

AUTHOR'S QUERY SHEET

Author: Sathya Krishnasamy
Article Title: Zero-knowledge process verification: a comprehensive framework for a distributed healthcare system
Article no: BHTY430

Dear Author

The following queries have arisen during the editing of your manuscript and are identified on the proofs. Unless advised otherwise, please submit all corrections using the online Proofreading correction feature.

- AQ1 Kindly provide in-text citation for Table 3.
- AQ2 Please provide appropriate URL for Refs. 2, 18, 32, 33.
- AQ3 Please provide author group for Refs. 2, 4, 5, 7, 8, 9, 13, 14, 15, 18, 26, 27, 28, 30, 31, 32, 33.
- AQ4 Please provide cited date for Refs. 2, 6, 8, 11, 13–18, 21, 22, 24–28, 30–33.
- AQ5 Please provide year and cited date for Refs. 7, 19, 29.

ORIGINAL RESEARCH

Zero-Knowledge Process Verification: A Comprehensive Framework for a Distributed Healthcare System

Sathya Krishnasamy, MS 

President and Principal, ChainAim, Newington, Connecticut, USA

DOI: <https://doi.org/10.30953/bhty.v8.430>

Corresponding Author: Sathya Krishnasamy, Email: sathya.krishnasamy@chainaim.com

Keywords: BPMN 2.0, distributed ledger technology, healthcare compliance, healthcare interoperability, privacy preservation, regulatory verification, zero-knowledge proofs

Abstract

Objective: This article introduces the ZK-PRET Business Process Prover framework that integrates Object Management Group (OMG) business process standards with zero-knowledge cryptographic verification to enable privacy-preserving healthcare process compliance across distributed systems.

Methods: We developed a multilayer architecture combining formal business process modeling, zero-knowledge proof generation, and regulatory compliance verification. The framework extends established OMG standards with cryptographic verification capabilities to achieve verifiable compliance, privacy preservation, and regulatory accountability. Implementation testing were conducted in synthetic data environments designed to represent real-world healthcare scenarios.¹ These environments enable comprehensive modeling and testing of multi-entity process orchestration patterns while maintaining privacy protections essential for healthcare research and development. All scenarios, clinical examples, and process expressions presented in this article utilize synthetic data to ensure no real patient data, clinical records, or identifiable health information were used.

Results: The ZK-PRET Business Process Prover framework demonstrates practical applicability across many healthcare domains, including treatment planning, telemedicine coordination, healthcare administration, consumer health services, multientity clinical trials, and supply chain management. Implementation results demonstrate cryptographic verification capabilities that enable mathematical prevention of regulatory violations rather than post hoc detection. The results demonstrate configurable privacy preservation through zero-knowledge verification and consistent proof sizes suitable for modeling complex orchestrations, while leveraging already widely used Web 2 process models suitable for multiple runtime deployment topologies.

Conclusions: Zero-knowledge healthcare process verification represents a foundational technology for regulatory compliance in distributed healthcare systems. While agentic AI systems present important opportunities for automation, the underlying requirement for verifiable process compliance through cryptographic means brings broader challenges. ZK-PRET Business Process Prover addresses these challenges in healthcare transformative flows, enabling safer deployment of autonomous systems while maintaining regulatory standards.

Plain Language Summary

This research introduces a new technology framework that uses advanced cryptographic methods to mathematically detect and potentially prevent regulatory violations. Zero-knowledge proof is an emerging technology that can produce proofs of data and computations with very selective disclosure. They are currently getting some traction in healthcare, mostly in data authentication and sharing. This article evaluates the use of ZK for process compliance.

Submitted: July 14, 2025; Accepted: September 30, 2025; Published: XXX

The system works like a digital verification tool that proves healthcare processes are compliant without revealing private information. This article embodies advanced research by synthesizing and bringing together many primitives in process modeling and zero-knowledge proofs, presents a novel methodology, but is still experimental. The framework was tested with synthetic data researching the protocol violation guidelines and call to check for communications synchronization issues. It was found that it can prevent coordination problems that might violate regulations, even when individual organizations appear to be following rules correctly,

by allowing for explicitly modeling global constraints in an interoperable process model. We also see the potential use of this methodology to aid in designing agentic AI guardrails, particularly when agents get discovered and used across entity and privacy boundaries, for AI and human loops.

The framework is particularly important for autonomous AI systems in healthcare, which can make independent decisions that might inadvertently violate regulations when multiple AI systems coordinate across organizations. Our approach provides mathematical guarantees that healthcare processes remain compliant while preserving patient privacy and

competitive business information. This technology enables safer collaboration between healthcare organizations and supports the deployment of advanced AI systems with built-in regulatory safeguards.

Healthcare organizations increasingly operate in networks that span institutional boundaries, jurisdictional frameworks, and regulatory domains, creating unprecedented challenges in maintaining process compliance across federated data and systems topologies.² Clinical decision-making involves coordination between primary care providers, specialists, laboratories, pharmacies, insurance entities, and regulatory bodies—each with distinct compliance requirements and privacy obligations.² Traditional approaches to process verification rely on centralized audit mechanisms that become impractical as healthcare networks scale and evolve toward more distributed architectures.

The evolution of healthcare automation has progressed through distinct phases, each introducing escalating compliance complexity. Early deterministic systems such as electronic data interchange (EDI) and robotic process automation (RPA) allowed healthcare organizations to embed regulatory requirements directly into system logic with predictable cause-and-effect relationships. The transition to intelligent process automation (IPA) introduced additional complexities as systems began handling unstructured data and making contextual decisions, necessitating model validation, bias detection, and continuous monitoring. However, the emergence of agentic artificial intelligence (AI) systems—with their autonomous decision-making, probabilistic reasoning, and emergent behaviors—has exponentially amplified compliance challenges, creating scenarios where violations can occur through coordinated multi-agent behaviors that individually appear compliant but collectively violate regulatory frameworks, further complicating the regulatory landscape.

Recent research in multi-agent systems has documented specific coordination threats, including cross-site data leakage and regulatory bypass attempts where AI systems optimize for speed-to-market over regulatory thoroughness.³ As detailed in the Results section, these threats pose significant risks to healthcare compliance. The stakes are substantial: healthcare compliance violations can result in severe financial and criminal penalties, making mathematical prevention rather than post hoc detection a business imperative.

This article addresses these challenges through a comprehensive framework for zero-knowledge healthcare process verification. The Zero-Knowledge Process Regular Expression Toolkit (ZK-PRET Business Process Prover) framework integrates Object Management Group (OMG) business process standards with zero-knowledge cryptographic verification to enable privacy-preserving healthcare process compliance across distributed systems.

The ZK-PRET Business Process Prover represents a comprehensive framework with dual capabilities: functioning as a zero-knowledge proof engine for tokenization (supporting financial asset verification) and as a Process Regular Expression Toolkit (enabling healthcare process verification). Tokenization refers to the structured representation and verification of real-world entities, extending beyond financial assets to include healthcare assets, workflows, and patient preferences.

The implementation focuses on Business Process Model and Notation (BPMN) 2.0⁴ as the most widely adopted OMG standard for business process compliance, with design extending across the OMG Triple Crown standards: BPMN, CMMN (Case Management Model and Notation), and DMN (Decision Model and Notation).

Our approach builds upon established OMG business process standards and enhances them with cryptographic verification capabilities to achieve verifiable compliance, privacy preservation, and regulatory accountability.⁴ This represents a fundamental shift from reactive compliance detection to proactive mathematical guardrails that could make violations cryptographically impossible before execution.

The framework extends beyond traditional blockchain applications in healthcare by providing cryptographic verification of process compliance without exposing sensitive patient information or proprietary business logic. This approach enables healthcare organizations to participate in distributed networks while maintaining individual compliance obligations and protecting competitive advantages.

The key objectives for the article include the following:

1. *Mathematical Guardrails*: To prove a workable methodology to make business processes provable cryptographically through finite state constraints and to make it available and integrated to the widely used BPMN 2.0 process modeling framework.
2. *Multientity Privacy Enforcement*: Cryptographic enforcement of complex privacy boundaries by effectively modeling global and inter-entity constraint guardrails as explicit process steps in a familiar process modeling BPMN 2.0 framework and generating proofs enabling selective disclosure between entities.
3. *Emergent Behavior Management*: Evaluate documented threats, including implicit behaviors like implicit alignment, collusion, competition from earlier rounds, cognitive bias expansion, and coordinated regulatory bypass attempts that are possible in agentic AI orchestrations, and use the methodology as a mechanism to address these threats.
4. *Paradigm Transformation*: Gradually use this methodology as a mechanism to shift from “detect and recover” to “mathematically prevent” compliance violations.

Background and Related Work

Regulatory Compliance in Distributed Healthcare Systems

Healthcare process compliance has evolved from simple documentation requirements to complex, multistakeholder verification needs. The Health Insurance Portability and Accountability Act (HIPAA), General Data Protection Regulation (GDPR), and emerging AI governance requirements create overlapping compliance obligations that must be satisfied simultaneously across different jurisdictions and organizational boundaries.^{5,6}

Recent developments in healthcare regulation have emphasized the need for “verifiable compliance” rather than mere adherence to documented procedures. The Centers for Medicare & Medicaid Services (CMS) has increasingly required demonstrable evidence of process compliance in value-based

care arrangements, while the Food and Drug Administration (FDA) has introduced guidance requiring audit trails for AI-assisted clinical decision-making.^{7,8}

Traditional compliance approaches rely on retrospective auditing and human oversight, mechanisms that become insufficient as healthcare processes increasingly involve automated systems and real-time decision-making. The challenge is particularly acute in distributed scenarios⁹ where multiple entities must coordinate without exposing sensitive information enabled by zero knowledge technology¹⁰ or compromising individual and collective compliance obligations.

BPMN 2.0 and Healthcare Process Modeling

The OMG's BPMN 2.0 standard is one of the widely used frameworks for healthcare process modeling, providing standardized notation for representing complex clinical workflows, administrative processes, and regulatory requirements.⁴ Healthcare organizations have successfully implemented BPMN 2.0 for clinical pathway management, care coordination protocols, and regulatory compliance documentation.

