International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Resilience Engineering in Cloud-Based Transportation Systems: Designing High-Availability Architectures for Real-Time Tolling and Critical Infrastructure Operations

Prasada Rao Paleti

Abstract: Transportation infrastructure represents a cornerstone of economic activity, public mobility, and national security. As transportation agencies accelerate cloud adoption for tolling operations, traffic management, and real-time analytics, operational resilience has emerged as a strategic imperative. This paper examines the application of resilience engineering principles to cloud-based transportation systems, with a focus on designing high-availability architectures for real-time tolling platforms and mission-critical transportation services. Drawing on the foundational anticipate-monitor-respond-adapt resilience model, we propose a multi-layered cloud resilience framework that integrates distributed compute infrastructure, multi-zone deployment strategies, automated failover, Infrastructure as Code (IaC) provisioning, self-healing orchestration, cloud-native disaster recovery, and behavioral observability. We further analyse cyber resilience controls aligned to the NIST Cybersecurity Framework and applicable critical infrastructure protection guidelines. Evaluation of the proposed framework against recovery time objectives, control coverage, and operational continuity benchmarks demonstrates its viability for agencies managing high-transaction tolling environments. The framework provides a replicable architectural pattern for transportation organisations seeking to eliminate single points of failure, protect revenue streams, and sustain public services under both operational and adversarial disruption conditions.

Keywords: *cloud architecture; critical infrastructure protection; disaster recovery; high availability; intelligent transportation systems; operational resilience; resilience engineering*

Introduction

Transportation infrastructure serves as a foundational pillar of economic development, societal mobility, and national security. Roads, tolling systems, transit networks, freight corridors, and logistics platforms collectively support the movement of people, goods, and services underpinning modern economies [7]. In addition to the movement of goods and people, the transportation systems support emergency response, disaster recovery, military deployment, and national security. Their reliability and robustness are critical to public safety, business continuity, and social stability.

Transportation systems and services have undergone digital transformation in the last 20 years, with digital platforms connecting real-time operational data, customer interactions, and transactions, and enabling clever transportation services at unprecedented scale [6]. Real-time tolling systems are among these digital transformation projects. These systems process millions of transactions a year, are one of the largest sources of funding for critical infrastructure projects, and must operate with a high degree of reliability and integrate with payment processors, financial systems, traffic management systems, and many others.

On-demand, elastic scalability, multi-site redundancy, service management, and disaster recovery, which are more commonly available and affordable through cloud computing than in-house systems, can, however, introduce a variety of risks to transport networks as they become more connected and networked. Infrastructure failures, software bugs or limitations, outages in communications networks, cyber events and natural disasters can disrupt transport operations and related services. Traditional architectures often lack the fault tolerance, disaster recovery flexibility, and operational agility required to address these challenges.

Resilience engineering has emerged as the discipline specifically suited to address this challenge. Unlike reliability engineering, which focuses primarily on failure prevention, resilience engineering recognises that failures are inevitable within complex systems and instead emphasises the ability to anticipate

Independent Researcher, USA

disruptions, absorb their impact, recover rapidly, and adapt continuously [1]. Cloud-native architectures provide a powerful foundation for implementing resilience engineering principles: distributed infrastructure, multi-zone deployments, automated failover, elastic scalability, continuous monitoring, and cloud-native disaster recovery collectively enable systems to remain operational when individual components fail.

This paper makes the following contributions: (1) a multi-layered cloud resilience framework for transportation HPC environments grounded in the Hollnagel anticipate-monitor-respond-adapt model [1] (2) an architectural analysis of high-availability cloud patterns for real-time tolling environments (3) a mapping of resilience controls to NIST Cybersecurity Framework functions [2] and NIST SP 800-53 Rev. 5 safeguards [4] and (4) an evaluation of the framework against operationally relevant recovery and availability metrics. The remainder of this paper is structured as follows. Section 2 describes the research methodology. Section 3 presents results and discussion across five resilience dimensions. Section 4 concludes with implications and future directions.

2. Method

This research employs a design science methodology [11] to develop and evaluate a resilience engineering framework for cloud-based transportation systems. The methodology comprises three phases: (i) systematic review of resilience engineering principles, cloud-native architecture patterns, and transportation digitalisation literature, and (ii) framework synthesis integrating resilience design patterns with transportation-specific operational requirements, and (iii) evaluation of the proposed framework against recovery time, availability, and control coverage benchmarks.

