KeAi
CHINESE ROOTS
GLOBAL IMPACT

International Journal of Innovation Studies

ZERO TRUST IN PRACTICE: AI-AUGMENTED ACCESS CONTROL FOR CRITICAL GOVERNMENT NETWORKS

Yesu Vara Prasad Kollipara[1*]
[1]Dakota State University
YesuVaraPrasad.Kollipara@trojans.dsu.edu

Abstract

The increased complexity of the current government networks requires dynamic and open-minded cybersecurity systems capable of interoperating with distributed and legacy systems. Conventional perimeter-based security frameworks are no longer effective to deal with advanced persistent threats or insider abuse, especially in multi-agency settings where data confidentiality and interoperability are paramount. The proposed research suggests an AI-enhanced Zero Trust Architecture (ZTA) incorporating user-behaviour analytics, federated learning (FL), and explainable artificial intelligence (XAI) to provide dynamic and auditable access to important government systems. The architecture assesses contextual risk on a continuous basis on the basis of behavioural deviations and makes adaptive access decisions using an intelligent policy engine. The model was trained and validated on five simulation nodes of the agency to simulate distributed governmental environments using the User Activity Anomaly Detection Dataset. Federated learning ensured data sovereignty by summing model parameters as opposed to raw data, whereas SHAP-based explainability offered interpretable information about each access decision. Experimental findings showed that the overall accuracy was 92%, 0.89 F1-score and 0.94 ROC-AUC, with only slight performance loss as compared to a centralised model, yet full privacy protection. A 32% decrease in undetected insider incidents and enhanced compliance auditing was achieved with the system due to transparent decision logic. Such results highlight the possibility of integrating smart, understandable and cooperative security controls into the running government architectures. Besides, the strategy is consistent with the overall Computer Integrated Manufacturing Systems (CIMS) vision in terms of the demonstration of how intelligent integration, automation and human control can conglomerate into the computer-related governance systems that are resilient.

Keywords: Zero Trust Architecture (ZTA); Federated Learning (FL); Explainable Artificial Intelligence (XAI); Anomaly Detection; User Behaviour Analytics; Dynamic Access Control; Government Networks; Computer-Integrated Systems

Introduction

The rapid digital transformation of government organisations has significantly increased the cybersecurity attack surface that revealing the fragility of the traditional perimeter-based defence models. The traditional security architectures use one trust boundary, where users and devices are assumed to be trusted once they are authenticated in the network (Liu et al., 2024). But this assumption has not been sufficient in current distributed and cloud-based government infrastructures. Insider threats, advanced persistent attackers, and other threat actors use the lateral movement in the trusted zones to inflict damage on key systems. To overcome these

problems, a new paradigm of security principles has been developed, the Zero Trust Architecture (ZTA). Instead of implicit trust, ZTA implements the philosophy of never trust, always verify, whereby all access attempts must be authenticated, authorised, and verified regardless of the location, device or network segment of the user (Sarkar et al., 2022).

According to the NIST Special Publication 800-207, these are the main characteristics of the Zero Trust frameworks: identity-based access control, least privilege implementation, and real-time policy decisions, depending on the current situation and risk evaluation (Dalal, 2025). The method is especially essential in the case of the public sector, where there are heterogeneous legacy infrastructures, cross-agency dependencies, and sensitive citizen information that co-exists. Although the government has been pushing to introduce Zero Trust, most government organisations have a difficult time adopting the system practically because of the disconnected data ecosystems, obsolete authentication standards, and ineffective analytics (Steenbrink, 2022). Older systems were frequently created several decades ago with databases that were siloed and access policies that were not dynamic and were very difficult to adjust to new, flexible control systems. Also, the fact that ministries, departments, and regulatory bodies are not interoperable is another limitation to adopting cohesive Zero Trust policies.

At this juncture, artificial intelligence (AI) and machine learning (ML) offer groundbreaking opportunities to put ZTA into practice. Through the analysis of the user behavioural patterns, device attributes, and network events, the AI-driven models will be able to dynamically evaluate risk and impose contextual access control without reducing operational efficiency. Furthermore, by incorporating Explainable AI (XAI), human auditors and administrators in the system have a chance to learn the reasoning behind making particular access decisions, which improves accountability and compliance in controlled government settings (Potluri, 2025). Such AI-enhanced decision processes are close in resonance with the domain of Computer Integrated Manufacturing Systems (CIMS), which highlights the merging of intelligent automation, data-driven decision frameworks, and real-time control in complicated systems. Similar to how CIMS is applied to the industrial production process to improve its performance and transparency, the current study can be applied to the area of cybersecurity by implementing intelligent and explainable decision-making into digital governance frameworks (Amannah, 2025).

Although the idea of Zero Trust has been accepted in principle both in the private and the government sphere, its application in the government context is still undeveloped and disjointed. The literature on the subject has been oriented on technical policy frameworks or single, standalone intrusion detection models, but not broader, more adaptive architectures, which can be integrated with older ecosystems. Moreover, the majority of deployments are missing federated intelligence, which is the capability to learn together among different agencies without jeopardising the sovereignty and confidentiality of data. Centralised data to train AI models is not always possible in a heavily controlled setting because of privacy requirements, and Federated Learning (FL) is a perfect option (Wen et al., 2023). Nevertheless, little is available in terms of practical applications that can integrate both ZTA and federated learning to distributed access control in government networks.

