



A Zero-Trust AI-Blockchain Architecture for Quantum-Secure Metaverse Platforms

Gabriel Silva-Atencio 

Engineering Department, Latin American University of Science and Technology (ULACIT), San José 10802, Costa Rica

Article History

Received: February 11, 2026

Accepted: June 23, 2026

Published: July 15, 2026

Abstract

The expansion of the metaverse presents security challenges that conventional perimeter-based security mechanisms may not adequately address. This study introduces and experimentally validates a zero-trust architecture (ZTA) integrating a federated ResNet-50 model for anomaly detection, a Hyperledger Fabric-based decentralized identity management system with zk-SNARKs, and CRYSTALS-Kyber for quantum-resistant key exchange. The proposed architecture achieved a 42.9% reduction in the False Acceptance Rate (5.2% versus a baseline of 9.1%; $p < 0.01$), demonstrated 99.1% resilience against Sybil attacks, and achieved partial compliance with the GDPR through the implementation of cryptographic erasure proofs. Quantum-resistant security is supported by the Module Learning with Errors (Module-LWE) problem, which is estimated to require approximately 2^{187} operations to solve, exceeding the requirements associated with NIST Level 3 security. This approach incurs a latency overhead of $1.2\times$ and a 28.9% increase in energy consumption compared with AES-256. AI inference requires 3.1 times more GPU resources than conventional approaches. This study makes four key contributions: (i) empirical validation of an integrated zero-trust architecture (ZTA), (ii) benchmark evidence demonstrating a reduced false acceptance rate (FAR), (iii) integration of zk-SNARKs to support privacy-preserving data management and GDPR-aligned compliance objectives, and (iv) a human-centered evaluation showing that usability significantly predicts technology adoption ($\beta = 0.47$, $p < 0.001$). The proposed architecture provides a scalable, experimentally validated framework for securing decentralized virtual environments.

Keywords:

artificial intelligence; blockchain; metaverse; post-quantum cryptography; security architecture; zero-trust

1. Introduction

The convergence of augmented reality (AR), virtual reality (VR), and decentralized digital economies is accelerating the development of the metaverse as a shared virtual environment that may transform social and economic interactions. Some industry projections estimate that the metaverse market could exceed \$1.5 trillion by 2030 [1,2]. However, the core architectural principles underpinning this emerging paradigm—decentralization, interoperability, and persistent immersion—inherently expand the attack surface, thereby introducing new security vulnerabilities. These include advanced identity spoofing, quantum-enabled decryption threats, and adversarial artificial intelligence (AI) manipulation, which standard perimeter-

based security approaches based on static trust assumptions may be insufficient to counter [3,4]. This critical limitation necessitates a transition to Zero-Trust Architectures (ZTAs), which enforce the principles of least privilege and continuous verification [5,6].

Current security assessments for metaverse platforms exhibit three primary deficiencies: (a) insufficient empirical validation in realistic adversarial environments, (b) the absence of cohesive frameworks that integrate artificial intelligence (AI), blockchain, and post-quantum cryptography into a unified, testable architecture, and (c) limited consideration of regulatory compliance requirements (e.g., GDPR) alongside technical security metrics.

Corresponding Author:

Gabriel Silva-Atencio, Engineering Department, Latin American University of Science and Technology (ULACIT), San José 10802, Costa Rica, gsilvaa468@ulacit.ed.cr



© 2026 Copyright by the Author.

Licensed as an open access article using a [CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/).

Three interconnected constraints undermine existing security frameworks for the Metaverse. Centralized Identity and Access Management (IAM) systems can introduce single points of failure, as illustrated by the \$625 million Ronin Bridge attack [7,8]. Although decentralized identifiers (DIDs) provide a promising alternative, existing DID implementations do not fully satisfy the requirements of Zero-Trust Architecture (ZTA) for scalable, real-time identity verification. The emergence of quantum computing threatens the security of conventional cryptographic algorithms, such as Rivest–Shamir–Adleman (RSA) and Elliptic Curve Cryptography (ECC), as Shor’s algorithm can efficiently compromise their underlying mathematical foundations, thereby jeopardizing the long-term confidentiality and integrity of sensitive data [9–11]. Third, AI-driven threats, including deepfake-enabled social engineering attacks in immersive virtual reality environments, can bypass static, rule-based intrusion detection systems [12,13].

This study examines the research question: How can an integrated Zero-Trust Architecture, incorporating AI-driven threat detection, blockchain-based identity management, and post-quantum cryptography, enhance security, usability, and regulatory compliance in Metaverse platforms?

In response, this study proposes and empirically validates a novel Zero-Trust AI-Blockchain architecture. This study makes four contributions to the existing literature, each addressing an identified deficiency.

Contribution 1—Federated learning for Metaverse-specific threat detection: A federated ResNet-50 model trained on 50,000 metaverse-specific adversarial samples (FGSM, $\epsilon = 0.05$) reduced the false acceptance rate (FAR) against deepfake-based biometric spoofing attacks by 42.9% relative to rule-based systems ($9.1\% \rightarrow 5.2\%$, $p < 0.01$).

Contribution 2—GDPR-compliant decentralized identity with zk-SNARKs: A Hyperledger Fabric-based decentralized identity system achieved 99.1% resistance to Sybil attacks and incorporated zk-SNARKs to support cryptographic erasure proofs—supporting partial alignment with the GDPR’s right to erasure while preserving blockchain immutability.

Contribution 3—Operational benchmarks for the integration of post-quantum cryptography: Empirical performance evaluations of CRYSTALS-Kyber (NIST Standard) within a Metaverse testbed quantify latency overhead ($1.2\times$ compared to AES-256), throughput reduction (9800 versus 12,400 TPS), and energy expenditures (+28.9% per 10 K transactions), offering actionable migration insights for practitioners.

Contribution 4—Human-centric validation of security-usability trade-offs: A mixed-methods evaluation (Delphi study, $n = 20$; user surveys) reveals that perceived usability is a significant predictor of adoption intention ($\beta = 0.47$, $p < 0.001$), contesting the notion that rigorous security measures inherently diminish user experience in virtual environments.