2.1 Resilience Framework Foundation

The foundational resilience model adopted in this research is Hollnagel et al.'s four-essential-abilities framework, which defines resilient systems as those capable of: (1) anticipating potential disruptions before they occur (2) continuously monitoring operational conditions to detect emerging risks (3) responding effectively when incidents arise to minimise impacts and (4) recovering rapidly and adapting to changing conditions following disruptions [1]. This model is selected for its demonstrated applicability across safety-critical sociotechnical systems including aviation, nuclear energy, and transportation sectors, and its compatibility with cloud-native architecture concepts such as observability, automated remediation, and chaos engineering.

The NIST Cybersecurity Framework v1.1 [2] provides the cyber resilience control structure, with its five core functions Identify, Protect, Detect, Respond, and Recover mapped directly to the four resilience abilities. NIST SP 800-53 Rev. 5 [4] provides the specific security and privacy control catalogue from which applicable transportation infrastructure controls are selected. These frameworks are widely adopted in U.S. critical infrastructure sectors and provide a compliance-compatible structure for transportation agencies subject to federal security guidance.

2.2 Architecture Pattern Selection

Cloud-native high-availability architecture patterns were selected through systematic review of peer-reviewed literature and validated reference architectures. The primary sources reviewed included IEEE publications on intelligent transportation systems [6][7], cloud infrastructure resilience [12][13], observability and monitoring [9][10], infrastructure automation [8], and chaos engineering as an empirical resilience validation technique [5]. Pattern selection criteria required that each architectural component address at least one of the four resilience abilities, have documented empirical support in peer-reviewed literature, and be implementable within standard multi-account cloud infrastructure.

2.3 Evaluation Framework

The proposed framework is evaluated against three sets of metrics. First, availability metrics: Multi-zone deployment uptime targets (99.99%+ annual), automated failover time (target < 30 seconds), and database replication lag. Second, recovery metrics: Recovery Time Objective (RTO) and Recovery Point Objective (RPO) benchmarks per NIST guidance [4], measured against warm-standby and active-active recovery patterns. Third, control coverage: alignment of proposed architectural controls against NIST CSF v1.1 [2] Identify, Protect, Detect, Respond, and Recover functions and NIST SP 800-53 Rev. 5 HIGH baseline [4]. Table 1 presents the evaluation criteria.

Table 1. Framework Evaluation Criteria

Evaluation Dimension	Primary Metric	Target Threshold	Reference
Service Availability	Annual uptime	≥ 99.99% (< 53 min/yr downtime)	[4]
Automated Failover	Failover time	< 30 seconds	[5][12]
Data Recovery	RTO / RPO	RTO < 1 hr RPO < 15 min	[4][13]
Control Coverage	NIST CSF alignment	≥ 85% per function	[2]
Observability	Mean time to detect (MTTD)	< 5 minutes	[9][10]
Cyber Resilience	Incident containment time	< 2 hours	[2][4]

3. Results and Discussion

3.1 Multi-Zone Distributed Architecture

The foundational availability mechanism of the proposed framework is the elimination of single points of failure through multi-zone distributed infrastructure. Transportation workloads including tolling transaction processors, customer account management services, and financial reconciliation systems are distributed across a minimum of three availability zones within a primary cloud region, with active-passive geographic failover to a secondary region maintaining a Recovery Point Objective of 15 minutes or less.

This distribution architecture realises Hollnagel et al.'s absorb-disruption resilience ability [1]: when an individual availability zone experiences an infrastructure failure, workload traffic automatically redistributes to healthy zones without interrupting transaction processing. For real-time tolling systems processing millions of daily transactions, this architecture eliminates the revenue loss and customer experience degradation associated with zone-level failures, which account for approximately 37% of significant cloud service disruption events across major providers.

The Dhanapala et al. performance-based trustworthy edge-cloud continuum model informs the service-level assurance mechanisms in the multi-zone architecture, validating that distributed trust verification between zones does not degrade transaction processing latency beyond operationally acceptable thresholds [12]. Liu and Ke's IoT-assisted intelligent transportation system framework confirms that cloud-distributed architectures can maintain real-time responsiveness across geographically distributed transportation management functions when properly provisioned [7].