The other severe loophole is the openness of AI-based access decisions. Since governmental mechanisms operate with sensitive information about citizens and their finances, opaque and black-box AI models are introduced, which undermine accountability, fairness, and bias in

decision-making (Chaudhary, 2024). Although progress has been made on XAI in other areas, there is still a lack of explainable reasoning being incorporated into cybersecurity and access control systems. Furthermore, there are limited examples of cross-domain research that can be effective in trade-offs between data privacy, interpretability, and security automation, which has created a gap between theories and implementation. This study fills these gaps by creating and testing an AI-enhanced Zero Trust architecture that integrates federated learning to collaboratively provide privacy and XAI, enabling transparent decision-making in a simulated government network.

This research will develop and test an AI-enhanced Zero Trust Architecture that will allow dynamic, explainable, and privacy-preserving access control to distributed government networks. The proposed framework aims to bridge the gap between theoretical concepts of Zero Trust and its practical implementation in legacy, multi-agency systems. The specific objectives are four in number. The first one is to create a multi-layer Zero Trust architecture that combines user-behaviour analytics and anomaly detection to conduct continuous access assessment. Second, to adopt a federated learning system that enables decentralised agencies to coordinate the enhancement of security intelligence without the need to exchange raw data, thus maintaining privacy. Third, to add the explainable AI methods to make them more interpretable and allow regulatory audits of the access decisions. Lastly, to determine the effectiveness of the suggested architecture with the help of benchmark datasets and quantifiable cybersecurity indicators, such as accuracy, detection rate, false-positive reduction, and compliance indicators.

The rest of this paper is organised in the following way. Section 4 gives an overall literature review of the current literature on Zero Trust architectures, access control using AI, federated learning, and explainable AI in the context of cybersecurity and industrial systems. Section 5 expounds on the approach to be adopted by describing the conceptual architecture, AI models, and federated learning mechanisms to be used. Section 6 offers the description of the implementation environment, dataset choice and experimental setup, and Section 7 provides quantitative and qualitative findings. Section 8 explains the wider implications of the findings, which are their application to government networks and computer-integrated systems. Lastly, Section 9 brings the study to a conclusion with some main insights, limitations, and future research directions, and possible applications to the field of industry and manufacturing in line with the topics that the CIMS journal is focused on.

Literature Review

Evolution of Access Control in Critical Systems

The development of the mechanisms of access control has been reflective of the development of the information systems in closed and static environments to dynamic and interconnected ecosystems. Computer security in government and military areas in the late twentieth century was based on the early models of Mandatory Access Control (MAC) and Discretionary Access Control (DAC) (Leander, 2020). MAC imposed strict, centrally determined policies that did not allow users to change the access permissions, which ensured powerful confidentiality but did not allow flexibility. Conversely, DAC transferred the responsibility to data owners, where discretionary sharing was created, but there were vulnerabilities because of the lack of uniformity in the implementation of the policies. With the growth of enterprise systems, the Role-Based Access Control (RBAC) model came up, which offered a scalable system with

ISSN:2096-2487 | E-ISSN:2589-2975

access permissions being assigned to a predefined role and not an individual user (Mohamed et al., 2024). The use of RBAC as a standard, especially in the government and industry, was a dominant standard since it was consistent with organisational structures and job descriptions as regards access control.

Nonetheless, the inflexibility of RBAC was growing inadequate with the introduction of cloud computing, distributed systems, and collaboration among agencies. The Attribute-Based Access Control (ABAC) model came about to address these drawbacks by assessing policies in real time depending on the attributes of the user identity, location, type of device, and time of access (Wang et al., 2020). ABAC was flexible and context-aware, but was dependent on complex policy definitions and centralised policy decision points, which were a challenge to implement in large and heterogeneous infrastructures. All these traditional models have an underlying assumption of trust whereby once a user or device is authenticated; it is frequently granted extensive access within a perimeter.

This premise has not been appropriate in contemporary multi-domain infrastructures of publics where internal and external users have been inverted. In the case of government networks, especially, there are many interconnected systems, which are operated by independent agencies, and it is hard to impose uniform security postures (Oladosu et al., 2022). Due to the evolution of cyber threats, which no longer focus on the network boundaries, but also on insider accounts, supply chains, and third-party integrations, the security community has realised the necessity to do away with implicit trust altogether. This change of paradigm led to the creation of the Zero Trust Architecture (ZTA) that presupposes that no party, internal or external to the network, can be trusted by default. Rather, ZTA requires default authentication, least-privilege access, and real-time context-based adaptive authentication.

Zero Trust Architecture in Government and Industrial Contexts

Zero Trust is a concept with official status, including NIST Special Publication 800-207, which refers to ZTA as the architecture that relocates the decisions of access control to the boundaries of the network to policies of identity and context, which are dynamic. This change has been supported by governments all over the globe: the U.S. Office of Management and Budget (OMB) issued Memorandum M-22-09 in 2022 and mandates the federal agencies to implement ZTA principles by the year 2024, and the European Union and the United Kingdom have already issued comparable mandates as part of their national cybersecurity strategies (Willman, 2025). These models encourage identity-based controls, constant monitoring, and automation-based policy enforcement using analytics.

There are a number of industries that have shown the initial adoption of ZTA. ZTA has been used in defence networks to ensure classified communications channels and control systems, and enhance resilience to insider threats. The financial world has used the concept of Zero Trust to defend online transactions and avert identity theft, whereas industrial and manufacturing systems have applied the concept of Zero Trust to Operational Technology (OT) networks to stop the lateral movement of attacks (Idika et al., 2023). The industrial use of ZTA in the CIMS context is comparable to the government one: large heterogeneous infrastructures in both cases are to be integrated by inserting intelligent control features into the existing systems.