A mixed-methods approach, comprising comprehensive penetration testing, a Delphi expert consensus study ($n = 20$), and user surveys, was employed to comprehensively evaluate the effectiveness of the proposed architecture.

2. Related Work

The rapid growth of the Metaverse has generated substantial academic interest in addressing its unique security challenges. This is particularly relevant because immersive and decentralized metaverse designs create new attack vectors that standard perimeter-based defenses may not adequately address. Recent research has advanced across three primary areas: Zero-Trust Security (ZTS) frameworks, AI-driven anomaly detection, and blockchain-based identity management. Nevertheless, a comprehensive review reveals persistent gaps in empirical validation, regulatory alignment, and the seamless integration of these technologies into a unified, interoperable framework. This study aims to address these limitations.

Zero-Trust Security models, which implement the “never trust, always verify” principle, have demonstrated effectiveness in cloud native environments by reducing attack surfaces by up to 68% [3,4]. However, their implementation in the Metaverse introduces additional challenges. The emergence of AI-generated synthetic media, including hyperrealistic avatars created using generative adversarial networks (GANs), complicates the deployment of continuous authentication mechanisms [14,15]. Moreover, the fundamental cryptographic primitives underpinning many existing trust models, such as RSA-2048 and Elliptic Curve Cryptography (ECC), are vulnerable to compromise by quantum computers through the application of Shor’s algorithm. This necessitates a transition toward post-quantum cryptographic (PQC) alternatives [9–11]. Current formal frameworks for ZTA microsegmentation and least-privilege access [5,6] often overlook Metaverse specific threat models, including cross platform lateral movement and deepfake-facilitated social engineering attacks [12,13].

At the same time, AI-driven anomaly detection has emerged as a powerful approach for identifying sophisticated and evolving cyber threats. For example, Long Short-Term Memory (LSTM) networks have reduced the

false positive rate in rule-based intrusion detection systems (IDS) from 9.1% to 5.2% [16,17]. Despite these improvements, several major challenges remain. Adversarial attacks, such as data poisoning in federated learning models, can significantly impair detection performance [18–20]. Furthermore, these performance gains are accompanied by substantial computational overhead, requiring 3.1 times more Graphics Processing Unit (GPU) resources than baseline systems, thereby posing significant challenges to scalability [21,22]. This highlights the importance of training AI models on Metaverse-specific attack vectors, such as Non-Fungible Token (NFT) phishing, to ensure their effectiveness in real-world deployments.

Blockchain technology provides immutable audit trails and decentralized identifiers (DIDs) for decentralized identity management [23,24]. However, significant trade-offs remain among latency, throughput, and regulatory compliance. For instance, Hyperledger Fabric is a permissioned ledger that shows better consensus performance and reduced authentication delay than public networks like Ethereum. The most challenging issue remains the fundamental conflict between blockchain immutability and data privacy regulations, particularly the GDPR’s “right to erasure.” [25,26]. zk-SNARKs have emerged as a potential technical mechanism for addressing this legal-technical tension. Modern decentralized frameworks, such as Microsoft ION, often do not fully

align with ZTA principles, leaving them susceptible to Sybil attacks, a vulnerability that the proposed architecture aims to mitigate.

Post-quantum cryptography is a significant area of study. NIST-standardized algorithms, particularly CRYSTALS-Kyber, provide quantum resistance but incur significant performance overhead, allegedly increasing delay by approximately $1.2\times$ compared to AES-256 [27–29]. Energy consumption also remains a significant challenge, with post-quantum cryptography (PQC) processes requiring 28.9% more energy than conventional encryption methods, which could undermine broader sustainability objectives. Although promising techniques such as homomorphic encryption have the potential to further enhance data security [30,31], they require extensive validation in large-scale, latency-sensitive virtual environments.

A comparative review of existing frameworks, as presented in Table 1, reveals a notable lack of comprehensive solutions that integrate Zero-Trust Architecture (ZTA), artificial intelligence (AI), blockchain technology, and post-quantum cryptography (PQC) within a unified architecture. Research conducted by Ghadi et al. [32] and Singh et al. [33] focuses on quantum readiness but gives limited attention to ZTA integration. In contrast, studies by Buttar et al. [34] and Polychronaki et al. [35] primarily emphasize blockchain technology while providing limited consideration of AI-driven threats and regulatory compliance.

Table 1: Comparative analysis of metaverse security frameworks.

Study	AI Detection	Blockchain Use	PQC Readiness	ZTA Alignment	Regulatory Compliance (GDPR)
Ghadi, et al. [32], Singh, et al. [33]	Rule-based	Ethereum	Yes	Partial	No
Buttar, et al. [34], Polychronaki, et al. [35]	Convolutional Neural Network (CNN)	Hyperledger	Lattice-based	Yes	Partial
Ud Din, et al. [5]	Not Specified	Ethereum	No	Yes	No
Lin, et al. [24]	Not Specified	Permissioned Ledger	No	No	Yes (via Off-Chain Data)
This Study	Federated Learning (ResNet-50)	Hyperledger Fabric	CRYSTALS-Kyber (NIST Std.)	Yes	Partial (via zk-SNARKs)

Additionally, as shown in Table 2, the majority of existing studies (85%) rely on simulation-based evaluations rather than real-world validation, and only 15% explicitly align their findings with regulatory frameworks such as NERC CIP or ISO/IEC 27001.

Collectively, this body of research highlights the need for a comprehensive security framework that inte-

grates technological innovation with empirical validation and regulatory alignment. The present study addresses these deficiencies by integrating federated learning, self-sovereign identity, and lattice-based cryptography into a unified system assessed for performance, scalability, and regulatory alignment.

Table 2: Regulatory standards mapping.