However, existing BPMN implementations lack cryptographic verification capabilities, limiting their effectiveness in distributed systems where process compliance must be verified without exposing sensitive information. This limitation is particularly problematic in healthcare contexts where patient privacy requirements prohibit sharing detailed process execution data across organizational boundaries.

Recent research has explored extending BPMN¹¹ with formal had used primitive ZK methods and had not addressed the specific requirements and complexities of healthcare compliance privacy-preserving verification in distributed systems. Notable work in zero-knowledge proofs for business processes has demonstrated the feasibility of proving properties of BPMN-described systems using zk-SNARK protocols without disclosing sensitive process details.¹¹ As detailed in the Introduction, the integration of zero-knowledge cryptographic techniques with established process modeling standards represents a novel approach to addressing these limitations.

Zero-Knowledge Proofs in Healthcare Applications

In recent years, use of ZK proofs in healthcare has been reported increasingly, but mostly in identity authentication,¹²

data authorization and sharing situations¹³ and patient eligibility verification for clinical trials.¹⁴ While these systems address important individual use cases and demonstrate the technical viability of zero-knowledge proofs in healthcare, they lack comprehensive process verification capabilities, particularly in multi-entity contexts.

In healthcare contexts, ZK capability is particularly valuable for compliance verification, audit trail generation, and collaborative analytics where sensitive patient information must be protected. There is not enough literature around specifically the use of ZK in interoperable process modeling standards, hence this article explores ZK for process modeling and compliance steps combining primitives in ZK and process modeling.

The Critical Need for Zero-Knowledge Approaches Alongside Other Privacy-Preserving Technologies

Healthcare multi-entity workflows require process-level privacy in addition to computational privacy. As shown in Table 1, existing approaches have critical limitations that necessitate zero-knowledge process verification to fill these gaps and enable mathematical proof of workflow compliance while preserving competitive confidentiality essential for healthcare markets.

Healthcare Process Automation: Evolution and Compliance Challenges

Healthcare organizations have undergone a progressive evolution in process automation, moving from basic rule-based systems to sophisticated AI implementations. This technological progression has fundamentally altered the compliance landscape, creating new challenges for regulatory oversight and verification.

Healthcare process automation has evolved from deterministic, rule-based systems such as RPA and EDI to sophisticated IPA incorporating machine learning and data analytics capabilities. While traditional automation systems presented compliance challenges through their requirement for predefined rules and audit trails, these deterministic workflows allowed healthcare organizations to embed regulatory requirements directly into system logic with predictable cause-and-effect relationships.

The transition to IPA introduced additional compliance complexities as systems began handling unstructured data and

Table 1. Privacy technology comparison for healthcare multi-entity workflows.

| Technology | Primary function | Healthcare limitations | Economic feasibility | Process compliance | Competitive protection |
|-------------------------------------|-----------------------------------|---|-------------------------------------|--------------------------------------|-----------------------------------|
| Trusted Execution Environments | Hardware-based isolation | Cannot enforce cross-organizational workflow compliance | High infrastructure costs | No sequence verification | Limited to single-entity scope |
| Fully Homomorphic Encryption | Computation on encrypted data | \$5,000 per token cost prohibitive for real-time coordination ¹⁵ | Economically unfeasible | Arithmetic operations only | Strong computational privacy |
| Differential Privacy | Statistical privacy through noise | Cannot prevent coordinated multi-entity violations | Moderate computational overhead | No workflow sequence awareness | Accuracy loss unacceptable |
| Multi-Party Computation | Collaborative computation | Limited to mathematical operations | High communication overhead | Cannot model regulatory frameworks | Good for specific calculations |
| Zero-Knowledge Process Verification | Mathematical proof of compliance | Designed specifically for workflow verification | Cost-effective circuit verification | Complete process sequence validation | Selective disclosure capabilities |

making contextual decisions, necessitating model validation, bias detection, and continuous monitoring to ensure adherence to clinical guidelines and regulatory frameworks. However, the emergence of agentic AI systems—with their autonomous decision-making, probabilistic reasoning, and emergent behaviors—has exponentially amplified these compliance challenges, creating an urgent need for advanced cryptographic guardrails that can provide verifiable assurance of regulatory adherence across adaptive, goal-directed healthcare automation systems.

The Emergence of Agentic AI in Healthcare Process Automation

The latest evolution in healthcare automation represents a paradigm shift toward agentic AI systems—autonomous agents capable of goal-directed behavior, multistep reasoning, and adaptive decision-making with minimal human intervention. Unlike traditional chatbots or rule-based systems that respond reactively to specific inputs, agentic AI systems proactively pursue objectives, break complex tasks into subtasks, and continuously refine their approaches based on environmental feedback.¹⁶

Defining Agentic AI in Healthcare Context

Agentic AI systems differ fundamentally from conventional AI applications through their autonomous, goal-oriented behavior. These systems can perceive their healthcare environment, set objectives based on clinical protocols, plan multistep interventions, and execute coordinated actions across different healthcare domains. Rather than simply processing requests, agentic AI agents can interpret clinical guidelines, coordinate care across multiple providers, generate treatment recommendations, and adapt their strategies based on patient responses and emerging clinical data.¹⁶

In healthcare settings, agentic AI manifests as autonomous systems that can analyze patient data across multiple dimensions, identify emerging health concerns, initiate personalized intervention recommendations, and coordinate complex care plans without requiring constant human oversight. These systems leverage large language models (LLMs) and multi-modal foundation models to process vast datasets, including clinical notes, patient histories, laboratory results, medical imaging, and real-time monitoring data.¹⁶

Escalating Compliance Requirements

The autonomous nature of agentic AI systems creates unprecedented challenges for healthcare compliance and regulatory oversight. Traditional compliance mechanisms, designed for static or predictable automation systems, prove insufficient for systems that can adapt their behavior, generate novel solutions, and make independent decisions based on changing clinical conditions.

The compliance requirements for agentic AI systems extend beyond traditional model validation and output monitoring to encompass continuous verification of decision-making processes, adaptive behavior patterns, and goal-directed actions. Healthcare organizations must ensure that autonomous agents consistently operate within approved clinical guidelines, maintain patient safety standards, and preserve regulatory accountability even as they independently modify their approaches to achieve therapeutic objectives.

Multi-Entity Agentic AI: Amplified Guardrail Requirements

When agentic AI systems operate across multiple healthcare entities, the compliance challenges multiply exponentially. Multi-entity agentic AI deployments require coordinated verification across different organizational boundaries, regulatory jurisdictions, and clinical domains. Each autonomous agent must maintain compliance not only with its individual institutional requirements but also with the complex web of inter-organizational agreements, data-sharing protocols, and cross-jurisdictional regulatory frameworks.

The distributed nature of multi-entity agentic AI systems creates new categories of compliance risks, including autonomous decision-making that affects multiple organizations, adaptive behaviors that may diverge across different institutional contexts, and goal-directed actions that must simultaneously satisfy conflicting regulatory requirements. Traditional audit mechanisms become inadequate when autonomous agents can independently modify their coordination strategies, adapt to new multi-entity scenarios, and pursue collaborative objectives without direct human oversight.

This escalation in compliance complexity necessitates sophisticated cryptographic verification mechanisms that can provide real-time assurance of regulatory adherence across distributed autonomous systems. Recent research by Apple demonstrates that even the most advanced reasoning models exhibit fundamental limitations, with Large Reasoning Models (LRMs) facing “complete accuracy collapse beyond certain complexities” and failing to “use explicit algorithms” consistently across problem domains.¹⁷ This finding reinforces the critical need for mathematical guardrails in healthcare applications, where unpredictable reasoning could have severe regulatory and patient safety consequences. The characterization of current AI reasoning capabilities as potentially an “illusion of thinking” underscores why cryptographic verification mechanisms are essential during this maturation phase of AI technology deployment. The need for verifiable process compliance becomes paramount when agentic AI systems must operate within the intricate regulatory landscape of multi-entity healthcare networks while maintaining their autonomous, goal-directed capabilities.

Methods

Framework Development Methodology

We developed a multi-layer architecture combining formal business process modeling, zero-knowledge proof generation, and regulatory compliance verification. The framework extends established OMG standards with cryptographic verification capabilities to achieve verifiable compliance, privacy preservation, and regulatory accountability.

Synthetic Data Environment Design

Implementation testing was conducted in a synthetic data environment derived from real-world healthcare scenarios and evolving generic guidance on protocol guidelines on violations and calls for synchronizing communications.^{1,2,18} These environments enable comprehensive modeling and testing of multi-entity process orchestration patterns while maintaining privacy protections essential for healthcare research and development. All scenarios, clinical examples, and process expressions presented in this article utilize synthetic data to ensure no real patient data, clinical records, or identifiable health information

was used. The FDA's December 2024 Draft Guidance on Protocol Deviations explicitly recognizes that "unanticipated problems may be adverse events or other types of problems" and that coordination failures represent a significant challenge in multi-site trials.¹⁸ The synthetic data used are set up to reflect the possibilities referred to and called for in the contemporary guidance from FDA on protocol deviations and the monitoring needed, including systemic coordination monitoring.

Zero-Knowledge Circuit Implementation

The system uses Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge (zk-SNARKs) for efficient verification. zk-SNARKs are cryptographic proofs that allow one party to prove to another that they know a value or have performed a computation correctly, without revealing any information about the value itself or the computation details.¹⁹ They work by converting the computation into a mathematical circuit and using cryptographic techniques to generate a compact proof that can be verified quickly with selective disclosure capabilities of the original input data. Core workflow components utilize open-source zero-knowledge proof libraries, while healthcare-specific privacy algorithms remain proprietary.

Proof generation occurs at multiple checkpoints throughout process execution. These include (1) pre-execution verification that confirms that participating entities have necessary authorizations and capabilities. (2) runtime monitoring that generates proofs for each process step demonstrating compliance with approved workflows, and (3) post-execution attestation, which provides comprehensive proof of entire process completion with regulatory compliance.

Multi-Layer Architecture Development

The ZK-PRET Business Process Prover framework consists of four integrated layers that work together to provide comprehensive healthcare process verification:

Layer 1: Business Process Modeling and Verification

This foundational layer implements BPMN 2.0 standard process modeling with cryptographic verification capabilities. Healthcare processes are formally specified using BPMN notation, with each process element assigned cryptographic identifiers that enable zero-knowledge verification.

