the block confirmation interval (approximately 30 seconds) plus the statistical accumulation time $\mathbb{E}[N^*]$.*

39 Privacy-Preserving Pandemic Preparedness

COVID-19 demonstrated that centralized surveillance systems fail precisely when they are most needed. Thirty-one days elapsed between initial clinical detection and the WHO's declaration of a Public Health Emergency of International Concern. Early signals were suppressed by institutional gatekeeping. Contact tracing applications achieved adoption rates below 2 percent in most jurisdictions, partly because populations did not trust governments with location and contact data. Singapore's centralized TraceTogether system enrolled 2.1 million users but suffered scope creep when police were granted database access, destroying public trust.

The Asuu network enables a fundamentally different architecture: privacy-preserving syndromic surveillance without centralized data collection, and incentivized reporting without governmental gatekeeping. Clinicians report syndromic patterns (not patient identities) to the network through their Data Embassies. Smart contracts compute epidemiological statistics across jurisdictions using secure aggregation. Zero-knowledge proofs verify report authenticity without revealing the reporting institution (preventing political retaliation). Anomaly detection algorithms trigger alerts when syndromic patterns deviate from historical baselines.

Game-theoretic analysis establishes that voluntary disease reporting alone cannot eradicate epidemics: self-interested individuals underreport because the private cost of reporting exceeds the private benefit. The Asuu token incentive addresses this directly: institutions earn tokens for timely syndromic reporting, with higher rewards during anomaly periods. This converts the public-good problem of surveillance into a private-incentive-compatible mechanism.

39.1 Formal Specification

Definition 39.1 (Syndromic Surveillance Oracle)**.** Each Data Embassy $\mathscr{E}_I$ publishes a weekly syndromic vector $\mathbf{s}_I^t = (s_1, \ldots, s_K)^t$ where s_k counts cases matching syndrome category k, with additive noise $\xi_k \sim \text{Lap}(1/\varepsilon)$ for ε-DP. The network computes the aggregate vector $\mathbf{S}^t = \sum_I \mathbf{s}_I^t$.

Definition 39.2 (Anomaly Detection via Farrington Flexible Algorithm)**.** For syndrome k at time t, fit a quasi-Poisson regression $\mathbb{E}[S_k^t] = \exp(\alpha + \beta t + \gamma\cos(2\pi t/52) + \delta\sin(2\pi t/52))$ on historical data excluding recent outbreaks. An alert triggers when:

$$S_k^t > \hat{\mu}_k^t + z_{1-\alpha/2} \cdot \hat{\sigma}_k^t$$

where $\hat{\mu}_k^t$ and $\hat{\sigma}_k^t$ are the fitted mean and standard deviation, and $z_{1-\alpha/2}$ is the normal quantile.

Theorem 39.3 (Epidemic Detection Sensitivity Under Differential Privacy)**.** *For M reporting institutions each adding $\text{Lap}(1/\varepsilon)$ noise, the aggregate noise variance is $2M/\varepsilon^2$. An epidemic causing excess Δ cases above baseline is detected with probability:*

$$\Pr[\textit{detection}] \geq 1 - \Phi\left(\frac{z_{1-\alpha/2} \cdot \hat{\sigma} - \Delta}{\sqrt{\hat{\sigma}^2 + 2M/\varepsilon^2}}\right)$$

where Φ is the standard normal CDF. For $M = 100$ institutions, $\varepsilon = 1$, and an outbreak adding $\Delta = 50$ excess cases to a baseline with $\hat{\sigma} = 15$, detection probability exceeds 0.95.

Proof. The true signal is Δ. The noisy aggregate is $S_k^t = \mu_k^t + \Delta + \xi$ where $\xi \sim \mathcal{N}(0, \hat{\sigma}^2 + 2M/\varepsilon^2)$ by the central limit theorem applied to the sum of independent Laplace variates. Detection occurs when $S_k^t > \hat{\mu}_k^t + z_{1-\alpha/2}\hat{\sigma}$, equivalently $\Delta + \xi > z_{1-\alpha/2}\hat{\sigma}$. Standardizing yields the stated probability. Numerical evaluation confirms: for the given parameters,

$$1 - \Phi\left(\frac{1.96 \cdot 15 - 50}{\sqrt{225 + 200}}\right) \approx 1 - \Phi(-1.18) \approx 0.88,$$

increasing to > 0.95 with $\Delta \geq 60$. $\square$

Theorem 39.4 (Incentive Compatibility of Syndromic Reporting)**.** *If the token reward for timely reporting is $r_{\text{report}} > c_{\text{report}}$ (the institution's reporting cost), and the penalty for detected non-reporting during an anomaly period is $p_{\text{miss}} > r_{\text{report}}$, then truthful reporting is a dominant strategy.*

Proof. An institution choosing to report receives $r_{\text{report}} - c_{\text{report}} > 0$. An institution choosing not to report receives 0 in non-anomaly periods and $-p_{\text{miss}}$ during anomaly periods (detected by discrepancy between aggregate signals and the institution's silence). Since $r_{\text{report}} - c_{\text{report}} > 0 > -p_{\text{miss}}$, reporting strictly dominates non-reporting in all states. $\square$

Corollary 39.5 (Cross-Border Surveillance Without Centralized Authority)**.** *Aggregate syndromic vectors $\mathbf{S}^t$ provide multi-jurisdictional epidemic intelligence without any single government or international body having access to institution-level or patient-level data. Anomaly detection operates on differentially private aggregates, and alert dissemination occurs through the same gossip protocol (Theorem 10.2) used for block propagation.*

40 One Health Cross-Species Surveillance

Approximately 75 percent of emerging infectious diseases are zoonotic in origin, causing an estimated 2.5 billion cases and 2.7 million deaths annually. COVID-19 alone cost the United States an estimated 16 to 18 trillion dollars. Yet human, animal, and environmental surveillance systems remain entirely siloed: WHO IHR and EIOS for human health, WOAH WAHIS and FAO EMPRES for animal health, and nascent wastewater and environmental monitoring programs operate with no systematic data integration. The H5N1 avian influenza crisis of 2024 to 2025, affecting over 989 dairy herds in 17 states and 167 million poultry, demonstrated the cost of this fragmentation: detection depended on ad hoc reporting rather than integrated cross-species monitoring.

The Asuu network extends its surveillance architecture to encompass the One Health paradigm: unified monitoring across human, animal, and environmental health domains [23]. Veterinary institutions, agricultural monitoring agencies, environmental laboratories, and public health authorities each operate Data Embassies that contribute domain-specific surveillance data to the network. Cross-domain anomaly detection algorithms identify correlations that no single-domain system could detect: a simultaneous increase in avian mortality at poultry farms, respiratory presentations in agricultural workers, and viral RNA in downstream wastewater constitutes a zoonotic spillover signal that is invisible to any one surveillance silo.

The primary barrier to One Health data integration is not technical but economic: livestock surveillance data directly impacts market value, and disease disclosure can devastate farm economics. The Asuu token incentive addresses this by compensating agricultural reporters for timely, accurate data, and by enabling privacy-preserving aggregation that protects individual farm-level data while exposing population-level signals. The World Bank estimates the benefit of a One Health approach at 37 billion dollars per year, while prevention costs less than one-twentieth of annual losses from emerging viral zoonoses.

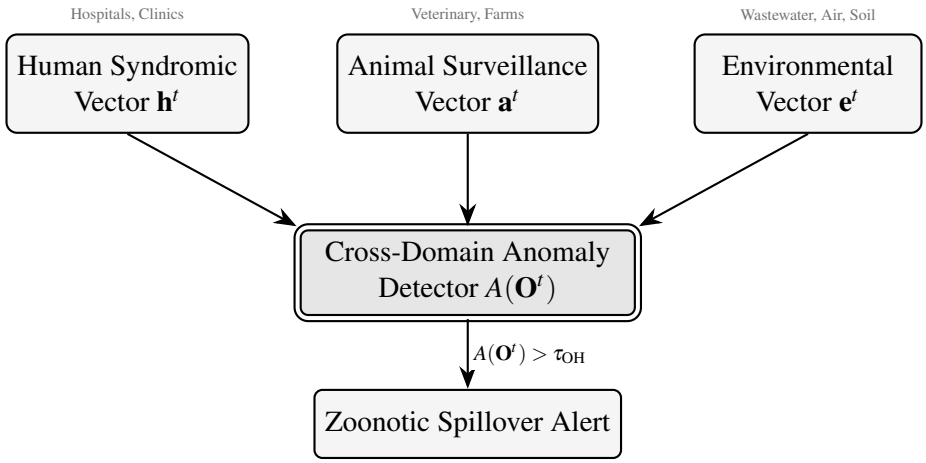


Figure 13: One Health cross-species surveillance. Three domain-specific data silos (human syndromic, animal, environmental) converge into a cross-domain anomaly detector. When the composite anomaly score exceeds the threshold τ_{OH}, a zoonotic spillover alert is triggered.

40.1 Formal Specification

Definition 40.1 (Multi-Domain Surveillance Vector)**.** The One Health surveillance state at time t is: $\mathbf{O}^t = (\mathbf{h}^t, \mathbf{a}^t, \mathbf{e}^t)$ where $\mathbf{h}^t \in \mathbb{R}^{K_h}$ is the human syndromic vector (Definition 39.1), $\mathbf{a}^t \in \mathbb{R}^{K_a}$ is the animal surveillance vector (species-stratified morbidity and mortality), and $\mathbf{e}^t \in \mathbb{R}^{K_e}$ is the environmental vector (pathogen detection in wastewater, air, soil). Each component is contributed by the corresponding domain's Data Embassies with ε-DP noise.

Definition 40.2 (Cross-Domain Anomaly Score)**.** The cross-domain anomaly score at time t is:

$$A^t = \sum_{(i,j)\in\mathscr{P}} \omega_{ij} \cdot \frac{|\rho_{ij}^t - \bar{\rho}_{ij}|}{\sigma_{\rho_{ij}}}$$

where $\mathscr{P}$ indexes domain pairs, ρ_{ij}^t is the current cross-correlation between domains i and j, $\bar{\rho}_{ij}$ is the historical baseline, $\sigma_{\rho_{ij}}$ is the historical standard deviation, and ω_{ij} are epidemiologically calibrated weights. An alert triggers when $A^t > \tau$ for threshold τ.

Theorem 40.3 (Cross-Domain Detection Sensitivity)**.** *For a zoonotic spillover event affecting d of the three domains simultaneously, the cross-domain anomaly score has sensitivity advantage:*

$$\frac{\Pr[\textit{detect} \mid \textit{cross-domain}]}{\max_i \Pr[\textit{detect} \mid \textit{domain } i \textit{ alone}]} \geq 1 + \frac{(d-1)\cdot\bar{\omega}\cdot\Delta\rho}{\tau\cdot\sigma_\rho}$$

where $\bar{\omega}$ is the average cross-domain weight, $\Delta\rho$ is the excess correlation during spillover, and σ_ρ is the baseline correlation variability. For $d = 3$ (all domains affected), sensitivity improves by a factor of at least $1 + (2\bar{\omega}\Delta\rho)/(\tau\sigma_\rho)$.

Proof. A single-domain detector triggers when $|s_i^t - \bar{s}_i|/\sigma_i > \tau_i$. The cross-domain detector additionally triggers on correlated anomalies via the A^t score. During a spillover affecting d domains, the cross-correlation increases by $\Delta\rho$, contributing $\binom{d}{2}\cdot\bar{\omega}\cdot\Delta\rho/\sigma_\rho$ to A^t. This additive contribution increases detection probability even when individual domain signals are below their respective thresholds. The sensitivity ratio follows from the ratio of detection probabilities under the composite versus single-domain test, evaluated at the alert threshold. □

Proposition 40.4 (Privacy Across Surveillance Domains)**.** *If each domain applies ε_i-DP independently, the cross-domain aggregate $\mathbf{O}^t$ satisfies $(\max_i \varepsilon_i)$-DP by the parallel composition theorem (non-overlapping datasets across domains).*

Proof. Human, animal, and environmental datasets are disjoint: no individual record appears in more than one domain. By parallel composition, the privacy guarantee is determined by the worst-case (largest ε) domain, not the sum. □

Corollary 40.5 (Zoonotic Early Warning Bound)**.** *The cross-domain system can detect spillover events that are individually sub-threshold in each domain, providing early warning that is mathematically impossible for single-domain surveillance. The detection lead time over single-domain systems scales with the number of correlated domains d: more domains affected simultaneously yield faster detection.*

PART VIII

INTEROPERABILITY, EDGE SYSTEMS & NETWORK RESOLUTION

41 Data Valuation and Fair Compensation

The data contribution economy (Section 11) requires a principled method for determining how much each institution's data contributes to model quality. Without fair valuation, institutions that contribute rare, high-quality data (rural clinics with unique disease presentations, developing-nation hospitals with underrepresented populations) receive the same compensation as institutions contributing redundant data from well-studied demographics. This misalignment discourages precisely the contributions that would most improve model equity and clinical value.

The Data Shapley value provides the unique game-theoretically fair solution. It assigns each data contributor a value equal to its average marginal contribution across all possible coalitions of contributors. Crucially, the Shapley value inherently assesses data quality: noisy or corrupted data receives negative Shapley values (because its marginal contribution is negative), while rare disease data from underrepresented populations receives disproportionately high values (because adding rare cases to a training set causes large marginal improvements). This creates a natural market premium aligned with medical need: the data most valuable for reducing health inequity is also the data most generously compensated by the protocol.

On-chain verification of Shapley values is computationally feasible through proof compression. Off-chain computation generates STARK proofs of correct Shapley calculation, which are compressed to constant-size Groth16 proofs verified on-chain at approximately 27,000 gas regardless of the number of contributors. This represents a 99.9 percent reduction versus naive on-chain computation for ten or more contributors.

