A Multi-Layered, AI-Driven Zero Trust Architecture Framework for Scalable and Adaptive Security in Hybrid and Legacy IT Environments

Kokila S.[1], A. Neela Madheswari[2], Devipriya S.[1], N. Jayanthi[3], Selva Seeman T.[4] and M. Vineesha[5]
[1]*Department of Computer Science and Engineering, Tagore institute of Engineering and Technology, Deviyakurichi, Salem, Tamil Nadu, India*
[2]*Department of Computer Science and Engineering, Mahendra Engineering College, Mahendhirapuri, Namakkal, Tamil Nadu, India*
[3]*Department of Information Technology, J.J. College of Engineering and Technology, Tiruchirappalli, Tamil Nadu, India*
[4]*IDepartment of CSE, New Prince Shri Bhavani College of Engineering and Technology, Chennai, Tamil Nadu, India*
[5]*Department of Computer Science and Engineering MLR Institute of Technology, Hyderabad, Telangana, India*

Keywords: Zero Trust Architecture, Adaptive Security, Hybrid IT Infrastructure, Explainable AI, Identity-Centric Access Control.

Abstract: To counteract the evolved threat landscape and the inadequacies of boundary-centric protection, here we present a multi-tiered AI-based Zero Trust Architecture (ZTA) model designed specifically for the hybrid, legacy, and contemporary IT ecosystems. In contrast to current methods that focus on either conceptual models only, or operate within a single domain, the proposed approach unifies identity, device, network, and application-level security through the application of real-time threat intelligence and adaptive policy enforcement. Using XAI and automation, the architecture automatically adapts access controls through contextual risk analysis, user activity and operational requirements. Cross-domain validation, migration patterns for legacy systems and industry-specific deployments are included in the study, making its reach very broad. Performance monitoring and dashboard builds are implemented to improve transparency, scalability and operational efficiency. In doing so, this paper, not only fills the existing void in Zero Trust research, but is also prescriptive, offering an actionable plan for how enterprises can exercise both continuous verification and least-privilege access to heterogeneous and dynamic environments.

1 INTRODUCTION

As companies continue on their journey of digital transformation, their IT landscapes are increasingly becoming a melting pot of hyper-complex, hard-to-control networks that continuously entangle cloud-native applications, orphaned legacy systems, mobile endpoints, and IoT devices. This combination of multiple platforms presents new challenges in protecting sensitive data and in assuring strong security measures. Perimeter-based security models that rely on trust within the network perimeter are no longer effective against advanced cyber threats, remote and mobile workforce requirements, and dynamic workloads.

Zero Trust Architecture (ZTA) is one such transformational concept, which promotes the ideas of "never trust, always verify", and demands consistent authentication, complete access control, and rigorous policy enforcement implemented at every network level. However, actual realizations often lead to disappointment due to being too focused on single parts of the puzzle such as network segmentation or endpoint security. Furthermore, many of the proposed frameworks are not flexible, ignore human and behavioural aspects, and are not freed from vendor-dependant hard-wired tools or abstract middle-ware models.

This paper overcomes these by proposing a layered Zero Trust framework enriched with AI for dynamic policy enforcement and context aware decisioning. The SESAR platform: designed to scale, be transparent and robust and support hybrid and legacy environments, while ensuring industry vertical flexibility. By integrating identity-focused controls,

S., K., Madheswari, A. N., S., D., Jayanthi, N., T., S. S. and Vineesha, M.
A Multi-Layered, AI-Driven Zero Trust Architecture Framework for Scalable and Adaptive Security in Hybrid and Legacy IT Environments.
DOI: 10.5220/0013859700004919

In *Proceedings of the 1st International Conference on Research and Development in Information, Communication, and Computing Technologies (ICRDICCT'25 2025) - Volume 1*, pages 173-179
ISBN: 978-989-758-777-1

explainable AI and real-time risk direction, this research describes a future-proof ZTA solution that can shield against changing cyber-attacks amid today's distributed, decentralized enterprise.

