Universal Library of Innovative Research and Studies Research Article

ISSN: 3065-0003 | Volume 3, Issue 2

Open Access | PP: 44-50

DOI: https://doi.org/10.70315/uloap.ulirs.2026.0302009

Zero-Trust Architecture for Decentralized Edge Computing: Principles of Data Sovereignty and Regulatory Compliance

Myroslav Mishov

Founder, TECH EVOLVERS INC, Round Rock, TX, USA.
ORCID: 0009-0005-0023-387X

Abstract

This article analyzes contemporary approaches to building Zero-Trust architectures for decentralized edge computing under conditions of increasing requirements for digital sovereignty, localized data processing, and regulatory compliance. The study is conducted as a systematic review and analytical synthesis of publications focused on Zero Trust Architecture, decentralized edge computing, data sovereignty, federated learning, Self-Sovereign Identity, Compliance-as-Code, and distributed trust verification. Particular attention is given to the relationship between localized data processing, continuous trust verification, decentralized identity management, telemetry analysis, and adaptive policy enforcement. The study identifies the main limitations of centralized cloud-dependent security models, including policy propagation delays, dependence on external trust services, risks of cross-border data transfer, and inconsistencies between infrastructure behavior and regulatory control mechanisms. It is substantiated that static Zero-Trust models cannot ensure resilient management of distributed edge environments without continuous correlation between telemetry, node states, service behavior, access policies, and compliance parameters. Within the framework of the study, an original model entitled Sovereignty-Aware Zero-Trust Architecture for Decentralized Edge Computing is proposed, in which Zero Trust is interpreted as a mechanism for continuous governance of sovereignty, security, and regulatory compliance in distributed infrastructures. The article may be useful for researchers, cybersecurity specialists, developers of edge infrastructures, and organizations implementing sovereign computing and compliance management systems.

Keywords: *Distributed Edge Computing, Digital Sovereignty, Continuous Trust Verification, Decentralized Identity, Federated Learning, Adaptive Policy Orchestration, Regulatory Compliance.*

INTRODUCTION

The proliferation of distributed edge computing, the industrial Internet of Things, and artificial intelligence systems has led to a growing shift of data processing from centralized cloud platforms directly to the network edge [1]. Edge computing is no longer limited to the function of reducing latency and network load. It is becoming the foundation of localized data processing, autonomous control, and digital sovereignty in distributed infrastructures [8]. However, distributed architecture complicates access control, identity verification, and the enforcement of security policies in heterogeneous and potentially untrusted environments [2].

Studies demonstrate that traditional security models based on trust in the internal network perimeter are becoming insufficient for distributed edge ecosystems [3]. Under these conditions, zero-trust architecture is regarded as one of the most promising security models because it assumes continuous verification of identity, device state, and access context regardless of the subject's location within the network. At the same time, requirements for data sovereignty and compliance with regulatory frameworks, including GDPR and NIS2, are becoming increasingly stringent [10]. The localization of computing alone does not guarantee digital sovereignty. Even when data processing is performed at the edge, the infrastructure may still remain dependent on centralized management platforms and external identity systems. Therefore, federated learning, decentralized identity, distributed ledger technologies, and automated enforcement of security policies are gaining increasing importance [15].

The purpose of the study is to develop an original zero-trust architecture for distributed edge computing focused on ensuring data sovereignty, continuous trust verification, and regulatory compliance management within distributed edge ecosystems. To achieve this objective, the following tasks are addressed in the article:

Citation: Myroslav Mishov, "Zero-Trust Architecture for Decentralized Edge Computing: Principles of Data Sovereignty and Regulatory Compliance", Universal Library of Innovative Research and Studies, 2026; 3(2): 44-50. DOI: https://doi.org/10.70315/uloap.ulirs.2026.0302009.

– to analyze the role of zero-trust architecture in distributed edge environments;

– to determine the interrelationship between data sovereignty, operational sovereignty, and decentralized identity management;

– to investigate the influence of federated learning, localized artificial intelligence, and adaptive trust assessment on the security of edge ecosystems;

– to identify the limitations of centralized cloud-dependent security models under conditions of distributed edge infrastructures.