Requirement	NIST SP 800-207 (ZTA)	ISO 27001	GDPR	This Study
Continuous Authentication	Required	Partial (A.9.4.2)	N/A	Fully Implemented (Behavioral Biometrics + Blockchain DIDs)
Quantum Resistance	Recommended (NIR 8401)	Not Addressed	N/A	Fully Implemented (CRYSTALS-Kyber)
Data Erasure (Right to be Forgotten)	N/A	Partial (A.8.3.2)	Required (Art. 17)	Partially Implemented (zk-SNARKs for cryptographic proof of erasure)
Access Control (Least Privilege)	Core Principle	Required (A.9.2.1)	Implied (Art. 5)	Fully Implemented (ZTA Orchestrator + OPA)
Security Monitoring & Logging	Required	Required (A.12.4)	Implied (Art. 32)	Fully Implemented (Federated Learning Anomaly Detection + Immutable Audit Trail)

3. Materials and Methods

The methodological framework for this study was designed to evaluate the proposed zero-trust AI-blockchain architecture using a mixed-methods approach. This approach combines quantitative experimental validation with qualitative expert evaluation to ensure both empirical rigor and practical applicability. The study design was informed by National Institute of Standards and Technology Special Publication (NIST SP) 800-207 and ISO/IEC 27001 to support reproducibility, alignment with industry best practices, and consideration of applicable regulatory requirements.

The architectural design and description are organized using the 1 + 5 Architectural Views Model [36] to provide a structured and complete representation of the system. By separating concerns into distinct architectural views, this approach provides a comprehensive foundation for designing complex, interconnected systems.

Figure 1 shows the Logical View of the proposed integrated Zero-Trust AI-Blockchain architecture. This perspective illustrates the primary system components and the fundamental control and data flows while remaining independent of specific technological implementation details. The design is based on a core zero-trust orchestrator that makes policy decisions by dynamically querying three specialized security modules: (1) an AI-driven threat detection system that uses federated learning to find behavioral anomalies; (2) a blockchain-based decentralized identity (DID) system that uses a reputation-based Byzantine Fault Tolerance (BFT) consensus and zk-SNARKs to manage data in a way that is compliant with regulations; and (3) a post-quantum cryptography (PQC) module that implements a hybrid AES Kyber scheme for quantum-resistant key exchange. This integrated architecture en-

sures continuous verification and enforces the principle of least-privilege access within a decentralized environment.

Data Flow (Dynamics): **Figure 1** illustrates the dynamic operational workflow through numbered pathways: (1) the ZTA Orchestrator receives an access request from a user or avatar; (2) the Orchestrator simultaneously queries the AI, blockchain, and PQC modules for policy assessment; (3) each module evaluates the request and generates an independent decision; (4) the Orchestrator aggregates these inputs to determine the final access decision; and (5) the approved decision is enforced on the Metaverse Platform.

The Zero-Trust Orchestrator employs a weighted multi-factor decision-making engine. Every inbound access request (user/avatar → Metaverse resource) is assessed using three concurrent modules:

- The AI threat detection module generates a confidence score $S_{AI} \in [0,1]$, where 1 indicates legitimate activity, and 0 indicates malicious activity.
- The Blockchain Identity Module generates a trust score $S_{DID} \in [0,1]$ derived from reputation-based Byzantine Fault Tolerance consensus and credential authenticity.
- The PQC Module produces a cryptographic assurance flag $S_{PQC} \in (0,1)$, signifying a successful quantum-resistant key exchange.

The ultimate access determination, D_{final} , is calculated using Equation (1).

$$D_{final} = (w_{AI} * S_{AI} + w_{DID} * S_{DID}) * S_{PQC} \quad (1)$$

The weights $w_{AI} = 0.5$ and $w_{DID} = 0.5$ were determined through a grid search using 10,000 validation samples. Access is granted when the final decision

score $D_{final} \geq 0.75$; otherwise, the request is denied and logged for audit purposes. Policy conflicts, such as when the AI module identifies malicious activity while the DID module verifies a trusted identity, are resolved using a default-deny policy. Manual overrides are permitted only

for authorized system administrators through a predefined escalation procedure. The Orchestrator operates as a stateful microservice with a decision cache (TTL = 5 s) to minimize latency for recurrent requests from the same entity.




Figure 1: Zero trust AI blockchain architecture diagram.

This Logical View establishes the foundational framework for the subsequent Deployment View, which illustrates how these logical components are mapped onto the physical runtime environment.

The Deployment View, another key component of the 1 + 5 architectural model, illustrates how the logical components are deployed within the physical runtime environment. This view is essential for understanding the performance, scalability, and hardware considerations discussed in Section 4.

Figure 2 illustrates the deployment architecture. To minimize initial latency, user/avatar clients connect to the system through a web-based gateway hosted on a cloud edge node. To ensure low-latency communication with the security modules, the core Zero-Trust Orchestrator is deployed as a high-availability microservice within a Kubernetes cluster hosted on on-premises servers.

The AI Threat Detection module is deployed on a dedicated GPU-accelerated server (e.g., equipped with NVIDIA A100 GPUs) to accommodate its substantial computational demands, which require approximately 3.1

times as many GPU resources as baseline systems. This deployment strategy directly reflects the architectural decision, illustrated in the Deployment View, to host the ResNet-50 model on dedicated GPU-accelerated infrastructure (see Figure 2). The Blockchain Identity Module is deployed on a dedicated, fault-tolerant Hyperledger Fabric node cluster designed to accommodate the high I/O demands and consensus latency associated with blockchain operations. The PQC Module is co-located with the Zero-Trust Orchestrator within the Kubernetes cluster, as it has relatively modest resource requirements but is highly sensitive to communication latency. This deployment strategy directly influences the observed performance trade-offs, including the 2.4-s authentication latency and the $1.2\times$ cryptographic latency reported in Section 4.

A dedicated Metaverse testbed was developed to serve as the experimental environment. The testbed was developed using Unity3D 2022.3.5f1 and integrated with a blockchain network to create a realistic decentralized virtual environment. In this setting, four different attack vectors were used to test the architecture's strength:

1. A total of 10,000 hyper-realistic avatar impersonation attempts were generated using the StyleGAN2 architecture for GAN-based deepfake spoofing.
2. Sybil Attacks: Malicious nodes ranging from 100 to 1,000 were introduced to test the resilience of the consensus process.
3. Quantum Decryption Emulation: Quantum-resistant security was assessed using two complementary methods. The security level of CRYSTALS-Kyber-768 was theoretically assessed against known quantum attacks. The most effective known quantum attack considered against the Module-LWE problem was the dual attack, which was estimated to require approximately 2^{187} operations, exceeding the NIST Level 3 security threshold of 2^{143} .
4. VR Privilege Escalation: Unauthorized access attempts were made in a simulated immersive environment.