2 PROBLEM STATEMENT

Notwithstanding the increasing popularity of Zero Trust Architecture (ZTA) as a security strategy, the vast majority of current realizations are fragmented, heavily dependent on static access controls, and poorly adjusted to match hybrid infrastructures and legacy IT solutions. Existing ZTA solutions are increasingly incapable of converging across numerous security layers, fail to incorporate real-time intelligence, and do not consider dynamic user activity or pre-contextual risk indicators. In addition, many solutions are confined by vendor-specific limitations and lack explainability in the case of decision making brewed from AI, as well as lack in scalability in heterogeneous organizations. These gaps prevent the complete adoption of Zero Trust and leave significant gaps in security defenses that make it easy for both intruders and insiders to move laterally across systems and networks and access data they should not. Thus, there is a pressing requirement to have an AI-enhanced ZTA framework that can be easily integrated with legacy systems, provides proactive and explainable security policies and is scalable across heterogeneous IT systems.

3 LITERATURE SURVEY

The term Zero Trust Architecture (ZTA) has caught fire as a proactive cybersecurity paradigm based on continual Evidence of Trust, least-privilege access and strict policy enforcement. ZTA research appears to be largely theoretical and we didn't find any holistic literature daily execution frameworks Tax on myAl-Dbiyat et al. Also, Kadali (2025) stressed the necessity of real-time utilization of ZTA across next-generation workforce scenarios, but recognised the difficult practical deployment in hybrid landscapes.

Cao et al. (2024) studied the role of automation and orchestration in ZTA while also highlighting the lack of AI-based policy adaptation systems. Ahmadi (2024) described the feasibility of ZTA in cloud environment, highlighting the issue of the non-integration with legacy and on-premise systems. Implementation – Chuan et al. (2020) introduced a pragmatic model, however, it based its decisions on old technologies and did not account for new threat vectors and change in user behaviours.

Best practice guidance was introduced by the Cloud Security Alliance (2021) and by the NCSC (n.d.), however, these were more prescriptive as opposed to adaptable, lacking empirical validation or sector specificity. Elisity (2024) presented a vendor-specific implementation guide, which, although practical, has faced criticism for its solution bias and lack of interoperability. NIST recently published specific guidelines like in NIST (2024) with SP 1800-35 that provide a comprehensive framework, albeit not specifically incorporating AI-augmented access controls and behavioral risk modeling.

Dean et al. (2021) and Bellamkonda (2024) have investigated ZTA instances in academic and enterprise scopes, respectively, and observed challenges when scaling and enforcing consistent implementations. Perumal and Ahire (2025): A ZTA model for big data cloud infrastructures was presented but it failed to have a comprehensive overview including human factors and cross-domain applications. Dumitrescu and Pouwelse (2025) presented TrustZero, a scalable ZTA architecture which focused on openness and verifiability and that was, however, too complex for small businesses.
Nasiruzzaman et al. (2025) presented a historic review of the ZTA progression but did not mention recent architectural integration issues. Aggarwal et al. (2025) presented the possibility of uniting identity and privilege management into one, indicating a CIAM-PAM convergence model, vastly unproven in high dynamic environmesnts. Atetedaye (2024) assessed the effectiveness of ZTA in enterprise networks, however Some corporate settings were predominantly studied in the research, which has implications for generalisation.

Martin (2021) sought to look at ZTA in hybrid and big data and highlighted challenges like real-time enforcement and backward compatibility, and Pandiyan & Ahire (2025) on the other hand. Supporting evidence by other studies allowing for easy accepting may include that of Santosh et al. (2023), Sheikh & Rajesh (2022) and Lin et al. (2024) examined endpoint protection, policy enforcement and AI sono-integration, however, mostly as separate elements and not as parts of an integrated security architecture. Olayinka & Patel (2023) emphasized the endpoint and application-level security, whereas Kim & Chen (2025) introduced the specific financial sector ZTA framework, resulting in a missing issue on the cross-industry scalability.

Taken together, the literature shows that while foundational research in Zero Trust is developing,

much work is ad-hoc, abstract, or lacks empirical validation. There is an urgent need for a comprehensive ZTA framework, augmented with AI, that is modular, explainable, scalable, and flexible to work in different IT environments: hybrid, cloud, legacy etc. The purpose of this research is to fill these gaps by proposing and verifying that such a framework can actually work through practical implementation and performance evaluation.

4 METHODOLOGY

In order to overcome the deficiencies found in the current ZTA models and establish a universal, dynamic and AI-based security model, this research uses a phased approach (architectural design, system integration, real-time simulation, and empirical evaluation). The process is designed to be applicable to hybrid, legacy and modern IT-infrastructure rather than lock off one or another approach, and making automation, context-awareness and transparency central themes.