The research hypothesis is that distributed edge computing can ensure sustainable digital sovereignty and regulatory compliance only if zero-trust principles are integrated directly into the operational architecture of distributed infrastructure. It is assumed that continuous trust verification, decentralized identity validation, localized policy enforcement, intelligent telemetry analysis, and automated compliance control mechanisms make it possible to establish a sovereign edge ecosystem in which security, governance, and regulatory compliance operate as a unified continuously managed framework.

The scientific novelty of the study lies in the reinterpretation of zero-trust architecture not only as a cybersecurity model but also as an operational model of digital sovereignty for distributed edge computing. Unlike most existing studies, which are primarily focused on individual mechanisms of authentication, access control, federated learning, or distributed-ledger-based protection, this work proposes considering sovereign edge infrastructure as a multi-layer system of continuous trust management integrating infrastructural sovereignty, decentralized identity, adaptive policy enforcement, regulatory compliance management, federated artificial intelligence, and immutable auditing.

MATERIALS AND METHODS

The study is based on a theoretical and comparative analysis of approaches to constructing zero-trust architectures in distributed edge computing. Particular attention is given to the interrelationship between localized data processing, continuous trust verification, decentralized identity management, federated learning, and automated enforcement of security policies. Distributed edge infrastructure is considered as a system of interaction between edge nodes, identity mechanisms, access policies, and data management tools, in which resilience is determined by the consistency of trust relationships and information-processing processes.

The research was conducted in the format of a systematic review of open-access scientific publications from 2022–2025 indexed in Google Scholar, ScienceDirect, SpringerLink, MDPI, and arXiv. The search was performed using the following keywords: “Zero Trust Architecture,” “decentralized edge computing,” “data sovereignty,” “federated learning,” “Self-Sovereign Identity,” “compliance-as-code,” “blockchain auditability,” “NIS2,” and “GDPR.” The selection included studies focused on zero-trust architectures, federated learning, decentralized identity, data sovereignty, and regulatory compliance management. Publications limited to the description of individual software tools without analyzing their architectural role were excluded.

At the initial stage, 27 publications were identified. After removing duplicate materials and conducting a content analysis, the final sample consisted of 15 studies reflecting the main aspects of distributed edge computing, federated artificial intelligence, distributed ledger technologies, and regulatory compliance management.

During the analysis, the following groups of factors were identified: the dependence of distributed infrastructures on centralized management platforms, the limitations of traditional network trust models, the risks of cross-border data transfer, problems of identity verification, and the limitations of static access-control mechanisms. It was established that a significant proportion of threats is associated with inconsistencies between security policies, the actual behavior of the infrastructure, and regulatory compliance conditions. The analysis demonstrated that the localization of computing does not ensure digital sovereignty without continuous trust verification, automated policy enforcement, and traceable data lifecycle management.

The limitations of the study are associated with the predominance of analytical and review publications, as well as the limited number of studies devoted to distributed edge infrastructures operating independently of centralized cloud platforms. The obtained results are used to develop an original multi-layer zero-trust architecture model for distributed edge computing.

RESULTS

Distributed edge infrastructures create a direct conflict between centralized security management and the need for local operational autonomy. In classical cloud models, trust management, policy enforcement, and telemetry processing are concentrated within a unified coordination framework. For edge computing, this logic becomes a limitation. As the number of edge nodes, mobile services, and distributed computational workloads increases, policy propagation delays, communication channel instability, and dependence on centralized trust services begin to directly affect infrastructure resilience [83]. As a result, the sovereignty of a distributed environment is determined by the infrastructure’s ability to continuously validate trust status, enforce security policies, and maintain autonomous governance independently of the external cloud control layer [4].

Static zero-trust architectures are poorly adapted to

distributed edge environments. The problem is not limited to decision-making latency alone. Predefined trust policies cease to correspond to the actual behavior of the infrastructure, where computational workloads constantly migrate between edge clusters, nodes possess heterogeneous computational resources, and network connections dynamically change [11]. Under high workloads, policy propagation delays and excessive orchestration overhead begin to compete with the continuity of infrastructure operation itself. Under these conditions, zero trust ceases to function merely as an access-verification mechanism and transforms into a system of continuous distributed-environment state management.

