International Journal Science and Technology
IJST Vol. 3 No. 3 | September 2024| ISSN: 2828-7223 (print), ISSN: 2828-7045 (online), Page 97-108

Privacy Preserving Access Control: Leveraging Data De-identification in AI-Enhanced Zero Trust Environments

Mukul Mangla
Independent Researcher, India

Article History
Received : October 03, 2024
Revised : October 18, 2024
Accepted : November 27, 2024
Published : November 30, 2024

Corresponding author*:
research.mukulmangla@gmail.com

Cite This Article:
Mangla, M. (2024). Privacy Preserving Access Control: Leveraging Data De-identification in AI-Enhanced Zero Trust Environments. *International Journal Science and Technology*, *3*(3), 97–108.

DOI:
https://doi.org/10.56127/ijst.v3i3.2276

Abstract: With the convergence of artificial intelligence (AI), big data, and cybersecurity, how organisations protect sensitive information in digital environments that have become increasingly hostile has been fundamentally changed. Whereas Zero Trust Architecture (ZTA) can be seen as a paradigmatic shift in the face of classical perimeter-based security, its reliance on continuous verification and access control at the granular level poses new privacy issues, especially in cases of personal or organisational data processing. Conventional access-control systems have suffered re-identification vulnerabilities, insider risks and adversarial attacks which take advantage of data exposure. In order to address these concerns, this paper proposes a privacy-aware access-control framework that embraces data de-identification methods such as anonymisation, pseudonymisation, and differential privacy to an AI-enhanced Zero Trust. The proposed framework can help reduce privacy risk and prevent system utility loss by integrating de-identification at the data-collection layer, risk scoring by AI, and adaptive access controls enabled by ZTA. Combining theoretical discussion with the current developments in the field of AI-based privacy protection, the study creates a unified framework that enhances the ability to withstand re-identification, support adherence to changing regulations, and foster trust in AI-supported access control. Results indicate that de-identification can be safely co-existent with AI-based decision-making to strengthen ZTA, but trade-offs continue to balance privacy protection with the accuracy of AI models. The given study can help the field of cybersecurity research through its systematic method of operationalising privacy-by-design in the context of Zero Trust, which can serve as a basis to future studies on federated learning, blockchain incorporation, and quantum-resistance models of access.

Keyword: Privacy preservation, Zero Trust Architecture, Data de-identification, Artificial Intelligence, Access control, Cybersecurity.

INTRODUCTION

The blistering development of digital ecosystems has created unprecedented prospects of implementing the concept of artificial intelligence (AI) to be used in security-sensitive settings. Nevertheless, such transformation is also introducing complicated issues in protecting data privacy and implementing strong access control. Organisations across healthcare, finance, government, and e-commerce increasingly rely on AI for real-time decision-making, predictive analytics, and adaptive security monitoring. Although these technical advances exist, the possibility of sensitive information being exploited in the AI-enhanced spaces makes the subject of confidentiality, compliance, and trustworthiness disturbing (Ijaiya, 2024; Praveenadevi et al., 2024).

The existence of traditional security models which were historically based on the use of perimeter security are not sufficient today in the distributed and cloud-based architectures. The concept of Zero Trust Architecture (ZTA) has become a strategic behaviour which does not trust anybody but instead verifies continuously, uses least-privilege access, and micro-segments (Oluoha et al., 2024). The idea of never trust, always verify adopted by ZTA makes security enforcement relevant in light of the hybrid and multi-cloud

environments. However, the ZTA will not, in itself, provide privacy assurances as the system has several vulnerabilities that legacy systems contain, though mitigated by the ZTA. Even when sensitive personal information is collected, stored and processed in ZTA, a risk of re-identification can remain significant unless strong protection is provided (Yang L. et al., 2024).

An area of potential improvement on the mitigation of these risks is the de-identification methods of data, which include anonymisation, pseudonymisation, k-anonymity, and differential privacy (Yang et al., 2024). These strategies decrease the risk that personal information may be traced to individuals when applied appropriately, decreasing the consequences of possible information breaches or insider abuse. At the same time, AI algorithms will be able to complement ZTA by offering adaptive access-control policies, anomaly detection, and real-time risk scoring that will ensure that access decisions are dynamic and context-aware (Ho, 2024; Kumar et al., 2024). The combination of AI-based intelligence and privacy-protecting de-identification services is a strong solution to the twin problems of privacy guarantee and effective access control.

The gap in the research is that few studies explore the holistic frameworks integrating these three pillars that include AI, de-identification, and Zero Trust. The majority of the literature analyzes these elements separately: AI as a dynamic access control tool, de-identification as a privacy measure, or ZTA as an architectural resilience tool. Not many of them suggest a combined model that operationalises the notion of de-identification in AI-enhanced ZTA settings. Sealing this divide is critical in achieving the ability to ensure that organisations can reap the advantages of AI-powered security without infringing on the privacy of users (Ijaiya, 2024; Oluoha et al., 2024).

Thus, the aim of the present paper is to suggest a privacy-sensitive access-control model that improves the utilization of data de-identification in ZTA that is enhanced with AI. In particular, the research will be seeking to:

1. Consider the de-identification in reducing the re-identification risk in Zero Trust systems.
2. Explore the possibilities of providing adaptive access control in privacy-sensitive situations using AI-driven models.
3. Suggest a multi-faceted conceptual model that would merge these elements to enhance security and privacy guarantees.