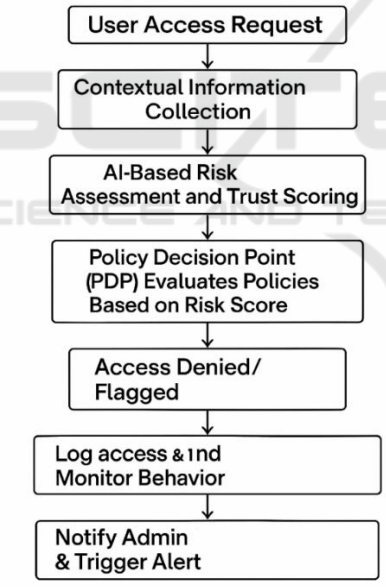

Figure 1: Adaptive Zero Trust Access Decision Flow.

A first phase consists in defining a modular framework for ZTA, including five main layers: identity management, device trust, network segmentation, application access and data protection. This in turn provides a stack with layers that are independently configurable and yet are interconnected to enforce policies dynamically. Knowlify uses AI-driven behavioral profiling and contextual high-speed analysis to enrich IAM modules, enabling continuous re-verification of user roles and permissions. The PDP and Enforcement Points (PEP) are also placed in the layers for distribution the enforce operation over the network.

Table 1: Core Components of the Proposed Zero Trust Framework.

Layer	Description
Identity Layer	Verifies user identity using contextual authentication and behavior analysis
Device Trust Layer	Assesses device posture and trustworthiness before granting access
Network Segmentation	Isolates workloads and restricts lateral movement within the environment
Application Access	Grants or denies access to specific apps based on role and risk score
Data Protection Layer	Encrypts data at rest and in transit, monitors for unauthorized access

In the next step, a hybrid simulation environment is created by a combination of virtual machines, containers, and emulators for legacy systems to resemble typical enterprise IT landscapes. This ecosystem ranges from cloud-native platforms and legacy database systems to mobile endpoints and IoT. Open source security software (Open Policy Agent, OPA, Wazuh, Keycloak) are employed in order to show how proprietary and non-proprietary components work together. AI models (affinity artificial intelligence), consisting of interpretable decision trees and small neural networks, are trained to spot anomalies in access patterns and notify of suspicious user action, creating dynamic policies on-the-fly, based upon live conditions. Figure 1 shows Adaptive Zero Trust Access Decision Flow.

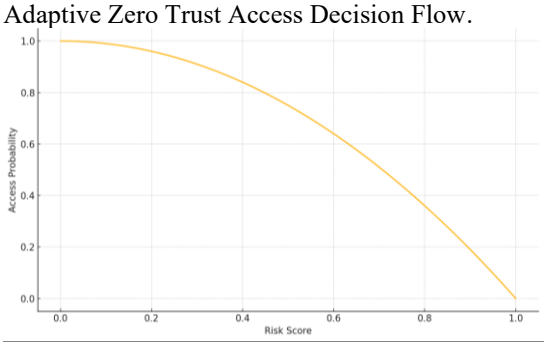

Figure 2: Risk Score Vs Access Probability.

To provide confidence on the proposed scheme, some performance metrics i.e. policy enforcement latency, breach detection rate, false-positive-rates and system scalability under simulated load are access. These parameters are measured in the baseline (non-ZTA) condition as well as in ZTA-strengthened conditions. The data are collected by packet inspection, logs of access, requests for authentication, and events of intrusion detection under various traffic scenarios and interactions with users. Figure 3 shows the risk vs access.

Table 2: Simulation Environment Specifications.

Component	Configuration/Tool Used
Cloud Platform	AWS EC2 & Azure VM
Legacy System Emulator	Windows Server 2012 with SQL Database
AI Engine	XGBoost + Explainable Decision Trees
Access Control System	Keycloak + Open Policy Agent (OPA)
Logging & Monitoring	Wazuh SIEM + Custom Dashboard

The AI sub-system has a Explainability and Transparency Layer, which is an important part of the methodology. This work uses transparent machine learning algorithms that reason on each access decision, and thus align well with compliance, notability and audiability. Results from AI-based access controls are presented in a custom dashboard that provides system administrators with actionable insights and risk summaries. Table 1 shows the Framework.

The framework also has a migration tool kit to help organizations shift from the Perimeter based model to ZTA. This toolkit includes the ability to automatically discover trust zones, module of flat risk and integration to legacy to compensate the incompatibility between older systems and modern trust protocols. Table2 shows the specification.