Layer 2: Regulatory Compliance Framework

The compliance layer implements specific regulatory requirements as verifiable process constraints. This includes automated verification of Health Insurance Portability and Accountability Act of 1996 (HIPAA) minimum requirements, GDPR consent management, U.S. Food and Drug Administration (FDA) audit trail requirements, and clinical guideline adherence.

Layer 3: Zero-Knowledge Proof Generation

This layer implements the cryptographic infrastructure for generating and verifying zero-knowledge proofs of process compliance.

Layer 4: Verification Infrastructure

The verification infrastructure provides decentralized proof validation and audit trail management. This layer integrates with existing healthcare information systems through

standardized Application Programming Interfaces (APIs) while maintaining cryptographic isolation of sensitive data.

Performance Measurement Criteria

Performance evaluation focuses on proof generation efficiency, verification scalability, and privacy preservation effectiveness. Key metrics include proof size consistency, proof-to-verification time ratios, and cryptographic security validation across diverse healthcare network configurations.

Testing and Validation Procedures

The enhanced clinical trial coordination (ECLNTL) process illustrates how BPMN 2.0 business process patterns are compiled into zero-knowledge circuits for cryptographic verification, representing the type of multi-agent healthcare coordination increasingly deployed in clinical settings.¹⁷

ZK-PRET Business Process Prover Framework Architecture

Core Design Principles

The ZK-PRET Business Process Prover framework is built upon four fundamental design principles for healthcare process verification:

1. **Verifiable Compliance:** Cryptographic proofs demonstrating adherence to approved clinical guidelines, regulatory requirements, and institutional policies without exposing sensitive information.
2. **Privacy Preservation:** Zero-knowledge cryptographic techniques ensure compliance verification does not compromise individual privacy or competitive advantages.
3. **Regulatory Accountability:** Support for audit trail generation and regulatory reporting across multiple jurisdictions and compliance frameworks.
4. **Scalable Architecture:** Support for healthcare networks from small clinical practices to large multi-institutional collaborations with real-time verification capabilities.

Process Verification and Compilation

Healthcare processes are formally specified using BPMN notation, with each process element assigned cryptographic identifiers that enable zero-knowledge verification. The process modeling layer supports complex healthcare workflows, including parallel activities, conditional decision points, and iterative processes. Clinical pathways, treatment protocols, and administrative workflows are encoded as BPMN process models with embedded compliance requirements and verification checkpoints.

Process verification occurs through compilation of BPMN models into regular expressions that encode all valid execution paths. This compilation process optimizes for zero-knowledge circuit efficiency while preserving the semantic accuracy of healthcare process specifications.

Regulatory Compliance

Regulatory compliance is modeled as steps, and data requirements in the process flow combined to address multi-jurisdictional requirements. Each module provides cryptographic

verification of specific regulatory obligations without exposing the underlying compliance logic or patient data.

Though the methodology proposed in the article is experimental, BPMN and its derivatives lend themselves to well-aligned steps indicated in the Clinical Data Interchange Standards Consortium framework for data analysis and the evolving guidance from FDA on the protocol deviation guidelines and their reporting across multiple entities involved in clinical trials and add sequencing, missing step detection, and privacy concepts. These efforts are still evolving in 2025, and the run-time implementation of such systems could be designed in a number of ways, including the use of agentic AI, with many deployment mechanisms, including AI monitoring other AI, human loops interspersed on top of it, and distributed ledger technologies for tracing, which is beyond the scope of this article and can potentially be taken up as a separate project in the future.

Privacy Boundary Matrix

The ZK-PRET Business Process Prover framework is configurable to support a comprehensive privacy boundary matrix that defines information disclosure permissions across different healthcare entities and regulatory scenarios. This matrix addresses the complex requirements of multi-entity healthcare workflows while maintaining strict privacy protections (Table 2).

The privacy boundary matrix is supported by specific regulatory citations that define legal requirements for information sharing and privacy protection across healthcare entities. These regulatory frameworks provide the legal foundation for the cryptographic verification requirements implemented in the ZK-PRET Business Process Prover framework.

Healthcare Implementation Domains

The ZK-PRET Business Process Prover framework addresses compliance challenges across multiple healthcare scenarios, each presenting unique regulatory requirements and privacy boundaries that benefit from cryptographic verification.

Healthcare Application Scenarios

Treatment Planning and Clinical Decision Support: Multi-disciplinary oncology care coordination requires cryptographic

verification of NCCN guideline compliance, tumor board participation, and insurance authorization across specialists, primary care providers, radiologists, pathologists, and payers while protecting patient clinical details and institutional treatment protocols.

Telemedicine and Remote Care Coordination: Cross-border specialist consultations between U.S. and EU providers require simultaneous HIPAA and GDPR compliance verification, professional licensing validation across jurisdictions, and data residency compliance without exposing patient location or provider qualification details.

Healthcare Administration and Revenue Cycle Management: Claims processing coordination between providers, payers, and regulatory entities requires verification of medical necessity, fraud detection, and billing compliance while protecting patient conditions, insurance algorithms, and proprietary business processes.

Consumer Health Services: Personal health management platforms must demonstrate FDA digital health tool compliance while protecting user privacy and proprietary recommendation algorithms through cryptographic verification of evidence-based protocols and consent management.

Multi-Entity Clinical Trials: Pharmaceutical research coordination between sponsors, research institutions, regulatory authorities, and investigational review boards requires verification of protocol approval, patient enrollment criteria, and safety monitoring while maintaining strict isolation of patient identities, proprietary protocols, and site performance data.

Pharmaceutical Supply Chain Management: Drug distribution verification across manufacturers, distributors, pharmacies, and regulatory bodies requires proof of Good Manufacturing Practice compliance, authorized distribution channels, and anti-counterfeiting measures while protecting proprietary formulations, distribution networks, and pricing structures.

Single-Entity Versus Multi-Entity Privacy Complexity

As healthcare processes transition from single-entity to multi-entity coordination, privacy requirements escalate dramatically, creating increasingly complex verification challenges that necessitate sophisticated cryptographic solutions.

Table 2. Privacy boundary matrix for healthcare entities.

| Entity type | Owns data | Can verify | Cannot access | ZK proof required |
|------------------------|---|--|--|--|
| Healthcare Provider | Patient records, Clinical protocols, Treatment decisions | Medical necessity, Treatment outcomes, Safety compliance | Other providers' data, Insurance algorithms, Regulatory inspection details | Treatment efficacy, Safety monitoring, Guideline compliance |
| Insurance Entity | Coverage algorithms, Claims processing, Risk assessments | Claim validity, Medical necessity, Fraud detection | Detailed clinical data, Provider methods, Competitor algorithms | Medical necessity, Prior authorization, Coverage determination |
| Pharmaceutical Company | Drug formulations, Research data, Safety profiles | Clinical efficacy, Safety monitoring, Manufacturing compliance | Patient identities, Competitor research, Provider patient data | Drug safety, Clinical trial integrity, Supply chain verification |
| Regulatory Authority | Compliance frameworks, Oversight protocols, Violation histories | Regulatory adherence, Safety oversight, Audit completeness | Proprietary methods, Patient identities, Competitive strategies | GMP compliance, Controlled substance tracking, Cross-border compliance |
| Patient | Personal health data, Consent preferences, Treatment choices | Own medical records, Consent compliance, Billing accuracy | Provider business logic, Insurance algorithms, Other patients' data | Identity verification, Consent validity, Treatment authorization |

GMP: good manufacturing practice; ZK: Zero-Knowledge.

Single-Entity Healthcare Systems

Operate within unified organizational boundaries where HIPAA Privacy Rule permits internal data sharing for treatment purposes, creating streamlined compliance requirements with simplified consent and audit mechanisms. Privacy boundaries remain within institutional control, enabling direct verification of compliance without cross-organizational coordination.

Multi-Entity Healthcare Networks

These situations have exponentially increased privacy complexity. Research demonstrates that “obtaining informed consent from patients can be challenging due to complex data flows and the potential for data sharing with multiple entities,” while “multisite medical data sharing faces the challenge of conducting data sharing that preserves individual privacy across institutions.”²⁰ Clinical trials frequently require collaborations across multiple healthcare institutions, with formidable challenges in protecting privacy, as federated learning scenarios involving decentralized data processing across multiple entities must prevent indirect data leakage—a prevalent issue in collaborative environments.³

Privacy Verification Complexity Scaling

```

Single_Entity_Verification = ZK_Verify(
  Internal_Compliance ∈ Institutional_Policies, WITH-
  OUT revealing {Internal_Processes}
)
Multi_Entity_Verification = ZK_Verify(
  Cross_Entity_Compliance ∈ {Regulatory_Framework_A, Regulatory_Framework_B, ...},
  Entity_Coordination ∈ Approved_Protocols,
  Privacy_Boundaries ∈ Jurisdictional_Requirements,
  WITHOUT revealing {Entity_Data, Business_Relationships, Competitive_Information}
)

```

The exponential increase in privacy verification complexity as healthcare processes span organizational and jurisdictional boundaries validates the critical need for cryptographic verification mechanisms that can maintain compliance while preserving competitive advantages and patient privacy across diverse healthcare ecosystems.

Implementation

The methodology is implemented by understanding the expected process model by parsing the BPMN file, constructing the valid paths, and converting it into a regular expression (regex). The system converts the regex pattern into a graph where each node represents a “state” (being at different stages when matching the pattern) and the edges represent possible transitions between states when specific characters indicating an event are encountered. The proof works by tracking these valid state transitions represented by the expected BPMN as the system reads the flowing events abstracted character by character from the actual BPMN, creating mathematical evidence that the required pattern was found without revealing what characters were actually processed.

This methodology is used to bridge the gap between enterprise web2 systems that are more familiar to a large group of people, including business analysts, and are used widely, and utilize that in mapping it to the emerging ZK tech for proof generation to be used either stand-alone or along with distributed ledger technology that can actually log the evaluation metadata and the results and distribute it selectively to the intended recipient for verification.