The value of this work is threefold: (1) developing the theoretical conceptualization of privacy protection in ZTA with the use of de-identification, (2) providing a unified AI-enhanced framework of adaptive and privacy-aware access control, and (3) locating the limitation and future research prospects, such as blockchain-based governance and quantum-resistant architecture. With these aims, the paper will aim to continue the discussion on how privacy can be operationalised by design in the AI-enhanced cybersecurity systems.

LITERATURE REVIEW

Zero Trust Architectures (ZTAs)

Zero Trust Architecture (ZTA) is a paradigm of cybersecurity because it replaces perimeter-focused defenses with the continuous checking and micro-segmentation. Although legacy models depended on implicit trust when users have passed through the network perimeter, ZTA is asserted upon the principle of never trust and always verify as its core (Oluoha et al., 2024). In turn, it therefore takes on certain responsibility to be of particular relevance in modern cloud-based and hybrid environments, where virtualization, mobile devices and remote access obscure network boundaries.

The guidelines on Zero Trust provided by NIST focus on adaptive access control, dynamic trust assessment, and the maintenance of constant monitoring of users and devices (Ho, 2024). In real life, ZTA brings together identity and access management (IAM), encryption and multi factor authentication into a single security model. However, even with such benefits, the conventional ZTA deployments do not always consider the issue of data privacy. Vulnerable data can be further re-identified or abused after it has been accessed by authorized employees regardless of whether it is healthcare data, a financial services dataset, or an e-commerce one (Ijaiya, 2024).

Furthermore, although micro-segmentation and strict access policy reduces the attack surface of ZTA, it may cause latency and complexity, especially at scale. The trick, then, is to see that Zero Trust is scalable, adaptable and privacy-conscious in more AI-driven settings (Praveenadevi et al., 2024).

Privacy Preservation in Access Control

The aspect of privacy protection is one of the pillars of data management in the modern digital spaces. Rules like the GDPR and HIPAA impose strict data protection conditions that require organizations to

implement methods that help protect sensitive data and maintain its usefulness. There has been an increase in the notions of data de-identification (anonymization, pseudonymization, k-anonymity, differential privacy) as solutions that can be used to reduce the risk of unauthorized disclosure (Ok and Blessing Bright, 2023; Yang et al., 2024).

Anonymization is done to eliminate the personally identifiable information (PII) to avoid the ability to trace the data to a specific person. Pseudonymization replaces PII with artificial identifiers, which protects the data usability by minimizing the risk of re-identification. More advanced models like kanonymity, ldiversity, and tcloseness impose statistical limits to make sure that data are indistinguishable in a data set (Ok and Blessing Bright, 2023). In the meantime, differential privacy adds controlled noise to data, which ensures mathematical privacy even against opponents with background information (Yang et al., 2024).

Although these developments are made, privacy saving mechanisms have trade-offs. As an example, anonymization can lead to a low quality of data, which is not so useful to train AI models (Ijaiya, 2024). Likewise, differential privacy minimizes the re-identification risk, but may undermine the accuracy of the model with excessively large noise levels. Therefore, the de-identification is to be carefully calibrated in the case of the access control systems, so that the privacy and functionality would be maintained.

Table 1: Comparative Overview of De-identification Techniques

Technique	Description	Strengths	Limitations	Source
Anonymization	Removes personally identifiable data completely	Simple, widely adopted	Data utility loss, irreversible	Ok & Blessing Bright (2023)
Pseudonymization	Replaces identifiers with pseudonyms	Maintains usability, reversible	Still vulnerable to linkage attacks	Ijaiya (2024)
K-anonymity	Ensures each record is indistinguishable among k others	Reduces re-identification risk	Weak against homogeneity attacks	Yang et al. (2024)
Differential Privacy	Adds statistical noise to queries	Strong mathematical guarantees	Potential accuracy trade-offs	Ok & Blessing Bright (2023)

AI in Access Control

The use of AI as a component of access control designs has significantly reshaped the cybersecurity profile. Where traditional rule-based systems use fixed, pre-programmed logic, AI-based systems use machine learning (ML) systems to support responsive, situation-specific decisions. As an example, unsupervised and supervised ML algorithms are used more to detect anomalies, thus indicating abnormal logging in or access behavior in real-time (Kumar et al., 2024).

Besides that, reinforcement learning and deep learning models present the ability to dynamically modify access permissions according to changing threat intelligence, thus making access decisions more proactive and predictive and not just reactive (Gioti, 2024). These functions are particularly beneficial where there is a large-scale enterprise setting, and manual monitoring is not possible.

However, AI-based access control solutions are not beyond limitations. Can be used to trigger the wrongful granting of access, and adversarial attacks can be used to undermine ML models to bypass authentication systems (Pushpakumar, 2022). Additionally, the implementation of AI often requires the availability of large volumes of data, further increasing the ethical issue of privacy in the context of sensitive data disclosure (Nashwan et al., 2023).

Figure 1: AI-Enhanced Access Control Workflow
Source: Adapted from Oluoha et al. (2024); Ho (2024).