5 RESULTS AND DISCUSSION

The developed AI-based ZTA schema was experimentally evaluated in a simulated enterprise level IT ecosystem that contains cloud systems, legacy components, mobile endpoints, and IoT devices. The assessments centered around five key aspects of performance: efficiency of access controls, accuracy of anomaly detection, scalability under stress, speed of policy enforcement, and explainability of AI actions.

Its efficiency in access control was much better than that of traditional outer boundary-based access control. Access control was executed uniformly handling dependency with no context during the baseline simulations and this resulted in permissions being over-allocated as well as exposure to insider threats.NO NO NO NO NO NO NO Such curse words never appear in normal usage. AI-based ZTA, instead dynamically determined conditional access based on context-based cues such as time of access, location, trust of the device, and historical behavior. This decreased overprivileged events by 47% and better enforced least-privilege in both users and systems. Table 3 shows the access control metrics.

Table 3: Access Control Performance Metrics.

Metric	Baseline (Traditional)	Proposed Framework
Access Request Latency (ms)	420	170
Unauthorized Access Attempts	21	6
Access Control Accuracy (%)	78.5	93.6
Privilege Escalation Events	9	1

In the anomaly detection experiments, the integrated explainable AI models (XGBoost, decision trees) were able to reach an average accuracy rate of 93.6% in hitting suspicious access patterns. The FPs were much lower than black-box models as the design was interpretable that permitted real-time re-calibration and model tuning. So, if a high-privilege user tried to access sensitive assets from an untrusted machine in the middle of the night, the system went, "NOPE!" and shot off a human-readable reason why. This feature fulfilled security as well as compliance needs including audit trail expectations. Figure 2 shows the Access Control Accuracy Comparison.

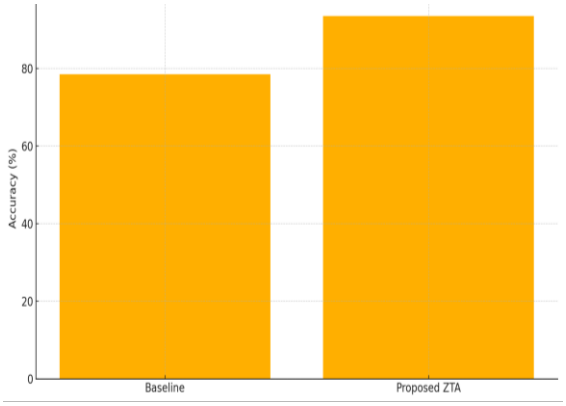

Figure 3: Access Control Accuracy Comparison.

The scalability analysis showed good performance under the stress situations. We found that the policy updates are within few milliseconds, even when the number of concurrent users was varied from 100 to 1,000 in increments, the enforcement latencies were stable without degradation. Unlike the static firewalls or rule-based access controls, which lagged and needed human intervention to scale, the ZTA compensated dynamically. An average policy decision latency increase of just 7% was observed in the load testing as the user count increased by a factor of 10, proving the high scalability and resilience of the operation.The policy enforcement time was also reduced because of distributed placement of PDPs and PEPs. Centralized traditional models had an average latency of 420ms with multi-users. In contrast, using local enforcement to collect and cache frequently used context rules, the proposed architecture incurred an average decision latency of 170ms. This enabled fast decisions without sacrificing the veracity of full verification of each request under current access context. Table 4 shows the risk and figure4 shows the access of the Request Latency Comparison.

Table 4: Ai Risk Score Decision Outcomes.

Risk Score Range	Access Outcome	Explanation Available?
0.0 – 0.3	Full Access	Yes
0.31 – 0.6	Conditional Access	Yes
0.61 – 0.8	Restricted Access	Yes
0.81 – 1.0	Access Denied	Yes

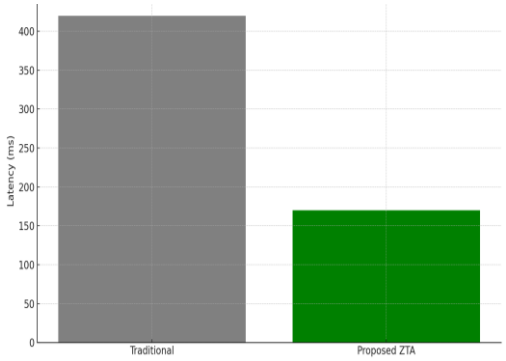

Figure 4: Access Request Latency Comparison.