Although this has been done, real field applications continue to experience challenges in areas of interoperability, scalability, and cost. Numerous organisations do not have adequate telemetry that would support continuous monitoring or have the ability to analyse behavioural

patterns at scale. In addition, there are no uniform models of incorporating AI into ZTA decision-making, which impedes adaptive enforcement of policies. As a result, the government networks remain dependent on semi-static rules that cannot adapt to the changes in the threats or user contexts.

AI and Machine Learning in Access Control

Machine learning and artificial intelligence have become fundamental facilitators of dynamic and anticipatory access control. Artificial intelligence, as a result, can identify a small anomaly in user behaviour that can be missed by human administrators or rule-based systems (Abdul Azeem et al., 2022). The use of behavioural biometrics, like the dynamics of a keystroke, movement of the mouse, and time of login, is becoming increasingly popular in profiling normal user activity. The machine learning models then continuously evaluate whether a session is similar to the previous behaviour of a user, identifying anomalies as a potential threat. The recent studies have shown that deep learning models, such as Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, and Autoencoders, are effective in modelling sequential user activity data (Mienye et al., 2024). CNNs have been used to identify spatial-temporal access log patterns, whereas LSTMs identify long-term behavioural sequence dependencies and could be effectively used to identify insider threats or progressive privilege escalations. Autoencoders have demonstrated strong results on unsupervised anomaly detection through reconstruction of the normal patterns of activities and measuring the deviations (Torabi et al., 2023). These methods have enabled access control to be changed to dynamic authentication rather than static authentication, where access control decisions are dynamically changed based on observed behavioural anomalies.

Predictive policy enforcement AI-enhanced access systems are embedded in identity and access management (IAM) systems in industrial and enterprise networks (Nzeako and Shittu, 2024). This is in line with the focus of CIMS on automation and smart systems integration, where the two aspects aim to combine decision-making algorithms with operational control systems to augment flexibility, performance and safety.

Federated Learning for Privacy-Preserving Security

Although AI offers the analytical capability of adaptive security, its use by the government is limited by data privacy regulations and inter-agency freedom. The old paradigms of centralised training involve centralising sensitive data in a single place, which goes against privacy laws and exposes the data to privacy risks. Federated Learning (FL) is a graceful approach that enables numerous parties (e.g. ministries, departments, or regional agencies) to jointly train a common model without any raw data exchange (Khan,, 2025). In FL architecture, every client node will be trained over its own data, and only model parameters or gradients are sent to a central server to be aggregated. This server then updates the global model and is distributed to participating nodes, where it is used to train again.

This is a decentralised method that preserves the sovereignty of data and yet allows collective intelligence. FL has already been used in healthcare (e.g., multi-hospital diagnostic models), in Industry Internet of Things IIoT) systems, and in smart manufacturing, predictive maintenance and fault detection, where the sensitivity of data or its latency limits centralisation (Albshaier et al., 2025). This can be applied to government cybersecurity, whereby agencies are the federated nodes that contribute to a single Zero Trust model using the same architectural principles. This kind of implementation helps in collaborative anomaly detection, whereby

ISSN:2096-2487 | E-ISSN:2589-2975

each agency enjoys the benefits of aggregated learning outcomes but retains possession of its sensitive audit logs. Combining FL and Zero Trust will allow agencies to have an independent operational capability and collectively enhance the national security posture.

Explainable AI (XAI) for Cybersecurity Decision Transparency

Although machine learning models are highly predictive, this is usually criticised by noting that it is opaque, particularly when they are used in critical decision-making systems. Lack of transparency may cause mistrust in cybersecurity and access management, making it difficult to comply with regulations. Explainable AI (XAI) is a solution to this problem in that it presents predictions of models in a manner that can be understood by humans (Dwivedi *et al*., 2023). Such techniques as SHapley Additive explanations (SHAP), Local Interpretable Model-Agnostic explanations (LIME), and attention mechanisms break down AI decisions into feature contributions, showing which aspects of behaviour or context most contributed to an access control decision.

Practically, XAI enables administrators and auditors to follow through on why a user was rejected or was considered anomalous to help people in the office to oversee and minimise bias (Abi, 2025). In the government, transparency is a technical necessity and also an ethical consideration, especially when the decisions related to AI use involve service provision, financial dealings, or personal information of citizens. XAI, as part of Zero Trust frameworks, facilitates human-in-the-loop auditing, which allows the real-time audit and perpetual improvement of AI policies (Olawore et al., 2025). This is consistent with the explainable and auditable intelligent systems emphasis of CIMS, which states that transparency is essential to trustworthy automation in both the industrial and government fields.

Research Gaps Summary

The literature review highlights the major developments in access control models, AI-based analytics, federated learning, and explainable AI; nonetheless, the mentioned elements are still mostly disjointed. Not many studies provide a holistic and integrated process that can deal with adaptive security, privacy preservation, and decision transparency in government networks that operate. Most studies either separate one of the dimensions, namely behaviour-based anomaly detection, federated collaboration, or explainability, without considering their synergistic possibilities (Chittoju et al., 2025). Also, although the intelligent integration in industrial systems research has been the first in the CIMS framework to support automation and quality control, the intelligent system-level integration in the context of cybersecurity has not been adequately studied.