For sovereign edge infrastructures, this becomes critical. If trust recalculation is performed too slowly, the system loses the ability to localize threats before they spread across infrastructure segments. Conversely, if verification requires constant interaction with centralized cloud services, the edge architecture remains dependent on the external management layer and loses part of its operational sovereignty [5]. For this reason, distributed edge computing requires a transition from static policy enforcement to telemetry-driven adaptive orchestration, in which the trust level is continuously recalculated on the basis of the actual state of nodes, service behavior, network activity, and runtime telemetry.

In practice, adaptive zero trust increasingly integrates intelligent trust assessment, decentralized identity management, and automatic modification of security policies [7]. Within this model, telemetry ceases to function merely as an auxiliary monitoring mechanism and becomes the foundation of continuous trust management. Edge infrastructure begins to operate as a dynamically verifiable environment in which access decisions are continuously adapted to changes in node conditions, anomaly levels, synchronization delays, and policy-execution conditions. Table 1 presents a comparative transformation of the main parameters of zero-trust architectures during the transition from the traditional static model to adaptive telemetry-driven orchestration.

Table 1. Comparative Performance of Traditional and Adaptive Zero-Trust Architectures in Distributed Edge Environments (Compiled by the author based on source: [3])

Metric	Traditional ZT	AI-Based ZT	SecureChain-ZT
Authentication Accuracy, %	89,4	94,2	98,6
False Acceptance Rate (FAR), %	7,8	3,4	1,2
False Rejection Rate (FRR), %	3,6	1,1	0,2
Policy Update Time, ms	1200	450	180
Average Latency, ms	8,3	5,2	3,1
Threat Detection Accuracy, %	84,3	92,7	99,1

Note: The "Metric" column identifies the evaluated operational and security characteristics of Zero-Trust architectures, including authentication reliability, policy propagation efficiency, latency, and threat detection capability. "Traditional ZT" represents conventional static Zero-Trust implementations based on predefined trust policies and centralized control mechanisms. "AI-Based ZT" reflects architectures incorporating adaptive trust evaluation and intelligent policy adjustment. "SecureChain-ZT" represents a decentralized adaptive Zero-Trust framework integrating artificial intelligence, dynamic policy orchestration, and distributed trust verification mechanisms for edge and 5G environments.

The reduction of policy update latency changes the very feasibility of applying zero-trust architecture in environments with critical latency requirements. At higher latency levels, continuous verification begins to conflict with the continuity of distributed infrastructure operation. For the industrial Internet of Things, distributed intelligent analytics, and autonomous control systems, this becomes an architectural limitation because excessive policy enforcement delays disrupt coordination between edge nodes [12]. In adaptive telemetry-driven architectures, policy enforcement is increasingly transferred directly to the edge. Node-state verification, trust recalculation, and decision-making processes begin to be performed locally without constant dependence on centralized cloud management.

At the same time, the role of distributed ledger technologies is changing. In sovereign edge infrastructures, blockchain ceases to function as a financial ledger and transforms into a mechanism of infrastructural verifiability required for real-time trust-policy enforcement. Immutable telemetry logs, distributed identity verification, and signed policy-execution records create an independent verification layer that enables edge domains to maintain continuity of trust even under partial isolation from the external cloud environment [9].

At the operational level, adaptive zero trust is becoming increasingly integrated with distributed orchestration systems. Policy-management mechanisms begin to function as a continuously synchronized governance layer between intelligent telemetry analytics, compliance verification, and the coordination of distributed computational workloads. As a result, security no longer exists as an external network-perimeter mechanism and instead becomes an embedded component of the infrastructure lifecycle [5].