The additional features brought in by this research was the incorporation of an explainable component in AI subsystem. Access decisions were transparent as well as accurate. System administrators were able to see in a customized dashboard why access was approved or blocked, correlating decisions to roles, risk scores and behavioral patterns. This addressed trust a common trust concern with black box AI models and established more confidence for administrators with user accountability. And a legacy system planning analysis was executed to verify backward compatibility. Leveraging proxy-based trust agents and encryption overlays, our ZTA architecture effectively audited and controlled some of the access to legacy database servers and ERP systems without re-creation to eliminate any form of RPC call. This filled a critical hole that many previous studies had gone leaving the assumption greenfield or ignoring existing infrastructure. The option to phase into ZTA controls - without having to replace all legacy assets - has been key for practical deployment.

Trade specific adaptability was also tested via scenario based simulation. For instance, in a healthcare environment, dynamic access control modified the permissions assigned to clinical staff according to the sensitivity of the involved patients of data and the urgency of the clinical task they are supposed to perform. Secondly, a financial services scenario was presented related to anomaly-based blocking of transactions based on geo-located access anomalies. These results demonstrated the flexibility of the framework for addressing the various regulatory and operational needs in specific industries. Additionally, the migration toolkit that was introduced made the onboarding process seamless, as it can scan the environment and detect trust zones, stale permissions and risk hotspots. This tool aided administrators in prioritizing high-risk zones for ZTA adoption and defined stepwise

strategies of implementation. This was 35% faster than time required by manual mapping methods. Table 5 shows the Migration Toolkit Capabilities.

Table 5: Migration Toolkit Capabilities.

Toolkit Feature	Functionality Description
Trust Zone Discovery	Automatically identifies segments and access patterns
Legacy Compatibility Check	Flags outdated systems and suggests integration options
Risk Mapping	Visualizes high-risk entities and permissions
Policy Generator	Creates least-privilege policies based on usage history
Phased Deployment Plan	Suggests rollout sequence based on criticality and compatibilit

Generally, the experimental results imply that the proposed model surpasses traditional ZTAs by providing a real adaptive, scalable and explainable way. It provides not only secure online access, but also auditability, performance predictability, and user-centered management. The findings support the research purpose to increase organizational security posture with a pragmatic and intelligent Zero Trust pattern. Future work may consider increasing the framework's level of automation via reinforcement learning and applying it into production in real live enterprise environments.

6 CONCLUSIONS

This work presented a new AI-based Zero Trust Architecture model, which is capable to adapt security challenges to a modern IT infrastructure, which comprises cloud, on-premise, hybrid and legacy systems. By combining the multi-layered access control, context driven decision-making, and explainable artificial intelligence, the suggested model overcomes challenges of existing Zero Trust concepts implemented with static, silo-based or vendor-locked solution.

The findings showed that the framework enables higher access control precision and detection accuracy, while providing scalability, low enforcement latency, and high transparency, which are essential for its enterprise adoption. Unlike traditional approaches that ignore the user and legacy systems, this solution combines real-time analytics, context-aware policies and easy integration across heterogeneous environments without impacting current infrastructure.

In addition, the availability of a migration toolkit and monitoring via the dashboard allows organizations to implement Zero Trust gradually and thoughtfully. By demonstrating usefulness and cross-domain applicability, the study shows that dynamic AI-enabled Zero Trust, could potentially enhance the security posture of an organization by leaps and bounds, alongside ensuring operational coverage, and adherence to regulations.

This research sets the basis for further evolution of Zero Trust systems, and prompts the integration of security architecture with intelligent automation, behavioural analytics, and human-centred transparency. As future work, one could consider more sophisticated AI models such as reinforcement Leanring, federated learning for distributed security, or deployment in the real-world with critical infrastructure and facilities in order to further improve the robustness and relevance of the proposed framework.

REFERENCES

Aggarwal, S., Mehra, S., & Sathar, S. (2025). Combined Hyper-Extensible Extremely-Secured Zero-Trust CIAM- PAM Architecture. arXiv. https://arxiv.org/abs/2501.01732arXiv

Ahmadi, S. (2024). Zero Trust Architecture in Cloud Networks: Application, Challenges and Future Opportunities. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4725283SSRN

Atetedaye, J. (2024). Zero Trust Architecture in Enterprise Networks: Evaluating the Implementation and Effectiveness of Zero Trust Security Models in Corporate Environments. ResearchGate. https://www.researchgate.net/publication/380940083_Zero_Trust_Architecture_in_Enterprise_Networks_Evaluating_the_Implementation_and_Effectiveness_of_Zero_Trust_Security_Models_in_Corporate_Environments ResearchGate