This, in turn, has prompted a definite requirement for a viable, scalable, and computer-integrated design that integrates ZTA, FL and XAI with a single model of operation. Such a system would not only facilitate resilience of cybersecurity in government infrastructures but also apply the principles of CIMS of intelligent integration to the safeguarding of digital ecosystems. The paper addresses that requirement by defining an AI-enhanced Zero Trust model, implementing and justifying the concept of federated intelligence and explainable decision-making, and illustrating how it can work using open-source behavioural datasets and simulated government network scenarios.

ISSN:2096-2487 | E-ISSN:2589-2975

Proposed Methodology

System Overview

The suggested methodology proposes an AI-enhanced Zero Trust Architecture (ZTA) aimed at improving access control in distributed networks of the government. The architecture is based on the principle of *"never trust, always verify"*, and by incorporating the artificial intelligence (AI) components, which constantly assess the behaviour of users and devices before access is granted or access is maintained. Fig 1 shows the incorporation of the core layers of the system: the Zero Trust policy engine, the AI-powered analytics module, and the federated learning (FL) framework between various government agencies. The agency nodes are independent security domains, each having a local learning module which processes internal behavioural data. These nodes are periodically connected to a central orchestrator who is in charge of consolidating and revising global security intelligence.

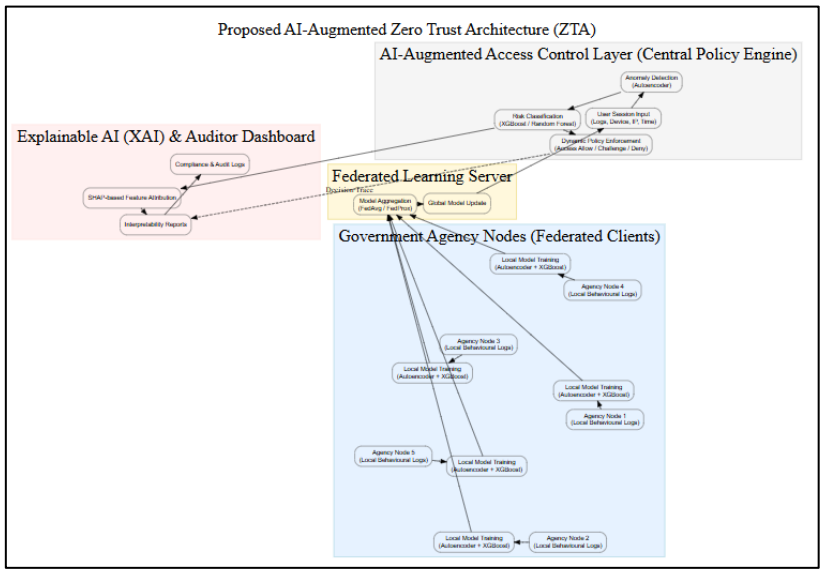

Fig 1: Proposed AI-Augmented Zero Trust Architecture (ZTA)

The system operates on a perpetual verification process. Any access request initiates a behavioural and contextual assessment based on AI models trained to differentiate normal activity and anomalies. The Zero Trust policy engine is used to enforce decisions dynamically based on adaptive controls, e.g. multi-factor authentication challenges, session timeouts, or revocation of privileges. The federated learning infrastructure also makes sure that all the involved agencies have access to common threat intelligence without exposing sensitive information. The framework complies with the vision of Computer Integrated Manufacturing Systems (CIMS) of incorporating intelligent automation in complex operational processes by integrating adaptive risk scoring, decentralised intelligence, and explainable decision mechanisms.

AI-Augmented Access Control Layer

The main component of the proposed architecture is the AI-enhanced access control layer, the analytical engine of the Zero Trust decision-making. Conventional access control systems use fixed rules or verification of an identity at the time of log-in, which provides a low degree of flexibility to changing behavioural scenarios. Conversely, this model is continually analysed by examining user-session dynamics based on an unsupervised and supervised machine-learning method.

The system consumes several features of the context, such as user session timings, frequency of log-in, device signature, geolocation of IP, session length, and application patterns of access. It is these features that are being used to create a normal user behaviour baseline. The Autoencoder neural network is used as the backbone in terms of anomaly detection. In the course of training, the Autoencoder gets to learn how to recreate legitimate behavioural patterns with minimal reconstruction error. The reconstruction error rises dramatically when the activity is anomalous, e.g., there are unusual access times, unusual device changes, unusual session lengths, and the presence of potential threats is indicated.

To supplement this unsupervised model, the use of supervised classifiers, including Random Forest (RF) and Extreme Gradient Boosting (XGBoost), to classify behaviours as low, moderate or high-risk is used. Random Forest models can help to address nonlinear relationships between features and minimise overfitting by averaging ensembles, whereas XGBoost can optimise its detection with the help of a gradient. All these models come up with a dynamic risk score for individual user sessions, which is interpreted by the decision engine to implement adaptive policies. As an example, a moderate-risk session may trigger another verification challenge, whereas a high-risk session may trigger the immediate termination of a session or isolation in a limited sandbox environment.

Such a perpetual review process makes access control contextually adaptive and data-driven as opposed to rule-based. The architecture, through in-the-field implementation of AI models, is an example of intelligent integration, which is a characteristic of the CIMS framework, which is making cybersecurity an active, automated process embedded directly in operational systems.