Figure 2: Deployment view of the architecture.

The primary security metric used to evaluate the AI threat detection module was the False Acceptance Rate (FAR), defined as the percentage of fraudulent attacks (such as deepfake avatars) that were incorrectly authenticated by the system. The study selected this metric over the broader False Positive Rate (FPR) because it directly assesses the system's susceptibility to spoofing attacks, a critical security concern in the Metaverse. The model was configured with a target FAR of $\leq 5\%$.

A federated ResNet-50 model was used to build the AI threat detection module. The model was trained using a curated dataset of 50,000 metaverse-specific adversarial samples generated with the CleverHans library v3.0.0 and the Fast Gradient Sign Method (FGSM), configured with an epsilon (ϵ) value of 0.05 to support realistic threat representation. Model training was conducted using TensorFlow Federated 2.8 with NVIDIA CUDA 11.7 acceleration. The primary security objective was to achieve a

FAR of $\leq 5\%$ for the detection of deepfake avatars. The secondary objective was to maintain an inference latency of ≤ 80 ms. The ResNet-50 architecture was selected due to its demonstrated effectiveness in image-based anomaly detection and its compatibility with federated learning environments.

The blockchain identification module was developed using W3C-compliant DIDs and Hyperledger Fabric 2.5. A reputation-based Byzantine Fault Tolerance (BFT) consensus mechanism was implemented across a network of 500 nodes. This number was selected based on preliminary scaling tests, which indicated an appropriate balance between Sybil resistance and transaction finality. To address the conflict between blockchain immutability and the GDPR's right to erasure, zk-SNARKs (implemented using the libsnark package) were incorporated to enable cryptographic proof of data deletion without modifying the blockchain.

The Open Quantum Safe Library v0.7 used the CRYSTALS-Kyber (NIST Standard #1) technique to protect keys in the cryptography layer. The study implemented a hybrid AES–Kyber encryption scheme using 768-bit keys to balance quantum-resistant security with operational efficiency. The objective was to maintain latency overhead below 1.2 times that of conventional AES 256 encryption.

A two-round Delphi assessment was performed with a panel of twenty cybersecurity professionals to qualitatively assess the architecture's viability and practical applicability. Panelists were selected based on having a minimum of two Scopus-indexed publications and demonstrated practical experience in Zero-Trust or blockchain deployments. To minimize potential geographic bias, the panel comprised participants from diverse regions, with 70% from North America, 20% from Europe, and 10% from Asia. Cohen's κ coefficient was used to quantify consensus, with values >0.75 indicating substantial agreement. Qualitative responses were analyzed using thematic analysis performed with NVivo 14.

The study systematically collected quantitative performance metrics. Prometheus monitoring was used to evaluate authentication latency, the Hyperledger Caliper benchmarking tool was employed to measure throughput (TPS), and Joulemeter and NVIDIA SMI were utilized to assess energy consumption. The experimental architecture incorporated baseline comparisons with Snort 3.0 and Ethereum Mainnet, following the measurement standards established in NIST Internal Report (IR) 8286 and ISO/IEC 25023.

All data collection procedures were conducted in accordance with Association for Computing Machinery

(ACM) and Institute of Electrical and Electronics Engineers (IEEE) ethical guidelines. All user data were pseudonymized in accordance with NIST SP 800-122 requirements and GDPR regulations, and expert responses were anonymized.

All penetration testing activities, including the 10,000 avatar impersonation attempts, Sybil attacks (500–10,000 malicious nodes), and VR privilege escalation attempts, were conducted exclusively within a controlled and isolated testbed environment specifically constructed for this research. No live or operational Metaverse platforms were tested. The testbed was hosted on dedicated servers owned by the affiliated institution. All testing complied with applicable data protection policies, and no real user data were exposed or compromised during any security assessment activity.

4. Results

The experimental assessment of the proposed zero-trust AI–blockchain architecture demonstrated statistically significant improvements in selected security and performance metrics, providing empirical evidence of its potential effectiveness in addressing metaverse-specific risks.

In the realm of AI-driven threat detection, the federated ResNet-50 model attained a False Acceptance Rate (FAR) of 5.2%, reflecting a 42.9% decrease from the 9.1% FAR baseline of rule-based systems ($p < 0.01$, one-way ANOVA $F(1, 198) = 29.4$). The model also reduced inference latency by 60%, from 4.5 ms to 1.8 ms, indicating a measurable performance improvement. The system identified 94% of GAN-generated deepfake avatars, representing a 31-percentage-point improvement over previous techniques [14,15]. Adversarial robustness testing showed 87% resistance to FGSM-based attacks, representing a 35-percentage-point improvement over baseline systems [18,20]. Figure 3 demonstrates that this enhanced level of protection required substantially greater computational resources, utilizing 3.1 times the GPU capacity of baseline configurations.

The blockchain-based identity module showed 99.1% resilience against Sybil attacks, indicating strong performance for decentralized identity management in virtual environments. The analysis revealed an anticipated performance trade-off: at the optimal node count of 500 ($R^2 = 0.89$), throughput decreased by 14%, while identity issuance latency remained stable at 2.4 s under peak load conditions (see Figure 4). The integration of zk-SNARKs provided a cryptographic mechanism for achieving partial compliance with the GDPR's "right to erasure," addressing a limitation present in previous blockchain-based implementations [25,26].




Figure 3: AI threat detection: performance vs. resource overhead.




Figure 4: Blockchain node scalability trade-offs: throughput, latency, and sybil resistance.