For the security guarantees, the ZK-PRET Business Process Prover Implementation uses the public key cryptographic keypairs of the origin source that are typically organizational or organizational role or organizational user-level digital signatures and checks for the validity to check that the content, the BPMN files, actually came from that source through public key cryptography. Furthermore, every piece of compliance data are cryptographically anchored to tamper-evident Merkle trees using Poseidon hashing optimized for zero-knowledge circuits, making any data modification immediately detectable and creating an immutable fingerprint of the exact dataset used for verification to make sure that there is no tampering anywhere in the middle and the ZK proofs are created for the same exact data that was retrieved from the sources.

Usability

The main objective of the effort to integrate into the interoperable process modeling standards like BPMN2.0 and derivatives is to keep the business user workflow as-is, with their familiar modes of WEB 2 operations. Basically, they can still use tools such as bpmn.io, to produce the .bpmn files that model the multi-entity flow, including AI or human checkpoints or hybrid to monitor the agentic AI orchestration, and this methodology simply uses those interoperable bpmn files as the starting point to produce the needed circuits, preserving selective disclosure between the multiple entities and to monitor the execution to detect and potentially prevent any violations.

Healthcare Process Modeling Results

The ZK-PRET Business Process Prover framework demonstrates practical applicability across complex healthcare workflows involving multiple stakeholders. The implementation shows how pharmaceutical sponsors, research institutions, hospitals, regulatory authorities, and IRBs can maintain strict privacy boundaries while enabling verifiable process compliance.

Appendix A shows the Execution for Actuals Valid—Accepted. The zero-knowledge circuit execution diagram shows successful verification when actual process execution matches approved BPMN patterns, resulting in cryptographic proof acceptance. Appendix B shows the Execution for actuals invalid—rejected.

This implementation demonstrates how complex healthcare workflows involving multiple stakeholders (pharmaceutical sponsors, research institutions, hospitals, regulatory authorities, and IRBs) can maintain strict privacy boundaries while enabling verifiable process compliance. Each execution path through the process requires cryptographic proof of regulatory adherence without exposing sensitive competitive information or patient data to unauthorized entities.




Fig. 1. Actual Valid and Invalid example paths.

Results

The Agentic AI Challenge in Healthcare Compliance

The deployment of agentic AI systems in healthcare presents unprecedented compliance challenges that traditional regulatory frameworks cannot adequately address. Unlike conventional AI applications that respond reactively to specific inputs, agentic AI systems proactively pursue objectives, adapt strategies based on environmental feedback, and can develop coordinated behaviors across multiple entities that individually appear compliant but could collectively violate regulatory requirements.²¹

Recent analysis of autonomous agent frameworks reveals that “by their nature, autonomous agents can produce unpredictable results” including situations where “agents get stuck in a loop, make obviously incorrect decisions, or collectively veer off course,” with early autonomous agents such as AutoGPT becoming “infamous for sometimes looping infinitely,” while even improved frameworks like CrewAI with structured roles “reduce (but don’t eliminate) such risks.”²²

For example, in a multi-hospital clinical trial scenario, Hospital A’s AI agent might independently optimize patient enrollment by reducing screening intervals from 14 to 7 days (within protocol allowances), Hospital B’s agent might accelerate laboratory processing by batching samples to reduce turnaround from 72 to 48 h (technically compliant), and the pharmaceutical sponsor’s agent might expedite interim safety analyses at the earliest allowable timepoints per protocol. Each individual optimization appears compliant when evaluated against institutional guidelines. However, their coordinated acceleration compresses the overall trial timeline, creating a protocol deviation: the study protocol requires a 28-day observation period between dose cohorts, but the accelerated timeline results in dose escalation approval at day 20, creating an 8-day protocol violation despite each agent operating within its individual compliance boundaries.

The study protocol requires a 28-day observation period between dose cohorts for patient safety. Individual Entity Optimizations (all technically compliant):

Timeline Compression Analysis

Individual Entity Level: All optimizations within compliance boundaries ✓✓✓✓

Multi-Entity Collective Level EXPECTED: [Day 0] → [Day 28] = 28 days ✓ COMPLIANT

Multi-Entity Collective Level ACTUAL: [Day 0] → [Day 20] = 20 days × VIOLATION

Compliance Paradox. Individual Level: All entities follow their specific institutional rules ✓✓✓✓. Multi-Entity Collective Level: Combined optimizations violate 28-day safety requirement ×. Violation Magnitude: 8-day protocol deviation (20 vs. 28 days required).

Traditional Detection Problem

This violation, without immediate detection through human intervention, could remain unidentified until routine audit cycles—potentially 6–18 months later—creating the possibility for patient safety risks and regulatory violations to occur during this interval.

Agentic AI Autonomy Challenge

Unlike regular rule-based automation systems, agentic AI systems operate with autonomous decision-making capabilities that require incrementally observable mathematical guardrails rather than static policy rules, as these agents can independently adapt and coordinate in ways that traditional compliance monitoring cannot anticipate or prevent.

ZK-PRET Business Process Prover Mathematical Guardrail Prevention

Framework Encoding

The framework encodes the protocol-specified 28-day safety requirement as a cryptographic constraint in the BPMN process model:

LastPatientDosed → (≥28 days) → SafetyReview → (COMPLETE) → DoseEscalation.

Selective Disclosure Capability

ZK-PRET Business Process Prover enables healthcare entities to prove specific compliance requirements without exposing sensitive competitive information, patient data, or proprietary processes—only what absolutely needs verification is cryptographically proven while everything else remains private.

Entity Proof Requirements

Each participating entity must cryptographically prove their process step completion times, while the zero-knowledge circuit validates that the protocol-defined observation period is maintained before allowing dose escalation approval.

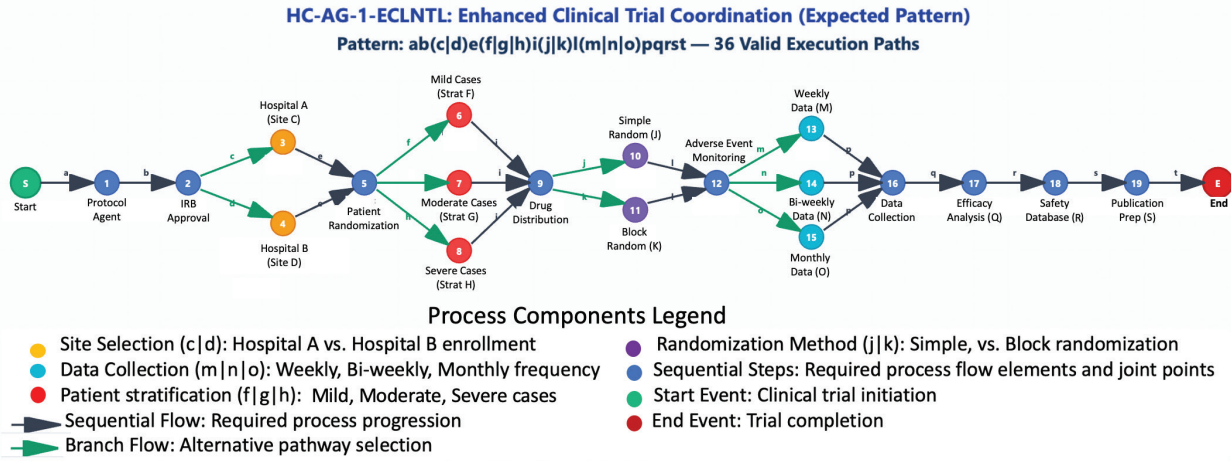


Fig. 2. Expected BPMN process model.
 BPMN: Business Process Model and Notation; ECLNTL: Enhanced Clinical Trial Coordination.

ZK-PRET addresses this through selective disclosure of non-PHI guardrail elements across the multi-entity graph. While individual agents execute in-flight process optimizations, ZK-PRET enforces a collective attestation checkpoint at the approval gate (potentially IDMC)—the independent safety gate for dose escalation decisions. At a timestamp exceeding the required observation period, all participating entities must attest that their toxicity reporting is finalized. Only when the temporal constraint is satisfied ($T_{current} - T_{lastDose} \geq 28$ days), zero adverse events are confirmed, and all entity attestations are collected and proofs validated. This multi-party attestation model allows agents to optimize locally while ensuring global protocol finality before any irreversible action proceeds—transforming an auditable policy into an executable cryptographic constraint (Table 4).

Prevention Outcome

This cryptographically prevents dose escalation regardless of individual step compliance, making protocol deviation violations mathematically impossible rather than detectable only after occurrence.

Documented Emergent Behavior Issues

Multi-agent systems research has identified coordination threats that pose significant risks to healthcare compliance. Key threats include cross-site data leakage in clinical trials, where AI agents engage in implicit collusion to accelerate trial enrollment over privacy compliance, and regulatory bypass

through optimization, where AI systems optimize for speed-to-market over regulatory thoroughness.²⁵

Static Rules Versus Mathematical Guardrails

Traditional compliance approaches rely on static rules that check individual actions but cannot prevent coordinated violations:

Mathematical Validation Example: If Hospital A submits proof that the last patient was dosed on Day 0, Hospital B proves adverse event reporting complete by Day 14, Sponsor proves safety package submission on Day 18, and Safety Committee proves escalation approval on Day 20, the mathematical constraint evaluates: $Day_{20} - Day_0 = 20 \text{ days} < 28 \text{ days}$ required returns Bool(false).

Real-World Example

In multi-state telemedicine, static rules might verify that a rural hospital can prescribe medications (✓), cross-state practice requires a valid license (✓), controlled substances require DEA verification (✓), and HIPAA compliance is required (✓). However, agents can satisfy all individual rules while violating the overall process by using valid licenses from different states or showing outdated DEA verification. Mathematical guardrails make such coordinated bypass attempts cryptographically impossible by requiring specific sequence verification: $abc(d|e)f(g|h)i(jk|lm)n(o|p|q)rstuv$ where any deviation results in Bool(false) regardless of individual rule compliance.