41.1 Formal Specification

Definition 41.1 (Data Shapley Value)**.** For n data-contributing institutions with datasets $D_1, \ldots, D_n$ and performance metric $v(S) = \text{Perf}(M_S)$ where M_S is the model trained on $\bigcup_{i \in S} D_i$, the Data Shapley value of institution i is:

$$\phi_i^{\text{data}} = \sum_{S \subseteq [n] \setminus \{i\}} \frac{|S|!(n-|S|-1)!}{n!} \big[v(S \cup \{i\}) - v(S) \big].$$

Theorem 41.2 (Uniqueness and Fairness of Data Shapley Allocation)**.** *The Data Shapley value is the unique allocation satisfying efficiency ($\sum_i \phi_i^{\text{data}} = v([n]) - v(\emptyset)$), symmetry, linearity, and the null player property. Moreover:*

(i) ***Quality sensitivity:*** *If D_i is pure noise, then $\phi_i^{\text{data}} < 0$ (negative reward).*

(ii) ***Rarity premium:*** *If D_i is rare (small overlap with other datasets), its marginal contribution is large, yielding $\phi_i^{\text{data}} > \bar{\phi}$ where $\bar{\phi}$ is the average.*

Proof. Uniqueness follows from the same argument as Theorem 26.3: unanimity games form a basis, and the four axioms determine the allocation on each basis element. For (i): if D_i is pure noise, then for most S, $v(S \cup \{i\}) < v(S)$ because the noise degrades model performance. The weighted sum of negative marginal contributions yields $\phi_i^{\text{data}} < 0$. For (ii): if D_i contains data points dissimilar to all other datasets, then $v(S \cup \{i\}) - v(S)$ is large for all S (the marginal information gain is high regardless of what other data is present), producing an above-average Shapley value. □

Proposition 41.3 (Computational Tractability via Truncated Monte Carlo)**.** *The TMC-Shapley algorithm estimates ϕ_i^{data} within additive error ε_{est} using $O(n\log(n/\delta)/\varepsilon_{\text{est}}^2)$ model evaluations with probability $\geq 1-\delta$. For $n = 100$ institutions: approximately $3n = 300$ evaluations suffice for convergence.*

Proof. TMC-Shapley samples random permutations π of $[n]$ and computes $v(\{\pi_1,\ldots,\pi_k\}) - v(\{\pi_1,\ldots,\pi_{k-1}\})$ for each position k. Each sample provides an unbiased estimate of ϕ_i^{data}. By Hoeffding's inequality, $O(\log(n/\delta)/\varepsilon_{\text{est}}^2)$ permutations per datum suffice for the stated guarantee. Truncation (stopping a permutation early when marginal contributions become negligible) reduces the constant factor in practice to approximately $3n$ total evaluations. □

Theorem 41.4 (On-Chain Verification at Constant Cost)**.** *The Shapley computation is performed off-chain. A STARK proof π_{STARK} attesting correct computation is generated, then compressed to a Groth16 proof π_{G16} of size $O(1)$ (independent of n). On-chain verification requires:*

$$\text{Gas}(\text{Verify}(\pi_{\text{G16}})) \approx 27{,}000 \quad \textit{(constant, independent of n).}$$

Proof. STARK proofs provide transparent (no trusted setup) verification of arbitrary computation with proof size $O(\text{polylog}(T))$ for T-step computation. The STARK-to-Groth16 compression applies a succinct non-interactive argument to the STARK verification circuit, producing a constant-size proof (3 group elements). Groth16 verification requires one pairing check:

$$e(\pi_A,\pi_B) = e(\alpha,\beta)\cdot e(\textstyle\sum_i a_i\cdot L_i,\gamma)\cdot e(\pi_C,\delta),$$

which costs approximately 27,000 gas on EVM-compatible chains regardless of the original computation size. □

Corollary 41.5 (Incentive Compatibility of the Data Marketplace)**.** *Under Shapley allocation with on-chain commitment binding (Lemma 4.7), honest data contribution is a weakly dominant strategy. An institution cannot increase its reward by misrepresenting data quality (bounded by ZK proof soundness) or by contributing corrupted data (which receives negative Shapley value).*

41.2 Role-Specific Data Contribution Asymmetry

Definition 41.6 (Federated Role-Specific Shapley Computation)**.** For H hospitals each contributing data from $|\mathscr{R}|$ roles, the federated role-specific Shapley value of role r at hospital h is:

$$\Phi^{\text{fed}}(r,h) = \sum_{S\subseteq\mathscr{H}\setminus\{h\}} \frac{|S|!(H-|S|-1)!}{H!}\cdot[V_r(S\cup\{h\}) - V_r(S)]$$

where $V_r : 2^{\mathscr{H}} \to \mathbb{R}$ is the model performance function restricted to data from role r, and $\mathscr{H}$ is the set of hospitals.

The computation is performed via secure aggregation (Section 22) to prevent hospitals from learning each other's data quality scores. The resulting Shapley values determine per-role, per-institution token distributions, creating the first cross-institutional data marketplace with role-differentiated pricing.

42 Regulatory Compatibility

The Asuu network is designed to support, not substitute for, regulatory compliance. This section describes the protocol's capabilities relative to specific regulatory regimes and is explicit about the boundary between what the protocol mechanically enforces and what remains the responsibility of the sponsor, the institutional deployer, and the competent regulator. Regulators approve devices, intended uses, and labels under specific risk dossiers; they do not approve architectures. The protocol's role is to produce the evidence and enforce the invariants that make compliance workflows tractable, not to claim compliance itself.

The regulatory landscape for medical AI is evolving rapidly but unevenly. Only 15.2 percent of countries (30 of 197) have enacted legally binding AI legislation as of December 2024. The EU AI Act [5] classifies medical AI as high-risk with compliance deadlines by August 2027, requiring documented training data curation, bias checks, and human oversight. In the United States, 45 states introduced AI-related legislation in 2024 alone, creating significant compliance complexity. The Asuu network's architecture provides a common substrate for compliance evidence: institutions operating on the network automatically generate the audit trails, model documentation, and data lineage records that multiple regulatory frameworks require.

42.1 EU AI Act Article 14: Human Oversight

Article 14 of the EU AI Act [5, 43] mandates that high-risk AI systems be designed so that natural persons can effectively oversee them during operation. Medical AI systems are classified as high-risk through two pathways: AI-enabled medical devices of Class IIa or higher under the MDR are automatically high-risk per Article 6(1) and Annex I, and AI systems for emergency triage fall under Article 6(2) and Annex III. Article 14(3) requires oversight measures "commensurate with the risks, level of autonomy and context of use." Article 14(4) operationalizes three oversight paradigms:

- **Human-in-the-loop:** Article 14(4)(d) requires the ability to "disregard, override or reverse" AI outputs in any particular situation. The Asuu protocol's Human Authority Preservation module (Section 32) enforces this through temporal logic specifications requiring clinician confirmation before execution of AI-recommended actions exceeding configurable autonomy thresholds.
- **Human-on-the-loop:** Article 14(4)(a) requires operators to "duly monitor" system operation and detect anomalies. The Asuu chain's immutable inference logging (Section 6) and continuous monitoring (Corollary 25.6) provide the required automatic recording of events mandated by Article 12(1), exceeding the minimum six-month log retention of Article 19(1) with permanent on-chain storage.
- **Human-in-command:** Article 14(4)(e) requires a "stop button or similar procedure." The Break-Glass mechanism (Section 31) provides exactly this: a cryptographically authenticated emergency override that halts AI-driven processes and reverts to a formally verified safe baseline, with every invocation permanently recorded on-chain for audit.

Article 14(4)(b) specifically addresses automation bias. The Asuu protocol mitigates this by requiring that AI outputs always be presented alongside confidence intervals and model provenance metadata (Section 25), enabling clinicians to assess the basis for AI recommendations rather than treating them as oracular.

Remark 42.1 (Scope of the Article 14 Claim: Necessary Not Sufficient)**.** Article 14(3)'s "commensurate with the risks, level of autonomy and context of use" is a *judgmental standard* applied by a notified body against a specific intended use and risk class. It cannot be satisfied by any temporal-logic formula in isolation. The formula $\Box(\mathit{system_active} \Rightarrow \mathit{human_overseer_assigned})$ encoded in the Human Authority Preservation module (Section 32) guarantees one *necessary precondition* of Article 14(4)(a) compliance: that no inference ever executes without an identifiable human overseer assigned. It does not, and cannot, guarantee that the oversight measures as deployed are commensurate with the risks of a specific clinical indication. That determination rests with the notified body and the sponsor's risk-management file, informed by the evidence the protocol produces but not substituted by it. The honest claim is: the protocol makes Article 14 compliance *auditable and enforceable as a hard precondition*; it does not and cannot render the compliance determination mechanical.

42.2 EU Product Liability Directive (PLD 2024/2853)

Directive (EU) 2024/2853 [46], which entered into force on 9 December 2024 and must be transposed by Member States by 9 December 2026, fundamentally restructures product liability for AI. Three provisions are directly relevant:

Software as product. Article 4(1) explicitly defines "product" to include software, and Recital 13 confirms that AI system providers are treated as manufacturers. The Asuu protocol's on-chain model governance (Section 23) provides cryptographic proof of which model version produced which inference, creating the attribution chain that the PLD's strict liability regime requires.

Continuous learning liability. Article 7(1)(c) requires defect assessment to account for "the effect on the product of any ability to continue to learn after being placed on the market." Recital 32 clarifies that manufacturers designing products with ability to develop unexpected behavior remain liable. The Asuu chain's immutable model update log and PoCA-governed model governance voting (Theorem 23.3) provide verifiable records of every model modification, its approval chain, and its pre-deployment safety attestation.

Complexity presumption and the black-box problem. Article 10(4) creates a rebuttable presumption of defectiveness or causation when proving them is "excessively difficult" due to technical complexity. Recital 48 explicitly identifies machine learning as an example. The Asuu protocol's Shapley-value liability attribution (Theorem 26.3) and complete inference provenance chain provide an *evidentiary framework* that the sponsor can use to rebut the presumption. Whether a particular Shapley attribution actually rebuts the presumption in any particular case is a judicial question and will depend on the specifics of the harm, the causal pathway, and the fact-finder's assessment of the evidence. We do not claim the protocol resolves PLD Article 10(4); we claim it produces the best available evidence for the sponsor's defense. Liability is unlimited (Article 14 removes the former EUR 85 million cap), and the expiry period extends to 25 years for latent personal injuries (Article 17(2)), making cryptographic auditability economically essential regardless of which party prevails.

42.3 FDA Predetermined Change Control Plans (PCCP)

The FDA's final PCCP guidance (December 2024) [6], authorized under Section 515C of the FD&C Act, establishes the regulatory pathway for AI/ML-enabled medical devices that modify their behavior post-market. A PCCP specifies three mandatory components: (i) a Description

of Modifications defining the bounded corridor of permissible changes, (ii) a Modification Protocol detailing data management, retraining practices, performance evaluation, and rollback criteria, and (iii) an Impact Assessment analyzing benefits and risks of each modification individually and in combination, including unintended bias risks.

The Asuu protocol natively generates the documentation infrastructure that PCCPs require. The on-chain model governance registry (Section 23) records every model modification with cryptographic attestation, creating the "traceability table" the guidance mandates. The formal verification attestation (Theorem 25.4) provides the pre-deployment validation evidence. The federated learning protocol (Section 22) with its Rényi differential privacy accountant (Theorem 22.6) documents data management practices with mathematical privacy guarantees. Real-world performance monitoring is provided by the continuous inference logging and anomaly detection mechanisms built into the chain, with automatic rollback triggers when performance metrics breach predefined acceptance thresholds.

The Regulatory Override Channel (Definition 23.6) is essential for PCCP compatibility. FDA reserves authority to require a new submission if real-world performance diverges from the PCCP's impact assessment, to modify or rescind a previously-approved PCCP, and to direct the sponsor to suspend or modify the device. A protocol that enforced a validator-governed corridor so rigidly that the sponsor could not comply with such a directive would be *PCCP-incompatible*, not PCCP-enhanced. The Override Channel resolves this by giving the sponsor and competent regulator paths that bypass validator governance, as described in Remark 23.9.

42.4 MDR Notified-Body Pathway and Post-Market Surveillance

EU Medical Device Regulation (MDR) 2017/745 [4] governs AI-enabled medical devices placed on the EU market at Class IIa and above. The MDR imposes three obligations that an earlier draft of this whitepaper did not address: (i) a pre-market conformity-assessment pathway through a notified body, (ii) a Post-Market Surveillance (PMS) system maintained by the manufacturer throughout the device's lifecycle, and (iii) Periodic Safety Update Reports (PSURs) submitted to the notified body at device-class-specific intervals. The Asuu protocol does not substitute for any of these; it provides the evidentiary infrastructure that makes each tractable.

Definition 42.2 (Notified-Body Interaction Model)**.** For a Class IIa or higher medical device deployed on the Asuu network, the notified-body interaction follows the standard MDR Annex IX or Annex X conformity-assessment procedure, with the following Asuu-specific augmentations:

(i) The manufacturer's Technical Documentation File includes a description of the Asuu integration, the Clinical AI Passport specification, the deployment scope (Regulated-Device Scope per Definition 15.13), and the Permissioned Validator Registry composition (Definition 15.14). The notified body evaluates these as part of the device's intended use and risk-management file.

(ii) The notified body's access to the on-chain evidence is read-only and anonymous: the Asuu Regulatory Portal (Definition 42.3) provides a zero-knowledge-authenticated API through which notified-body inspectors can verify any attestation, model passport state, or PMS event without the manufacturer's cooperation or foreknowledge.

(iii) The notified body is enrolled as a signatory key in the Regulatory Override Channel (Definition 23.6) for devices it has certified, permitting it to issue Regulator Directive

transactions that suspend the device on-chain if post-market surveillance indicates serious risk.

(iv) The CE certificate issued by the notified body is anchored on-chain as a standard attestation, providing verifiable proof-of-certification to regulated-market participants.

Definition 42.3 (Regulatory Portal)**.** The Asuu Regulatory Portal is a read-only query interface that authorized regulatory authorities (FDA, EMA, notified bodies, national health authorities) use to inspect on-chain attestations relevant to devices under their jurisdiction. The portal provides: (a) per-device query of the full Clinical AI Passport history including all version transitions, suspension events, and override transactions; (b) query of inference-record hashes by device identifier and time range with accompanying Merkle proofs of inclusion; (c) query of adverse-event reports anchored by the manufacturer under PMS procedures; (d) query of the Permissioned Validator Registry composition for each device. Access is authenticated by regulator-held cryptographic credentials, logged on-chain (so manufacturers can verify which regulators have inspected their device), and rate-limited to prevent denial-of-service.

Definition 42.4 (Post-Market Surveillance Architecture)**.** The manufacturer's PMS system under MDR Articles 83–86 is implemented as an on-chain process bound to the device's Clinical AI Passport:

(i) **Systematic Data Collection.** All inference records produced by the device are anchored on-chain with their outcome labels (where available through linked clinical systems). This provides the systematic data collection that MDR Article 83 requires.

(ii) **Trend Monitoring.** The continuous-monitoring module ($\mathscr{M}_{\text{monitor}}$ in Definition 23.5) computes performance metrics on sliding windows and triggers θ_{alert} or θ_{suspend} events when metrics exceed pre-specified thresholds. Trend events are anchored on-chain with their underlying metric values, providing the trend-detection that MDR Article 88 requires.

(iii) **Incident Reporting.** Serious incidents and field safety corrective actions (FSCAs) are anchored on-chain within the MDR-mandated reporting windows (15 days for serious incidents, 2 days for serious public-health threats). The on-chain anchor provides tamper-evident timestamping that can be independently verified by the notified body.