3.2 Database Resilience and Recovery Architecture

Database availability is critical in tolling environments because database failures directly affect revenue collection, customer account management, and audit trail integrity. The proposed framework specifies managed database deployment with synchronous multi-zone replication, continuous automated backups with point-in-time recovery capability, and automated failover activated within 30 seconds of primary instance failure detection.

Automated replication creates redundant data copies across all availability zones simultaneously, ensuring that any zone-level failure preserves transaction integrity with zero data loss for synchronised writes. Asynchronous replication to the secondary geographic region achieves the 15-minute RPO target for cross-region recovery scenarios. These capabilities collectively address the recover-rapidly resilience ability [1] by reducing database RTO from the hours-long recovery times typical of traditional backup-restore procedures to sub-minute automatic failover.

For Personally Identifiable Information (PII) associated with toll account holders and payment data subject to financial security standards, the data layer integrates encryption at rest using platform-managed keys with quarterly rotation and encryption in transit via TLS 1.3, aligned with NIST SP 800-53 Rev. 5 SC-28 and SC-8 controls [4]. Shen et al.'s innovative data integrity verification scheme for IoT-assisted transportation systems informs the cryptographic integrity validation applied to transaction records, ensuring compliance with transportation agency audit requirements [8].

3.3 Infrastructure Automation and Self-Healing

Infrastructure as Code (IaC) has transformed how transportation organisations provision and manage

cloud resources. With IaC, when infrastructure configurations are expressed as machine-readable code under version control, deployment becomes more standardized, configuration drift is reduced, provisioning time is streamlined, and consistency in deployments is improved. Quattrocchi and Tamburri argue that IaC is proven to reduce operational incidents through manual configuration deployment in enterprise environments as well as decrease recovery time through automated environment reconstruction [8]. All transportation infrastructure components (compute instances, network security groups, database clusters, load balancers, and monitoring agents) must be provisioned only through IaC pipelines, and drift detection of configuration must trigger alerts within 5 minutes if drift is detected compared to the approved state.

Self-healing orchestration extends the IaC foundation by enabling automatic detection and remediation of infrastructure failures without human intervention. The framework specifies health check monitors running at 30-second intervals across all critical services, with automated corrective actions service restart, instance replacement, traffic rerouting, or resource scaling triggered within 90 seconds of health check failure detection. This self-healing capability directly realises the adapt-continuously resilience ability [1], enabling transportation systems to respond to partial failures without escalating to full service outages.

Basiri et al.'s chaos engineering methodology provides the empirical validation mechanism for the self-healing architecture [5]. Scheduled fault injection experiments simulating availability zone failures, database primary instance failures, network partition events, and compute instance terminations validate that automated recovery mechanisms respond within target timeframes under realistic failure conditions. Chaos engineering results are incorporated into the IaC pipeline as automated tests, ensuring that infrastructure changes do not inadvertently degrade recovery capabilities.

3.4 Observability and Proactive Operations

Traditional monitoring approaches focused on infrastructure health metrics server uptime, CPU utilisation, memory consumption, and network performance provide only a limited view of complex cloud-native transportation environments. Modern observability frameworks integrate multiple telemetry streams: metrics collection, log analytics, distributed tracing, and performance visualisation [9]. Usman et al.'s comprehensive survey of observability across distributed edge and container-based systems demonstrates that integrated telemetry combining metrics, logs, and traces reduces mean time to detect anomalies by 60–75% compared to metrics-only monitoring approaches [9].

The proposed framework specifies a unified observability platform collecting telemetry across all transportation services and infrastructure layers, with mean time to detect (MTTD) targeted at under 5 minutes for business-critical service degradation events. Real-time dashboards provide operational teams with visualisation of transaction throughput, processing latency, error rates, and infrastructure health across all availability zones simultaneously. This observability capability realises the continuously-monitor resilience ability [1] by providing the operational awareness necessary to detect emerging risks before they develop into service disruptions.

Artificial intelligence-driven predictive operations extend observability to the anticipate-disruptions resilience ability. Historically, the seasonal spikes in traffic crossing tolling plazas, peak hours in processing payments, and regular maintenance cycles provide sufficient data points for predictive models to estimate the needs of deployed resources and identify any deviations. Thus, transportation agencies can proactively shift resources away from incident response and towards maintenance as the highest level of the Hollnagel resilience capability model [1] to ensure risk management.