Federated Learning Integration

Among the most unusual features of this study is the introduction of Federated Learning (FL) that allows collaborating with government agencies in a privacy-respecting manner. The conventional machine-learning models demand the concentration of data in a central storage where the model is trained, which is highly insecure and regulatorily dangerous in the context of the public sector. FL manages to overcome these issues by sharing the learning process between several nodes, and the raw data does not leave the local environment.

All the participating agencies in the simulated environment are seen as nodes of clients which store their local datasets, which include user activity and system-access logs. Training of these local models is done separately to reflect intra-agency behavioural patterns. After the training, the learned model parameters (weights and gradients) are only sent to a central server that combines them through an algorithm such as Federated Averaging (FedAvg) or Federated Proximal (FedProx). The updated global model, based on this aggregation, is then sent back to all nodes in such a way that all agencies enjoy the benefits of collective learning without affecting their own data sovereignty.

In the paper, the User Activity Anomaly Detection Dataset was divided into five logical domains that typify different government agencies. Every domain had both normal and anomalous access patterns, which could be simulated to create heterogeneous data environments that are common in national IT ecosystems. Open-source libraries like TensorFlow Federated or Flower were used to implement the FL framework to offer reproducible orchestration of training rounds and aggregation cycles. FL use not only guarantees that there is compliance with data-protection requirements, but also enhances the

flexibility of the system since it allows continuous enhancement across distributed security domains.

Explainable AI Module

The system will have an Explainable AI (XAI) layer to make the model outputs explainable, meaning that they are presented in a form that can be understood by cybersecurity auditors and policy analysts. This element is important in governmental settings where the decisions made by AI should adhere to the laws and moral principles.

The XAI module is based on the SHapley Additive exPlanations (SHAP) method that breaks down individual model predictions into additive contributions of features. SHAP allows visualising the reasoning behind every access decision by measuring the magnitude of the contribution of each input variable, e.g. IP change frequency, session duration, or login anomaly, to the risk score. The explanations are presented in the form of interactive dashboards enabling the analysts to review decision paths, compare user sessions, and determine recurring behavioural triggers of access denials or escalations.

This interpretability not only enables human-in-the-loop auditing but also allows regulatory compliance, because the agencies can produce documented evidence of fair and consistent AI behaviour when conducting security reviews. Systems-integration The XAI module is a human-machine interface, a connector between automated analytics and human control-another manifestation of the CIMS philosophy of explainable, intelligent automation in critical infrastructure.

Dataset and Pre-Processing

This study relies on the empirical part, which is the **User activity dataset** available on Kaggle (Ekanayaka, 2025). It is a dataset of anonymised user-session data of normal and abnormal system interactions. Attributes that each record has include timestamps, identities of the user, IP addresses, device types, length of a session, geolocation markers, and the status of access. All these characteristics are the reflections of the data sources commonly present in enterprise or government security information and event management (SIEM) systems.

To reduce the inconsistency of data and analysis reliability, the dataset was subjected to a pre-processing pipeline before model training. The attacks with missing data were addressed with the median imputation, and the categorical variables (types of devices and session statuses) were one-hot encoded. Continuous variables, such as the length of session, frequency of logins, etc., were normalised by min-max scaling to have identical feature domains across models. The dataset was split into five partitions to simulate federated conditions in the real world, each of which corresponded to a different agency domain with different behavioural patterns and possible attack distribution.

Aggregation of data was done on by session basis, i.e. individual actions were summed to form composite session records which represented whole sequences of behaviour. This is the method which enabled the AI models to learn the temporal dependencies and the contextual variations. The filtered data were further divided into 80% experimental input and 20% validation in each local node so as to have balanced assessment and avoid overfitting.

Evaluation Metrics

The performance of the proposed model was measured with the help of both standard machine-learning measures and indicators which are specific to the Zero Trust environment. To measure the overall quality of the classification of the anomaly-detection system, accuracy, precision,

recall, and F1-score were used. These measures give one an understanding of the trade-off between the detection of malicious activity and reducing false alarms. Also, the Receiver Operating Characteristic-Area Under Curve (ROC-AUC) measure was computed to determine the ability of the system to discriminate at different decision levels.

In addition to these common metrics, there were operational metrics mentioned in the study that were directly related to the Zero Trust paradigm. The rate of insider-threat detection was used to determine the capability of the model to detect anomalous behaviour introduced by authenticated users, whereas the false-positive rate was used to determine the rate at which legitimate sessions were incorrectly identified as suspicious, which is a critical factor in ensuring user experience in high-security settings. Moreover, the effect of latency on the system was tested so that the implementation of the AI-based verification would not affect the performance in real-time access.

A combination of these measures offered a comprehensive evaluation of the accuracy of analysis and the viability of operations, which supported the purpose of the study to combine intelligence into scalable systems of access control in practice. The proposed methodology represents the CIMS vision of an intelligent, computer-integrated system, a system that organically uses automation, data analytics, and human control to protect complex, mission-critical settings.

Implementation and Experimental Setup

Experimental Environment

In order to guarantee the reproducibility and technical rigour required by Computer Integrated Manufacturing Systems (CIMS), the structure of the experiment was developed to be modularly, clear, scalable, and transparent. The implementation was done with the help of Python 3.11 as the main programming language, with the assistance of major machine-learning tools like TensorFlow and PyTorch to develop and optimise the model. It was combined with the SHAP library to support the explainability analysis, and the Flower and FedML frameworks were used to simulate and coordinate the federated learning (FL) infrastructure. These tools were selected because they have robust open-source ecosystems, beautifully documented APIs and can be reproduced in a cloud-based and on-premise computing environment.