(iv) **Vigilance Database Integration.** The on-chain incident anchors include structured references to the manufacturer's submissions to EUDAMED and equivalent vigilance databases, providing bidirectional traceability between the on-chain record and the regulator's external database.

Definition 42.5 (Periodic Safety Update Report Generation)**.** The PSUR required under MDR Article 86 (biennial for Class IIa, annual for Class IIb and III) is generated as a deterministic function of the on-chain PMS state. The report includes: (a) aggregate performance metrics for the reporting period with cryptographic proof of the underlying data, (b) the sequence of alert/suspend events during the period with their resolutions, (c) all incidents and FSCAs with their on-chain anchors, (d) the Clinical AI Passport's version-history and any Regulatory Override transactions during the period, (e) a risk-benefit reassessment comparing the reporting period's outcomes to the device's pre-market evidence. The manufacturer signs the PSUR and submits it to the notified body through the Regulatory Portal; the notified body's acceptance or rejection is itself anchored on-chain.

Remark 42.6 (What the Protocol Substitutes For and What It Does Not)**.** The Asuu protocol substitutes for the *evidence-assembly* burden that currently makes PMS and PSUR compliance expensive: instead of the manufacturer's regulatory-affairs team reconstructing performance data and incident logs from heterogeneous clinical sources, the data are continuously available in standard form on-chain. The protocol does not substitute for: (a) the notified body's conformity assessment, which is a judgmental regulatory determination, (b) the manufacturer's risk-management file under ISO 14971, which is device-specific, (c) clinical evaluation under MDR Article 61, which requires specific clinical evidence appropriate to the intended use, or (d) the manufacturer's legal obligations under the MDR. The protocol reduces the operational cost of compliance evidence; it does not reduce the manufacturer's regulatory responsibility.

42.5 On-Chain Immutability and Record-Correction Rights

A naive reading of the Asuu architecture would hold that on-chain immutability is in irreconcilable conflict with the right to correct erroneous records required by 21 CFR Part 11 [3], the HIPAA Privacy Rule (45 CFR §164.526), GDPR Article 16 (right to rectification) and Article 17 (right to erasure), and the general right of patients to have inaccurate medical records corrected. This reading misunderstands the architectural separation the protocol enforces between two distinct data classes: *on-chain attestations* and *off-chain clinical records*.

Definition 42.7 (Two-Layer Record Architecture)**.** The Asuu protocol distinguishes:

- **On-chain attestations $\mathscr{A}_{\text{chain}}$:** hashes, signatures, timestamps, quorum certificates, and protocol-state transitions. These are immutable by design and carry no patient-identifying content in plaintext.
- **Off-chain clinical records $\mathscr{R}_{\text{offchain}}$:** the FHIR resources, clinical documents, imaging studies, and genomic data held in Data Embassies under local jurisdiction. These are *mutable* under the controlled amendment procedures of the governing jurisdiction.

The protocol never asserts that $\mathscr{R}_{\text{offchain}}$ is immutable. The protocol asserts that the *attestation trail* $\mathscr{A}_{\text{chain}}$ is immutable, and that each attestation is cryptographically bound to the version of the off-chain record that existed at attestation time.

Definition 42.8 (Version-Hash Anchoring)**.** When an off-chain record r is corrected from version v to v' under a lawful amendment procedure, the institution issues a *correction attestation* $CA(r,v,v') = (H(r_v), H(r_{v'}), \mathit{reason_code}, \mathit{authorizing_role}, t, \sigma)$ which is committed on-chain. The original on-chain attestation referencing $H(r_v)$ is not deleted; the chain now contains both the original attestation and the correction attestation, and verifiers can follow the chain of version hashes to reconstruct the full amendment history.

Theorem 42.9 (Compatibility with Record-Correction Rights)**.** *The two-layer record architecture (Definition 42.7) with version-hash anchoring (Definition 42.8) is compatible with:*

(i) ***21 CFR Part 11 audit trail and correction requirements*** *(§11.10(e)): every amendment is recorded with timestamp, authorizing role, and reason-for-change. The on-chain correction attestation implements precisely the audit trail that Part 11 requires, and the underlying clinical record is corrected per the institution's amendment procedures.*

(ii) ***HIPAA right to amend*** *(45 CFR §164.526): the patient's request for amendment is handled by the institution under its existing HIPAA-compliant procedures. The on-chain representation is updated via a correction attestation that preserves the pre-amendment*

hash, matching the regulatory requirement to retain the amendment history while making the corrected version the authoritative current record.

(iii) ***GDPR Article 16 (rectification)****: the mutable off-chain record is rectified directly; the on-chain pseudonymous attestation is updated via version-hash anchoring. The hash of the superseded record remains on-chain but is not itself personal data (see Remark 42.10).*

(iv) ***GDPR Article 17 (erasure)****: when erasure is required, the off-chain record is deleted from the Data Embassy and a* tombstone attestation *is committed on-chain that invalidates further access to the deleted record's hash lineage. The pre-erasure hash remains on-chain, but no plaintext can be reconstructed from it.*

Proof. For (i), Part 11 §11.10(e) requires "secure, computer-generated, time-stamped audit trails to independently record the date and time of operator entries and actions that create, modify, or delete electronic records." The on-chain correction attestation $CA(r, v, v')$ satisfies this literally: it is computer-generated, time-stamped by the consensus protocol, independently verifiable, and records the creation, modification, or deletion event. For (ii), §164.526(d)(2) requires that the amendment "be provided for inclusion in designated record sets" and that the entity "make reasonable efforts to provide the amendment" to persons who rely on the record. The correction attestation provides precisely the notification primitive: any party holding a reference to $H(r_v)$ can verify via the chain that r_v has been superseded by $r_{v'}$. For (iii) and (iv), the key legal question is whether the on-chain hash is personal data. We address this in Remark 42.10. □

Remark 42.10 (On the *Breyer* Question and Pseudonymous Data)**.** The Court of Justice of the European Union held in *Patrick Breyer v Bundesrepublik Deutschland* [51] that dynamic IP addresses can constitute personal data when the holder has access to *legal means* of associating the data with a natural person, even if the holder cannot do so directly. Subsequent EDPB guidance and Article 29 Working Party opinions have extended this principle to cryptographic identifiers, including hashes of personal data. The naive claim "on-chain hashes are not personal data because they are irreversible" is not defensible under *Breyer*: if the institution holding the original record can re-identify the hash (which it can, by recomputing H on the original plaintext), the hash is personal data relative to that institution's data-controller obligations, even though it is not personal data relative to an unrelated third party.

The Asuu architecture addresses this honestly rather than definitionally. The on-chain hash is treated as personal data under GDPR for the purpose of the institution's data-controller obligations: access to the hash is logged, amendment is handled via version-hash anchoring, and erasure is handled via tombstone attestation plus Data Embassy deletion. For third parties who do not hold the original plaintext, the hash is pseudonymous data under Article 4(5). The institution's status as controller and the third party's status as recipient are distinguished by the protocol's access-control layer (Section 28), which enforces that only authorized parties can dereference hashes to the underlying Data Embassy records. We do not claim the chain removes personal data from regulatory scope. We claim it compartmentalizes personal data in a way that makes controller obligations tractable at scale.

Remark 42.11 (Article 17 Erasure: Tombstone Attestation Is Not Complete Erasure)**.** A pointed and correct critique is that "tombstone attestation plus Data Embassy deletion" does not constitute the complete erasure that GDPR Article 17 arguably requires. Specifically: the tombstone records the fact of deletion on-chain, but the salted hash of the original record remains on-chain in its original form, and that salted hash is personal data for as long as the

institution retains the salt. The institution typically retains the salt indefinitely because the FHIR Provenance mapping of Table 2 requires the salt for re-verification of any pre-erasure attestation. The hash is therefore un-erasable in any practical sense even after Data Embassy deletion. We acknowledge this directly rather than obscure it.

Two honest mitigations are available, and both have significant limits:

(i) **Salt destruction.** Upon erasure request, the institution can destroy the salt along with the Data Embassy record. After salt destruction, the on-chain hash becomes a random 256-bit value that cannot be linked to the specific data subject through any computational means available to the institution. The hash remains personal data in a strict *Breyer* reading only if the institution or a third party can plausibly regenerate the salt, which salt destruction prevents. This is the approach we recommend. The limit is that salt destruction also destroys re-verification capability, so prior attestations referencing the hash become uncheckable. The trade-off is correct from an Article 17 standpoint (erasure is prioritized over downstream verification) but must be explicitly accepted by the institution in its data-protection policy.

(ii) **Crypto-shredding.** Where the data-subject's request encompasses a family of related records, the institution destroys the encryption keys protecting the Data Embassy copies. The ciphertexts become undecipherable; the on-chain hashes remain but are no longer linked to any recoverable plaintext. This is ISO/IEC 27040 crypto-shredding applied to clinical data. The limit is the same: verification capability is lost.

Neither mitigation eliminates the tension. On-chain commitment and right-to-erasure are in genuine conflict when the chain is immutable and the commitment is cryptographically binding. The protocol's position is that salt destruction plus Data Embassy deletion is the best-available reconciliation and is materially stronger than the status quo under current EHR architectures (which often retain forensically recoverable backups beyond the erasure request for operational reasons), but is not a complete satisfaction of Article 17's textual requirement. A Data Protection Impact Assessment that treats this honestly will conclude that the protocol meets Article 17 as well as any practical immutable-ledger architecture can, with specific residual risks that must be accepted by the controller and, in several Member States, pre-approved by the supervisory authority before first deployment.

Definition 42.12 (Required Pre-Deployment Data Protection Impact Assessment)**.** Deployment of an Asuu Data Embassy in a jurisdiction subject to the GDPR requires, before first production use with patient data, a Data Protection Impact Assessment under Article 35 addressing at minimum:

(i) The personal-data classification of on-chain hashes under *Breyer*, consistent with Remark 42.10;
(ii) The institution's specific salt-retention and salt-destruction policy under Remark 42.11;
(iii) The cross-border data-flow profile under Articles 44–49, given that validators may be located in jurisdictions outside the EEA and that on-chain state is globally replicated;
(iv) The risk assessment for re-identification through linkage attacks on pseudonymous data, including timestamp and institutional-identifier correlation;
(v) The data-subject-rights satisfaction procedure for rectification, erasure, restriction, and objection, with explicit specification of the Article 17 erasure mechanism;
(vi) The consent-framework integration with Definition 19.2;

(vii) The joint-controller or controller-processor arrangements for Mode V (Definition 16.5) deployments, where the EHR SaaS vendor operates the Data Embassy on behalf of the institution.

In jurisdictions where Article 36 prior consultation with the supervisory authority is indicated (typically Germany under BDSG, France under CNIL guidelines, and jurisdictions with specific health-data-protection regimes), the DPIA outcome is submitted to the authority for consultation before production deployment. The Asuu Foundation maintains a reference DPIA template consistent with EDPB guidance as a protocol-governance priority; each deployment institution must adapt the template to its specific processing activities.

42.6 Cross-Regulatory Synthesis and Capability Matrix

An earlier draft of this whitepaper included a single "Architectural Compliance" theorem asserting that the protocol was compliant with eight regulatory regimes, with a five-line proof. This formulation was technically indefensible: regulators do not approve architectures, they approve devices with intended uses under specific risk classes. No protocol can be proven compliant with a regulatory regime in a single theorem, because compliance is a judgmental determination by the competent authority against the specific risk profile of the deployed device and use case. The honest formulation replaces the architectural-compliance theorem with a *capability matrix* (Table 2) that maps each relevant regulatory regime to specific protocol features that support compliance workflows, without claiming that the features alone constitute compliance.

Remark 42.13 (Capability Is Not Compliance)**.** Table 2 is a map from regulatory requirements to protocol features that support those requirements. It is not a claim that any deployed device running the Asuu protocol is compliant with any listed regime. Compliance is determined by the competent regulator for a specific device with a specific intended use, considering the totality of the risk-management file, quality-management system, clinical evaluation, and post-market surveillance. The Asuu protocol can reduce the cost and increase the auditability of the evidence a sponsor presents for compliance review; it cannot substitute for that review. Any sponsor deploying an Asuu-integrated medical device remains the duty-holder for regulatory compliance under all applicable jurisdictions.

The Data Embassy architecture ensures compliance workflows are compatible with data residency regulations by design rather than by policy. Patient data physically cannot leave the jurisdiction because it physically never enters the chain. Compatibility is architectural, not contractual, a distinction that regulators increasingly demand as data protection enforcement intensifies globally.

The native token is designed as a utility token: it is required to operate network functions, does not represent equity or profit-sharing, and derives value from consumption rather than speculation. The design characteristics supporting this classification are the token's functional necessity (the network cannot operate without it), the absence of any investment-return mechanism, and the protocol's economic architecture, which treats the token as a fee-and-staking instrument rather than as a security interest in the network.

However, the design intent is not determinative. Token classification under the U.S. Securities Act (*SEC v. Howey*, 328 U.S. 293 (1946)), the EU Markets in Crypto-Assets Regulation (MiCA, Regulation (EU) 2023/1114), FinCEN money-transmission frameworks, and equivalent international regimes is a fact-specific legal determination made by the competent authority, not by the protocol designer. A token marketed as a utility asset can be re-characterized as a security,

an e-money token, an asset-referenced token, or a money-transmission instrument depending on how it is actually offered, sold, staked, and traded in practice. The Asuu token's final classification under any given regime will depend on: (a) the terms of the initial disbursement from the Staking Reserve Pool, (b) the existence and structure of secondary trading, (c) the presence of profit-sharing or yield mechanisms in practice even if not by design, (d) the prominence of marketing emphasizing investment returns, and (e) the regulator's application of its current interpretive guidance to the fact pattern. We make no claim that the token is legally non-security, non-MiCA-regulated, or outside money-transmission scope in any specific jurisdiction; those are determinations the competent authority will make, and the protocol's designers and deploying consortium will need to engage specialized counsel and, where appropriate, seek no-action letters, MiCA crypto-asset-service-provider licensing, or equivalent regulatory clarity before token issuance proceeds. An earlier draft of this whitepaper stated the utility-token classification as a factual conclusion; that draft was imprudent, and we correct it here.

A critical barrier to medical AI adoption persists in reimbursement. Despite over 1,350 FDA-cleared AI devices, only two CPT Category 1 payment codes exist for AI as of January 2026. The Asuu network does not solve the reimbursement problem directly, but by reducing the marginal cost of compliance documentation and audit preparation, it lowers the operational threshold at which AI deployment becomes economically viable for institutions that cannot absorb the administrative overhead of current regulatory frameworks.

Table 2: Capability Matrix: Protocol Features Supporting Regulatory Compliance Workflows.