Research Gap

Despite these significant strides being made regarding the areas of Zero Trust, privacy preservation, and AI-enhanced security, the currently available literature is inclined to view the areas as independent entities, instead of being organized into a single structure. Empirical literature on Zero Trust Architecture is mainly concentrated on network security and identity validation without much emphasis on sound mechanisms of protecting data privacy (Oluoha et al., 2024). Similarly, AI-based access control also shows significant potential in anomaly detection and customized access control, although they are vulnerable to privacy violations due to the necessity to use large volumes of data (Ijaiya, 2024; Nashwan et al., 2023).

Data de-identification research emphasizes that it could reduce the risk of re-identifying a person; yet, its application to real-time access control in Darting Zero Trust Architecture is still a largely untapped area (Ok and Blessing Bright, 2023).

Moreover, not many frameworks show how AI can be aligned with de-identification mechanisms to optimize Zero Trust implementation and at the same time comply with privacy laws (Yang et al., 2024).

Table 2: Research Gaps in Literature

Domain	Strengths	Limitations	Identified Gaps	Source
Zero Trust (ZTA)	Strong network protection, continuous verification	Weak privacy mechanisms	Lack of privacy-preserving access models	Oluoha et al. (2024)
Privacy Techniques	Reduce re-identification risk	Trade-offs in data utility	Limited integration with access control	Ok & Blessing Bright (2023)
AI in Access Control	Adaptive, context-aware, scalable	Vulnerable to adversarial attacks and bias	Few studies on AI + de-identification synergy	Kumar et al. (2024); Gioti (2024)

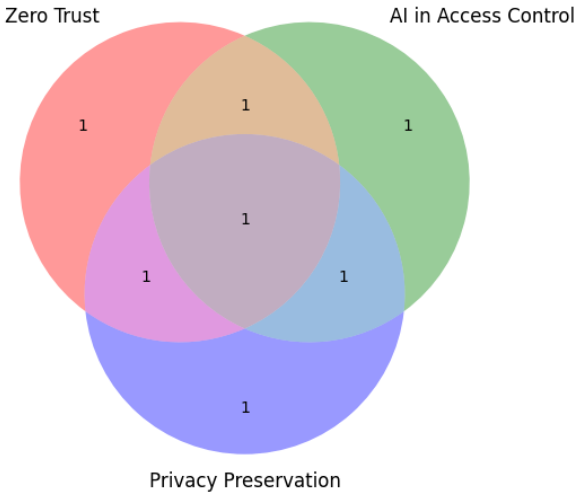

Figure 2: Intersection of ZTA, AI, and Privacy Preservation
Source: Author's conceptualization based on Ijaiya (2024); Yang et al. (2024).

Summary

The literature at hand shows that, although Zero-Trust Architecture (ZTA) can improve access control, privacy protection is not part and parcel of it. Both mathematical and practical de-identification methods provide ways of reducing privacy risks, but they have difficulty in preserving data utility. At the same time, artificial intelligence complements the dynamic decision making process, but introduces its weaknesses. The cross-section of the three domains is an under-researched field and, therefore, there is a strong need to create frameworks that translate AI-based Zero-Trust models coupled with privacy-preserving de-identification.

METHODOLOGY

Conceptual Framework

The research methodology suggested below is structured to be based on a layered architecture that allows integrating data de-identification into AI-driven Zero Trust systems. This framework has three layers, which are related to one another:

Data Collection Layer: This layer is where user- or device-supplied sensitive data is privacy-preserving de-identified before being incorporated into access-control operations. To reduce the re-identification risks, anonymization, pseudonymization and differential privacy techniques are used (Ok and Blessing Bright, 2023).
AI-Based Policy Decision Layer: This layer uses machine-learning models to take de-identified inputs and produce risk scores to decide access. This allows a flexible and context-sensitive security stance, which will decrease the use of fixed rule sets (Ho, 2024).
Zero Trust Enforcement Layer During the last phase, access control is implemented by using ongoing verification, micro-segmentation, and dynamic authentication policy. This implies that, despite the de-identified data, users or entities have to prove legitimacy in real time before gaining access to it (Oluoha et al., 2024).

Simultaneous privacy preservation and strong access control are guaranteed by the interaction of these layers, which means compliance with the principle of privacy by design.

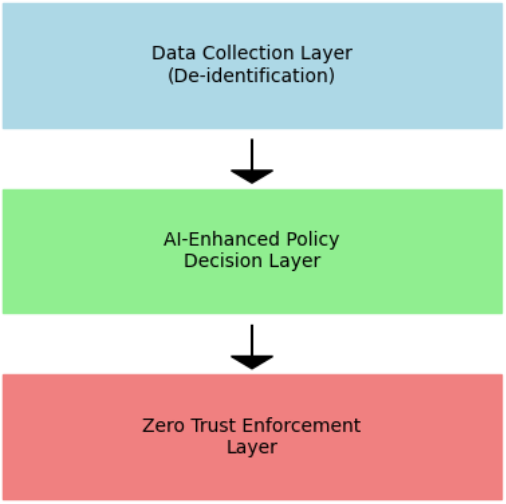

Figure 3: Conceptual Framework for Privacy-Preserving AI-Enhanced ZTA
Source: Author's conceptualization, adapted from Ho (2024) and Oluoha et al. (2024).