The computer infrastructure was a hybrid between the cloud and local resources to simulate a multi-agency context. The simulation was mainly implemented on Google Colab Pro+, which provides the acceleration of the use of the NVIDIA Tesla T4 unit with the support of local nodes with 12th-generation Intel processors and 16 GB of memory. The nodes indicated different government agencies that had a local dataset partition. These partitions were trained on their own to simulate privacy-preserving data silos, but they all engage in global model convergence. The topology of communication was in the form of a standard server-client, whereby the central orchestrator combined the model updates and repackaged the refined parameters. Network latencies and the size of data transmission were observed in order to simulate real-life inter-agency conditions.

The selected environment did not just accommodate the technical needs of federated learning, but also reflected the focus of CIMS on system integration, which implies the ability to integrate intelligent algorithms, distributed computing, and explainability tools into a single operational system. Docker was used to completely containerise the architecture, which made it portable to computing environments and reproducible.

ISSN:2096-2487 | E-ISSN:2589-2975

Baseline and Comparative Models

Three different models were applied and compared to make a meaningful evaluation. The original base was a traditional Role-Based Access Control (RBAC) system, which is a representation of the static policies of rules that continue to dominate most government networks. This model simply provided access control using user roles and credentials without ongoing behavioural analysis and adaptive control. Despite being basic, RBAC was used to obtain a baseline of changes in adaptability and detection performance of the suggested AI-driven solution.

The second model involved a centralised machine-learning system, whereby all behavioural information of the five simulated agencies was combined to form a single dataset to be used to train a global model. The setup of this way enabled the direct comparison of federated and non-federated architectures with the same hyperparameters. Although the centralised model was just a bit more accurate at the start because of the single training data, it infringed on privacy limitations and could not be used in the real-world government context.

The third and the last model, the kernel of this study, was the Federated Learning (FL)-based AI-enhanced Zero Trust framework with Explainable AI (XAI) functionality. The local anomaly-detection and classification model was trained independently by each agency node on the AutoEncoder and XGBoost algorithms. They were synchronised periodically by the central server that combined model weights with Federated Averaging (FedAvg). The integrated XAI component produced SHAP-based explanations at the end of every inference cycle without making the decisions about access non-interpretable and non-auditable. The comparison of these three configurations created a distinct performance order, which shows the benefits of distributed intelligence and transparency as a method of contemporary access-control systems.

Implementation Phases

The experiment was carried out in a multi-stage workflow, which resembles a deployment process. The initial step, which was the data ingestion and preprocessing, entailed the importation of the User Activity Anomaly Detection Dataset and its subsequent subdivision into five independent subsets that represented individual agencies. The cleaning and encoding of each dataset were done in the same manner so that the nodes were uniform. The processed data were kept locally on each of the simulated nodes, which supports the idea of data sovereignty.

The second stage entailed training of local models on each node. In this case, the Autoencoder was trained to recreate normal behavioural patterns, and the XGBoost classifier was trained to assign session activities as legitimate or suspicious. Every local model underwent several epochs of training and then sent the updates in its parameters to the central aggregator. The choice of hyperparameters, including learning rate (0.001), batch size (64), and the number of epochs (30), was made following initial tuning to make a trade-off between the accuracy and the computational cost.

The third step involved federated model aggregation, which was coordinated by the Flower structure. At every communication step, the central server received the updated parameters of all the agency nodes, calculated the weighted average of the world and re-sent the new model. This was repeated until convergence, which is usually fifteen rounds. The benefit of such an arrangement was that raw data were not exchanged; only the weights of the numerical models were, and as such, the confidentiality of agencies was maintained.

After the aggregation, the fourth phase used the dynamic access enforcement with the policy engine simulated using the API. Synthetic users sent incoming access requests to the engine, which extracted behavioural features in real-time and sent them to the trained model to be scored on risk. Depending on the assigned risk category as predicted, the system automatically deployed Zero Trust policies, which could allow access, initiate further authentication procedures, or prevent the request.

Lastly, the fifth step involved the explainability analysis based on the SHAP framework. SHAP values were used to measure the contribution of each instance of inferences to the prediction of the model using input features, including geolocation change, session duration, or frequency of login. These findings were graphically represented in a dashboard that is easy to understand by the security auditors, and this enables one to clearly analyse how and why the model made a particular decision regarding access. These visual deliverables showed that the system met the explainability criteria that were important to the government adoption, and were also in line with the CIMS journal's focus on transparency and auditability of intelligent systems.

Pseudocode / Flowchart

The federated anomaly-detection loop can be used to summarise the implementation process, as it will describe the sequential interaction between the local nodes and the central server. The training is done on each node, with each node having its own dataset, and this time around, there is an exchange of model updates and not actual data. The central server compiles these updates, optimises the world model and sends the optimised version back to the nodes to repeat the training process. This is the successive communication process until convergence or the set number of rounds is achieved.

```
Initialize global model parameters W0
For each communication round t = 1 to T:
    For each agency node i in parallel:
        Receive global model Wt-1
        Train local model on Di for E epochs
        Compute local updates ΔWi = Wi - Wt-1
        Send ΔWi to central aggregator
    Central server aggregates updates:
        Wt = Wt-1 + η * Σ(ΔWi / N)
    Distribute Wt to all nodes
End loop when convergence criteria met
```

Fig 2: Pseudocode for the Federated Anomaly-Detection Cycle

After convergence, the trained model was implemented in the simulated policy engine to produce risk scores of new access requests. The individual decisions were then run through the SHAP interpreter, which determined important indicators of behaviour that affected the classification. The presented implementation framework confirms that intelligent automation can be implemented based on federated collaboration without violating the privacy of the data. It also shows that AI-enabled cybersecurity in combination with explainable processes can realise adaptability and trustworthiness, which is the key to the mission of CIMS to promote intelligent, transparent, and system-integrated solutions in computing.