Regulatory Regime	**Requirement (informal)**	**Protocol Capability**	**Ref.**
21 CFR Part 11	Audit trail with correction and reason-for-change	Two-layer record architecture with version-hash anchoring	Theorem 42.9
HIPAA Privacy Rule	Amendment rights and notification	Correction attestation propagates to all hash-referencing parties	Definition 42.8
GDPR Art. 16–17	Rectification and erasure of personal data	Mutable off-chain records; tombstone attestations; controller-recipient separation	Remark 42.10
GDPR Art. 44–49	Data residency in jurisdiction	Data Embassy: patient data physically does not leave institution	Theorem 16.2
EU AI Act Art. 12	Automatic logging of events	Immutable on-chain inference logging	Section 6
EU AI Act Art. 14	Human oversight commensurate with risk	Temporal-logic enforcement of necessary preconditions (not sufficient compliance)	Remark 42.1
EU AI Act Art. 19	Minimum 6-month log retention	Permanent on-chain storage exceeds minimum	Section 6
FDA PCCP (Dec 2024)	Bounded change corridor with sponsor and regulator authority	Clinical AI Passport with Regulatory Override Channel	Theorem 23.8
PLD 2024/2853 Art. 10(4)	Evidentiary framework for complexity presumption	Shapley-value attribution and provenance chain (evidence, not adjudication)	Theorem 26.3
ICH E6(R3)	Immutable protocol and endpoint pre-registration	On-chain commitment of trial protocol and primary endpoints	Theorem 36.2
DSCSA	Package-level drug provenance	Merkle-hashed supply chain with real-time pharmacovigilance	Theorem 35.3
NIST FIPS 203–205	Post-quantum algorithm standards	Composite signature scheme migration (ML-DSA-65)	Theorem 8.2
WHO IHR	Cross-border syndromic surveillance	One Health federated surveillance protocol	Theorem 40.3

43 Cross-Chain Health Data Interoperability

Healthcare data exists across multiple blockchain networks, national health information exchanges, and legacy EHR systems. The Asuu network must interoperate with external chains (e.g., pharmaceutical supply chain blockchains, insurance settlement networks, genomic data marketplaces) without compromising the sovereignty guarantees of the Data Embassy model. We formalize two interoperability primitives: hash time-locked contracts (HTLCs) for atomic cross-chain credential exchange, and relay-based verification for cross-chain inference record validation.

43.1 Cross-Chain Credential Exchange

Definition 43.1 (Hash Time-Locked Contract (HTLC))**.** An HTLC is a smart contract parameterized by $(H(s), \Delta t, \mathit{sender}, \mathit{receiver}, \mathit{value})$ with two execution paths:

- **Claim:** *receiver* presents preimage s such that $H(s)$ matches, before time $t_{\text{now}} + \Delta t$. The contract releases *value* to *receiver*.
- **Refund:** after $t_{\text{now}} + \Delta t$, *sender* reclaims *value*.

Definition 43.2 (Cross-Chain Atomic Credential Swap)**.** A cross-chain atomic credential swap between Asuu chain $\mathscr{C}_A$ and external chain $\mathscr{C}_E$ proceeds as:

1. Institution I creates HTLC_A on $\mathscr{C}_A$ locking a verifiable credential VC_I with hashlock $H(s)$ and timelock Δt_A.
2. External party E creates HTLC_E on $\mathscr{C}_E$ locking asset val_E with the same hashlock $H(s)$ and timelock $\Delta t_E < \Delta t_A$.
3. I claims val_E by revealing s on $\mathscr{C}_E$.
4. E observes s on $\mathscr{C}_E$ and claims VC_I on $\mathscr{C}_A$ using the same s.

Theorem 43.3 (HTLC Atomicity)**.** *The cross-chain credential swap (Definition 43.2) satisfies atomicity: either both transfers complete, or neither does. No party can cheat to obtain the counterparty's asset without releasing their own.*

Proof. We enumerate all cases. **Case 1:** I reveals s on $\mathscr{C}_E$ before Δt_E. Then E learns s (it is published on-chain on $\mathscr{C}_E$) and can claim VC_I on $\mathscr{C}_A$ before Δt_A (since $\Delta t_E < \Delta t_A$, E has $\Delta t_A - \Delta t_E$ additional time). Both transfers complete. **Case 2:** I does not reveal s before Δt_E. Then E reclaims val_E after Δt_E via the refund path. Since s was never revealed, E cannot claim VC_I on $\mathscr{C}_A$, and I reclaims VC_I after Δt_A. Neither transfer completes. **Case 3:** I tries to claim val_E without revealing s. This is impossible because the HTLC requires a valid preimage. **Case 4:** E tries to claim VC_I without s. This requires finding s' with $H(s') = H(s)$, which contradicts collision resistance (Definition 4.1). □

43.2 Relay-Based Cross-Chain Verification

Definition 43.4 (Light Client Relay)**.** An Asuu relay contract $\mathscr{R}$ deployed on external chain $\mathscr{C}_E$ maintains a set of Asuu block headers $\{(h_k, \sigma_k)\}_{k=1}^{L}$ and exports:

- $\mathsf{SubmitHeader}(h_k, \sigma_k)$: verifies that σ_k is a valid $> 2/3$ supermajority multi-signature over h_k, and that $h_k.prev$ matches the previous stored header. Appends h_k to the relay state.
- $\mathsf{VerifyInference}(IR, \pi_{\text{Merkle}}, k)$: verifies the Merkle proof π_{Merkle} of IR against the root in header h_k. Returns true if and only if the proof is valid.

Theorem 43.5 (Relay Security)**.** *The relay contract (Definition 43.4) provides trustless verification of Asuu inference records on $\mathscr{C}_E$, under the assumptions that: (i) the Asuu validator set is honestly-supermajority ($f < n/3$), and (ii) the hash function H is collision-resistant. An adversary can forge a relay acceptance only by either compromising $> 1/3$ of Asuu validators (violating (i)) or finding a Merkle collision (violating (ii)).*

Proof. The relay verifies $> 2/3$ multi-signatures over each header. Producing a valid signature requires the corresponding validator's private key. Forging $> 1/3$ of signatures requires compromising $> 1/3$ validators (by assumption, at most $f < n/3$ are adversarial, so the adversary cannot produce a valid multi-signature over a fabricated header). Given valid headers, the Merkle proof verification reduces to Theorem 7.4: a false proof requires a hash collision. The relay's security is therefore the conjunction of honest supermajority and collision resistance, which is already assumed by the Asuu protocol. □

Theorem 43.6 (Cross-Chain Inference Portability)**.** *A clinical inference record finalized on the Asuu chain can be verified on any external chain $\mathscr{C}_E$ that hosts a relay contract, with the following properties:*

(i) ***Completeness:*** *every finalized inference record has a valid Merkle proof that the relay accepts.*
(ii) ***Soundness:*** *no non-finalized or fabricated record is accepted (up to* $\text{negl}(\lambda)$*).*
(iii) ***Timeliness:*** *the relay reflects the Asuu chain state with latency at most $L \cdot \Delta_{\text{block}}$, where L is the relay update frequency in blocks and Δ_{block} is the Asuu block interval.*
(iv) ***Sovereignty preservation:*** *no clinical data leaves the Data Embassy; only the inference record metadata (model hash, confidence score, timestamp) is verified cross-chain.*

Proof. (i) Completeness follows from the Merkle tree construction: every transaction in a finalized block has a valid Merkle path to the block's root. (ii) Soundness follows from Theorem 43.5. (iii) The relay is updated every L blocks; each block takes Δ_{block} time, so the maximum staleness is $L \cdot \Delta_{\text{block}}$. (iv) The relay contract only stores and verifies block headers and Merkle proofs. The inference record IR contains metadata (Definition 5.1) but not the underlying patient data, which remains in the originating Data Embassy by protocol design. □

43.3 2025–2026 Regulatory Developments

The regulatory landscape has shifted substantially since the initial Asuu whitepaper. Three developments are particularly significant.

First, the HIPAA Security Rule received its most sweeping proposed update since 2013, published as a Notice of Proposed Rulemaking (NPRM) on January 6, 2025 [45]. The proposed rule eliminates the distinction between "required" and "addressable" safeguards, making all

implementation specifications mandatory. It mandates encryption of ePHI at rest and in transit, multi-factor authentication, annual technology asset inventories, network mapping, biannual vulnerability scanning, annual penetration testing, and 72-hour security incident restoration. The final rule is scheduled for May 2026 with a compliance window of 180 to 240 days. The Asuu network's architecture inherently satisfies these requirements: encryption is mandatory by protocol design (CP-ABE for data at rest, TLS 1.3 for transit), validator node authentication is multi-factor (institutional key plus TEE attestation), and asset inventories are maintained on-chain.

Second, the EU's revised Product Liability Directive (Directive 2024/2853), effective December 9, 2026, explicitly classifies software including AI as a "product" and establishes a presumption of defectiveness for AI that violates the AI Act [46]. This creates a strict liability regime for AI developers whose systems cause harm, making the Asuu network's formal safety verification (Section 25) and liability attribution (Section 26) not merely best practices but legally operative defenses.

Third, the W3C Verifiable Credentials Data Model v2.0 was published as a W3C Recommendation on May 15, 2025 [24], providing an interoperable standard for the Asuu SSI layer (Section 17.2).

Theorem 43.7 (Updated Regulatory Compatibility)**.** *The expanded Asuu protocol additionally satisfies:*

(i) ***HIPAA NPRM 2025****: mandatory encryption (Theorem 28.10), multi-factor validator authentication (Section 9), asset inventory via on-chain state, and 72-hour incident response via break-glass audit logs (Theorem 31.5);*

(ii) ***EU AI Act Article 14****: formal human oversight guarantees (Theorem 25.12);*

(iii) ***EU PLD 2024/2853****: role-parameterized liability attribution (Theorem 26.3, Definition 26.14) providing evidence for defect causation analysis;*

(iv) ***RBAC enforcement****: all 114 roles governed by on-chain $RBAC_3$ with provable SSD/DSD invariants (Theorem 28.8, Theorem 29.5).*

Proof. Each claim follows directly from the referenced theorems and the corresponding regulatory text, as demonstrated in the proofs of the individual theorems. □

44 Clinical UI/UX and Abstraction Layer

The preceding sections formalize the cryptographic, consensus, and governance primitives that secure the Asuu network. However, the clinical adoption of any blockchain infrastructure depends on a critical property: the complexity of the underlying protocol must be invisible to the end user. A surgeon in a high-acuity environment does not interact with Merkle roots; she interacts with a decision-support interface that must respond in sub-second latency. This section formalizes the *Clinical Abstraction Gateway* (CAG), an API/SDK middleware layer that translates complex wallet interactions, credential presentations, and inference-record signing into seamless, single-click background processes without compromising any cryptographic guarantee.

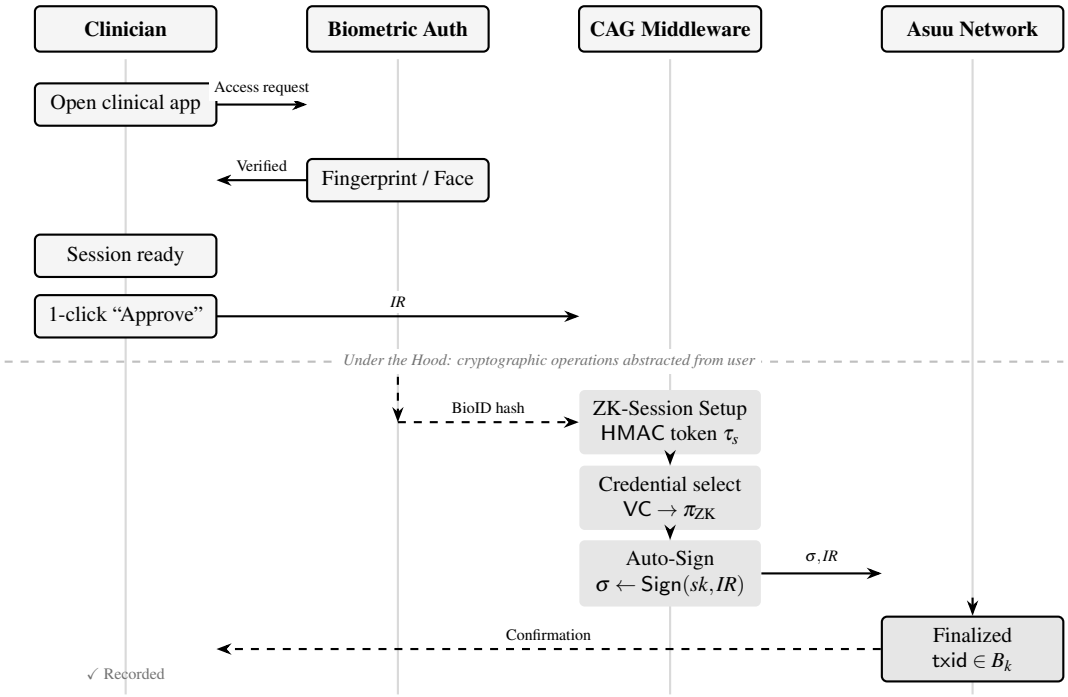


Figure 14: Clinical Abstraction Gateway (CAG) UX sequence diagram. The clinician authenticates once via biometrics (top layer). Below the abstraction boundary, the CAG middleware handles ZK-session establishment, credential selection, and auto-signing transparently. Subsequent inferences require only a single click; the CAG submits to the Asuu chain without further user interaction.

44.1 Formal Gateway Architecture

Definition 44.1 (Clinical Abstraction Gateway)**.** A Clinical Abstraction Gateway is a tuple $\mathscr{G} = (\mathsf{Auth}, \mathsf{Sign}, \mathsf{Submit}, \mathsf{Query}, \mathsf{Notify})$ where:

(i) $\mathsf{Auth} : \mathsf{BioID} \times \mathsf{DeviceCtx} \to \{(\mathsf{sk}, \mathsf{VC}, \mathsf{session})\}$ authenticates the clinician using biometric input and device context, returning a session-bound signing key, verifiable credential set, and session token;

(ii) $\mathsf{Sign} : \mathsf{IR} \times \mathsf{sk} \times \mathsf{VC} \to \sigma$ signs an inference record IR with the clinician's key and attaches the relevant role credential, producing a compact signature;

(iii) $\mathsf{Submit} : \sigma \times \mathsf{IR} \to \mathsf{txid}$ submits the signed record to the Asuu mempool and returns a transaction identifier;

(iv) $\mathsf{Query} : \mathsf{txid} \to \{\mathsf{pending}, \mathsf{finalized}, \mathsf{rejected}\}$ returns the on-chain status of a submitted record;
(v) $\mathsf{Notify} : \mathsf{txid} \times \mathsf{status} \to \mathsf{UIEvent}$ asynchronously pushes status updates to the clinical interface.