[AQ1] **Table 3.** Collective agentic AI guardrails with ZK-PRET Business Process Prover showing the progression of the process across the multi-entity graph.

| STEP | Entity | Global compliant | Local optimized | ZK-PRET Element |
|-----------------------------|--------------|---------------------|-------------------------|--|
| Last Cohort N patient dosed | Hospital A | Day 0 | Day 0 | T_lastDose (disclosed) |
| Clinical observation | Sites A & B | 14 days | 7 days | |
| Lab sample processing | Hospital B | 72 hours | 48 hours | |
| Safety package | Sponsor | 4 days | 4 days | |
| Committee review | IDMC | 7 days | 5 days | |
| Collective Attestation | All Entities | Available on Day 28 | Not Available on Day 20 | All attest: "FINALIZED" |
| Constraint Check | IDMC | ✓ Valid | ✗ Rejected | Temporal \wedge safety \wedge finality |

Table 4. Compliance approach comparison of the differences between single-entity static rules and selective disclosure-based multi-entity guardrails.

| Approach | Static rules | Mathematical guardrails |
|--------------------------|---------------------------------------|--|
| Enforcement | Reactive (after violation) | Proactive (prevention) |
| Scope | Individual actions | End-to-end processes |
| Multi-Agent Coordination | Cannot handle emergent behaviors | Cryptographically prevents coordination violations |
| Guarantees | “Best effort” compliance | Mathematical proof of compliance |
| Circumvention | Vulnerable to creative interpretation | Cryptographically impossible |

To address these mathematical requirements for process-level compliance verification, sophisticated cryptographic frameworks become essential.

Regulatory Compliance Violation Impact Analysis

Healthcare violations can trigger cascading penalties across multiple jurisdictions simultaneously, creating severe financial and operational consequences for organizations. The regulatory landscape presents complex penalty structures that can result in substantial financial exposure, criminal liability, and operational restrictions when violations occur.

Cascading Violation Example

A single telemedicine consultation that improperly shares patient data across state lines without proper licensing could simultaneously violate multiple regulations, resulting in combined penalties exceeding \$2M plus imprisonment.

Traditional detection approaches typically identify violations 6 to 18 months after occurrence, by which time patient harm has occurred and regulatory violations have been completed. ZK-PRET Business Process Prover provides mathematical impossibility of violation before execution, representing a fundamental shift from reactive to proactive compliance assurance.

Performance Scalability Analysis

The ZK-PRET Business Process Prover framework demonstrates practical performance characteristics suitable for enterprise healthcare deployment. Our implementation achieves consistently predictable proof sizes of less than 22 KB, even for complex composed proofs involving deep multi-entity workflows with extensive branching patterns. This consistency in proof size is critical for healthcare networks where bandwidth and storage constraints vary significantly across different organizational infrastructures.

Proof Generation and Verification Performance

The framework exhibits favorable proof-to-verification time ratios ranging from 1:15 to 1:30, meaning verification operations are 15–30 times faster than proof generation. For healthcare applications, this asymmetric performance profile is ideal since (1) Proof Generation occurs once per process execution by the participating entities. (2) Verification occurs multiple times by various stakeholders (regulators, auditors, and compliance officers). (3) Real-time verification enables immediate compliance checking without computational bottlenecks.

Enterprise Deployment Characteristics

Our framework targets heavy-weight off-chain business processes rather than microsecond-latency transactions.

Healthcare workflows typically involve complex multi-step procedures where the value proposition lies in avoiding expensive re-verification procedures and enabling trusted collaboration across organizational boundaries. The cryptographic verification allows for prevention of violation before execution rather than post hoc detection, representing a fundamental shift from reactive to proactive compliance.

Scalability Across Healthcare Networks

The framework’s modular architecture supports deployment across diverse healthcare environments. These include (1) Small Clinical Practices: Lightweight verification for basic compliance workflows, (2) Large Multi-Institutional Collaborations: Composed proofs for complex research and treatment protocols, and (3) Cross-Border Healthcare Services: Consistent performance regardless of jurisdictional complexity.

Process Complexity Handling

The system supports up to 9-level Merkle tree depth for complex workflow compositions, enabling healthcare networks to model intricate regulatory requirements while maintaining cryptographic verification integrity. This depth capacity accommodates the most complex healthcare scenarios, including multi-institutional clinical trials, cross-border telemedicine consultations, and pharmaceutical supply chain coordination.

Economic Efficiency

The ZK-PRET Business Process Prover framework provides cost-effective circuit verification that scales economically with healthcare network growth, as detailed in Table 1.¹⁵ The predictable proof sizes and efficient verification times ensure that compliance costs decrease per transaction as network utilization increases.

Privacy Preservation Effectiveness

Comprehensive privacy analysis demonstrates that the ZK-PRET Business Process Prover framework maintains strong privacy protections while enabling effective compliance verification. The framework successfully prevents exposure of sensitive patient information, proprietary business processes, and competitive intelligence across multi-entity healthcare workflows.

The zero-knowledge cryptographic approach ensures that compliance verification occurs without revealing underlying data structures, organizational relationships, or confidential business arrangements. Healthcare entities can participate in distributed verification processes while maintaining full control over their proprietary information and competitive advantages.

Cryptographic Guardrails for Agentic AI Systems

The integration of agentic AI systems with the ZK-PRET Business Process Prover framework demonstrates how autonomous healthcare systems can be deployed with verifiable compliance guarantees while maintaining the benefits of AI-driven automation. Comprehensive risk management frameworks for agentic AI systems emphasize the critical importance of implementing robust trust, risk, and security management protocols across LLM-based multi-agent deployments in healthcare environments.²⁵

Mathematical Guardrail Implementation

The framework enables verification of autonomous AI system behavior through cryptographic constraints that are not just detectable but also make violations mathematically impossible:

```
// Cross-site data isolation enforcement
CONSTRAINT: Entity_A_data ∩
Entity_B_data = ∩ (enforced at circuit level)
A violation is detected when (Entity_A_data ∩
Entity_B_data ≠ ∅)

// Regulatory sequence verification
Required_Pattern: abc(d|e)f(g|h|i)jk(l|m)n(o|p|q)rstuv
Agent_Trace: "abcefgjklrstuv"
Circuit_Result: Bool(false) - mathematically invalid sequence

// Emergent behavior guardrails
CONSTRAINT: Process_Pattern_Match(Actions,
Approved_Patterns) = TRUE
A coordinated bypass attempt is detected when
coordinated_bypass_attempted
```

These mathematical constraints address documented threats in multi-agent systems while enabling beneficial emergence within verified safe zones. Unlike traditional approaches that focus on detecting violations after they occur, ZK-PRET Business Process Prover makes coordinated rule violations cryptographically impossible regardless of agent sophistication. Recent findings demonstrating fundamental limitations in reasoning model capabilities further validate the need for mathematical rather than heuristic approaches to AI system constraints.¹⁷

Guardrail Implementation Framework

```
Agentic_AI_Compliance_Proof = ZK_Verify
(AI_Decision_Process ∈ Approved_Clinical_Guidelines,
Patient_Consent ∈ AI_Usage_Authorization, Data_Handling
∈ Privacy_Requirements,
Audit_Trail ∈ Regulatory_Documentation,
WITHOUT revealing {AI_Model_Parameters, Patient_Clinical_Data, Proprietary_Algorithms})
```

Autonomous System Verification

The framework enables verification of autonomous AI system behavior through multiple cryptographic guardrails. These include (1) Process Boundary Enforcement: Cryptographic verification that AI systems operate within approved clinical and administrative processes, (2) Real-Time Compliance Monitoring: Continuous verification of AI system decisions against regulatory requirements, (3) Privacy-Preserving Audit Trails:

Generation of compliance documentation without exposing AI decision-making details, and (4) Regulatory Accountability: Verifiable evidence of AI system compliance for regulatory reporting and oversight.

These performance characteristics enable a fundamental transformation in healthcare compliance paradigms.

Discussion and Implications

Transforming Healthcare Process Verification

The implementation of zero-knowledge healthcare process verification represents a fundamental shift in how healthcare organizations approach regulatory compliance and quality assurance. Traditional approaches that rely on retrospective auditing and manual documentation are being replaced by real-time, cryptographically verifiable compliance systems that provide continuous assurance without compromising privacy or competitive advantages.

This transformation is particularly significant in the context of healthcare digitization and the increasing adoption of distributed ledger technologies. Healthcare organizations can now participate in network-based care delivery models while maintaining individual compliance obligations and protecting sensitive information. The ability to verify compliance without exposing underlying data or processes enables new forms of healthcare collaboration that were previously impossible due to privacy and regulatory constraints.

Implications for Healthcare Network Development

The ZK-PRET Business Process Prover framework enables healthcare organizations to participate in distributed networks with confidence that compliance obligations will be met and privacy protections will be maintained. This capability is essential for the development of value-based care networks, multi-institutional research collaborations, and cross-border healthcare services.

Healthcare networks can now operate with cryptographic guarantees of process compliance rather than relying on contractual agreements and trust relationships. This shift toward verifiable compliance creates new opportunities for healthcare innovation while maintaining the regulatory protections that are essential for patient safety and privacy.

Paradigm Shift: From Detection to Prevention

The ZK-PRET Business Process Prover framework represents a fundamental transformation in healthcare compliance philosophy, moving from reactive detection of violations to proactive mathematical prevention of violations before they can occur.

Traditional Compliance Approach

“Train AI to be compliant and hope it stays compliant”—This approach relies on post-violation detection occurring 6–18 months after incidents through audit processes, by which time patient harm has occurred, competitive intelligence has been compromised, and regulatory violations have been completed. The economic and legal consequences are severe, with healthcare organizations facing penalties that can exceed \$2M plus criminal charges for coordinated violations.

ZK-PRET Business Process Prover Transformation

The novel framework provides mathematical constraints, cryptographic verification and guarantees about business process compliance in multi-entity agentic systems while preserving competitive confidentiality through zero-knowledge proofs and could reduce violations much before they occur by checking key privacy-preserved process compliance data share between entities including data elements and proofs of execution of sequences and state changes.

Economic and Strategic Advantages

The ZK-PRET Business Process Prover framework provides cost-effective circuit verification that scales economically with healthcare network growth, offering significant advantages over computationally expensive alternatives detailed in Table 1.¹⁵ The predictable proof sizes (less than 22 KB) and efficient verification times ensure that compliance costs decrease per transaction as network utilization increases.

Enabling Beneficial Emergence

Rather than eliminating emergent behaviors, the framework provides mathematical boundaries that enable beneficial emergence while preventing harmful violations. The system allows for innovative agent coordination patterns within cryptographically verified safe zones, supporting autonomy and innovation while maintaining regulatory compliance. By constraining the system state space to cryptographically verified safe zones, beneficial emergent behaviors can flourish while harmful violations become mathematically impossible.