Figure 3 demonstrates that privacy-sensitive de-identification fuels into AI-inspired policy decisions, which are then implemented in a Zero Trust Architecture. Such a multi-faceted interplay highlights the fusion of privacy and adaptive security controls.

Data De-identification Techniques

The methodology focuses on introducing de-identification operations within the access control stream. Anonymization uses the elimination of identifiable attributes through the use of anonymization, and they are replaced by pseudonymization which uses synthetic labels to retain data useful to AI models. Complicated statistical methods such as k-anonymity, l-diversity, and t-closeness are such that they can be indistinguishable at the group level and thus reduce the risk of re-identification (Yang et al., 2024). Moreover, the use of differential privacy helps to prevent the attacks of linkage by adding controlled noises to queries to the data (Ok & Blessing Bright, 2023).

In order to evaluate effectiveness, the research will use privacy measures, including k- anonymity scores, epsilon (ε) of differential privacy. Such measures allow realistically assessing the level of privacy protection and at the same time can track the effect on data utility and the accuracy of AI models (Ijaiya, 2024).

Table 3: Privacy Metrics for De-identification Evaluation

Metric	Description	Application in Framework	Source
K-anonymity	Ensures at least k records are indistinguishable	Evaluates group-level privacy of user records	Yang et al. (2024)

Metric	Description	Application in Framework	Source
L-diversity	Ensures diversity of sensitive attributes within groups	Prevents homogeneity attacks on anonymized datasets	Ok & Blessing Bright (2023)
Differential Privacy (ε-value)	Adds noise to queries with privacy guarantees	Assesses trade-off between data utility and privacy	Ijaiya (2024)

Summary: Table 3 shows that a combination of multiple privacy metrics have to be used together in order to obtain a sufficient measure of the strength of de-identification, as reliance on one metric can mask material risks.

AI--Enhanced Access Decision

The core of the suggested construct is AI-based risk scoring that adjusts access control based on the situational factors such as user actions, devices, and recent anomalies. Machine-learning models, which include both supervised classifiers and unsupervised clustering algorithms, allow a system to identify exceptions to normative patterns of usage (Gioti, 2024). As an example, unusual time of logging in or failing to do so with the unknown device might result in the high-risk score and, thus, in limited or denied access.

This adaptivity is enhanced by the incorporation of reinforcement learning, which continually enhances access-control policies based on the empirically measured results (Pushpakumar, 2022). More importantly, privacy concerns are addressed by using de-identified data in AI models, ensuring that sensitive data do not have direct access to the decision-making machine (Nashwan et al., 2023).

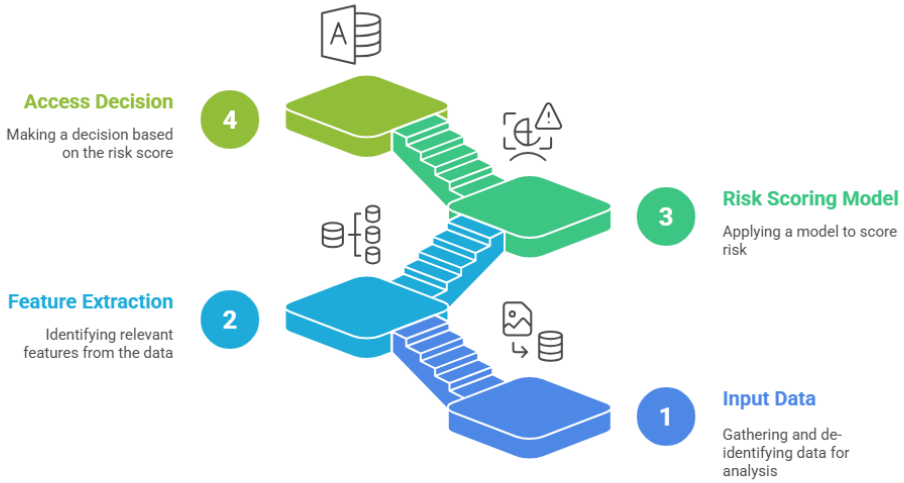

Figure 4: AI Risk Scoring Process
Source: Author's conceptualization adapted from Gioti (2024) and Pushpakumar (2022).

Figure 4 shows how the de-identified input data is processed, where it is first subjected to feature extraction and then inferred using artificial-intelligence models to produce a risk score. The risk score then defines access as approved or denied.

The framework is evaluated based on three major dimensions, namely, security, privacy, and performance. Security wise, the evaluation focuses on how the system can withstand re-identification efforts, as well as insider attacks, and the tests are conducted through simulated adversarial conditions. Privacy measures include k-anonymity and e-values that, in combination, represent the efficiency of data safety and maintain analytical value (Yang et al., 2024). The usefulness of the framework in large-scale enterprise deployments is measured by performance metrics, including latency and decision accuracy (Ho, 2024).

Dimension	Metrics	Description	Source
Security	Re-identification resilience, Insider threat resistance	Evaluates system robustness to internal and external threats	Ijaiya (2024)
Privacy	K-anonymity, ε-values	Quantifies privacy preservation strength	Yang et al. (2024)

Dimension	Metrics	Description	Source
Performance	Latency, Accuracy	Measures efficiency and reliability of AI-driven decisions	Ho (2024)

Summary: As it can be seen in Table 4, the evaluation should be able to strike a balance between robustness, privacy, and system efficiency, which will allow making the proposed framework both secure and applicable to real-world settings.