The post-quantum cryptography layer evaluation indicated that CRYSTALS-Kyber may be operationally feasible within the tested environment. The system maintained its asserted security level of 2^{187} quantum operations against the Module-LWE challenge, in accordance with NIST Level 3 requirements. Timing analysis revealed no significant difference between legal and hostile inputs (t -test $p = 0.43$), demonstrating resilience against basic side-channel assaults, but it added $1.2\times$ delay compared to AES-256 (6.2 ms vs. 5.1 ms). A hybrid AES-

Kyber implementation further improved performance by reducing latency to 5.7 ms ($p < 0.01$). The energy consumption analysis revealed that power usage increased by 28.9% (58 W vs. 45 W per 10K transactions) compared with conventional encryption. However, the hybrid AES-Kyber implementation reduced energy consumption to 51 W, representing a 12.1% decrease compared with pure Kyber. These findings highlight an important consideration for the sustainability of quantum-resistant cryptographic implementations (see Table 3).

Table 3: Cryptographic performance and energy efficiency.

Algorithm	Latency (ms)	Throughput (TPS)	Power (W per 10 K Tx)
AES-256	5.1	12,400	45
CRYSTALS-Kyber	6.2	9800	58
Hybrid AES-Kyber	5.7	11,200	51

Triangulation of qualitative data from the Delphi survey indicated a 92% expert agreement on the critical importance of AI–blockchain integration for Metaverse security, as indicated by a mean Likert score of 4.6/5 (± 0.5) for architectural integrity. The user survey analysis identified usability as the strongest predictor of adoption intention ($\beta = 0.47$, $p < 0.001$). In contrast, privacy concerns received a score of 4.2/5 (± 0.8). These human-factor findings demonstrate that the framework achieves an effective balance between stringent security requirements and user experience considerations.

5. Discussion

The empirical evidence suggests that the proposed zero-trust AI–blockchain architecture may advance security frameworks for metaverse environments. The 42.9% decrease in the FAR ($p < 0.01$) indicates improved resistance to biometric spoofing attempts in decentralized immersive settings. This improvement directly addresses a major limitation of traditional rule-based systems, which have struggled to detect emerging threats such as GAN-generated deepfakes and NFT phishing schemes [14,15]. The model’s 87% resistance to FGSM adversarial attacks, representing a 35-percentage-point improvement over conventional systems [18,20], further demonstrates its enhanced reliability. However, this enhanced level of protection incurs substantial computational costs, as demonstrated by a $3.1\times$ increase in GPU resource utilization. This increase is a direct consequence of deploying the ResNet-50 model on a dedicated GPU-accelerated server, as illustrated in the Deployment View. This finding aligns with previous studies highlighting the trade-offs associated with optimizing federated learning [21,22]. It also establishes a clear direction for future research on hardware acceleration and improving edge computing efficiency.

The blockchain identity module’s 99.1% resistance to Sybil attacks validates the efficacy of Hyperledger Fabric combined with a reputation-based Byzantine Fault Tolerance consensus mechanism for decentralized authentication in the Metaverse. The observed scalability-latency trade-off, exemplified by a 2.4-s authentication delay at the optimal node configuration (500 nodes, $R^2 = 0.89$),

empirically validates theoretical predictions concerning the performance limitations of distributed ledger technologies under security-intensive workloads [34,35]. A significant conceptual advancement is the effective use of zk-SNARKs to achieve partial GDPR compliance, offering a cryptographic solution to the persistent legal conflict between blockchain immutability and regulatory data-deletion requirements [25,26]. This critical challenge has been overlooked in 85% of previous frameworks, as identified in the literature review.

The experimental assessment of the quantum-resistant cryptographic layer provides a practical evaluation of post-quantum cryptography (PQC) operational readiness within a Metaverse testbed. The $1.2\times$ latency overhead of CRYSTALS-Kyber-768 compared to AES-256 (6.2 ms against 5.1 ms) corresponds with NIST reference forecasts, validating its practical applicability in latency-sensitive virtual settings [28,29]. The 28.9% increase in energy consumption (58 W compared with 45 W per 10,000 transactions) challenges prevailing assumptions regarding the efficiency of lattice-based encryption and necessitates a thorough reassessment of sustainability claims within the PQC literature [9,11].

Regarding the validation of quantum security, the proposed architecture does not provide empirical evidence of quantum resistance through Shor’s algorithm simulation, as such an approach would be theoretically inappropriate. Shor’s algorithm is not applicable to the Module Learning With Errors (MLWE) problem that underlies CRYSTALS Kyber. Quantum-resistant security is supported by established cryptographic assumptions; the most recognized quantum attack targeting Kyber-768 (the dual attack) requires approximately 2^{187} operations, exceeding the security requirements of NIST Level 3 standards (2^{143}). Implementation-side-channel testing (100,000 timing and power-analysis traces) revealed no exploitable leakage (t -test $p > 0.05$). These results establish Kyber’s security as mathematically derived rather than empirically “validated” against quantum attacks, a distinction that is essential for precise academic discourse.

The hybrid AES-Kyber solution, demonstrating a latency reduction of 5.7 ms ($p < 0.01$), exemplifies a practical balance between quantum resilience and operational efficiency. This result highlights the

significance of hybrid cryptographic frameworks during transitional migration phases, when systems must concurrently accommodate classical and post-quantum primitives.

The Delphi study and user surveys provided qualitative insights that complemented the technical findings with critical human-factor perspectives. The strong expert consensus (92% agreement; Likert score of $4.6/5 \pm 0.5$) regarding the architectural integrity of the AI-blockchain integration supports its conceptual validity. The significant correlation between perceived usability and adoption intention ($\beta = 0.47, p < 0.001$) challenges the traditional zero-sum assumption that security enhancements must be achieved at the expense of user experience [37,38]. These findings suggest that user-centric design is a critical enabler rather than a secondary consideration for implementing complex security systems in consumer-focused virtual platforms.

The interpretation of these data must be situated within the methodological limitations of the investigation. Reliance on reference implementations and theoretical cryptographic analysis may not fully represent the capabilities of future quantum adversaries [39,40]. Although multi-engine validation using Unreal Engine mitigated concerns regarding platform specificity, the generalizability of performance benchmarks across the diverse technical landscape of the Metaverse, including mobile AR platforms and multiple rendering engines, remains a challenge [14,15]. The geographical composition of the Delphi panel (70% North American experts) may have introduced potential cultural bias, particularly in the evaluation of GDPR-related compliance considerations.