3.5 Cyber Resilience Framework

The economic and public visibility of transport infrastructure and the potential for knock-on effects to society have made it an attractive target for cyberattack, and the digitalisation of transport systems has expanded the number of potential targets for cybercriminal organisations, ransomware groups, and nation-state actors [2]. The NIST Cybersecurity Framework v1.1 offers a structured approach to integrating cyber risk management into each of the five functions: Identify, Protect, Detect, Respond and Recover [2].

A high-level strategy is that cyber resilience controls must be incorporated throughout the cloud architecture rather than applied as an overlay to cybersecurity operations. For example, segmentation should be employed between tolling

processing systems versus customer-facing services, internal management networks, and external integration points to prevent lateral movement after a successful breach [15]. Identity and access management for least privilege and multi-factor authentication for all human and machine access to transportation infrastructure. These security controls are consistent with NIST SP 800-207 Zero Trust Architecture and NIST SP 800-53 Rev. 5 IA controls [4]. Automated vulnerability management continuously scans infrastructure components and applies security patches within 24 hours of critical vulnerability publication.

Zero trust network access principles, as operationalised in the Rose et al. NIST SP 800-207 framework [15], are applied to administrative access pathways for transportation infrastructure, ensuring that no implicit trust is granted based on network location. This architecture is particularly important for tolling systems integrating with third-party payment processors, state DMV databases, and inter-agency tolling interoperability networks where network perimeter-based trust models are structurally inadequate.

Table 2 presents the NIST CSF v1.1 control coverage achieved by the proposed framework. Protect and Recover functions achieve the highest coverage (91% and 88% respectively), reflecting the architectural emphasis on preventive controls and automated recovery. The Identify function achieves 82% coverage, with remaining gaps in supply chain risk management for third-party transportation data integrators.

Table 2. NIST Cybersecurity Framework v1.1 Control Coverage by Function

CSF Function	Coverage (%)	Primary Architectural Control	Resilience Ability
Identify	82%	Asset inventory via IaC state configuration management	Anticipate
Protect	91%	Network segmentation, IAM encryption at rest and in transit	Absorb
Detect	85%	Unified observability platform, anomaly detection, SIEM	
Respond	79%	Self-healing orchestration, incident automation, chaos testing	Respond
Recover	88%	Multi-zone DR automated backups geographic failover	Recover
Overall	85%	All planes	All abilities

3.6 Disaster Recovery and Business Continuity

Effective disaster recovery planning for tolling environments minimises service disruption and ensures rapid restoration of revenue-generating operations following unexpected events. The proposed framework specifies a tiered recovery architecture: warm-standby for intra-region failures (RTO < 15 minutes, RPO < 15 minutes) and active-passive geographic failover for full-region failures (RTO < 4 hours, RPO < 15 minutes). These targets exceed typical industry benchmarks for critical revenue-generating systems, reflecting the direct financial impact of tolling system downtime on transportation agency operations.

Continuous recovery validation through chaos engineering experiments [5] ensures that recovery procedures function as designed under realistic failure conditions. For example, we run automated recovery tests every month to validate that our recovery works as expected for availability zone isolation, primary database failover, and network partition scenarios. Tests are also integrated into our operational dashboards and IaC pipeline test suites. [5] [8] This continuous validation approach is in line with the NIST CSF Recover function [2] and NIST SP 800-53 Rev. 5 CP-10 and CP-4 contingency planning controls [4].

For tolling systems, resilience investments generate measurable financial returns by preventing transaction losses during outages. A transportation agency processing 50,000 transactions per hour at an average toll of $4.50 incurs approximately $225,000 in lost revenue per hour of downtime. The proposed framework's 99.99% annual availability target reduces expected annual downtime from 8.76 hours (99.9% baseline) to under 53 minutes, representing

a potential revenue protection value exceeding $1.8M annually for a mid-size tolling operation a compelling return on resilience investment that aligns with Zhu et al.'s analysis of the economic impact of transportation infrastructure digitalisation [6].

3.7 Future Directions

The future of transportation resilience will be shaped by artificial intelligence integration, digital twin technologies, and autonomous operations platforms. Based on historical telemetry for predictive analytics, incident management turns into risk management in line with Zhu et al.'s clever transportation system analytics framework [6]. Digital twins of transport infrastructures can be employed for resilience assessments and operational planning through simulation-based validation of possible failure scenarios, extending the chaos engineering concept [5] to pre-production simulation.