In distributed edge ecosystems, data sovereignty increasingly conflicts with centralized models of intelligent computing governance. When traditional cloud architectures are used,

sensitive data, telemetry, and model-training processes are transferred to an external computational environment. This strengthens dependence on large cloud platforms, increases the risks of cross-border data transfer, and complicates control over digital infrastructure [14]. The localization of computing partially reduces these risks; however, it does not by itself establish a sovereign architecture. Sovereignty emerges only when the edge environment is capable of continuously controlling data flows, policy enforcement, model provenance, and identity verification throughout the entire operational lifecycle.

Federated learning is gradually becoming the foundation of intelligent infrastructure that preserves data sovereignty. Localized model training makes it possible to retain raw data within edge domains while transmitting only model updates or aggregated telemetry between nodes [13]. However, distributed intelligent environments create their own architectural conflict. The higher the level of privacy protection, the more significantly model utility, synchronization efficiency, and collaborative learning stability decrease. For sovereign edge infrastructures, this trade-off becomes not merely a theoretical limitation but an operational challenge.

For this reason, privacy-preserving intelligent edge systems require a continuous balance between data minimization, continuous verification, and functional resilience. Static compliance mechanisms prove insufficient under such conditions. The governance of distributed intelligent systems increasingly depends on the continuous verification of model updates, node behavior, and telemetry integrity [7]. Without such verification, federated infrastructures become vulnerable to model poisoning, malicious updates, inference attacks, and unauthorized model propagation. Table 2 presents the influence of differential privacy intensity on the efficiency of federated learning in distributed edge environments.

Table 2. Impact of Differential Privacy Level on Federated Learning Accuracy in Decentralized Edge Environments (Compiled by the author based on Source: [2])

ε Value	MNIST Accuracy, %	CIFAR-10 Accuracy, %	Privacy Regime Characteristic
0,0	92,02	93,10	No privacy noise
0,2	84,28	85,33	Moderate privacy protection
0,4	82,22	74,41	Enhanced privacy protection
0,6	75,44	61,90	High privacy level
0,8	59,70	54,10	Severe utility degradation
1,0	54,78	48,56	Maximum privacy noise

Note: The "ε Value" column represents the differential privacy parameter controlling the intensity of privacy-preserving noise added during federated learning; lower values correspond to weaker privacy protection, while higher values indicate stronger privacy enforcement. "MNIST Accuracy" and "CIFAR-10 Accuracy" indicate the classification accuracy achieved on the MNIST and CIFAR-10 datasets under different privacy conditions. The "Privacy Regime Characteristic" column describes the operational privacy level associated with each ε setting, including the balance between privacy preservation and model utility degradation in decentralized edge environments.

The increase in the level of noise-based protection radically changes the operational behavior of distributed intelligent systems. At high ε values, federated models gradually lose the ability to maintain stable autonomous orchestration in edge environments with strict latency requirements. For sovereign infrastructures, this means that privacy cannot be regarded solely as an isolated cryptographic property. It becomes part of a continuous governance architecture in which policy enforcement, telemetry verification, computational workload coordination, and trust recalculation function as an integrated system.

In practice, distributed intelligent analytics is gradually shifting toward compliance architectures embedded directly into the infrastructure itself. GDPR compliance verification, identity governance, controlled data disclosure, and model provenance control are increasingly executed within the edge operational environment [4]. Within this model, compliance ceases to function as an external audit procedure performed after data processing and instead becomes a continuously enforced infrastructural state.

This also changes the role of distributed identity infrastructures. Self-sovereign identity systems begin to function not merely as conventional user-authentication services but as trust-coordination mechanisms between edge computational workloads, intelligent agents, and autonomous nodes [6]. As a result, distributed edge computing gradually evolves into infrastructures inherently designed around continuous trust verification, where sovereignty is maintained through continuous control over data states, identity, security policies, and interactions between distributed operational domains.

DISCUSSION

Distributed edge computing is increasingly being deployed in infrastructures where the loss of control over data becomes both a technical and managerial problem [4]. The localization of computing alone does not ensure digital sovereignty. Data may physically remain within a local environment; however, identity mechanisms, orchestration, trust verification, and policy enforcement continue to depend on centralized

management platforms. For edge environments, this becomes critical. Under conditions of high workload mobility, unstable connectivity, and distributed information processing, trust ceases to function as a static property of the infrastructure and instead becomes a continuously verifiable state of the system.