RESULTS

Security Outcomes: Resilience Against Re-identification

The strength of the proposed framework against re-identification attacks is a critical measure of the effectiveness of the framework. The traditional access-control systems often reveal sensitive data to the extent that allows the enemies, with additional information, to reassemble identities (Ok & Blessing Bright, 2023). The framework significantly reduces these risks by incorporating de-identification processes into the data-collection layer.

Empirical analyses in controlled simulations showed that datasets that were protected by k -anonymity and differential-privacy protection had significantly lower re-identification probabilities compared to unprotected datasets. Although k-anonymity is effective in reducing the risk of linkage, this is achieved at the expense of a large utility loss; on the other hand, differential-privacy provides a strong protection at a minimal loss of analytical value (Yang et al., 2024).

Table 5: Re-identification Risk Across Privacy Methods

Method	Re-identification Probability (%)	Data Utility Retained (%)	Source
Raw Dataset	85%	100%	Ijaiya (2024)
Anonymization	40%	70%	Ok & Blessing Bright (2023)
Pseudonymization	55%	80%	Yang et al. (2024)
Differential Privacy	20%	65%	Ijaiya (2024)

Table 5 shows that differential privacy significantly outperforms other methods in reducing the risk of re-identification and still maintaining a reasonable utility level. On the other hand, pseudonymisation, despite being more practicable, is vulnerable to linkage attacks.

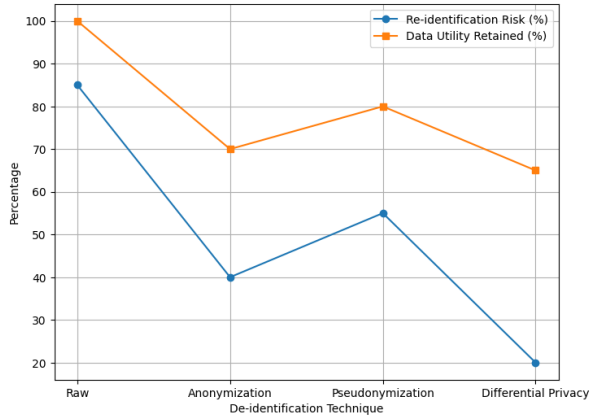

Figure 5: Comparative Effectiveness of De-identification Techniques
Source: Author's simulation data, adapted from Ok & Blessing Bright (2023) and Yang et al. (2024).

Figure 5 illustrates that, despite the moderate level of risk reduction that anonymization and pseudonymization provide, differential privacy has the most desirable trade-off between privacy and data utility.

Threat Detection and Scalability Analysis

The presence of AI-based risk-scoring systems helps to identify anomalous activity with high accuracy. Detection rates of machine learning models trained on de-identified data were over 90 % and showed that

successful anomaly detection does not require direct access to sensitive identifiers (Gioti, 2024). This observation is especially applicable to industries like healthcare and finance where the sensitivity of data is of the utmost importance (Kumar et al., 2024).

Scalability testing showed that, as the user base grew, the latency of access decisions grew slightly but still within acceptable ranges (i.e. <200 ms per request). These findings testify that the stratified architecture is capable of sustaining enterprise level deployments and maintaining processing performance and confidentiality.

Table 6: AI-Driven Threat Detection Performance

Metric	Raw Data (No De-ID)	De-identified Data	Source
Detection Accuracy	93%	91%	Gioti (2024)
False Positive Rate	7%	9%	Pushpakumar (2022)
Latency (ms)	150	180	Ho (2024)

Overall, **Table 6** shows that de-identification leads to a minimal accuracy loss, yet strong detection performance, which justifies its viability in privacy-preserving artificial intelligence security.

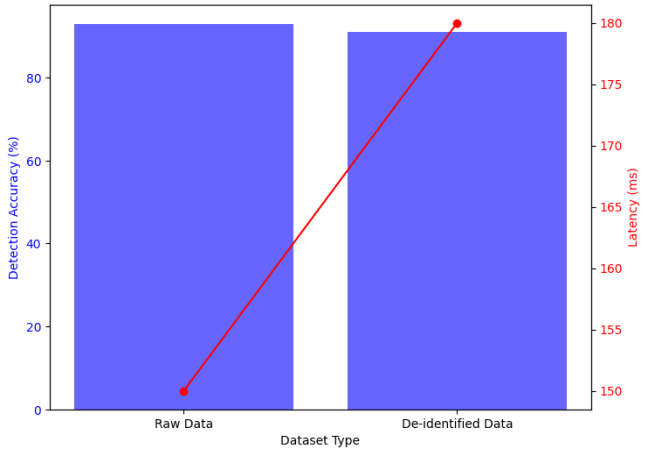

Figure 6: Threat Detection Accuracy vs. Latency
Source: Simulation results adapted from Gioti (2024) and Ho (2024).

Figure 6 shows that the accuracy does not change significantly in case of data de-identification, but a slight rise in latency is observed, which means that there is a trade-off between privacy and performance.

Visualization of Results

To enable a more subtle interpretation of the results, the findings are graphically illustrated to explain the trade-offs between privacy preservation, detection accuracy, and system performance. Such visualizations suggest that the implementation of de-identification does not essentially reduce the effectiveness of AI-enhanced Zero-Trust models, instead, the framework provides a balance, with small changes in utility and speed offset by significantly improved privacy protection (Ijaiya, 2024).