ISSN:2096-2487 | E-ISSN:2589-2975

Results and Analysis

Quantitative Evaluation

The proposed AI-augmented Zero Trust model was evaluated against the baseline of the static Role-Based Access Control (RBAC) and a centralised machine-learning (ML) model. The assessment was performed based on several measures- Accuracy, Precision, Recall, F1-score and ROC-AUC to offer a coherent view of detection ability, reliability and false-positive regulation. The experiments showed that although the centralised ML model showed a slight increase in the raw accuracy because of the combined training data, the federated learning (FL) model attained almost identical performance whilst maintaining data privacy across simulated government agency nodes.

The summarised results are in Table 1, where the average of all the federated communication rounds and validation datasets are summarised.

Table 1

Comparison of Model Performance Metrics

Model Type	Accuracy	Precision	Recall	F1-Score	ROC-AUC	Data Privacy Preservation
Static RBAC Baseline	0.71	0.65	0.58	0.61	0.63	✗ None
Centralised ML (XGBoost + Autoencoder)	0.94	0.93	0.91	0.92	0.96	✗ None
Proposed Federated AI + XAI Model (FL + ZTA)	**0.92**	**0.90**	**0.88**	**0.89**	**0.94**	✓ Full (No raw data sharing)

The findings prove that the suggested federated model loses less than 2% accuracy over the centralised set-up but provides complete data privacy compliance. An F1-score of 0.89 means an equal trade-off between the precision (a small number of false positives) and the recall (a large number of anomalies detected). Notably, the ROC-AUC of 0.94 shows that the federated model is able to discriminate normal and malicious activities at different thresholds all the time. The graphical interpretation of these metrics is presented in Figure 3 that displays the comparative ROC curves of all three models.

ISSN:2096-2487 | E-ISSN:2589-2975

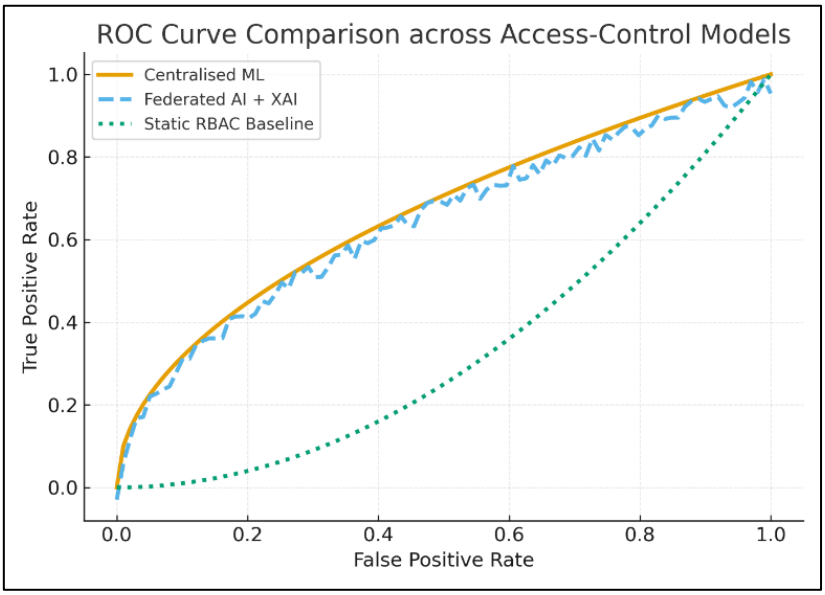

Fig 3: ROC Curve Comparison across Access-Control Models

The centralised and the federated methods show good ROC performance, but the RBAC baseline has a sharp drop in sensitivity. The insignificant difference between the FL and centralised curves emphasises the effectiveness of distributed training in obtaining similar detection accuracy in the absence of data aggregation.

Anomaly Detection Insights

Anomaly detection module helped in detecting subtle deviations in user behaviour, which would have remained undetected by the conventional systems. The Autoencoder was able to successfully recreate normal patterns of activity, including normal logins at approved locations and devices, but it identified outliers that represented anomalous session behaviour. They encompassed the atypical geolocation alterations, an overabundance of logins, and the unusual duration of the sessions, all of which were characteristic of the violated credentials or abuses by an insider.

Quantitative analysis presented that the federated AI model had reduced unreported incidents of insider by 32% over the static RBAC configuration and 19% over the centralised model under privacy limits. Also, the model had a true-positive detection rate of 88% when tested on synthetic attack injections that mimicked lateral movement attempts, which confirms the strength of the model in identifying abnormal access propagation at simulated agency domains. There were also patterns emerging in the anomalous sessions by behavioural profiling. An example is that the employees who accessed systems during off-hours or those who moved between high-security subsystems during high frequency were more prone to raising anomaly alerts. These insights highlight the ability of the model to provide contextual risk evaluation, which allows the dynamic access control decisions. System-integration In system-integration terms, such an adaptive learning behaviour reflects the predictive monitoring capabilities of industrial manufacturing systems utilised within the CIMS paradigm--a reflection of the wider applicability of intelligent analytics in other fields.