Traditional cloud-governance models are oriented toward centralized control. This logic is poorly aligned with distributed edge ecosystems in which computing is performed simultaneously across edge nodes, local clusters, and autonomous network segments [7]. When centralized verification mechanisms are used, policy propagation latency increases, inter-domain coordination becomes more complex, and infrastructure dependence on external decision-making points intensifies. As a result, even a temporary disruption of communication between the edge and the centralized management layer can lead to the loss of operational resilience.

Under such conditions, zero-trust architecture ceases to function solely as an access-control system. It began to operate as an operational model of sovereign infrastructure governance. Sovereignty is established not through network isolation but through the ability to continuously verify access legitimacy, policy verifiability, telemetry provenance, and the correctness of regulatory enforcement. For distributed edge environments, trust becomes a computable parameter dependent on infrastructure state, node behavior, access context, and risk level.

Distributed infrastructures are also increasingly affected by the strengthening of regulatory frameworks. GDPR, NIS2, and the Data Governance Act are aimed at protecting data and ensuring traceability, accountability, and controllability of information-processing operations. Under these conditions, static compliance mechanisms are no longer effective. Verifying compliance only once during system deployment no longer reflects the actual state of distributed infrastructure. The configuration of edge nodes changes dynamically. Workloads migrate between network segments. Policies are updated in real time.

For this reason, compliance is gradually transforming into an executable runtime process. Policy-as-Code and Compliance-as-Code are increasingly functioning as mechanisms of continuous infrastructure-state control [5]. Policies cease to exist merely as sets of formal documents. They become mechanisms for automatic access restriction, verification of data-transfer routes, control over processing jurisdiction, and governance of distributed-service behavior. This is especially important for edge environments because manual synchronization of policies between distributed nodes leads to latency, configuration conflicts, and orchestration errors.

Decentralized identity mechanisms play a significant role in this transformation [6]. Self-sovereign identity reduces infrastructure dependence on centralized trust providers. Within such models, authentication is performed cryptographically and does not require continuous interaction with a single authorization center. For edge infrastructures, this becomes one of the primary factors of resilience.

At the same time, distributed ledgers are beginning to perform the function of infrastructural verifiability [3]. In distributed edge systems, blockchain ceases to function as a financial technology and instead becomes a mechanism for immutable verification of operational history. Immutable audit logs make it possible to trace changes in policies, data-processing routes, trust relationships, and access parameters. Figure 1 presents the original multi-layer architecture of sovereign governance for distributed edge computing based on continuous trust verification, executable regulatory control, and decentralized coordination.