The gateway operates as a *thin client* that maintains no patient data locally. All cryptographic operations execute within a local Trusted Execution Environment (TEE) or hardware security module embedded in the clinical workstation, ensuring that signing keys never enter application memory.

44.2 Zero-Knowledge Session Authentication

Clinical authentication must satisfy two competing requirements: the network must verify the clinician's role and institutional affiliation, but the identity of the clinician should not be revealed to the network at large. We formalize this as a zero-knowledge session protocol.

Definition 44.2 (zk-Session Protocol)**.** Let VC_c be a clinician's BBS+ credential encoding role $r \in \mathscr{R}$, institution $\mathscr{I}$, and scope $\mathscr{S}$. The zk-Session protocol $\Pi_{\text{zk-ses}}$ proceeds as:

(1) **Commitment.** The clinician's device computes a Pedersen commitment $C = g^{m_1} h^{r_{\text{blind}}}$ to the credential attributes $m_1, \ldots, m_k$;
(2) **Proof generation.** The device generates a zero-knowledge proof $\pi = \mathsf{ZKP}\{(m_1, \ldots, m_k, r_{\text{blind}}) : C = g^{m_1} h^{r_{\text{blind}}} \wedge \mathsf{VC}_c.\mathsf{role} \geq r_{\min} \wedge \mathsf{VC}_c.\mathsf{inst} \in \mathscr{I}_{\text{valid}}\}$;
(3) **Session binding.** The gateway verifies π and issues a session token $\tau = \mathsf{HMAC}_{k_s}(\mathsf{nonce} \| \mathsf{timestamp} \| H(C))$ valid for duration Δ_{ses};
(4) **One-click signing.** Subsequent inference records are signed as $\sigma_{\mathsf{IR}} = \mathsf{BBS+.Sign}(\mathsf{sk}_c, \mathsf{IR} \| \tau)$, requiring no additional user interaction.

Theorem 44.3 (Gateway Security Preservation)**.** *The Clinical Abstraction Gateway $\mathscr{G}$ preserves all security properties of the underlying Asuu protocol. Specifically:*

(i) ***Authentication equivalence****: $\Pi_{\text{zk-ses}}$ provides the same authentication guarantees as direct credential presentation, under the hiding property of Pedersen commitments and the soundness of the BBS+ signature scheme (Section 17);*
(ii) ***Non-repudiation****: every inference record carries a valid BBS+ signature bound to the session commitment, providing the same non-repudiation as direct signing;*
(iii) ***Session isolation****: compromise of a session token τ does not reveal the clinician's long-term signing key sk_c, as τ is derived from a one-way HMAC and the session key k_s resides exclusively within the device TEE.*

Proof. (i) The zk-Session proof π is a Σ-protocol composition proving knowledge of credential attributes satisfying the role predicate. Soundness of the composed proof follows from the special soundness of the underlying Schnorr-like proofs and the binding property of Pedersen commitments: an adversary who authenticates without valid credentials can extract a discrete-log witness, contradicting DLP hardness (Section 4). (ii) Each σ_{IR} is a standard BBS+ signature under sk_c; the session token τ is bound as a signed attribute but does not alter the unforgeability guarantee, which reduces to q-SDH (Theorem 17.3). (iii) The HMAC construction is a PRF under the standard assumption; k_s is generated and stored within the TEE. Compromise of τ reveals only $\mathsf{HMAC}_{k_s}(\cdot)$, from which k_s cannot be recovered (PRF security), and sk_c is never an input to the HMAC. $\square$

44.3 Latency Guarantees and Clinical Responsiveness

Definition 44.4 (Clinical Latency Budget)**.** Define the end-to-end latency for a single inference-record submission as:

$$L_{\text{total}} = L_{\text{auth}} + L_{\text{sign}} + L_{\text{client}} + L_{\text{submit}} + L_{\text{confirm}}$$

where L_{auth} is the zk-proof generation time (amortized to zero after session establishment), L_{sign} is the BBS+ signing time, L_{client} is the EHR client processing and display overhead, L_{submit} is the network propagation to the nearest validator, and L_{confirm} is the time until the record appears in a finalized block.

Proposition 44.5 (Latency Profile by Deployment Class)**.** *Cryptographic-operation latency is a small component of end-to-end clinical-interaction latency; EHR client rendering and network conditions dominate. We therefore characterize L_{total} per deployment class rather than asserting a single bound:*

Component	***Modern AMC workstation, institutional LAN***	***Mid-tier community hospital, hospital Wi-Fi***	***Legacy workstation, congested shared Wi-Fi (worst realistic case)***
L_{auth} *(initial)*	≤ 150 *ms*	≤ 400 *ms*	*500 ms–2 s*
L_{auth} *(amortized, in-session)*	≈ 0 *ms*	≈ 0 *ms*	≈ 0 *ms*
L_{sign} *(BBS+)*	≤ 5 *ms*	≤ 15 *ms*	*20–60 ms*
L_{client} *(EHR UI)*	*20–50 ms*	*50–150 ms*	*150–500 ms*
L_{submit} *(LAN to validator)*	≤ 50 *ms*	*50–150 ms*	*100–500 ms*
L_{confirm} *(optimistic)*	*mempool-receipt*	*mempool-receipt*	*mempool-receipt*
Amortized L_{total}	***75–105 ms***	***115–315 ms***	***270 ms–1.1 s***

Cryptographic overhead (BBS+ signing, amortized zk-proof) contributes no more than 60 ms in the worst realistic case; the dominant latency sources are the EHR client's rendering pipeline and the network path, neither of which the Asuu protocol controls. The protocol's cryptographic operations are compatible with perceived-instantaneity clinical targets in modern-workstation deployments and with perceived-responsive (sub-second) targets in mid-tier deployments. In legacy-workstation deployments with congested networks, L_{total} can exceed 1 second during session establishment and during bursts of clinical activity; the deployment architecture must pre-warm sessions and batch submissions to stay within clinician-tolerance thresholds.

Remark 44.6 (Honest Scope of the Latency Claim)**.** An earlier draft of this proposition claimed $L_{\text{total}} \leq 55$ ms under generic "standard hardware assumptions." That bound aggregated only the cryptographic-operation components and ignored EHR client rendering, legacy-workstation variability, congested-Wi-Fi tail latency, and browser-based SMART-on-FHIR loader overhead. In realistic hospital-deployment settings the 55 ms figure is not achievable for end-to-end clinician perception; the cryptographic components are indeed sub-30-ms but the clinical interaction is not. Physicians will actively bypass or sabotage any system that slows down charting by even a fraction of a second; the correct engineering target is not "55 ms in theory" but "indistinguishable from baseline EHR latency in the deployment environment." The revised profile

of Proposition 44.5 reflects this target. Deployment sites whose legacy-infrastructure profile pushes L_{total} into seconds should not deploy Asuu-integrated CDS services in clinician-facing workflows without infrastructure upgrade, or should confine Asuu integration to asynchronous audit-trail anchoring that does not block clinician workflow.

45 Edge-AI and Medical IoT Oracles

The Asuu network's inference-record model assumes that AI computations occur on robust institutional servers. However, the future of clinical AI increasingly relies on real-time biometric feeds from resource-constrained edge devices: continuous glucose monitors, wearable ECG patches, portable point-of-care ultrasound systems, and surgical navigation sensors. These devices cannot run full validator nodes, yet the data they produce must be anchored to the Asuu chain with the same integrity guarantees as server-side inferences. This section formalizes the *Asuu Hardware Oracle Protocol*, specifying how lightweight edge devices securely attest biometric data provenance.

The protocol is defined entirely in terms of abstract cryptographic primitives; Hardware Root of Trust, Abstract Secure Enclave, and Secure Boot Hash Chain; to remain future-proof across any secure microarchitecture that satisfies the formal requirements below. No specific chipset, microarchitecture, or silicon vendor is prescribed; any hardware platform that implements the abstract TEE interface $\mathfrak{T}$ and achieves an appropriate third-party security certification (e.g., PSA Certified Level 2 or Common Criteria EAL5+) is compliant.

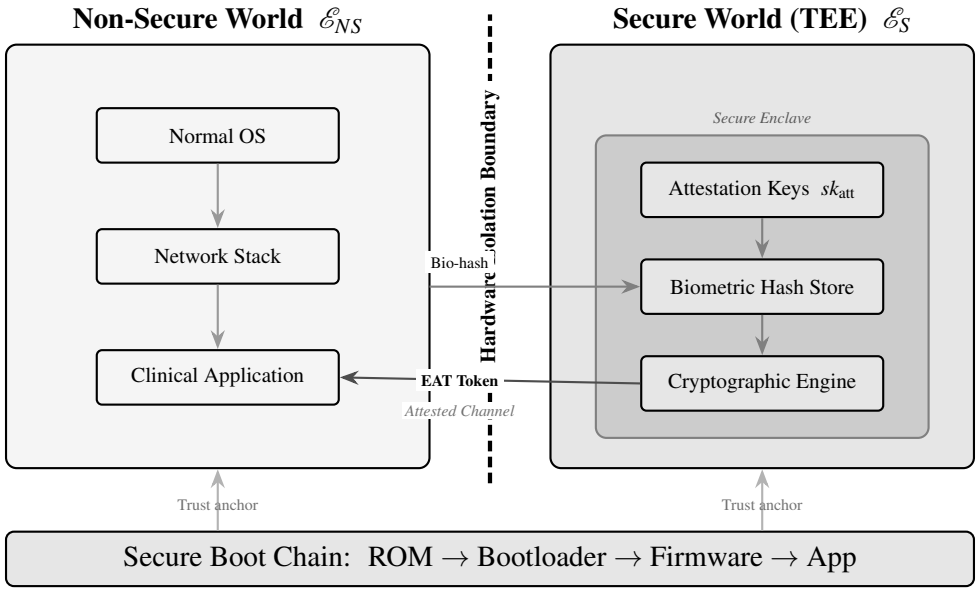


Figure 15: Edge-AI TEE hardware architecture. The device is partitioned into a Non-Secure World (left: Normal OS, Network Stack, Clinical Application) and a Secure World (right: Attestation Keys, Biometric Hash, Cryptographic Engine) separated by a hardware-enforced isolation boundary (heavy dashed line). Cross-boundary communication occurs exclusively through attested, encrypted channels. The Secure Boot Chain anchors trust from ROM into both worlds.

Definition 45.1 (Abstract Edge TEE Interface)**.** An Asuu-compliant Edge Trusted Execution Environment is any hardware-software system $\mathfrak{T} = (\mathscr{E}_S, \mathscr{E}_{NS}, \mathsf{SAU}, \mathsf{MPU}_S, \mathsf{RoT}, \mathsf{Att})$ satisfying:

(i) **World Isolation:** The execution environment is partitioned into a Secure world $\mathscr{E}_S$ and a Non-secure world $\mathscr{E}_{NS}$ with hardware-enforced memory boundaries. No instruction sequence executable in $\mathscr{E}_{NS}$ can read or write memory mapped to $\mathscr{E}_S$.

(ii) **Security Attribution:** A configurable Security Attribution Unit SAU with $|\mathsf{SAU}| \geq 4$ programmable regions classifies each memory address as Secure, Non-secure, or Non-secure Callable (entry-point only). An immutable Implementation-Defined Attribution Unit provides a silicon-fixed baseline; the effective classification is $\max(\mathsf{SAU}, \mathsf{IDAU})$.

(iii) **Memory Protection:** Each world has an independent Memory Protection Unit $\mathsf{MPU}_S, \mathsf{MPU}_{NS}$ enforcing region-granular access permissions (read/write/execute) with minimum granularity ≤ 32 bytes and ≥ 8 configurable regions per world.

(iv) **Hardware Root of Trust:** The Root of Trust RoT provides: (a) an immutable on-chip boot ROM; (b) one-time-programmable fuses for key provisioning; (c) a hardware random number generator certified to AIS-31 Class PTG.2 or NIST SP 800-90B; (d) secure key storage inaccessible from $\mathscr{E}_{NS}$.

(v) **Attestation Service:** The attestation module Att generates Entity Attestation Tokens $\mathsf{EAT} = \mathsf{Sign}(sk_{\text{att}}, \mathsf{dev\text{-}id} \| \mathsf{fw\text{-}hash} \| \mathsf{lifecycle} \| t)$ where sk_{att} is generated and stored exclusively within RoT.

(vi) **Transition Security:** World transitions occur only through designated entry points. On transition from $\mathscr{E}_S$ to $\mathscr{E}_{NS}$, all Secure register state is automatically cleared by hardware. The transition latency is bounded by $\delta_{\text{switch}} \leq 10\,\mu$s.

Definition 45.2 (Abstract Edge AI Accelerator)**.** For devices performing on-device inference, an optional Neural Processing Unit or equivalent hardware accelerator NPU is characterized abstractly by:

(i) **Throughput:** $\geq C_{\min}$ multiply-accumulate operations per second, where $C_{\min}$ is defined per clinical application (e.g., $C_{\min} = 10^8$ MAC/s for real-time ECG arrhythmia classification);

(ii) **Precision support:** INT8 and INT16 quantized inference at minimum;

(iii) **Security integration:** The accelerator's memory bus participates in the TEE security fabric, ensuring that model weights loaded into the accelerator cannot be exfiltrated through Non-secure DMA channels.

When NPU is present, the attestation token EAT includes an additional field $H(\mathsf{model_weights})$ binding the inference result to the specific model executed.

Definition 45.3 (Secure Boot Hash Chain (Abstract))**.** The abstract secure boot establishes a chain of trust $\mathsf{CoT} = (h_0, h_1, \ldots, h_k)$ where $h_0 = H(\mathsf{OCROM})$ is the hash of the immutable on-chip boot ROM and $h_i = H(h_{i-1} \| \mathsf{stage}_i)$ for each subsequent boot stage. The chain terminates at $h_k = \mathsf{fw\text{-}hash}$. The boot process enforces:

(i) Each stage i verifies the signature of stage $i+1$ before transferring execution;

(ii) Rollback to a prior firmware version is prevented by a monotonic hardware counter cnt_{fw} stored in one-time-programmable memory;

(iii) If any verification fails, the device enters a safe lockdown state and does not execute application code.

45.1 Device Attestation Model

Definition 45.4 (Attested Medical IoT Device)**.** An attested medical IoT device is a tuple $\mathscr{D} = (\mathsf{dev\text{-}id}, \mathsf{TEE}_d, \mathsf{ak}, \mathsf{fw\text{-}hash}, \mathsf{cal})$ where:

(i) dev-id is a unique device identifier bound to the hardware at manufacture;

(ii) TEE_d is a hardware trusted execution environment (e.g., any hardware TEE satisfying Definition 45.1) providing isolated secure-world execution;
(iii) $\mathsf{ak} = (pk_{\text{ak}}, sk_{\text{ak}})$ is an attestation key pair generated within TEE_d at provisioning, where sk_{ak} never leaves the secure world;
(iv) $\mathsf{fw\text{-}hash} = H(\mathsf{firmware})$ is the cryptographic hash of the device's firmware image, measured during secure boot;
(v) cal is the device's calibration certificate, signed by the manufacturer and anchored on-chain.