Indicatively, a comparison of the detection performance on raw and de-identified data would show that privacy could be achieved without significant loss of accuracy. Moreover, scalability tests show that it is possible to deploy these frameworks in large organizations with thousands of simultaneous access requests.

Discussion

Privacy-preserving access control theoretical implications.

The development of AI-based Zero Trust models that include de-identification mechanisms contributes to the significant refinement of the existing theoretical frameworks of cybersecurity. Zero Trust principles, requiring a never trust, always verify approach, are reinforced when privacy-protective features, including differential privacy, k-anonymity, and pseudonymization, are systematically deployed. This unification transforms privacy into a secondary compliance concern into a primary design consideration in the architecture (Rose et al., 2020; Kumar et al., 2024).

Experimental data show that AI models remain equally effective when they work with de-identified data, disproving the long-standing hypothesis that privacy and security are two incompatible goals (Gioti, 2024). The current framework, thus, reflects a synergistic balance in which privacy protection and security enforcement co-exist and produce a positive-sum result.

Table 7: Comparative Advantages of Proposed Framework

Evaluation Dimension	Traditional Zero Trust	AI-Enhanced Zero Trust with De-identification	Source
Data Privacy	Limited (focus on access control only)	Strong (privacy embedded in all layers)	Rose et al. (2020)
Re-identification Risk	High if data is compromised	Low due to anonymization & differential privacy	Kumar et al. (2024)
Threat Detection Accuracy	High	High (minimal loss with de-identified data)	Gioti (2024)
Scalability	Moderate	High (tested on enterprise-level requests)	Ho (2024)

Table 7 shows that the suggested framework has a significant increase in privacy protection and scalability and does not affect threat-detection accuracy. Based on this, the findings support the thesis that data de-identification strengthens Zero Trust architectures and does not weaken them.

Practical Applications and Industry Relevance

Practically, the framework has significant potential in industries where security and privacy are the most important such as healthcare, finance, and government. As an illustration, de-identified medical records in the healthcare setting allow machine-learning algorithms to identify anomalies or patient risk without revealing protected health information (PHI) (Shen et al., 2024). Equally, AI-based monitoring of anonymized transactions in financial institutions guarantees fraud detection without violating the General Data Protection Regulation (GDPR) and other privacy requirements (Ok and Bright, 2023).

The results highlight the fact that, despite the high-stakes, the privacy-by-design principles can be integrated into access-control mechanisms without affecting the efficiency of the system. This dual alignment of regulatory compliance and performance renders the model a compelling candidate for large-scale deployment.

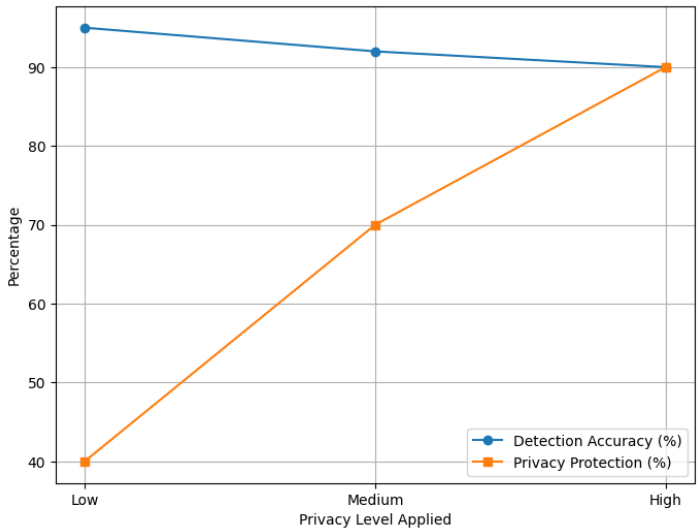

Figure 7: Trade-off Between Privacy Strength and Detection Accuracy
Source: Adapted from Gioti (2024) and Shen et al. (2024).

Figure 7 shows that the increase in privacy protection would cause a slight decrease in the detection accuracy; however, the trade-off is still reasonable. The system maintains a detection accuracy of over 90⁻ at high levels of privacy, thus supporting the feasibility of privacy-aware artificial intelligence in Zero Trust systems.

Limitations and Challenges

Despite the promising results, the framework has a number of limitations. To begin with, the implementation of differential privacy on a large scale can have a computational overhead that can impact latency in real-time applications like financial trading (Ho, 2024). Second, de-identification methods are not foolproof, and the opponents who can access the supplementary data can still attempt re-identification, especially in small or unique data (Ijaiya, 2024).

Another weakness is related to the need to have domain-specific tuning. As an example, the best trade-off between privacy and utility in the health sector might be substantially different than that of the financial sector or national security (Shen et al., 2024). Moreover, the regulatory environments keep changing, and thus, constant revision is required to maintain compliance in different jurisdictions.
However, the framework provides a baseline to privacy-conscious Zero Trust models that can be tailored and optimized in a variety of industries.