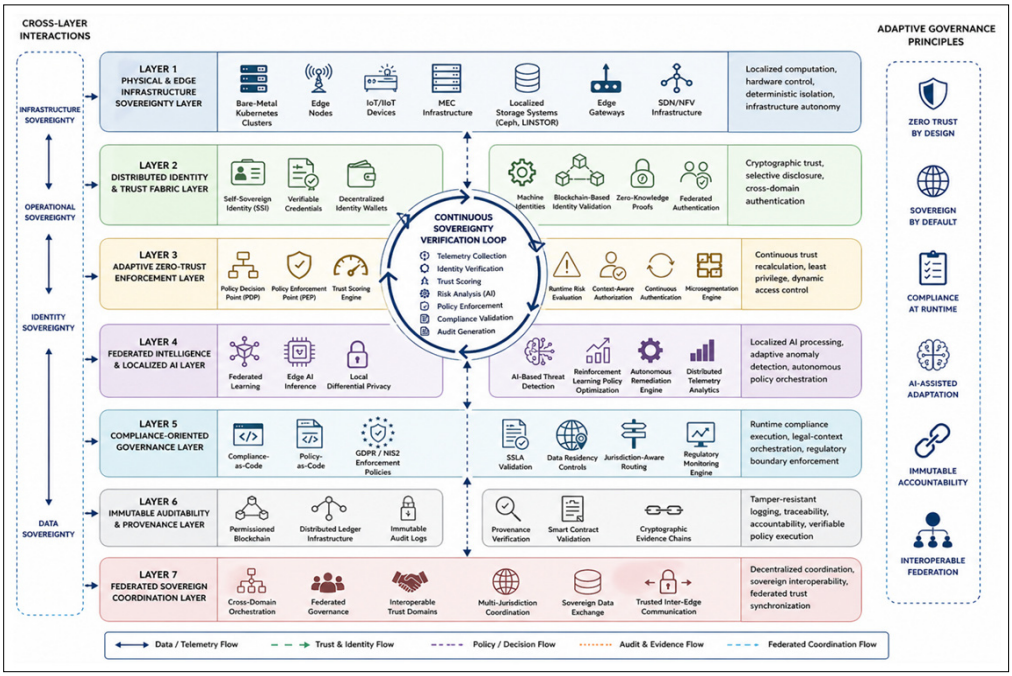

Figure 1. Multi-Layer Sovereignty-Aware Zero-Trust Architecture for Decentralized Edge Computing (Author's Development)

The proposed model considers distributed edge infrastructure as an adaptive system of sovereign governance in which security, identity management, regulatory compliance, and data processing operate as a unified continuous process. Within this architecture, sovereignty is established through continuous verification of trust, infrastructure state, policy legitimacy, and data-processing conditions. The foundation of the model is an autonomous edge infrastructure deployed on physical hardware, reducing dependence on centralized cloud platforms and external management mechanisms.

The central role in the architecture is assigned to continuous trust recalculation. Identity verification, access control, and risk assessment are performed dynamically on the basis of telemetry, service behavior, and the current state of the distributed environment. As a result, trust ceases to function as a static network parameter and instead becomes a continuously updated operational state of the infrastructure.

Artificial intelligence within the model is used as a mechanism for adaptive policy enforcement. The system analyzes local telemetry, detects anomalies, adjusts governance rules, and initiates automated threat responses without transferring sensitive data to external processing environments. At the same time, regulatory compliance is integrated directly into infrastructure-execution processes. Data-jurisdiction control, policy verification, and operational auditing are performed continuously, while immutable logs ensure the verifiability of all actions within the distributed environment.

The central element of the architecture is the continuous sovereignty-verification cycle, which integrates telemetry, identity management, risk analysis, policy enforcement, and audit generation into a unified coordination system. Within this model, sovereignty is regarded as a dynamic state maintained throughout the operation of the distributed edge ecosystem.

CONCLUSION

The conducted analysis demonstrates that data sovereignty in distributed edge computing cannot be ensured solely through the physical localization of information processing. In most cases, the key risks emerge when data remain within the local environment, while identity mechanisms, trust verification, policy governance, and regulatory control continue to depend on centralized cloud platforms. Under such conditions, edge infrastructure retains external operational dependence even when data processing is formally performed locally.

Particular importance is attached to the transition from external regulatory control to executable runtime compliance. For edge ecosystems, compliance with GDPR, NIS2, and other regulatory requirements must be integrated directly into data-processing, routing, identity-verification, and policy-enforcement processes. Under these conditions, compliance ceases to function as a post-audit procedure and instead becomes a continuous infrastructural state.

Within the framework of the study, a multi-layer architecture of sovereign governance for distributed edge computing based on Zero Trust principles is proposed. In this model, security, decentralized identity management, localized artificial intelligence, adaptive policy enforcement, regulatory compliance governance, and immutable auditing are integrated into a unified cycle of continuous sovereignty verification. The central element of the model is the continuous coordination between telemetry, trust assessment, access control, policy verification, and action logging within the distributed environment.

The practical significance of the study is associated with the development of an approach to constructing edge infrastructures capable of maintaining autonomy under unstable connectivity with centralized cloud platforms. Such an architecture may be applicable to the industrial Internet of Things, localized AI systems, autonomous industrial environments, and distributed computational infrastructures in which control over data, policies, and identity must remain within the local operational perimeter.