Despite these limitations, the study offers several contributions. It offers empirical substantiation for the use of Zero-Trust concepts in decentralized virtual environments, therefore addressing a deficiency highlighted in recent IEEE and ACM literature [41,42]. It offers practical benchmarks for integrating AI, blockchain, and PQC, which have traditionally been investigated independently within theoretical contexts [32,33]. Moreover, it presents verified human-factor models that effectively harmonize security and usability goals, which have historically been seen as fundamentally conflicting [43]. The methodological rigor, aligned with NIST SP 800-207 and ISO 27001 standards, establishes a reproducible framework for future security research in this emerging field. It provides a strong foundation for both academic investigation and practical implementation to address evolving cyber-physical threats in next-generation virtual environments.

6. Conclusions

This study introduced and empirically evaluated a zero-trust AI-blockchain architecture for metaverse environments, integrating federated learning for deepfake-resistant threat detection, a Hyperledger Fabric-based decentralized identity system with zk-SNARKs for GDPR-compliant erasure proofs, and CRYSTALS-Kyber for post-quantum key exchange. Experimental results demonstrated a statistically significant reduction in false acceptance rates, strong resilience against Sybil attacks, and quantifiable trade-offs in latency, throughput, and energy consumption. A mixed-methods human-centric evaluation further revealed that perceived usability strongly predicts adoption intention. Future research should focus on cross-platform generalization beyond Unity3D, hardware-accelerated optimization of federated learning to reduce GPU overhead, and longitudinal analysis of evolving attack patterns in live Metaverse deployments.

List of Abbreviations

ACM	Association for Computing Machinery
AES	Advanced Encryption Standard
AI	Artificial Intelligence
ANOVA	Analysis of Variance
AR	Augmented Reality
BFT	Byzantine Fault Tolerance
CNN	Convolutional Neural Network
CUDA	Compute Unified Device Architecture
DID	Decentralized Identifier
ECC	Elliptic Curve Cryptography
FAR	False Acceptance Rate
FGSM	Fast Gradient Sign Method
FPR	False Positive Rate
GAN	Generative Adversarial Network
GDPR	General Data Protection Regulation
GPU	Graphics Processing Unit
IAM	Identity and Access Management
IDS	Intrusion Detection System
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
ION	Identity Overlay Network
IR	Internal Report
ISO	International Organization for Standardization
LSTM	Long Short-Term Memory
MLWE	Module-Learning with Errors
NERC CIP	North American Electric Reliability Corporation Critical Infrastructure Protection

NFT	Non-Fungible Token
NIST	National Institute of Standards and Technology
OPA	Open Policy Agent
OQS	Open Quantum Safe
PQC	Post-Quantum Cryptography
ResNet	Residual Neural Network
RSA	Rivest–Shamir–Adleman
StyleGAN	Style Generative Adversarial Network
TPS	Transactions Per Second
VR	Virtual Reality
W3C	World Wide Web Consortium
ZTA	Zero-Trust Architecture
ZTS	Zero-Trust Security
zk-SNARK	Zero-Knowledge Succinct Non-Interactive Argument of Knowledge

Author Contributions

The author confirms that he was solely responsible for the conceptualization, methodology, validation, investigation, resources, data curation, writing—original draft, writing—review & editing, visualization, and project administration of the article. The author has read and agreed to the published version of the manuscript.

Availability of Data and Materials

The data produced in this investigation are accessible in the following repositories. A total of 50,000 adversarial samples tailored for the Metaverse (FGSM, $\epsilon = 0.05$) and a dataset of 10,000 deepfake avatars created by StyleGAN2 are available at <https://doi.org/10.5281/zenodo.20098586>. Anonymized raw Delphi survey results and Prometheus performance data (authentication delay, throughput, power consumption) may be obtained from the corresponding author upon reasonable request and contingent upon institutional ethical clearance. No proprietary data from other parties were used.

Ethical Committee Approval and Consent to Participate

This study, involving a Delphi expert panel ($n = 20$) and voluntary user surveys, was conducted in compliance with the ethical principles of the Declaration of Helsinki (2024 revision) and the Code of Nuremberg. In accordance with Costa Rica’s Regulatory Law of Biomedical Research (Law 9234), which applies to research involving human subjects, the study was determined to be exempt from full review by an institutional ethics committee. This exemption was justified because the research involved minimal risk, employed anonymous survey methodologies, and did not involve the collection of biological samples or sensi-

tive personal data that could compromise participants’ dignity or autonomy. Informed consent was obtained from all participants. Written consent was obtained from the Delphi panelists, while implied consent was obtained through voluntary survey completion after participants reviewed an information sheet outlining the study purpose, confidentiality measures, and their right to withdraw without consequences. The Delphi consultations and surveys were deemed to be of minimal risk, and the procedures for obtaining consent were appropriate for the nature of the research, as stipulated by national regulations.

Human Rights

This study involving human participants (Delphi expert panel, $n = 20$; user surveys) was conducted in accordance with Costa Rica’s Regulatory Law of Biomedical Research (Law 9234) and the ethical principles of the Declaration of Helsinki (2024 revision).

Conflicts of Interest

The author declares no conflicts of interest.

Funding

The study did not receive any external funding and was conducted using only institutional resources.

Acknowledgments

The author would like to thank all those involved in the work who made it possible to achieve the objectives of the research study.

AI Declaration

During the preparation of this work, the author used the AI-assisted tool Grammarly (Version 1.0) for grammar and spell-checking, and ChatGPT (GPT-4) for structuring the literature review and assisting with the initial drafting of the introductory sections. After using these tools, the author reviewed and edited the content as necessary and takes full responsibility for the final content of the publication. No generative AI was used to create figures, images, or graphical elements, nor for the collection or primary analysis of data.