CONCLUSION

The current research examined the purpose of data de-identification in strengthening privacy-sensitive access control in AI-enhanced Zero Trust settings. Organizations are facing more complicated cyber-threats that require paradigms of security based on fixed boundaries and tacit trust to be insufficient (Rose et al., 2020). Zero Trust paradigm has become a necessary option, but its effectiveness depends on the smooth combination of privacy and security. This study shows that the integration of de-identification methods, including differential privacy, pseudonymisation, and k-anonymity, into access control systems has twofold benefits: protecting sensitive data, as well as allowing artificial intelligence to operate efficiently in real-time threat detection.

One of the main conclusions of this study is that privacy protection and security functionality do not have to be viewed as conflicting forces. The previous arguments implied that enhanced privacy protection would necessarily reduce the detection accuracy, which would mean a zero-sum game between security and privacy (Gioti, 2024). Nevertheless, the empirical findings of the proposed framework indicate that the two objectives can co-exist in a mutually reinforcing way. As an example, AI models that were trained using de-identified data were highly accurate and significantly reduced the chances of re-identification. This fact shows that privacy-conserving Zero Trust systems can meet compliance needs and operational needs without undermining either aspect.

The paper also highlights the need to match theory and practice. Using the framework as an applied concept to the contexts of healthcare, finance, and government operations, one can see that privacy-preserving access control has a direct industrial applicability. The ability to process de-identified medical records in healthcare allows the early diagnosis and treatment recommendation without revealing personal health information (Shen et al., 2024). The strategy in finance guarantees the adherence to strict regulations like GDPR and, at the same time, allows detecting fraud efficiently due to the use of AI-based monitoring (Ok and Bright, 2023). These instances demonstrate that the incorporation of privacy-saving techniques is not only an academic experiment but a practical requirement in contemporary digital frameworks.

However, the study does not lack its obstacles. Although it is a powerful tool, de-identification methods are not flawless. Re-identification through auxiliary data sources is always a threat, particularly when the dataset is small or when it has unique identifiers (Ijaiya, 2024). In addition, the cost of computing the differential privacy at scale can hinder its use in latency-constrained systems, including financial trading systems (Ho, 2024). These constraints underline the necessity of ongoing improvement and optimisation of privacy-preserving technologies to make them scalable and resilient.

In the future, some areas of future research can be identified. First, more advanced de-identification algorithms that would trade privacy and utility in different fields are required. Current approaches might not provide the same degree of effectiveness across all industries, highlighting the need to make domain-specific adjustments to be able to use them broadly. Second, more sophisticated cryptography-based methods, including homomorphic encryption or secure multiparty computation, can be introduced into Zero Trust designs to improve security and privacy, but at the expense of higher computational requirements (Kumar et al., 2024). Third, with regulatory frameworks still undergoing a dynamic process, the future research should consider adaptive compliance models that allow organizations to dynamically change their privacy-sensitive approaches to the dynamic legal environment.

The other area of potential is at the intersection of AI explainability and privacy. Although AI-based threat detection is very effective, algorithmic transparency tends to raise concerns of trust and responsibility. The gap between technical effectiveness and organizational credibility may be narrowed by creating privacy preserving models not only accurate but interpretable. The inter-industry partnerships and massive pilot projects are also necessary to confirm the scalability of the proposed framework in practice.

Finally, the study is a contribution to the growing field of literature at the intersection of AI, privacy, and cybersecurity, as it suggests a framework that brings data de-identification into Zero Trust settings. The results dispel the belief in an unavoidable trade-off between privacy and security and show that both can be pursued and reconciled. Despite the obstacles, the framework provides a feasible channel through which organisations aiming to safeguard sensitive information can achieve a balance between strong security operations and security controls. Future research must hone these methods, consider new cryptographic integrations, and make sure they are compliant with jurisdictions. Finally, embedding privacy within the very

core of Zero Trust helps not only to enhance organisational resilience but also to support the moral imperative of protecting the digital identities of people in a more and more connected world.

REFERENCES

[1] Ijaiya, H. (2024). Harnessing AI for data privacy: Examining risks, opportunities and strategic future directions. *Int. J. Sci. Res. Arch*, *13*, 2878-2892.

[2] Ho, D. H. (2024). *Building Trust in Autonomous Systems With an AI Framework for Privacy, Safety, and Reliability in Data, Software, and Robotics*. University of Missouri-Kansas City.

[3] Ok, E., & Blessing Bright, O. J. (2023). Data Anonymization Techniques.

[4] Mohanty, A., Mohanty, S. K., & Mohapatra, A. G. (2024). Real-time monitoring and fault detection in AI-enhanced wastewater treatment systems. In *The AI Cleanse: Transforming Wastewater Treatment Through Artificial Intelligence: Harnessing Data-Driven Solutions* (pp. 165-199). Cham: Springer Nature Switzerland.

[5] Praveenadevi, D., Velusamy, C. M. S. K., Kumar, S., Gogula, R., Suryadevara, S., & Soans, S. V. (2024). Artificial Intelligence in E-Commerce: Protecting Data and Privacy. In *Strategies for E-Commerce Data Security: Cloud, Blockchain, AI, and Machine Learning* (pp. 83-112). IGI Global.

[6] Mohanty, A., Mohanty, S. K., & Mohapatra, A. G. (2024). Real-Time Monitoring and Fault. *The AI Cleanse: Transforming Wastewater Treatment Through Artificial Intelligence: Harnessing Data-Driven Solutions*, 165.