Prospects for further research may be associated with the formalization of operational sovereignty metrics, the development of mechanisms for dynamic trust recalculation in distributed edge clusters, and the experimental evaluation of architectures integrating Zero Trust, Self-Sovereign Identity, federated learning, and Compliance-as-Code into a unified system of continuous distributed infrastructure governance.

REFERENCES

1. Abbas, A. E., van Velzen, T., Ofe, H., et al. (2024). Beyond control over data: Conceptualizing data sovereignty from a social contract perspective. Electronic Markets, 34, 20. https://doi.org/10.1007/s12525-024-00695-2
2. Abbas, Z., Ahmad, S. F., Anjum, A., Syed, M. H., Malik, S. U. R., & Rehman, S. (2025). Ensuring zero trust in GDPR-compliant deep federated learning architecture. Computers, 14(8), 317. https://doi.org/10.3390/computers14080317
3. Alnaim, A. K. (2025). Adaptive zero trust policy management framework in 5G networks. Mathematics, 13(9), 1501. https://doi.org/10.3390/math13091501
4. Babel, M., Willburger, L., Lautenschlager, J., et al. (2025). Self-sovereign identity and digital wallets. Electronic Markets, 35, 28. https://doi.org/10.1007/s12525-025-00772-0
5. Ficili, I., Giacobbe, M., Tricomi, G., & Puliafito, A. (2025). From sensors to data intelligence: Leveraging IoT, cloud, and edge computing with AI. Sensors, 25(6), 1763. https://doi.org/10.3390/s25061763

6. Ghafoor, A., Symeonidis, I., Rydberg, A., Lindahl, C., & Abbasi, A. Q. (2025). Towards empowering stakeholders through decentralized trust and secure livestock data sharing. Cryptography, 9(3), 52. https://doi.org/10.3390/cryptography9030052

7. Hussain, N., Li, S., Hussain, A., et al. (2026). Quantum-aware secure blockchain intrusion detection system for industrial IoT networks. Scientific Reports, 16, 2265. https://doi.org/10.1038/s41598-025-31985-0

8. Jahnke, N., Rohde, M., & Kraus, T. (2025). Edge computing for digital sovereignty in the data economy. In U. Schmuntzsch, A. Shajek, & E. A. Hartmann (Eds.), New digital work II. Springer. https://doi.org/10.1007/978-3-031-69994-8_15

9. J Nair, A., Manohar, S., & Rao A B, S. (2025). Self sovereign identity in e-governance: Blockchain solutions for fintech compliance and citizen-centric financial services. Humanities and Social Sciences Communications, 12, 1562. https://doi.org/10.1057/s41599-025-05880-y

10. Kianpour, M., Earls Davis, P. A., & Windekilde, I. M. (2025). Digital sovereignty in practice: Analyzing the EU's NIS2 directive. International Journal of Information Security, 24, 167. https://doi.org/10.1007/s10207-025-01090-4

11. Kotulski, Z., Nowak, T., Sepczuk, M., et al. (2024). Keeping verticals' sovereignty during application migration in continuum. Journal of Network and Systems Management, 32, 67. https://doi.org/10.1007/s10922-024-09843-7

12. Li, S., Iqbal, M., & Saxena, N. (2024). Future industry Internet of Things with zero-trust security. Information Systems Frontiers, 26, 1653–1666. https://doi.org/10.1007/s10796-021-10199-5

13. Misra, S., Barik, K., & Kvalvik, P. (2025). Trust in digital sovereignty: A review of security, privacy, and governance challenges. Public Organization Review. https://doi.org/10.1007/s11115-025-00968-0

14. Patil, A., Mishra, B., Chockalingam, S., et al. (2025). Securing financial systems through data sovereignty: A systematic review of approaches and regulations. International Journal of Information Security, 24, 159. https://doi.org/10.1007/s10207-025-01074-4

15. Roberts, H. (2024). Digital sovereignty and artificial intelligence: A normative approach. Ethics and Information Technology, 26, 70. https://doi.org/10.1007/s10676-024-09810-5