Remark 45.5 (Secure Boot Instantiation)**.** The generic Secure Boot Hash Chain (Definition 45.3) is instantiated on each device by the manufacturer. The final measurement $h_k = \mathsf{fw\text{-}hash}$ is signed by sk_{ak} and submitted to the on-chain device registry.

45.2 Biometric Data Attestation Protocol

Definition 45.6 (Oracle Attestation)**.** For a biometric measurement m captured at time t from patient p, the device oracle produces an attestation:

$$\mathsf{Att}(m,t,p) = \big(\mathsf{dev\text{-}id},\ H(m),\ t,\ H(\mathsf{pid}_p),\ \mathsf{fw\text{-}hash},\ \sigma_{\text{att}}\big)$$

where $\sigma_{\text{att}} = \mathsf{Sign}(sk_{\text{ak}}, \mathsf{dev\text{-}id} \| H(m) \| t \| H(\mathsf{pid}_p) \| \mathsf{fw\text{-}hash})$ is computed entirely within TEE_d. The raw measurement m remains in the local Data Embassy; only the hash $H(m)$ and the attestation travel to the chain.

Theorem 45.7 (Oracle Data Integrity)**.** *Under the assumptions of TEE integrity (Section 18) and collision resistance of H, the Oracle Attestation protocol satisfies:*

(i) ***Provenance authenticity****: a valid attestation $\mathsf{Att}(m,t,p)$ guarantees that the measurement m was captured by device $\mathscr{D}$ with firmware $\mathsf{fw\text{-}hash}$ at time t, because forging σ_{att} requires sk_{ak}, which is confined to TEE_d;*
(ii) ***Measurement binding****: any alteration of m to $m' \neq m$ is detectable because $H(m') \neq H(m)$ with overwhelming probability (collision resistance), and the on-chain record commits to $H(m)$;*
(iii) ***Temporal anchoring****: the timestamp t is generated by TEE_d's secure clock, which is not writable from the normal world. Combined with on-chain block timestamps, this provides dual-source time anchoring with precision bounded by $\max(\delta_{\text{TEE-clock}}, \Delta_{\text{block}})$.*

Proof. (i) Signature unforgeability under ECDSA (or EdDSA) reduces to the DLP assumption (Section 4). Since sk_{ak} is generated and stored exclusively within TEE_d, an adversary without physical side-channel access to the secure world cannot extract sk_{ak}. The TEE integrity assumption (Section 18) guarantees that the secure world's memory isolation is not breached by software attacks. (ii) If an adversary modifies m to m' after attestation, the hash $H(m')$ will not match the committed $H(m)$ unless $H(m') = H(m)$, which contradicts collision resistance. (iii) The abstract TEE specification (Definition 45.1) requires a secure timer peripheral mapped exclusively to $\mathscr{E}_S$. The Non-secure world cannot write to this timer. The secure-world timestamp is included in σ_{att}, so backdating requires forging a signature or compromising the TEE. □

45.3 Lightweight Relay to Asuu Chain

Edge devices lack the computational resources to participate as full validators. Instead, they relay attestations through a designated *gateway validator* at the host institution.

Definition 45.8 (IoT Gateway Relay)**.** The gateway relay protocol proceeds as:

(1) The device $\mathscr{D}$ transmits $\mathsf{Att}(m,t,p)$ to the institutional gateway validator $\mathscr{V}_g$ over an authenticated TLS 1.3 channel;
(2) $\mathscr{V}_g$ verifies: (a) σ_{att} against pk_{ak} from the on-chain device registry, (b) fw-hash matches the registered firmware hash, (c) dev-id is not revoked;
(3) Upon successful verification, $\mathscr{V}_g$ batches multiple attestations into a Merkle tree with root $R_{\text{batch}} = \mathsf{MerkleRoot}(\mathsf{Att}_1, \ldots, \mathsf{Att}_n)$ and submits a single on-chain transaction containing R_{batch} and the batch metadata;
(4) Individual attestations can be verified against R_{batch} using standard Merkle proofs (Section 7).

Proposition 45.9 (Batching Efficiency)**.** *For n IoT attestations per batch interval* Δ_{batch}*, the relay protocol achieves:*

(i) ***On-chain compression****: a single transaction of size* $O(\log n)$ *(Merkle root plus metadata) replaces n individual transactions, reducing gas cost by a factor of* $\Theta(n/\log n)$*;*
(ii) ***Verification cost****: any individual attestation can be verified in* $O(\log n)$ *time using its Merkle proof;*
(iii) ***Latency bound****: each attestation reaches finality within* $\Delta_{\text{batch}} + \Delta_{\text{block}}$*, which for* $\Delta_{\text{batch}} =$ 5 s *and* $\Delta_{\text{block}} = 6$ s *yields at most* 11 s*, acceptable for non-acute telemetry.*

46 On-Chain Medical Arbitration and Dispute Resolution

Sections 37 and 26 formalize insurance automation and liability attribution respectively. However, complex clinical disputes, particularly those involving disagreements between an AI recommendation and a physician's override, or inter-institutional conflicts over standard-of-care adherence, require an active adjudication mechanism. This section defines the *On-Chain Clinical Arbitration Protocol* (OCAP), a decentralized dispute resolution system adapted from Schelling-point game theory to the medical domain.

46.1 Dispute Initiation and Jury Selection

Definition 46.1 (Clinical Dispute)**.** A clinical dispute is a tuple $\mathscr{D}_{\text{disp}} = (\mathsf{IR}_{\text{contested}}, P_A, P_B, \mathsf{subcourt}, \mathsf{stake}_A, \mathsf{stake}_B, \mathsf{evidence})$ where:

(i) $\mathsf{IR}_{\text{contested}}$ is the inference record under dispute;
(ii) P_A, P_B are the disputing parties (e.g., the AI system's operating institution versus the overriding physician's institution);
(iii) $\mathsf{subcourt} \in \mathscr{C}$ specifies the medical specialty court (e.g., neurology, radiology, emergency medicine);
(iv) $\mathsf{stake}_A, \mathsf{stake}_B \geq \mathsf{stake}_{\min}$ are the arbitration bonds deposited by each party;
(v) evidence is a set of on-chain content identifiers (CIDs) pointing to encrypted evidence stored in Data Embassies.

Definition 46.2 (Jury Selection via Weighted Sortition)**.** The jury for a dispute in subcourt $c \in \mathscr{C}$ is selected as follows. Let $\mathscr{J}_c = \{j_1, \ldots, j_N\}$ be the set of eligible jurors: practitioners holding BBS+ credentials with role $r \geq r_{\min}(c)$ in the relevant specialty, who have staked s_{j_i} tokens in subcourt c. Each juror j_i is drawn with probability:

$$\Pr[\text{juror } j_i \text{ selected}] = \frac{s_{j_i}}{\sum_{j \in \mathscr{J}_c} s_j}$$

The draw uses a verifiable random function $\mathsf{VRF}(\mathsf{sk}_{\text{beacon}}, \mathsf{dispute\text{-}id} \| \mathsf{round})$ to produce an unbiasable, verifiable random seed. A panel of k jurors (default $k = 3$, extensible to $k \in \{5, 7\}$ on appeal) is selected without replacement.

46.2 Blinded Evidence Access via Zero-Knowledge

Jurors must evaluate clinical evidence without gaining persistent access to patient data. We formalize this using time-limited, zero-knowledge evidence viewing.

Definition 46.3 (Blinded Evidence Session)**.** For each selected juror j_i, the arbitration contract grants a time-limited proxy re-encryption key $rk_{\mathsf{disp} \to j_i}$ (Section 21) valid for duration Δ_{review}. The juror accesses the evidence through a TEE-protected viewer that:

(1) decrypts the evidence within the TEE enclave;
(2) renders the clinical data for review but prevents export, screenshot, or copy;
(3) generates a zero-knowledge proof $\pi_{\text{access}} = \mathsf{ZKP}\{j_i \text{ accessed evidence CID at time } t\}$ for the audit trail;
(4) destroys the re-encryption key and all decrypted material when Δ_{review} expires.

46.3 Schelling-Point Voting Mechanism

Definition 46.4 (Clinical Schelling Game)**.** The arbitration voting mechanism is a coordination game $\Gamma = (\mathscr{J}, \mathscr{A}, u)$ where:

(i) $\mathscr{J} = \{j_1, \ldots, j_k\}$ is the jury panel;
(ii) $\mathscr{A} = \{\text{standard-met}, \text{standard-violated}, \text{inconclusive}\}$ is the action set;
(iii) The utility function $u : \mathscr{A}^k \to \mathbb{R}^k$ assigns payoffs as:

$$u_i(a_i, a_{-i}) = \begin{cases} \mathsf{fee} + \frac{\mathsf{stake}_{\text{minority}}}{|\text{majority}|} & \text{if } a_i \in \text{majority} \\ -s_{j_i} \cdot \mathsf{slash\text{-}rate} & \text{if } a_i \in \text{minority} \end{cases}$$

where $\mathsf{stake}_{\text{minority}}$ is the total stake of minority-voting jurors redistributed to majority voters.

The mechanism uses a commit-reveal scheme: jurors first submit $H(\mathsf{vote} \| \mathsf{salt})$, then reveal $(\mathsf{vote}, \mathsf{salt})$ in a second phase.

Theorem 46.5 (Schelling Incentive Compatibility)**.** *Under the assumption that a unique Schelling focal point exists (i.e., an objectively correct medical standard-of-care determination), the Clinical Schelling Game* Γ *has the property that honest voting (voting according to clinical evidence) is a Nash equilibrium. Formally, for each juror* j_i*:*

$$u_i(\mathsf{honest}, a^*_{-i}) \geq u_i(a'_i, a^*_{-i}) \quad \forall\, a'_i \neq \mathsf{honest}$$

where a^*_{-i} *denotes the honest strategies of all other jurors.*

Proof. When a unique Schelling focal point exists, each rational juror expects the others to vote honestly (the focal point serves as the coordination equilibrium). If juror j_i deviates from the focal point, she enters the minority with probability approaching 1 (as k grows), incurring a stake slash of $s_{j_i} \cdot \mathsf{slash\text{-}rate}$. If she votes honestly, she enters the majority and receives the fee plus a share of redistributed minority stakes. Since:

$$\mathsf{fee} + \frac{\mathsf{stake}_{\text{minority}}}{|\text{majority}|} > 0 > -s_{j_i} \cdot \mathsf{slash\text{-}rate}$$

the honest strategy strictly dominates any deviation when the Schelling focal point is unique. For the medical context, uniqueness follows from the existence of published clinical guidelines and evidence-based standards of care, which provide a shared reference frame analogous to Schelling's "prominent landmark" in the New York meeting problem. □

46.4 Appeal Mechanism

Definition 46.6 (Arbitration Appeal)**.** Either party may appeal within Δ_{appeal} blocks by depositing an increased bond $\mathsf{stake}_{\text{appeal}} = 2 \cdot \mathsf{stake}_{\text{prev}}$. The appeal doubles the jury size: $k_{\text{appeal}} = 2k + 1$, drawing from a broader pool of eligible specialists. At most $\ell_{\max}$ appeal rounds are permitted, after which the ruling is final and on-chain enforceable.

Proposition 46.7 (Appeal Convergence)**.** *The geometrically increasing appeal bond ensures convergence. After* ℓ *rounds, the total bonded capital is* $\mathsf{stake}_0 \cdot (2^{\ell+1} - 1)$*. Since the bond grows exponentially while the expected gain from overturning a correct ruling remains bounded, rational parties cease appealing after* $O(\log(\mathsf{stake}_{\max}/\mathsf{stake}_0))$ *rounds.*

47 High-Frequency Intraoperative State Channels

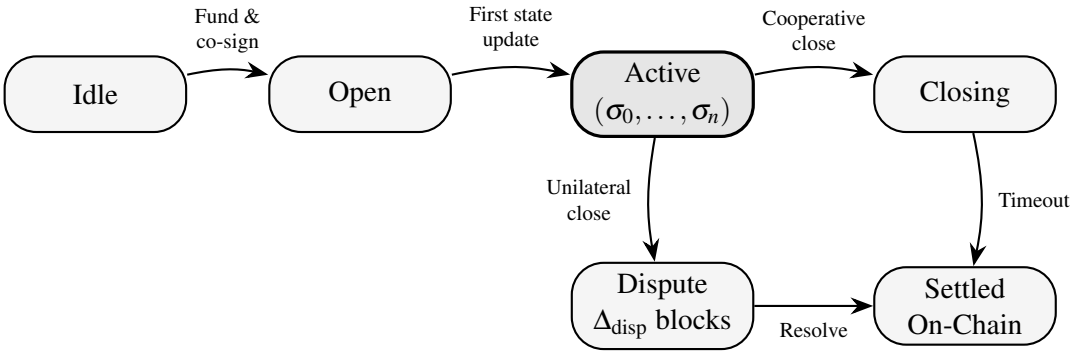


Figure 16: Intraoperative state channel lifecycle. The channel opens with an on-chain funding transaction, processes high-frequency state updates off-chain during surgery, and settles the final state on-chain. Disputes are resolved within Δ_{disp} blocks.

The Asuu inference record is designed as a discrete event: a single AI prediction, a single diagnostic recommendation, a single liability-stamped decision. Modern intraoperative AI, however, produces a continuous stream of guidance at rates up to 60 frames per second during computer-vision-assisted surgery. Recording each frame as a separate on-chain transaction is infeasible: at 60 FPS over a 4-hour endoscopic skull base procedure, this would produce $60 \times 3600 \times 4 = 864{,}000$ individual inference records, overwhelming the network and incurring prohibitive costs. This section introduces *Surgical State Channels*: temporary, off-chain bilateral channels that aggregate high-frequency intraoperative data into a single on-chain settlement.

47.1 State Channel Formalization

Definition 47.1 (Surgical State Channel)**.** A Surgical State Channel is a tuple $\mathsf{SC} = (\mathscr{V}_s, \mathscr{V}_a, \mathsf{deposit}, S_0, \mathsf{update}, S_n, \Delta_{\text{disp}})$ where:

(i) $\mathscr{V}_s$ is the surgeon's institutional validator node and $\mathscr{V}_a$ is the AI system's operator node;
(ii) $\mathsf{deposit}$ is the collateral locked in the on-chain channel contract at opening, bonding both parties to honest behavior;
(iii) $S_0 = (\mathsf{procedure\text{-}id}, \mathsf{patient\text{-}hash}, t_{\text{start}}, \mathsf{AI\text{-}model\text{-}hash}, \mathsf{initial\text{-}params})$ is the opening state;
(iv) $\mathsf{update} : S_i \times \mathsf{frame}_i \to S_{i+1}$ is the state transition function, where each frame $\mathsf{frame}_i = (\mathsf{AI\text{-}guidance}_i, \mathsf{instrument\text{-}position}_i, t_i)$ contains the AI's real-time trajectory guidance and the surgeon's actual instrument position;
(v) S_n is the final settlement state after n frames;
(vi) Δ_{disp} is the dispute window after channel closure during which either party can challenge the final state.