References

- [1] P. B. Lowry, W. F. Boh, S. Petter, and J. M. Leimeister, “Long live the metaverse: Identifying the potential for market disruption and future research,” *J. Manage. Inf. Syst.*, vol. 42, no. 1, pp. 3–38, 2025. [\[CrossRef\]](#)
- [2] J. Piñeiro-Chousa, M. Á. López-Cabarcos, V. VitorRomero, and A. Pérez-Pérez, “Evolution and trends of the metaverse in business and management:

- A bibliometric analysis,” *Rev. Manag. Sci.*, vol. 19, no. 1, pp. 197–222, 2025. [CrossRef]
- [3] A. S. George, “Emerging trends in AI-driven cybersecurity: An in-depth analysis,” *Partners Univ. Innov. Res. Publ.*, vol. 2, no. 4, pp. 15–28, 2024. [CrossRef]
- [4] Y. Huang, Y. J. Li, and Z. Cai, “Security and privacy in metaverse: A comprehensive survey,” *Big Data Min. Anal.*, vol. 6, no. 2, pp. 234–247, 2023. [CrossRef]
- [5] I. Ud Din, K. Habib Khan, A. Almogren, M. Zareei, and J. A. Pérez Díaz, “Securing the metaverse: A blockchain-enabled zero-trust architecture for virtual environments,” *IEEE Access*, vol. 12, pp. 92337–92347, 2024. [CrossRef]
- [6] M. Tsai, S. Lee, and S. W. Shieh, “Strategy for implementing of zero trust architecture,” *IEEE Trans. Rel.*, vol. 73, no. 1, pp. 93–100, 2024. [CrossRef]
- [7] M. G. K. Hemdani, “Cryptocurrencies and the dark web: A gateway to money laundering,” in *Cyber-crime Unveiled: Technologies for Analysing Legal Complexity*, M. Chawki and A. Abraham, Eds., Cham, Switzerland: Springer, 2025, pp. 217–247. [CrossRef]
- [8] M. Zhang, X. Zhang, Y. Zhang, and Z. Lin, “Security of cross-chain bridges: Attack surfaces, defenses, and open problems,” in *Proc. 27th Int. Symp. Res. Attacks Intrusions Def.*, 2024, pp. 298–316. [CrossRef]
- [9] D. J. Bernstein, “Post-quantum cryptography,” in *Encyclopedia of Cryptography, Security and Privacy*. Cham, Switzerland: Springer, 2025, pp. 1846–1847. [CrossRef]
- [10] G. Olaoye, “Post-quantum cryptography: Evaluating and standardizing quantum-resistant algorithms,” *Authorea*, 2025. [CrossRef]
- [11] O. Alnaseri, Y. Himeur, S. Atalla, and W. Mansoor, “Complexity of post-quantum cryptography in embedded systems and its optimization strategies,” in *Proc. 2025 Int. Wireless Commun. Mobile Comput. (IWCMC)*, May 2025, pp. 776–781. [CrossRef]
- [12] A. Ali, K. F. Khan Ghouri, H. Naseem, T. R. Soomro, W. Mansoor, and A. M. Momani, “Battle of deep fakes: Artificial intelligence set to become a major threat to the individual and national security,” in *Proc. 2022 Int. Conf. Cyber Resil. (ICCR)*, Oct. 2022, pp. 1–5. [CrossRef]
- [13] D. Shin, “Conclusion: Misinformation and AI—How algorithms generate and manipulate misinformation,” in *Artificial Misinformation: Exploring Human-Algorithm Interaction Online*. Cham, Switzerland: Springer, 2024, pp. 259–277. [CrossRef]
- [14] D. Bendig and M. Charlet, “Opportunities and challenges of blockchain for multi-sided platforms,” *Electron. Markets*, vol. 35, no. 1, 2025, Art. no. 25. [CrossRef]
- [15] H. Meng, J. Ding, H. Wang, Z. Zhang, X. Yao, and H. Ning, “Blockchain enabled metaverse: Development and applications,” *Tsinghua Sci. Technol.*, vol. 30, no. 4, pp. 1552–1582, 2025. [CrossRef]
- [16] P. Khan, M. Z. Islam, and S. Hossain, “AI-powered cybersecurity: Revolutionizing business threat detection and response,” *Am. J. Smart Technol. Sol.*, vol. 4, no. 1, pp. 37–48, 2025. [CrossRef]
- [17] M. A. M. Farzaan, M. C. Ghanem, A. El-Hajjar, and D. N. Ratnayake, “AI-powered system for an efficient and effective cyber incidents detection and response in cloud environments,” *IEEE Trans. Mach. Learn. Commun. Netw.*, vol. 3, pp. 623–643, 2025. [CrossRef]
- [18] Y. Bian, X. Zhang, G. Luosang, D. Renzeng, D. Renqing, and X. Ding, “Federated learning and semantic communication for the metaverse: Challenges and potential solutions,” *Electronics*, vol. 14, no. 5, 2025, Art. no. 868. [CrossRef]
- [19] L. Sun, Z. Zhang, and G. Muhammad, “Generative learning-based personalized federated learning for metaverse data security: Proposing two new frameworks,” *IEEE Syst. Man Cybern. Mag.*, vol. 11, no. 2, pp. 4–13, 2025. [CrossRef]
- [20] Y. Djenouri, A. Nabil Belbachir, A. Belhadi, T. Michalak, and G. Srivastava, “Next-gen metaverse security through intrusion detection enhanced by transformers and GANs,” *IEEE Internet Things J.*, vol. 12, no. 12, pp. 20640–20651, 2025. [CrossRef]
- [21] M. Z. Aloudat, M. Barhamgi, E. Yaacoub, and D. Aoun, “Security in metaverse markets: Challenges and solutions—A comprehensive review,” *Expert Syst.*, vol. 42, no. 8, 2025, Art. no. e70094. [CrossRef]
- [22] A. K. Lamba and S. Pal, “Metaverse security, privacy issues and cyber security opportunities & challenges,” in *Exploring the Metaverse*. Amsterdam, The Netherlands: Elsevier, 2025, pp. 123–138. [CrossRef]
- [23] K. Chang and O. Seneviratne, “Zero-knowledge proof framework for identity verification and interoperable payments on the decentralized web,” *ACM Trans. Web*, 2026. [CrossRef]
- [24] I.-C. Lin, I. Yeh, C.-C. Chang, J.-C. Liu, and C.-C. Chang, “Designing a secure and scalable data sharing mechanism using decentralized identifiers (DID),” *CMES-Comp. Model. Eng. Sci.*, vol. 141, no. 1, pp. 809–822, 2024. [CrossRef]
- [25] P. Ghaly, H. Gjermundrød, and I. Dionysiou, “Privacy-preserving “Right to Be Forgotten” in configurable blockchain tokenization: A cryptographic audit framework,” in *2025 7th Conf. Blockchain Res. Appl. Innov. Netw. Serv. (BRAINS)*, Zurich, Switzerland, 2025, pp. 1–8. [CrossRef]
- [26] M. Pilishvili, “Permacrisis: blockchain’s plan to fix the normative challenges in EU data protection law,” *Jog Állam Polit. Jog-És Polit. F.*, vol. 17, no. 3, pp. 81–92. 2025. [CrossRef]
- [27] M. Bisheh-Niasar, R. Azarderakhsh, and M. Mozaffari-Kermani, “Instruction-set accelerated implementation of CRYSTALS-Kyber,” *IEEE Trans.*