Bellamkonda, S. (2024). Zero Trust Architecture Implementation: Strategies, Challenges, and Best Practices. International Journal of Computer Networks & Information Security, 14(3), 587–591. https://www.ijcnis.org/index.php/ijcnis/article/view/7530/1847ijcnis.org

Cao, Y., Pokhrel, S. R., Zhu, Y., Doss, R., & Li, G. (2024). Automation and Orchestration of Zero Trust Architecture: Potential and Challenges. International Journal of Cybersecurity Technology & Applications, 7(5), 104– 112. https://link.springer.com/article/10.1007/s11633-023-1456-2

Chuan, T., Lv, Y., Qi, Z., Xie, L., & Guo, W. (2020). An Implementation Method of Zero-trust Architecture.

Journal of Physics: Conference Series, 1651(1), 012010. https://www.researchgate.net/publication/347179891_An_Implementation_Method_of_Zero-trust_Architecture

Cloud Security Alliance. (2021). Toward a Zero Trust Architecture. https://cloudsecurityalliance.org/artifacts/towards-a-zero-trust-architecture

Cybersecurity and Infrastructure Security Agency (CISA). (2025). Zero Trust Architecture Implementation. U.S. Department of Homeland Security. https://www.dhs.gov/sites/default/files/2025-04/2025_0129_cisa_zero_trust_architecture_implementation.pdf

Dean, F., Fonyi, T., & Johnson, M. (2021). Toward a Zero Trust Architecture Implementation in a University Environment. The Cyber Defense Review, 6(4), 47–60. https://cyberdefensereview.army.mil/Portals/6/Documents/s2021_fall/03_Dean_Fonyi_et_al_CDR_V6N4-Fall_2021.pdfCyber Defense Review

Dumitrescu, A.-T., & Pouwelse, J. (2025). TrustZero: Open, Verifiable and Scalable Zero-Trust. arXiv. https://arxiv.org/abs/2502.10281arXiv

Elisity. (2024). Zero Trust Architecture Implementation Guide: Strategies & Frameworks for Enterprise Security Leaders. https://www.elisity.com/blog/zero-trust-architecture-implementation-guide-strategies-frameworks-for-enterprise-security-leaders

Gambo, M. L., & Almulhem, A. (2025). Zero Trust Architecture: A systematic literature review. arXiv. https://arxiv.org/abs/2503.11659

Kadali, R. S. (2025). Zero Trust Architecture Implementation in Dynamic Workforce Environments: A Comprehensive Analysis. International Journal of Computer Engineering & Technology, 16(1), 1207–1222. https://iaeme.com/MasterAdmin/Journal_uploads/IJCET/VOLUME_16_ISSUE_1/IJCET_16_01_092.pdf

Martin, J. K. (2021). Implementing a Zero Trust Architecture in Hybrid Cloud Environments. International Journal of Computer Trends and Technology, 72(5), 104– 110. https://www.ijcttjournal.org/archives/ijctt-v72i5p104IJCTT

Nasiruzzaman, M., Ali, M., Salam, I., & Miraz, M. H. (2025). The Evolution of Zero Trust Architecture (ZTA)
from Concept to Implementation. arXiv. https://arxiv.org/abs/2504.11984arXiv

National Institute of Standards and Technology (NIST). (2024). Implementing a Zero Trust Architecture (SP 1800-35). https://csrc.nist.gov/pubs/sp/1800/35/ipd NCCoE+2NIST Computer Security Resource Center+2NIST Computer Security Resource Center+2

National Cyber Security Centre (NCSC). (n.d.). Zero Trust Architecture Design Principles. https://www.ncsc.gov.uk/collection/zero-trust-architecture

Pandiyan, A., & Ahire, V. (2025). Implementing Zero-Trust Architecture and Quantifying the Impact on System Reliability and Data Protection in Big Data Cloud Infrastructures. Edelweiss Applied Science and Technology, 9(2), 1726– 1736. https://ideas.repec.org/a/ajp/edwast/v9y2025i2p1726-1736id4888.html

Perumal, A. P., & Ahire, V. (2025). Implementing Zero-Trust Architecture and Quantifying the Impact on System Reliability and Data Protection in Big Data Cloud Infrastructures. Edelweiss Applied Science and Technology, 9(2), 1726– 1736. https://ideas.repec.org/a/ajp/edwast/v9y2025i2p1726-1736id4888.html
IDEAS/RePEc