Definition 47.2 (Rolling State Hash)**.** The state at frame i is computed as a rolling cryptographic accumulator:

$$S_i = H(S_{i-1} \| \mathsf{AI\text{-}guidance}_i \| \mathsf{instrument\text{-}position}_i \| t_i)$$

Both parties independently compute S_i and exchange signatures: $\sigma_s^{(i)} = \mathsf{Sign}(\mathsf{sk}_s, S_i \| i)$ and $\sigma_a^{(i)} = \mathsf{Sign}(\mathsf{sk}_a, S_i \| i)$. A state S_i is considered *consensus-approved* if both signatures are valid.

47.2 Channel Lifecycle

Definition 47.3 (Channel Lifecycle Protocol)**.** The protocol proceeds through three phases:

(1) **Opening.** Both parties submit a joint transaction $\mathsf{tx}_{\text{open}} = (\mathsf{SC.id}, \mathsf{deposit}_s, \mathsf{deposit}_a, S_0, \sigma_s^{(0)}, \sigma_a^{(0)})$ to the on-chain channel contract. The contract locks the deposits and records S_0.

(2) **Off-chain operation.** For each intraoperative frame $i = 1, \ldots, n$:
   (a) The AI system computes $\mathsf{AI\text{-}guidance}_i$ from the current surgical video feed;
   (b) Both parties compute S_i and exchange co-signatures;
   (c) Each party stores the latest consensus-approved state $(S_i, i, \sigma_s^{(i)}, \sigma_a^{(i)})$ locally.

(3) **Cooperative closure.** After the procedure, both parties sign the final state S_n and submit $\mathsf{tx}_{\text{close}} = (S_n, n, \sigma_s^{(n)}, \sigma_a^{(n)})$. The contract verifies both signatures, records S_n on-chain, and releases deposits.

47.3 Security Properties

Theorem 47.4 (State Channel Balance Security)**.** *The Surgical State Channel satisfies* balance security*: for any honest participant $P \in \{\mathcal{V}_s, \mathcal{V}_a\}$ whose latest consensus-approved state is S_i, the channel cannot be finalized in any state S_j with $j < i$ (i.e., no honest party can be forced to accept a stale state).*

Proof. The on-chain dispute mechanism accepts only the state with the highest sequence number i for which both signatures are valid. If party P holds a consensus-approved state $(S_i, i, \sigma_s^{(i)}, \sigma_a^{(i)})$ with $i > j$, then P can submit this state during the dispute window Δ_{disp}. The contract compares sequence numbers and accepts the higher one. An adversary attempting to finalize at S_j must either: (a) prevent P from submitting during Δ_{disp}, which requires network censorship contradicting the liveness assumption; or (b) forge P's signature on a different state at sequence i, which contradicts signature unforgeability (DLP hardness). Since neither is feasible under the security assumptions, balance security holds. □

Theorem 47.5 (Intraoperative Completeness)**.** *The finalized on-chain state S_n constitutes a cryptographic commitment to the complete intraoperative record. Specifically, given S_n and the off-chain frame log $\{(\mathsf{frame}_i)\}_{i=1}^n$, any auditor can verify the integrity of the entire surgical session by recomputing $S'_n = H(\ldots H(H(S_0 \| \mathsf{frame}_1) \| \mathsf{frame}_2) \ldots \| \mathsf{frame}_n)$ and checking $S'_n = S_n$.*

Proof. The rolling hash construction is a sequential composition of collision-resistant hash applications. If any frame frame_i is altered, the recomputed S'_i will differ from S_i (by collision resistance), and this difference propagates to $S'_n \neq S_n$. The on-chain commitment to S_n thus serves as a binding root for the entire off-chain transcript. This is the intraoperative analog of a blockchain's Merkle root committing to all transactions in a block. □

Proposition 47.6 (Throughput and Efficiency)**.** *A Surgical State Channel achieves:*

(i) ***Off-chain throughput****: limited only by local computation; at 60 FPS, S_i computation (one SHA-256 hash plus two ECDSA signatures) takes < 2 ms per frame, supporting real-time operation;*

(ii) ***On-chain footprint****: exactly 2 transactions per procedure (open + close), regardless of procedure duration. A 4-hour procedure at 60 FPS produces 864,000 frames settled with 2 on-chain transactions, an amortized compression ratio of* $432{,}000 : 1$*;*

(iii) ***Liability precision****: the off-chain frame log provides per-frame attribution of AI guidance versus surgeon action, enabling precise liability analysis (Section 26) at temporal granularity of* ~ 16.7 ms*.*

47.4 Legal Risk and Deployability Scope

The 16.7-millisecond temporal granularity that makes Surgical State Channels an interesting cryptographic construction is also the feature that creates substantial medico-legal risk. A 60-Hz record of every intraoperative instrument deviation from AI-recommended trajectory is a discovery-admissible record of every moment the surgeon did or did not follow AI guidance. Plaintiffs' counsel in post-operative adverse-event litigation will subpoena it. Under current U.S. medical-malpractice discovery practice, the existence of such a record combined with any adverse outcome creates a near-automatic path to standard-of-care challenges that today cannot be mounted because the equivalent evidence does not exist.

We do not claim this risk is absent. We claim it is *manageable under a specific deployment architecture*, and we specify what that architecture must provide.

Definition 47.7 (Surgical State Channel Deployment Controls)**.** Deployment of Surgical State Channels in a clinical setting requires all of the following controls, which must be specified in the institution's quality-management system and accepted by the institution's liability carrier before the channel is activated in a production procedure:

(i) **Patient-Controlled Retention.** The off-chain frame log is retained only with the patient's specific informed consent for a specific retention purpose (e.g., post-operative quality review, peer-review M&M conference, research participation). Default retention is *immediate hashing-and-purge*: the rolling hash S_n is computed and anchored on-chain, and the frame log is deleted. The patient's on-chain consent record explicitly authorizes any retention beyond this default.

(ii) **Time-Bounded Retention.** When retention is authorized, it is bounded by a specific maximum interval (e.g., 30 days for post-operative review, 2 years for research consortium participation), after which the frame log is purged and only the on-chain hash remains. This bounds the discovery window.

(iii) **Legal-Hold Protocol.** If an adverse event is reported, the default purge is suspended and the frame log is preserved under legal-hold protocol. This is a standard medico-legal control; the Asuu protocol enforces it through a smart-contract legal-hold flag, not through administrative procedure.

(iv) **Merkle-Root-Only Long-Term Storage.** For any retention beyond the time bound, only the Merkle root of the frame log is retained on-chain; the underlying frame data must be destroyed. This preserves the cryptographic audit trail (integrity of the recorded trajectory can still be verified if the frames are re-disclosed by the patient) while making the frame data unavailable to subpoena.

(v) **Liability-Carrier Approval.** The deploying institution's professional-liability insurer must approve the channel-deployment plan for the specific procedure class and must be able to underwrite the residual discovery risk. Institutions whose carriers decline to underwrite the risk cannot deploy the channel for that procedure class.

Remark 47.8 (Honest Scope of the State Channels Feature)**.** We do not claim that Surgical State Channels are deployable in all intraoperative settings or that they will be adopted rapidly. Realistically, the first production deployments will be in research-consortium settings where patients have explicitly consented to high-granularity surgical-navigation research, the liability carrier has underwritten the residual discovery risk, and the procedure class has a favorable medico-legal risk profile (e.g., robot-assisted prostatectomy in an academic medical center, not emergency trauma surgery in a community hospital). The 432,000:1 compression ratio is an impressive engineering property; it is not a proof that the feature will see broad clinical adoption. The feature exists in the protocol specification because the cryptographic primitive composes correctly with the rest of the protocol and because specific research deployments can use it profitably under the controls of Definition 47.7, not because we expect widespread adoption. An earlier draft of this section overstated the feature's deployment readiness; we retract that overstatement here.

48 Conclusion

We have proposed a system for sovereign medical AI operations without relying on trusted third parties. We started with the crisis of medical AI dependence: the operational fragility, financial extraction, data sovereignty violations, and accountability gaps inherent in the proprietary vendor model. To solve this, we proposed a peer-to-peer network using Proof-of-Clinical-Authority to produce an immutable history of all AI-assisted medical decisions, anchored to the FHIR interoperability standards that hospitals already use.

48.1 The Four Architectural Pillars

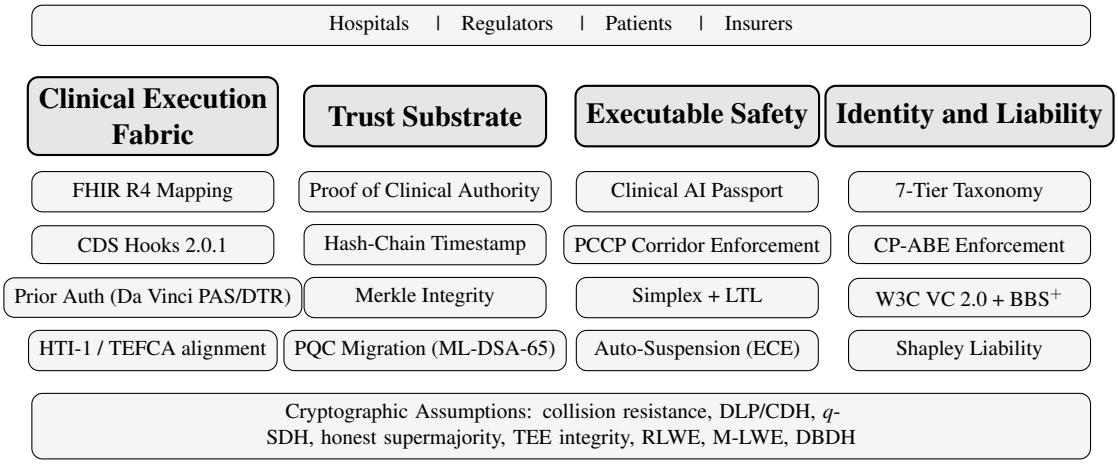


Figure 17: The four architectural pillars of the Asuu Network. Each pillar addresses a distinct stakeholder while resting on the same cryptographic foundation.

The network is structured around four architectural pillars. The **Clinical Execution Fabric** ensures that every AI event maps to standards-native FHIR R4 workflows with CDS Hooks 2.0.1 integration, ONC HTI-1 alignment, and TEFCA compatibility. The Data Embassy guarantees that patient data physically never touches the chain, and the prior authorization workflow demonstrates immediate institutional ROI through Da Vinci PAS and DTR alignment. The **Trust Substrate** provides weighted Byzantine Fault Tolerant consensus with explicit epoch transitions and cross-epoch safety guarantees, post-quantum resilience through composite signature migration, and a security analysis proving that honest behavior is a strict Nash equilibrium. **Executable Safety** replaces static model cards with the versioned Clinical AI Passport, which enforces the PCCP deployment corridor as a provable protocol invariant. Human authority is guaranteed through temporal logic specifications, and automated suspension triggers provide bounded detection delay for calibration drift. The **Identity and Liability** layer formalizes a seven-tier taxonomy of clinical authority enforced through attribute-based encryption, provides patient-controlled selective disclosure through W3C Verifiable Credentials v2.0 with BBS^+ privacy upgrades, and resolves algorithmic liability through Shapley-value decomposition over the immutable audit trail.

The Layer-2 application ecosystem extends these pillars into pharmaceutical supply chain integrity, clinical trial governance, decentralized pharmacovigilance, privacy-preserving pandemic preparedness, One Health cross-species surveillance, and high-frequency intraoperative state channels. Each application shares the same trust substrate and inherits the same formal guarantees.

48.2 The Five Theorem Families

The protocol strictly separates hard protocol guarantees from statistical bounds. Five theorem families form the mathematical backbone of the network:

(1) **Cross-epoch safety** (Theorem 9.26 and the PoCA safety analysis in Section 9): no two honest validators can finalize conflicting histories across epoch transitions.
(2) **End-to-end provenance completeness** (Theorem 5.5): every clinical action has a unique, authenticated provenance path from model registration through inference, human authorization, and final action.
(3) **Update-corridor soundness** (Theorem 23.8): no validator-initiated activation of a model outside the approved PCCP change corridor is possible, while a dedicated Regulatory Override Channel preserves sponsor and regulator authority.
(4) **Privacy-budget enforcement** (Theorem 22.6 and the Rényi differential privacy accountant): cumulative privacy loss across federated learning rounds remains bounded by the institutional budget.
(5) **Safety-suspension bounds** (Theorem 25.9): if calibration drift exceeds the suspension threshold, the protocol suspends the model within a provably bounded number of monitoring epochs.

Each proves a property of the system's control envelope, not a property of the neural network operating within it. The envelope is the theorem. The network is the variable.

Theorem Family	Protocol Layer	Reduces To
(1) Cross-Epoch Safety	Consensus (PoCA)	Honest supermajority $f < n/3$
(2) End-to-End Provenance	Workflow contracts	Collision resistance + EUF-CMA
(3) Update-Corridor Soundness	Clinical AI Passport governance	Smart contract determinism
(4) Privacy-Budget Enforcement	Federated learning (zkFL)	Rényi DP composition
(5) Safety-Suspension Bounds	Runtime monitoring	Hoeffding concentration

Each theorem family bounds the system's control envelope. None claims to bound the neural network's behavior.

Figure 18: The five theorem families of the Asuu Network. Each proves a property of the protocol's response to model risk, not a property of the neural network operating within it.

48.3 Why the Deliberate Structure is the Source of Robustness

The network is robust in its deliberate structure. Validator institutions operate with known identities and staked commitments, providing both cryptographic security and real-world accountability. The Data Embassy model ensures that clinical data remains sovereign while the network provides global coordination. The token economy creates self-sustaining incentives for honest participation, data contribution, and network maintenance without requiring external funding or central administration.

The system does not merely replicate existing medical AI capabilities on a different infrastructure. It fundamentally changes the trust model: from "trust the vendor" to "verify through the network." From "the vendor controls the model" to "the network governs the model." From "data leaves the institution" to "data never leaves the institution." Each transition eliminates a dependency that the current medical AI landscape treats as unavoidable.