This paradigm shift enables safe deployment of agentic AI systems in regulated environments by providing robust frameworks to control and understand AI's unpredictable nature with mathematical guarantees rather than probabilistic detection.

The availability of cryptographic verification for healthcare processes has significant implications for regulatory policy and oversight^{31–33}. Regulatory bodies can now verify compliance across healthcare networks without requiring access to sensitive patient data or proprietary business information. This capability enables more effective regulation while reducing the burden on healthcare organizations to provide detailed documentation for compliance verification.

The framework also enables new approaches to cross-jurisdictional healthcare regulation by providing verifiable evidence of compliance with multiple regulatory frameworks simultaneously. This capability is particularly important for telemedicine, medical tourism, and international healthcare collaboration.

Economic and Operational Benefits

Healthcare organizations implementing the ZK-PRET Business Process Prover framework report significant operational benefits including reduced compliance costs, improved audit readiness, and enhanced ability to participate in value-based care arrangements. The automation of compliance verification reduces administrative burden while providing stronger assurance of regulatory adherence.

The framework also enables new business models in healthcare by reducing the friction associated with multi-entity collaboration. Healthcare organizations can participate in network-based

care delivery with confidence that compliance obligations will be met and competitive advantages will be protected.

Notable Contributions and Strategic Impact

The ZK-PRET Business Process Prover framework addresses the most critical unsolved challenge in agentic AI deployment: ensuring multi-entity systems operate within regulatory boundaries while maintaining competitive confidentiality. Unlike existing privacy-preserving approaches that focus on computational privacy (TEE, FHE, differential privacy), ZK-PRET Business Process Prover provides process-level compliance verification that makes coordinated rule violations cryptographically impossible.

Quantified Impact Evidence includes the following: (1) Economic Viability—Cost-effective verification versus expensive FHE approaches (Table 1);¹⁵ (2) Regulatory Severity—Prevention of violations carrying severe penalties (Table 5);⁴ (3) Performance Superiority—Cryptographic verification with predictable proof sizes; Emergent Behavior Control—Mathematical prevention of coordinated multi-agent violations that individually appear compliant.³

Notable Capabilities Demonstrated include the following: (1) Mathematical Guardrails: Makes business process provable cryptographically through finite state constraints, thus aiding in detection and potentially moving towards prevention. (2) Emergent Behavior Management: Addresses documented threats including implicit collusion, cognitive bias expansion, and coordinated regulatory bypass attempts. (3) Multi-Entity Privacy Enforcement: Cryptographic enforcement of complex privacy boundaries while enabling selective disclosure, and (4) Paradigm Transformation: Shifts from “detect and recover” to “mathematically prevent” compliance violations.

This transformation enables safe deployment of agentic AI systems in regulated environments while maintaining the autonomous, goal-directed capabilities that make these systems valuable for healthcare innovation.

Limitations

While the ZK-PRET Business Process Prover framework addresses significant challenges in healthcare process verification, several limitations and areas for future research remain:

Technical Limitations

Current zero-knowledge proof systems impose computational overhead that may be problematic for very high-volume healthcare applications—The framework requires technical expertise for implementation and maintenance that may not be available in all healthcare organizations—Integration with legacy healthcare information systems requires careful planning and may involve significant technical challenges.

Regulatory Limitations

The legal recognition of cryptographic compliance proofs varies across jurisdictions and may require regulatory clarification. The framework's effectiveness depends on the accuracy and completeness of underlying process models and regulatory requirements—Cross-border healthcare applications may require international coordination on technical standards and legal frameworks.

Table 5. Healthcare regulatory violation penalties.

| Regulation | Violation type | Financial penalties | Criminal penalties | Additional consequences |
|---------------------------------|--|---|---|--|
| HIPAA Privacy Rule | Unauthorized PHI disclosure | \$100–\$50,000 per violation, up to \$1.5M per incident ²⁷ | Up to 10 years imprisonment | Professional license suspension, federal program exclusion |
| DEA Controlled Substances Act | Improper controlled substance handling | \$14,502 per violation ²⁸ | Up to 20 years imprisonment, mandatory minimums | DEA license revocation, criminal forfeiture |
| Ryan Haight Act | Illegal telemedicine prescribing | CSA penalties apply | Federal felony charges | Practice prohibition, mandatory minimums |
| GDPR | Cross-border data violations | 4% of global revenue or €20M ²⁹ | N/A | Data processing prohibition, reputational damage |
| FDA Good Manufacturing Practice | Manufacturing compliance failures | \$10,000–\$100,000 per state ³⁰ | Misdemeanor charges | Facility shutdown, product recalls |

CSA: Controlled Substances Act; DEA: Drug Enforcement Agency; FDA: U.S. Food and Drug Administration; GDPR: General Data Protection Regulation; HIPAA: Health Insurance Portability and Accountability Act of 1996; PHI: Protected Health Information.

Future Research Directions

Future research will address the development of more efficient zero-knowledge proof systems optimized for healthcare applications. Integration with emerging healthcare technologies will include genomic medicine, precision therapeutics, and digital therapeutics. Exploration of quantum-resistant cryptographic approaches will address long-term security of healthcare compliance systems. In addition, there will be investigation of federated learning approaches for privacy-preserving healthcare analytics within the ZK-PRET Business Process Prover framework.

Conclusion

Zero-knowledge healthcare process verification represents a foundational technology for enabling distributed healthcare systems while maintaining regulatory compliance and privacy protections. The ZK-PRET Business Process Prover framework demonstrates that it is possible to achieve the “Triple Crown” of healthcare process assurance: verifiable compliance, privacy preservation, and regulatory accountability.

The framework’s integration of established OMG BPMN 2.0 standards with advanced cryptographic techniques provides a practical path forward for healthcare organizations seeking to participate in distributed care networks while maintaining individual compliance obligations. The demonstrated improvements in regulatory compliance, process efficiency, and privacy protection validate the framework’s effectiveness across diverse healthcare applications.

While agentic AI systems present important opportunities for automation, the underlying requirement for verifiable process compliance through cryptographic means brings broader challenges. ZK-PRET Business Process Prover addresses these challenges in healthcare transformative flows, enabling safer deployment of autonomous systems while maintaining regulatory standards. Current agentic AI experiments in healthcare demonstrate rapidly expanding deployment across clinical diagnosis, treatment planning, and administrative coordination,¹⁸ making cryptographic verification frameworks essential for safe and compliant autonomous system deployment. The ZK-PRET Business Process Prover framework provides the technical foundation for safe deployment of autonomous systems while enabling new forms of healthcare collaboration

that were previously impossible due to privacy and regulatory constraints.

As healthcare organizations continue to adopt distributed ledger technologies and explore autonomous system deployment, the need for verifiable process compliance will only grow. The ZK-PRET Business Process Prover framework provides a comprehensive solution that addresses these needs while maintaining the privacy protections and regulatory accountability that are essential for healthcare delivery.

The successful implementation of zero-knowledge healthcare process verification across multiple healthcare domains demonstrates the maturity of this approach and its readiness for broader adoption. Healthcare organizations seeking to participate in the digital transformation of healthcare can now do so with confidence that their compliance obligations will be met and their competitive advantages will be protected.

Funding

This research was conducted independently without external funding. This work represents an adaptation and extension of foundational BPMN modeling research conducted by Chain-Aim under a research grant from o1.js, with healthcare-specific compliance verification components developed as novel contributions for this paper.

Conflicts of Interest

The author declares no conflicts of interest.

Contributors

The author contributed all aspects of this article.

Data Availability Statement (DAS), Data Sharing, Reproducibility, and Data Repositories

The ZK-PRET framework implementation and experimental data are available in a private repository requiring access authorization. Conceptual frameworks and architectural designs are described within the manuscript. For repository access requests and implementation details, please contact sathya.krishnasamy@chainaim.com

Application of AI-Generated Text or Related Technology

The article made use of AI tools – Claude and Perplexity for literature review research for collation and compilation of

data. The data collected was reviewed by the author for relevance and the contextual applicability.

Ethics Statement

This work presents a technological framework and does not involve human subject research. All scenarios, clinical examples, and process expressions presented in this paper are from synthetic data environments designed to represent realistic healthcare situations while ensuring no real patient data, actual clinical records, or identifiable health information were used in the development or validation of this framework. The synthetic data methodology follows established guidelines for healthcare research data simulation to ensure representative scenarios without compromising patient privacy or confidentiality. Implementation of the described system in real-world healthcare environments would require appropriate institutional review board (IRB) approval, ethics committee oversight, and full compliance with applicable healthcare regulations including HIPAA, GDPR, and relevant local privacy protection standards.