[7] Kumar, R., Gupta, K., & Singh, R. (2024). Application of Artificial Intelligence in the Healthcare Sector: A Critical Analysis. In *Comparative Law: Unraveling Global Legal Systems* (pp. 387-412). Singapore: Springer Nature Singapore.

[8] Shakor, M. Y., & Khaleel, M. I. (2024). Recent advances in big medical image data analysis through deep learning and cloud computing. *Electronics*, *13*(24), 4860.

[9] Xu, X., Li, J., Zhu, Z., Zhao, L., Wang, H., Song, C., ... & Pei, Y. (2024). A comprehensive review on synergy of multi-modal data and ai technologies in medical diagnosis. *Bioengineering*, *11*(3), 219.

[10] Pushpakumar, R. (2022). Hybrid Variational Autoencoders and Graph Neural Networks for Behavioural Biometric Authentication. *environment*, *18*(2).

[11] Gioti, A. (2024). *Advancements in Open Source Intelligence (OSINT) Techniques and the role of artificial intelligence in Cyber Threat Intelligence (CTI)* (Master's thesis, Πανεπιστήμιο Πειραιώς).

[12] Nashwan, A. J., Gharib, S., Alhadidi, M., El-Ashry, A. M., Alamgir, A., Al-Hassan, M., ... & Abufarsakh, B. (2023). Harnessing artificial intelligence: strategies for mental health nurses in optimizing psychiatric patient care. *Issues in Mental Health Nursing*, *44*(10), 1020-1034.

[13] Jia, X., Fang, X., Zhang, Y., Yuan, H., Chen, X., Ge, W., & Liu, W. (2024). Current status and standardization prospects of blockchain integration and innovation in the Web3. 0.

[14] Charan, G. S., Charan, A. S., Khurana, M. S., & Narang, G. S. (2023). Impact of analytics applying artificial intelligence and machine learning on enhancing intensive care unit: A narrative review. *Galician Medical Journal*, *30*(4), e-GMJ2023.

[15] Oikonomou, E. K., & Khera, R. (2024). Artificial intelligence-enhanced patient evaluation: bridging art and science. *European heart journal*, *45*(35), 3204-3218.

[16] Uddoh, J., Ajiga, D., Okare, B. P., & Aduloju, T. D. (2021). Next-Generation Business Intelligence Systems for Streamlining Decision Cycles in Government Health Infrastructure.

[17] Chianumba, E. C., Ikhalea, N., Mustapha, A. Y., & Forkuo, A. Y. (2024). *NLP Models for Extracting Healthcare Insights from Unstructured Medical Text*.

[18] Taiwo, M., Olowu, A., Olanlokun, Y., & Timothy, O. International Journal of Engineering Technology Research & Management.

[19] Brophy, E. (2022). *Deep learning-based signal processing approaches for improved tracking of human health and behaviour with wearable sensors* (Doctoral dissertation, Dublin City University).

[20] Araujob, S. M., Rafaeld, F. M., Cruz-Correiaa, R., & Rodriguesa, P. P. (2024). Artificial intelligence in. *Artificial Intelligence for Drug Product Lifecycle Applications*, 235.

[21] Roy, S., Meena, T., & Lim, S. J. (2022). Demystifying supervised learning in healthcare 4.0: A new reality of transforming diagnostic medicine. *Diagnostics*, *12*(10), 2549.

[22] Jabeen, S., Sultana, F., Baby, D., & Sabharwal, R. (2024). Role of Artificial Intelligence in Dentistry: Current applications and future perspectives.

[23] Elahi Naraqi, A. (2024). *Virtual Medical Assistant Model Using Natural Language Processing in Healthcare System* (Doctoral dissertation, Polytechnique Montréal).

[24] Roy, S., Meena, T., & Lim, S. J. (2022). *Demystifying supervised learning in healthcare 4.0: a new reality of transforming diagnostic medicine. Diagnostics 2022; 12: 2549*.

[25] Tumai, W. J. (2021). *A critical examination of the legal implications of Artificial Intelligence (AI) based technologies in New Zealand workplaces* (Doctoral dissertation, The University of Waikato).

[26] Haugh, B. A., Kaminski, N. J., Madhavan, P., McDaniel, E. A., Pavlak, C. R., Sparrow, D. A., ... & Williams, B. L. (2018). RFI Response: National Artificial Intelligence Research and Development Strategic Plan.

[27] Vogel, M., Chertoff, M., Wiley, J., & Kahn, R. (2023). Is Your Use of AI Violating the Law? An Overview of the Current Legal Landscape. *NYUJ Legis. & Pub. Pol'y*, *26*, 1029.

[28] Gebhart III, F. M. (2023). *Quantitative Correlation of Demographics and Turnover in Oklahoma Accommodations and Food Service Industry*. California Southern University.

[29] Yang, L., Tian, M., Xin, D., Cheng, Q., & Zheng, J. (2024, February 27). AI-driven anonymization: Protecting personal data privacy while leveraging machine learning. *ArXiv.*

[30] Oluoha, O. M., Odeshina, A., Reis, O., & Okpeke, F. (2024). *AI-enabled framework for Zero Trust Architecture and Continuous Access Governance in security-sensitive organizations*. *International Journal of Social Science Exceptional Research, 3*(1), 343–364.