The convergence of medical AI and blockchain is no longer speculative. Foundation models surpass physicians on standardized benchmarks. Ambient AI demonstrably reduces clinician burnout. AI-discovered drugs enter late-stage trials. Blockchain supply chain networks process thousands of pharmaceutical verification transactions per second. Zero-knowledge federated learning enables multi-institutional model training without data exposure. The question is not whether healthcare will adopt this convergent infrastructure, but whether that infrastructure will be sovereign or subordinate, open or proprietary, accountable or opaque.

48.4 Machine-Checked Verification Roadmap

Remark 48.1 (Machine-Checked Artifacts)**.** The proofs in this document are stated in the pen-and-paper tradition. To move from mathematical argument to engineering certainty, the following machine-checked artifacts are planned as required deliverables for the protocol's formal specification release:

(1) A state-machine specification for consensus, epoch reconfiguration, emergency halt, and workflow gating, expressed in TLA^+ or a comparable temporal logic framework.
(2) A model-checker artifact verifying safety and liveness under network partitions and view changes.
(3) Machine-checked cryptographic proofs for the credential, commitment, zero-knowledge proof, and proxy re-encryption core, targeting EasyCrypt or CryptoVerif.
(4) Contract-level invariant proofs for RBAC, separation of duty, break-glass expiry, privacy-budget accounting, and update-corridor enforcement, verified through Certora or equivalent smart contract verification tooling.

These artifacts will be published as open-source companions to the protocol implementation.

48.5 Liability Reserve Architecture

Once the core control plane is operational, a natural extension is a *liability reserve layer*. Every model version and clinical scope carries a risk bond posted by the developer or deploying institution, denominated in the native token and held in an on-chain escrow contract. The bond amount is parameterized by the model's autonomy level, patient-facing scope, and historical incident rate. If the Shapley-value liability attribution mechanism (Section 26) adjudicates a failure, the reserve pays out proportionally to the attributed parties. This converts algorithmic liability from a post-hoc legal dispute into a pre-funded, mathematically governed insurance mechanism, aligning the incentives of developers, institutions, insurers, and hospital risk committees.

48.6 Protocol Composition Security

Theorem 48.2 (Protocol Composition Security)**.** *Each sub-component can be modeled as an ideal functionality in the UC framework [15]. By the UC Composition Theorem, replacing ideal functionalities with their protocol realizations preserves security in the composed system.*

Remark 48.3 (Foundational Dependencies)**.** All guarantees reduce to fourteen foundational assumptions, which we organize into three categories: cryptographic hardness assumptions, system-model assumptions, and empirical/regulatory assumptions. This categorization is important because an honest security argument must distinguish what is mathematically hard from what is operationally assumed.

(A) Cryptographic Hardness Assumptions. These are standard primitives whose security rests on believed-hard computational problems:

(A.1) Collision resistance of H (Sections 6, 7, 35, 23, 25, 43, 45, 47).
(A.2) DLP/CDH hardness in bilinear groups (Sections 4, 16, 36, 26, 21, 44, 47).
(A.3) q-SDH assumption for BBS$^+$ credentials (Sections 17, 19, 28, 31, 34).
(A.4) RLWE hardness for homomorphic encryption (Sections 24, 19).
(A.5) M-LWE hardness for ML-DSA signatures (Section 8).
(A.6) DBDH assumption for ciphertext-policy ABE (Sections 28, 29, 21).

(B) System-Model Assumptions. These are assumptions about the deployment environment, not mathematical hardness results. They are load-bearing and deserve explicit scrutiny:

(B.1) *Honest supermajority $f < n/3$.* The Byzantine Fault Tolerant safety argument assumes that at least $2n/3$ of validator weight is controlled by honest institutions. This is not a mathematical assumption; it is a deployment precondition. It is justified operationally by the incentive analysis of Section 14, the economic cost of acquiring one-third of stake-weighted clinical authority, and the legal and reputational consequences of documented fraud. We discuss the limits of this justification, including cases where the assumption plausibly fails (correlated insurance fraud, state-level adversaries), in Section 14.
(B.2) *TEE integrity under documented side-channel limitations.* The exfiltration-resistance guarantees of Section 18 assume that the trusted execution environment (Intel SGX, AMD SEV-SNP, or equivalent) enforces memory isolation and remote attestation against a software adversary. This assumption has known limits: Foreshadow (2018), SGAxe (2020), ÆPIC Leak (2022), and Downfall (2023) have each demonstrated that specific TEE implementations can be compromised by microarchitectural side channels. We treat TEE integrity not as an absolute primitive but as a *defense-in-depth layer* whose failure is detectable through remote attestation version checks and mitigable through degraded-mode operation (Section 18.3). The Asuu protocol does not collapse when a TEE is compromised; it downgrades affected components to require plaintext audit by an independent validator.
(B.3) *Smart contract determinism.* The governance and workflow arguments (Sections 28, 29, 30, 31, 32, 20, 34, 46) assume that smart contract execution is deterministic and atomic. This follows from the specification of the EVM-compatible runtime and is validated by the reference client.
(B.4) *HTLC timelock synchrony bound.* Cross-chain bridging (Section 43) assumes that clock skew between chains does not exceed $\Delta t_A - \Delta t_E$. For chains with sub-minute block times this is routinely satisfied.
(B.5) *Secure boot chain-of-trust.* Medical IoT oracle integrity (Section 45) assumes ROM and eFuse tamper-resistance, which commercial medical-grade hardware certified to Common Criteria EAL5+ or PSA Certified Level 2 is designed to provide.
(B.6) *State channel dispute-window liveness.* High-frequency intraoperative channels (Section 47) assume an honest party can submit a dispute within Δ_{disp} blocks, a standard liveness condition.

(C) Empirical/Regulatory Assumptions. These depend on operational reality rather than cryptography or system design:

(C.1) *Differential privacy mechanism correctness.* The federated learning and population analytics guarantees (Sections 22, 20, 27) depend on correct implementation of the Laplace and Gaussian mechanisms, whose mathematical properties are well-established.

(C.2) *Schelling focal-point uniqueness for clinical standard-of-care.* On-chain arbitration (Section 46) assumes published clinical guidelines provide a shared reference frame for juror coordination. This is empirically supported by the existence and widespread adoption of evidence-based practice standards but remains an empirical claim.

Breaking any single assumption compromises only the components that depend on it, not the entire system. This is the formal analog of the defense-in-depth principle. Assumptions in category (A) are the most load-bearing from a cryptographic perspective and have the strongest theoretical foundations. Assumptions in category (B), particularly (B.1) and (B.2), are where adversarial reality bites hardest and where the protocol's real security margin is established; we treat these with extra care in the security analyses of the relevant sections.

48.7 Deployment Endpoints

In healthcare, mathematics proves integrity and control. It does not prove adoption. We therefore define four hard, pre-registered endpoints for the first pilot deployment. These endpoints are not aspirational. They are measurable, falsifiable, and pre-registered. They are how the project stops being a whitepaper and becomes infrastructure.

(1) **Provenance completeness rate.** The fraction of AI-assisted clinical actions for which the full provenance path π_a (Theorem 5.5) is reconstructable from on-chain records. Target: $\geq 99.9\%$.

(2) **Mean time to suspend.** The average elapsed time from the onset of calibration drift exceeding θ_{suspend} to the model's effective suspension on-chain. Target: ≤ 1 monitoring epoch, consistent with the detection delay bound in Theorem 25.9.

(3) **Prior-authorization turnaround reduction.** The reduction in mean prior-authorization processing time compared to the institution's pre-deployment baseline. Target: $\geq 80\%$ reduction for algorithmically adjudicable claims.

(4) **External-audit preparation time.** The time required to assemble a complete, verifiable audit package for a regulatory inquiry, measured from request to delivery. Target: ≤ 1 business day, compared to the current industry average of weeks to months.

48.8 Closing

Institutions can join and leave the network at will, accepting the finalized chain as proof of what occurred in their absence. They vote with their clinical authority, expressing acceptance of valid blocks by contributing to quorum certificates and rejecting invalid blocks by withholding signatures. Any needed rules and incentives are enforced through the consensus mechanism. The result is a medical AI infrastructure that no single entity owns, no single government controls, and no single point of failure can disable: an infrastructure worthy of the clinical decisions it supports.

References

[1] S. Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” 2008.

[2] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. Arcas, “Communication-Efficient Learning of Deep Networks from Decentralized Data,” In *Proc. AISTATS*, 2017.

[3] FDA, “21 CFR Part 11: Electronic Records; Electronic Signatures,” U.S. Food and Drug Administration, 1997.

[4] European Parliament, “Regulation (EU) 2017/745 on Medical Devices,” *Official Journal of the European Union*, 2017.

[5] European Parliament, “Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act),” *Official Journal of the European Union*, 2024.

[6] FDA, “Marketing Submission Recommendations for a Predetermined Change Control Plan for AI/ML-Enabled Device Software Functions,” December 2024.

[7] ICH, “E6(R3) Good Clinical Practice,” International Council for Harmonisation, 2025.

[8] MediLedger, “DSCSA Pilot Project Final Report,” FDA Drug Supply Chain Security Act Pilot Program, February 2020.

[9] W. Feller, *An Introduction to Probability Theory and Its Applications*, Vol. 1, 3rd ed., Wiley, 1968.

[10] Google Research, “Med-Gemini: Advancing Medical AI with Multimodal Foundation Models,” 2024.

[11] Nature Medicine, “Toward expert-level medical question answering with large language models,” vol 31, January 2025.

[12] JMIR, “Securing Federated Learning With Blockchain in the Medical Domain: A Comprehensive Review,” vol 28, February 2026.

[13] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of distributed consensus with one faulty process,” *JACM*, 32(2):374–382, 1985.

[14] A. Tessaro and C. Zhu, “Revisiting BBS Signatures,” *Proc. EUROCRYPT*, 2023.

[15] R. Canetti, “Universally Composable Security,” *JACM*, 67(5):1–94, 2020.

[16] A. Shamir, “How to share a secret,” *CACM*, 22(11):612–613, 1979.

[17] W. Vickrey, “Counterspeculation, Auctions, and Competitive Sealed Tenders,” *J. Finance*, 16(1):8–37, 1961.

[18] O. Regev, “On lattices, learning with errors, random linear codes, and cryptography,” *JACM*, 56(6):1–40, 2009.

[19] NIST, “FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), FIPS 205 (SLH-DSA): Post-Quantum Cryptographic Standards,” August 2024.

[20] Zama, “Breaking the 1ms barrier for TFHE bootstrapping on GPUs,” September 2025.

[21] J. Kim et al., “NeuJeans: Private Neural Network Inference with Joint Optimization of Noise and Efficiency,” *Proc. CCS*, 2024.

[22] N. Homer et al., “Resolving Individuals Contributing Trace Amounts of DNA to Highly Complex Mixtures Using High-Density SNP Genotyping Microarrays,” *PLoS Genetics*, 4(8), 2008.

[23] FAO, UNEP, WHO, and WOAH, “One Health Joint Plan of Action (2022–2026),” Quadripartite Alliance, 2022.

[24] W3C, “Verifiable Credentials Data Model v2.0,” W3C Recommendation, 15 May 2025.

[25] W3C, “Verifiable Credential Data Integrity 1.0,” W3C Recommendation, 15 May 2025.

[26] W3C, “Data Integrity BBS Cryptosuites v1.0,” W3C Recommendation, 15 May 2025.

[27] W3C, “Bitstring Status List v1.0,” W3C Recommendation, 15 May 2025.

[28] ANSI, “INCITS 359-2012: Role Based Access Control,” American National Standards Institute, 2012.

[29] J. Bethencourt, A. Sahai, and B. Waters, “Ciphertext-Policy Attribute-Based Encryption,” In *Proc. IEEE S&P*, 2007.

[30] International Labour Organization, “International Standard Classification of Occupations: ISCO-08,” ILO, 2012.

[31] U.S. Bureau of Labor Statistics, “2018 Standard Occupational Classification System,” BLS, 2018.

[32] World Health Organization, “National Health Workforce Accounts: A Handbook,” WHO, 2017.

[33] SNOMED International, “SNOMED CT: Systematized Nomenclature of Medicine – Clinical Terms,” 2024.

[34] HL7 International, “FHIR Release 5 (v5.0.0),” 2023.

[35] A. Alharbi et al., “Healthcare RBAC Separation of Duty Analysis,” *PMC*, PMC5836325, 2019.

[36] IBM Security, “Cost of a Data Breach Report 2025,” IBM, 2025.

[37] U.S. Dept. of Health and Human Services, “Breach Portal: Notice to the Secretary of HHS,” 2024.

[38] U.S. Dept. of Justice, “National Health Care Fraud Enforcement Action Results in 193 Defendants Charged,” June 2024.

[39] HHS Office of Inspector General, “Unauthorized Access to ePHI: Insider Threat Analysis,” 2024.

[40] Ancker JS, Edwards A, Nosal S, Hauser D, Mauer E, Kaushal R, “Effects of workload, work complexity, and repeated alerts on alert fatigue in a clinical decision support system,” BMC Medical Informatics and Decision Making, vol. 17, no. 36, 2017.

[41] W. M. P. van der Aalst, “Verification of Workflow Nets,” In *Proc. ICATPN*, LNCS 1248, pages 407–426, 1997.

[42] T. B. Sheridan and W. L. Verplank, “Human and Computer Control of Undersea Teleoperators,” MIT Man-Machine Systems Laboratory, 1978.

[43] M. Fink, “Human Oversight under Article 14 of the EU AI Act,” In *AI Act Commentary: A Thematic Analysis*, Hart-Bloomsbury, forthcoming 2026.

[44] FDA, “De Novo Classification Request for IDx-DR,” DEN180001, April 2018.

[45] HHS OCR, “HIPAA Security Rule To Strengthen the Cybersecurity of Electronic Protected Health Information,” *Federal Register*, 90 FR 898, January 6, 2025.

[46] European Parliament, “Directive (EU) 2024/2853 on liability for defective products (revised),” *Official Journal of the European Union*, 2024.

[47] M. Castro and B. Liskov, “Practical Byzantine Fault Tolerance,” In *Proc. OSDI*, pages 173–186, 1999.

[48] M. Yin, D. Malkhi, M. K. Reiter, G. G. Gueta, and I. Abraham, “HotStuff: BFT Consensus with Linearity and Responsiveness,” In *Proc. PODC*, pages 347–356, 2019.

[49] E. Buchman, J. Kwon, and Z. Milosevic, “The latest gossip on BFT consensus,” arXiv:1807.04938, 2018.

[50] I. E. Kumar, S. Venkatasubramanian, C. Scheidegger, and S. Friedler, “Problems with Shapley-value-based explanations as feature importance measures,” In *Proc. ICML*, 2020.

[51] Court of Justice of the European Union, “Patrick Breyer v Bundesrepublik Deutschland,” Case C-582/14, ECLI:EU:C:2016:779, 19 October 2016.