- Circuits Syst. I Reg. Pap.*, vol. 68, no. 11, pp. 4648–4659, 2021. [[CrossRef](#)]
- [28] M. Abbasi, F. Cardoso, P. Váz, J. Silva, and P. Martins, “A practical performance benchmark of post-quantum cryptography across heterogeneous computing environments,” *Cryptography*, vol. 9, no. 2, 2025, Art. no. 32. [[CrossRef](#)]
- [29] M. A. G. de la Torre, I. A. M. Sandoval, G. T. F. de Abreu, and L. H. Encinas, “Post-quantum wireless-based key encapsulation mechanism via CRYSTALS-Kyber for resource-constrained devices,” *IEEE Access*, vol. 13, pp. 66714–66725, 2025. [[CrossRef](#)]
- [30] G. Tu, W. Liu, T. Zhou, X. Yang, and F. Zhang, “Concise and efficient multi-identity fully homomorphic encryption scheme,” *IEEE Access*, vol. 12, pp. 49640–49652, 2024. [[CrossRef](#)]
- [31] O. Bernard, M. Joye, N. P. Smart, and M. Walter, “Drifting towards better error probabilities in fully homomorphic encryption schemes,” in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Cham, Switzerland: Springer, 2025, pp. 181–211. [[CrossRef](#)]
- [32] Y. Y. Ghadi et al., “A hybrid AI-blockchain security framework for smart grids,” *Sci. Rep.*, vol. 15, no. 1, 2025, Art. no. 20882. [[CrossRef](#)] [[PubMed](#)]
- [33] A. Singh, N. Dhanda, and K. K. Gupta, “A hybrid AI-blockchain framework for secure and resilient network infrastructure,” in *Proc. 2025 6th Int. Conf. Intell. Commun. Technol. Virtual Mobile Netw. (ICICV)*, Jun. 2025, pp. 1574–1580. [[CrossRef](#)]
- [34] A. M. Buttar, M. A. Shahid, M. N. Arshad, and M. A. Akbar, “Decentralized identity management using blockchain technology: Challenges and solutions,” in *Blockchain Transformations: Navigating the Decentralized Protocols Era*. Cham, Switzerland: Springer, 2024, pp. 131–166. [[CrossRef](#)]
- [35] M. Polychronaki, M. G. Xevgenis, D. G. Kogias, and H. C. Leligou, “Decentralized identity management for metaverse-enhanced education: A literature review,” *Electronics*, vol. 13, no. 19, 2024, Art. no. 3887. [[CrossRef](#)]
- [36] T. Górski, “The 1+5 architectural views model in designing blockchain and IT system integration solutions,” *Symmetry*, vol. 13, no. 11, 2021, Art. no. 2000. [[CrossRef](#)]
- [37] A. M. Albarrak, “Integration of cybersecurity, usability, and human-computer interaction for securing energy management systems,” *Sustainability*, vol. 16, no. 18, 2024, Art. no. 8144. [[CrossRef](#)]
- [38] B. Naqvi and K. Smolander, “Practitioners’ perspectives on and prospects for usable security,” *Computer*, vol. 57, no. 10, pp. 66–74, 2024. [[CrossRef](#)]
- [39] K. F. Hasan et al., “A framework for migrating to post-quantum cryptography: Security dependency analysis and case studies,” *IEEE Access*, vol. 12, pp. 23427–23450, 2024. [[CrossRef](#)]
- [40] C. Näther, D. Herzinger, S. L. Gazdag, J. P. Steghöfer, S. Daum, and D. Loebenberger, “Migrating software systems toward post-quantum cryptography—A systematic literature review,” *IEEE Access*, vol. 12, pp. 132107–132126, 2024. [[CrossRef](#)]
- [41] A. Agarwal, R. Ramachandra, S. Venkatesh, and S. R. M. Prasanna, “Biometrics in extended reality: A review,” *Discov. Artif. Intell.*, vol. 4, no. 1, 2024, Art. no. 81. [[CrossRef](#)]
- [42] M. El-Hajj, “Cybersecurity and privacy challenges in extended reality: Threats, solutions, and risk mitigation strategies,” *Virtual Worlds*, vol. 4, no. 1, 2025, Art. no. 1. [[CrossRef](#)]
- [43] S. Roychowdhury, B. K. Singh, A. Das, and A. Choudhury, “Challenges and solutions for balancing usability and security for users with physical difficulties in medical cyber-physical systems,” in *Cyber-Physical Systems Security: A Multi-disciplinary Approach*. Singapore: Springer, 2025, pp. 151–177. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The views expressed in this article are those of the author(s) and do not necessarily reflect the views of the publisher or editors. The publisher and editors assume no responsibility for any injury or damage resulting from the use of information contained herein.