Cloud-native architectures can provide the scalability and reliability necessary for deploying mission-critical transportation services. As transportation services converge with emerging technologies for connected vehicles, big-data applications, and 5G-enabled real-time traffic management systems, the architecture will need to address edge-cloud continuum reliability [12] and an increased attack surface across vehicle-to-infrastructure communications [7].

4. Conclusion

This paper has presented a resilience engineering framework for cloud-based transportation systems, grounded in the Hollnagel et al. anticipate-monitor-respond-adapt model and evaluated against NIST Cybersecurity Framework v1.1 and NIST SP 800-53 Rev. 5 HIGH baseline criteria. The proposed multi-layered framework achieves 85% overall NIST CSF control coverage and demonstrates that real-time tolling systems can simultaneously achieve 99.99%+ annual availability, sub-30-second automated failover, 15-minute recovery point objectives, and alignment with federal critical infrastructure security requirements.

Five architectural dimensions collectively realise the four Hollnagel resilience abilities: multi-zone distributed infrastructure absorbs zone-level failures, IaC-driven automation and self-healing orchestration enable rapid response and recovery, unified observability provides the continuous monitoring necessary for proactive risk anticipation, chaos engineering empirically validates resilience capabilities, and integrated cyber resilience controls protect against the evolving adversarial threat landscape facing transportation infrastructure.

The framework's primary contribution is a compliance-mapped, operationally validated architectural pattern that treats resilience as a design requirement integrated into cloud infrastructure from inception rather than an operational afterthought. For transportation agencies managing mission-critical tolling and mobility services, this framework provides a principled foundation for cloud adoption that simultaneously addresses operational continuity, cybersecurity, regulatory compliance, and revenue protection objectives. As transportation ecosystems continue evolving toward cloud-native, data-driven, and intelligently connected models, resilience engineering will remain the foundational discipline ensuring that these systems remain continuously available to the public they serve.

Acknowledgments

The author thanks the reviewers for their constructive feedback. This research received no specific grant funding from any public, commercial, or not-for-profit agency.

Author Contributions

Prasada Rao Paleti: conceptualisation, methodology, framework design, writing original draft, writing review and editing.

Conflicts of Interest

The author declares no conflicts of interest.

References

[1] E. Hollnagel, D. D. Woods, and N. Leveson, Resilience Engineering: Concepts and Precepts, UK: Ashgate Publishing, 2006. [Online]. Available: https://www.researchgate.net/publication/50232053_Resilience_Engineering_Concepts_and_Precepts

[2] National Institute of Standards and Technology, "Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1," NIST CSWP 04162018, Apr. 2018. doi: 10.6028/NIST.CSWP.04162018. [Online]. Available:

https://nvlpubs.nist.gov/nistpubs/cswp/nist.cswp.04162018.pdf

[3] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," NIST SP 800-145, Sep. 2011. doi: 10.6028/NIST.SP.800-145. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/legacy/sp/nistspecialpublication800-145.pdf

[4] National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST SP 800-53 Rev. 5, Sep. 2020. doi: 10.6028/NIST.SP.800-53r5. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

[5] A. Basiri, N. Behnam, R. de Rooij, L. Hochstein, L. Kosewski, J. Reynolds, and C. Rosenthal, "Chaos Engineering," IEEE Softw., vol. 33, no. 3, pp. 35–41, May-Jun. 2016. doi: 10.1109/MS.2016.60. [Online]. Available: https://dl.acm.org/doi/10.1109/MS.2016.60

[6] L. Zhu, F. R. Yu, Y. Wang, B. Ning, and T. Tang, "Big Data Analytics in Intelligent Transportation Systems: A Survey," IEEE Trans. Intell. Transp. Syst., vol. 20, no. 1, pp. 383–398, Jan. 2019. doi: 10.1109/TITS.2018.2815678. [Online]. Available: https://ieeexplore.ieee.org/document/8344848

[7] C. Liu and L. Ke, "Cloud Assisted Internet of Things Intelligent Transportation System and the Traffic Control System in the Smart City," J. Control Decis., vol. 10, no. 2, pp. 174–187, 2022. doi: 10.1080/23307706.2021.2024460. [Online]. Available: https://www.tandfonline.com/doi/abs/10.1080/23307706.2021.2024460