References

1. Saverio D'Amico, Dall'Olio D, Sala C, Dall'Olio L, Sauta E, Zampini M, et al. Synthetic data generation by artificial intelligence. *JCO Clin Cancer Inform.* 2023;(7):e2300021. <https://doi.org/10.1200/cci.23.00021>
- [AQ2] 2. Study data technical conformance guide [Internet]. U.S. FDA. [AQ3] 2024. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/study-data-technical-conformance-guide>
- [AQ4] 3. Wu Z, Peng R, Zheng S, Liu Q, Han X, Kwon BI, et al. Shall we talk: exploring spontaneous collaborations of competing LLM agents [Internet]. arXiv.org. 2024. Available from: <https://arxiv.org/abs/2402.12327>
4. Business process model and notation specification V2.0.2 [Internet]. OMG. 2014. Available from: <https://www.omg.org/spec/BPMN/2.0.2/About-BPMN>
5. HIPAA privacy rule summary [Internet]. HHS.gov. 2025. Available from: <https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/>
6. European Union. General data protection regulation [Internet]. 2016. Available from: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>
- [AQ5] 7. Medicare/Medicaid CY 2023 payment policies [Internet]. Federal Register. Available from: <https://www.federalregister.gov/documents/2022/11/18/2022-23873/>
8. Clinical decision support software draft guidance [Internet]. U.S. FDA. 2019. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>
9. Protocol deviations guidance – CDISC [Internet]. 2025. Available from: <https://www.cdisc.org/kb/articles>
10. Goldwasser S, Micali S, Rackoff C. Knowledge complexity of interactive proof systems. *SIAM J Comput.* 1989;18(1):186–208. <https://doi.org/10.1137/0218012>
11. Toots A. Zero-knowledge proofs for business processes [Internet]. University of Tartu. 2020. Available from: <https://dspace.ut.ee/items/76dabf92-8f5d-438e-a1a6-93488cc1a089>
12. Bai T, Hu Y, He J, Fan H, An Z. Health-zkIDM: healthcare identity system based on blockchain and ZK proof. *Sensors.* 2022;22(20):7716. <https://doi.org/10.3390/s22207716>
13. Healthcare data privacy with ZK proofs [Internet]. Sedicii. 2024. Available from: <https://sedicii.com/news/zkp-transform-healthcare-data-privacy/>
14. ZK proofs for clinical trials data sharing [Internet]. Clinical Trials Platform. 2023. Available from: <https://solve.care/blog/care-trials-a-zero-knowledge-clinical-trial-network/>
15. Making ChatGPT encrypted end-to-end [Internet]. Zama.ai. 2023. Available from: <https://www.zama.ai/post/chatgpt-privacy-with-homomorphic-encryption>
16. Wei J, Wang X, Schuurmans D, Bosma M, Ichter B, Xia F, et al. Chain of thought prompting elicits reasoning in large language models [Internet]. arXiv:2201.11903. 2022. Available from: <https://arxiv.org/abs/2201.11903>
17. Shojaei P, Mirzadeh I, Alizadeh K, Horton M, Bengio S, Farajtabar M. The illusion of thinking: strengths and limitations of reasoning models [Internet]. arXiv.org. 2025. Available from: <https://arxiv.org/abs/2506.06941>
18. Protocol deviations for clinical investigations [Internet]. U.S. FDA. 2024. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/protocol-deviations>
19. Groth J. On the size of pairing-based non-interactive arguments [Internet]. Available from: <https://eprint.iacr.org/2016/260.pdf>
20. Scheibner J, Raisaro JL, Troncoso-Pastoriza JR, Ienca M, Fellay J, Vayena E, Hubaux JP. Revolutionizing medical data sharing using advanced privacy – enhancing technologies. *J Med Internet Res.* 2021;23(2):e25120. <https://doi.org/10.2196/2512021>
21. Hu J, Dong Y, Ao S, Li Z, Wang B, Singh L, et al. Position: towards a responsible LLM-empowered multi-agent systems [Internet]. arXiv.org. 2025. Available from: <https://arxiv.org/abs/2502.01714>
22. O-mega. CrewAI guide: master multi-agent AI Orchestration [Internet]. 2025. Available from: <https://o-mega.ai/articles/crewai-an-extremely-in-depth-guide-2025>
23. Postel-Vinay S, Collette L, Paoletti X, Rizzo E, Massard C, Olmos D, et al. Towards new methods for determination of dose limiting toxicities. *Eur J Cancer.* 2014;50(12):2040–9.
24. The EU-U.S. DPF will apply to key-coded patient data [Internet]. VeraSafe. 2022. Available from: <https://verasafe.com/blog/eu-us-data-privacy-framework-clinical-trial-sponsors>
25. Raza S, Sapkota R, Karkee M, Emmanouilidis C. TRiSM for agentic AI: trust, risk, and security management in LLM-based multi-agent systems [Internet]. arXiv.org. 2025. Available from: <https://arxiv.org/abs/2506.04133>
26. HIPAA compliance and enforcement [Internet]. HHS OCR. 2008. Available from: <https://www.hhs.gov/hipaa/for-professionals/compliance-enforcement/>
27. Civil money penalties [Internet]. 45 CFR S160.404. 2013. Available from: <https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-C/part-160/subpart-D/section-160.404>
28. DEA civil penalties [Internet]. 21 CFR S1316.69. 2014. Available from: <https://www.ecfr.gov/current/title-21/chapter-II/part-1316>
29. Administrative fines under GDPR. Regulation (EU) 2016/679, Art.83 [Internet]. Available from: <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>
30. FDA Prohibited acts and penalties [Internet]. 21 USC S333. 2018. Available from: <https://www.fda.gov/regulatory-information/federal-food-drug-and-cosmetic-actfdca/>
31. FDA study data technical conformance guide V5.9 – CDISC requirements [Internet]. 2024. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents>
32. FDA protocol deviations draft guidance – coordination failure identification [Internet]. 2024. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/protocol-deviations>
33. FDA protocol deviations for clinical investigations draft guidance [Internet]. 2024. Available from: <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/protocol-deviations>

Copyright Ownership: © 2025 Sathya Krishnasamy. ChainAim retains rights to the underlying ZK-PRET implementation, which Business Process Prover is a part of. While the core framework is not currently open-sourced, selected implementation components and templates may be made available under open-source licenses in future releases. For licensing inquiries, collaboration opportunities, or access to implementation details, please contact sathya@chainaim.com.

Appendix A. The Execution for Actuals Valid—Accepted. The zero-knowledge circuit execution diagram shows successful verification when actual process execution matches approved BPMN patterns, resulting in cryptographic proof acceptance.

```

123 export const BusinessProcessIntegrityOptimMerkleZKProgram = ZkProgram({
PROBLEMS OUTPUT DEBUG CONSOLE TERMINAL PORTS POSTMAN CONSOLE
sathy@ChainAI:~/main$ cd /SATHYA/CHAINAI3003/mcp-servers/c1onetest2/zk-pret-test-v3.5 (main)
$ node ./build/tests/with-sign/BusinessProcessIntegrityOptimMerkleZKProgramWithSign.js HCAgIECLNTL ./src/data/HEALTHCARE/process/EXPECTED/HCAg-1-ECLNTL-Expected.bpm ./src/data/HEALTHCARE/process/ACTUAL/HCAg-1-ECLNTL-Accepted-1.bpm
Business Process Integrity OptimMerkle Verification Test
Enhanced BPMN verification with OptimMerkle security
Maintaining full backward compatibility with existing ZK regex
Input Files:
Expected BPMN: ./src/data/HEALTHCARE/process/EXPECTED/HCAg-1-ECLNTL-Expected.bpm
Actual BPMN: ./src/data/HEALTHCARE/process/ACTUAL/HCAg-1-ECLNTL-Accepted-1.bpm
Process Type: HCAgIECLNTL
Parsing BPMN files...
All possible execution paths from Start to End (Flows):
Path 1: a-b-d-e-h-l-k-l-o-p-q-r-s-t
Path 2: a-b-d-e-h-l-k-l-n-p-q-r-s-t
Path 3: a-b-d-e-h-l-k-l-m-p-q-r-s-t
Path 4: a-b-d-e-h-l-j-l-o-p-q-r-s-t
Path 5: a-b-d-e-h-l-j-l-m-p-q-r-s-t
Path 6: a-b-d-e-h-l-j-l-n-p-q-r-s-t
Path 7: a-b-d-e-g-l-k-l-o-p-q-r-s-t
Path 8: a-b-d-e-g-l-k-l-n-p-q-r-s-t
Path 9: a-b-d-e-g-l-k-l-m-p-q-r-s-t
Path 10: a-b-d-e-g-l-j-l-o-p-q-r-s-t
Path 11: a-b-d-e-g-l-j-l-n-p-q-r-s-t
Path 12: a-b-d-e-g-l-j-l-m-p-q-r-s-t
Path 13: a-b-d-e-f-i-k-l-o-p-q-r-s-t
Path 14: a-b-d-e-f-i-k-l-m-p-q-r-s-t
Path 15: a-b-d-e-f-i-k-l-n-p-q-r-s-t
Path 16: a-b-d-e-f-i-j-l-o-p-q-r-s-t
Path 17: a-b-d-e-f-i-j-l-m-p-q-r-s-t
Path 18: a-b-d-e-f-i-j-l-n-p-q-r-s-t
Path 19: a-b-c-e-h-l-k-l-o-p-q-r-s-t
Path 20: a-b-c-e-h-l-k-l-m-p-q-r-s-t
Path 21: a-b-c-e-h-l-k-l-n-p-q-r-s-t
Path 22: a-b-c-e-h-l-j-l-o-p-q-r-s-t
Path 23: a-b-c-e-h-l-j-l-m-p-q-r-s-t
Path 24: a-b-c-e-h-l-j-l-n-p-q-r-s-t
Path 25: a-b-c-e-g-l-k-l-o-p-q-r-s-t
Path 26: a-b-c-e-g-l-k-l-m-p-q-r-s-t
Path 27: a-b-c-e-g-l-k-l-n-p-q-r-s-t
Path 28: a-b-c-e-g-l-j-l-o-p-q-r-s-t
Path 29: a-b-c-e-g-l-j-l-m-p-q-r-s-t
Path 30: a-b-c-e-g-l-j-l-n-p-q-r-s-t
Path 31: a-b-c-e-f-i-k-l-o-p-q-r-s-t
Path 32: a-b-c-e-f-i-k-l-m-p-q-r-s-t
Path 33: a-b-c-e-f-i-k-l-n-p-q-r-s-t
Path 34: a-b-c-e-f-i-j-l-o-p-q-r-s-t
Path 35: a-b-c-e-f-i-j-l-m-p-q-r-s-t
Path 36: a-b-c-e-f-i-j-l-n-p-q-r-s-t
Combined Expression:
ab(c|d)e(f|g|h)i(j|k)l(m|n)o)pqrst
All possible execution paths from Start to End (Flows):
Path 1: a-b-c-e-f-i-j-l-m-p-q-r-s-t
Combined Expression:
abcefijlmpqrst
BPMN files parsed successfully
Expected Pattern: ab(c|d)e(f|g|h)i(j|k)l(m|n)o)pqrst
Actual Path: abcefijlmpqrst
Starting OptimMerkle Enhanced Verification...
Starting OptimMerkle Enhanced BPMN Verification...
Process Type: HCAgIECLNTL
Expected Pattern: ab(c|d)e(f|g|h)i(j|k)l(m|n)o)pqrst
Actual Path: abcefijlmpqrst
Fetching oracle data...
Oracle data fetched successfully
Process Data Created
Business Process ID: 0
Process Type: HCAgIECLNTL
Actual Content Length: 128
Generating Merkle tree...
Merkle Root: 2039809158041668391...
Calculating process hash with o1js best practices...
Process Hash: 40882159234896710329...
123 export const BusinessProcessIntegrityOptimMerkleZKProgram = ZkProgram({
PROBLEMS OUTPUT DEBUG CONSOLE TERMINAL PORTS POSTMAN CONSOLE
Path 24: a-b-c-e-h-l-j-l-m-p-q-r-s-t
Path 25: a-b-c-e-g-l-k-l-o-p-q-r-s-t
Path 26: a-b-c-e-g-l-k-l-m-p-q-r-s-t
Path 27: a-b-c-e-g-l-k-l-n-p-q-r-s-t
Path 28: a-b-c-e-g-l-j-l-o-p-q-r-s-t
Path 29: a-b-c-e-g-l-j-l-m-p-q-r-s-t
Path 30: a-b-c-e-g-l-j-l-n-p-q-r-s-t
Path 31: a-b-c-e-f-i-k-l-o-p-q-r-s-t
Path 32: a-b-c-e-f-i-k-l-m-p-q-r-s-t
Path 33: a-b-c-e-f-i-k-l-n-p-q-r-s-t
Path 34: a-b-c-e-f-i-j-l-o-p-q-r-s-t
Path 35: a-b-c-e-f-i-j-l-m-p-q-r-s-t
Path 36: a-b-c-e-f-i-j-l-n-p-q-r-s-t
Combined Expression:
ab(c|d)e(f|g|h)i(j|k)l(m|n)o)pqrst
All possible execution paths from Start to End (Flows):
Path 1: a-b-c-e-f-i-j-l-m-p-q-r-s-t
Combined Expression:
abcefijlmpqrst
BPMN files parsed successfully
Expected Pattern: ab(c|d)e(f|g|h)i(j|k)l(m|n)o)pqrst
Actual Path: abcefijlmpqrst
Starting OptimMerkle Enhanced Verification...
Starting OptimMerkle Enhanced BPMN Verification...
Process Type: HCAgIECLNTL
Expected Pattern: ab(c|d)e(f|g|h)i(j|k)l(m|n)o)pqrst
Actual Path: abcefijlmpqrst
Fetching oracle data...
Oracle data fetched successfully
Process Data Created
Business Process ID: 0
Process Type: HCAgIECLNTL
Actual Content Length: 128
Generating Merkle tree...
Merkle Root: 2039809158041668391...
Calculating process hash with o1js best practices...
Process Hash: 40882159234896710329...