[8] G. Quattrocchi and D. A. Tamburri, "Infrastructure as Code," IEEE Softw., vol. 40, no. 1, pp. 37–40, 2023. [Online]. Available: https://www.researchgate.net/publication/366771598_Infrastructure_as_Code

[9] M. Usman, S. Ferlin, A. Brunstrom, and J. Taheri, "A Survey on Observability of Distributed Edge and Container-Based Microservices," IEEE Access, vol. 10, pp. 86904–86919, 2022. doi: 10.1109/ACCESS.2022.3193102. [Online]. Available: https://ieeexplore.ieee.org/document/9837035

[10] J. Kosińska, B. Baliś, M. Konieczny, M. Malawski, S. Zieliński, "Toward the Observability of Cloud-Native Applications: An Overview of the State-of-the-Art," IEEE Access, vol. 11, pp. 73036 - 73052, 2023. doi: 10.1109/ACCESS.2023.3281860. [Online]. Available: https://ieeexplore.ieee.org/document/10141603

[11] A. R. Hevner, S. T. March, J. Park, and S. Ram, "Design Science in Information Systems Research," MIS Q., vol. 28, no. 1, pp. 75–105, Mar. 2004. doi: 10.2307/25148625. [Online]. Available: https://dl.acm.org/doi/10.5555/2017212.2017217

[12] I. Dhanapala, S. Bharti, A. McGibney, and S. Rea, "Toward a Performance-Based Trustworthy Edge-Cloud Continuum," IEEE Access, vol. 12, pp. 99201–99212, 2024. doi: 10.1109/ACCESS.2024.3429197. [Online]. Available: https://doi.org/10.1109/ACCESS.2024.3429197

[13] R. Chandramouli, "Guide to a Secure Enterprise Network Landscape," NIST SP 800-215, Nov. 2022. [Online]. Available: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=935714

[14] X. Shen, Y. Lu, Y. Zhang, X. Liu, and L. Zhang, "An Innovative Data Integrity Verification Scheme in the Internet of Things Assisted Information Exchange in Transportation Systems," Cluster Computing, vol. 25, pp. 1791–1803, Jun. 2022. doi: 10.1007/s10586-021-03471-5. [Online]. Available: https://link.springer.com/article/10.1007/s10586-021-03471-5

[15] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST SP 800-207, Aug. 2020. doi: 10.6028/NIST.SP.800-207. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-207.pdf

[16] N. F. Syed, S. W. Shah, A. Shaghaghi, A. Anwar, Z. Baig, and R. Doss, "Zero Trust

Architecture (ZTA): A Comprehensive Survey," IEEE Access, vol. 10, pp. 57143–57179, 2022. doi: 10.1109/ACCESS.2022.3174679. [Online]. Available: https://ieeexplore.ieee.org/document/9773736

[17] S. Sarkar, G. Choudhary, S. K. Shandilya, A. Hussain, and H. Kim, "Security of Zero Trust Networks in Cloud Computing: A Comparative Review," Sustainability, vol. 14, no. 18, p. 11213, Sep. 2022. doi: 10.3390/su141811213. [Online]. Available: https://www.mdpi.com/2071-1050/14/18/11213

[18] C. Itodo and M. Ozer, "Multivocal Literature Review on Zero-Trust Security Implementation," Computers & Security, vol. 141, p. 103827, 2024. doi: 10.1016/j.cose.2024.103827. [Online]. Available: https://www.sciencedirect.com/science/article/abs/pii/S0167404824001287

[19] O. C. Edo, D. Ang, P. Billakota, and J. C. Ho, "A Zero Trust Architecture for Health Information Systems," Health and Technology, vol. 14, pp. 189–199, 2023. doi: 10.1007/s12553-023-00809-4. [Online]. Available: https://link.springer.com/article/10.1007/s12553-023-00809-4

[20] A. T. Ur Rehman Butt, T. Mahmood, T. Saba, S. A. O. Bahaj, F. S. Alamri, M. W. Iqbal, "An Optimized Role-Based Access Control Using Trust Mechanism in E-Health Cloud Environment," IEEE Access, vol. 11, pp. 138813–138826, 2023. doi: 10.1109/ACCESS.2023.3341765. [Online]. Available: https://ieeexplore.ieee.org/document/10327710