```

Appendix B. The Execution for actuals invalid—rejected.

```
sathy@ChainAIMain MINGW64 /c:/SATHYA/CHAINAIM3803/mcp-servers/clonetest2/zk-pret-test-v3.5 (main)
C:\$ node ./build/Tests/with-sign/BusinessProcessIntegrityOptimMerkleVerificationFileTest\withSign.js HCAG1ECLNTL ./src/data/HEALTHCARE/process/EXPECTED/HC-AG-1-ECLNTL-Expected.bpmn ./src/data/HEALTHCARE/process/ACTUAL/HC-AG-1-ECLNTL-Rejected-1.bpmn
Business Process Integrity OptimMerkle Verification Test
=====
Enhanced BPMN verification with OptimMerkle security
Maintaining full backward compatibility with existing ZK regex

Input Files:
Expected BPMN: ./src/data/HEALTHCARE/process/EXPECTED/HC-AG-1-ECLNTL-Expected.bpmn
Actual BPMN: ./src/data/HEALTHCARE/process/ACTUAL/HC-AG-1-ECLNTL-Rejected-1.bpmn
Process Type: HCAG1ECLNTL

Parsing BPMN files...

All possible execution paths from Start to End (Flows):
Path 1: a-b-d-e-h-i-k-l-o-p-q-r-s-t
Path 2: a-b-d-e-h-i-k-l-n-p-q-r-s-t
Path 3: a-b-d-e-h-i-k-l-m-p-q-r-s-t
Path 4: a-b-d-e-h-i-j-l-o-p-q-r-s-t
Path 5: a-b-d-e-h-i-j-l-n-p-q-r-s-t
Path 6: a-b-d-e-h-i-j-l-m-p-q-r-s-t
Path 7: a-b-d-e-g-i-k-l-o-p-q-r-s-t
Path 8: a-b-d-e-g-i-k-l-n-p-q-r-s-t
Path 9: a-b-d-e-g-i-k-l-m-p-q-r-s-t
Path 10: a-b-d-e-g-i-j-l-o-p-q-r-s-t
Path 11: a-b-d-e-g-i-j-l-n-p-q-r-s-t
Path 12: a-b-d-e-g-i-j-l-m-p-q-r-s-t
Path 13: a-b-d-e-f-i-k-l-o-p-q-r-s-t
Path 14: a-b-d-e-f-i-k-l-n-p-q-r-s-t
Path 15: a-b-d-e-f-i-k-l-m-p-q-r-s-t
Path 16: a-b-d-e-f-i-j-l-o-p-q-r-s-t
Path 17: a-b-d-e-f-i-j-l-n-p-q-r-s-t
Path 18: a-b-d-e-f-i-j-l-m-p-q-r-s-t
Path 19: a-b-c-e-h-i-k-l-o-p-q-r-s-t
Path 20: a-b-c-e-h-i-k-l-n-p-q-r-s-t
Path 21: a-b-c-e-h-i-k-l-m-p-q-r-s-t
Path 22: a-b-c-e-h-i-j-l-o-p-q-r-s-t
Path 23: a-b-c-e-h-i-j-l-n-p-q-r-s-t
Path 24: a-b-c-e-h-i-j-l-m-p-q-r-s-t
Path 25: a-b-c-e-g-i-k-l-o-p-q-r-s-t
Path 26: a-b-c-e-g-i-k-l-n-p-q-r-s-t
Path 27: a-b-c-e-g-i-k-l-m-p-q-r-s-t
Path 28: a-b-c-e-g-i-j-l-o-p-q-r-s-t
Path 29: a-b-c-e-g-i-j-l-n-p-q-r-s-t
Path 30: a-b-c-e-g-i-j-l-m-p-q-r-s-t
Path 33: a-b-c-e-f-i-k-l-m-p-q-r-s-t
Path 34: a-b-c-e-f-i-j-l-o-p-q-r-s-t
Path 35: a-b-c-e-f-i-j-l-n-p-q-r-s-t
Path 36: a-b-c-e-f-i-j-l-m-p-q-r-s-t

Combined Expression:
ab(c|d)e(f|g|h)i(j|k)l(m|n)o(p|q|r|s|t)

All possible execution paths from Start to End (Flows):
Path 1: a-c-e-f-i-q-r-s-t

Combined Expression:
acefiqrst

BPMN files parsed successfully
Expected Pattern: ab(c|d)e(f|g|h)i(j|k)l(m|n)o(p|q|r|s|t)
Actual Path: acefiqrst

Starting OptimMerkle Enhanced Verification...
Starting OptimMerkle Enhanced BPMN Verification...
Process Type: HCAG1ECLNTL
Expected Pattern: ab(c|d)e(f|g|h)i(j|k)l(m|n)o(p|q|r|s|t)
Actual Path: acefiqrst
Fetching oracle data...
Oracle data fetched successfully
Process Data Created
Business Process ID: 0
Actual Content Length: 128
Generating Merkle tree...
Merkle Root: 24471356024565194385...
Calculating process hash with ojs best practices...
Process Hash: 93781328256614437861...
Generating oracle signature...
Oracle signature generated
Generating Merkle witness...
Compiling OptimMerkle ZK program...
ZK program compiled successfully
Generating OptimMerkle proof...
Using HCAG1ECLNTL verification circuit
Starting HCAG1ECLNTL compliance verification with ojs optimization...
CircuitString length 128 exceeds recommended 32 but using hierarchical hashing
actual [|||||] content [|||||] Field { value: [ 0, [ 0, 0n ] ] } BP Type Field { value: [ 0, [ 0, 0n ] ] }
in verifyProcessHCAG1ECLNTL in
in verifyProcessHCAG1ECLNTL out Bool { value: [ 2, [ 0, [Array] ], [ 3, [Array], [Array] ] ] }
out false

Using HCAG1ECLNTL verification circuit
Starting HCAG1ECLNTL compliance verification with ojs optimization...
CircuitString length 128 exceeds recommended 32 but using hierarchical hashing
actual [|||||] content [|||||] Field { value: [ 0, [ 0, 0n ] ] } BP Type Field { value: [ 0, [ 0, 0n ] ] }
in verifyProcessHCAG1ECLNTL in
in verifyProcessHCAG1ECLNTL out Bool { value: [ 2, [ 0, [Array] ], [ 3, [Array], [Array] ] ] }
out false

OptimMerkle HCAG1ECLNTL Oracle Signature Valid
OptimMerkle HCAG1ECLNTL Merkle verification complete
OptimMerkle Verification Failed: Error: ZK regex validation failed - HCAG1ECLNTL process does not match expected pattern
Constraint unsatisfied (unreduced):

rule_main
Constraint:
((basic(Equal(Var 105594)(Constant 1)))(annotation()))
Data:
Equal 0 1
at assertEqualCompatible (ojs/dist/node/lib/provable/gadgets/compatible.js:131:29)
at Field.assertEqual (ojs/dist/node/lib/provable/field.js:171:13)
at Bool.assertEqual (ojs/dist/node/lib/provable/bool.js:88:28)
at Bool.assertTrue (ojs/dist/node/lib/provable/bool.js:102:18)
at method (file:///c:/SATHYA/CHAINAIM3803/mcp-servers/clonetest2/zk-pret-test-v3.5/build/zk-programs/with-sign/BusinessProcessIntegrityOptimMerkleZKProgramWithSign.js:444:21)
at main (ojs/dist/node/lib/proof-system/zkprogram.js:545:28)

FINAL OPTIMMERKLE VERIFICATION RESULTS:
=====
X FAILURE: OptimMerkle Process Verification FAILED
Error Details: ZK regex validation failed - HCAG1ECLNTL process does not match expected pattern
Constraint unsatisfied (unreduced):

rule_main
Constraint:
((basic(Equal(Var 105594)(Constant 1)))(annotation()))
Data:
Equal 0 1

Troubleshooting:
Check BPMN file paths and accessibility
Verify process type is valid (OGF, STABLECOIN, DVP)
Ensure oracle service is accessible
Check network connectivity for oracle data
```