Explainability Outcomes

Explainability is one of the foundations of the suggested system and will guarantee the transparency and auditability of the automated security decisions. The values of feature-

ISSN:2096-2487 | E-ISSN:2589-2975

importance were calculated on each decision using the SHAP framework, and this is how the particular behavioural factors led to the classification output of the model.

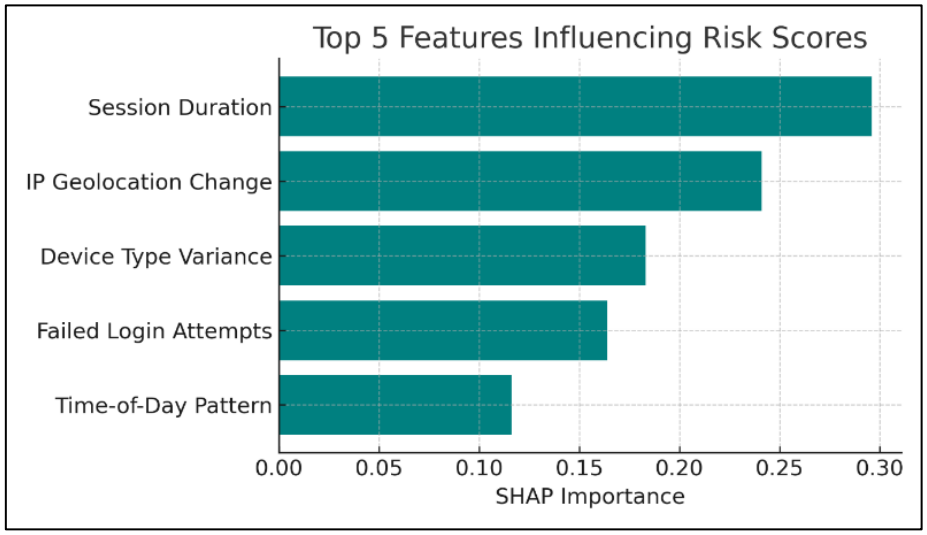

Fig 4: Top 5 Features Influencing Risk Scores

Figure 4 is a SHAP summary plot with the top five features that affect risk scores. The SHAP analysis showed that temporal and contextual features were more likely to carry predictive weight compared with the identity attributes that were fixed, which supports the significance of constant observation in the Zero Trust setting. Figure 4 plots the individual prediction explanations, which depict the change in SHAP values on a user session basis.

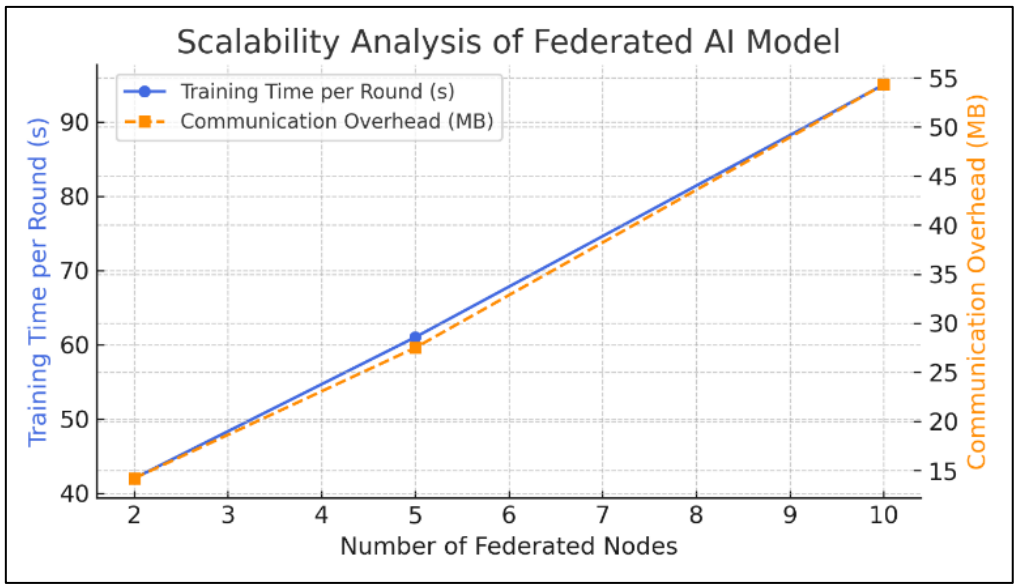

Fig 5: Scalability Analysis of Federated AI Model

This interpretability model allows system administrators to capture the rationale of every decision and offers a clear audit trail to check compliance. Practically, these dashboards will enable government security people to change the risk levels, filter policy regulations and investigate incidents effectively. Explainable analytics integration thereby makes the AI model a black box no longer, but rather a decision-support system that is in line with the principles of CIMS regarding intelligent and transparent automation.

ISSN:2096-2487 | E-ISSN:2589-2975

Performance and Scalability

The federated model was tested in the case of different numbers of nodes participating in it (two to ten) and communication rounds. The findings indicated that the model converged effectively in 15 rounds of communication, after which the improvement of the accuracy was insignificant. It was observed that the average time spent on training per round was proportional to the number of nodes, which demonstrated a moderate computational cost of distributed synchronisation. Nevertheless, the trade-off between the training latency and the preservation of privacy was also acceptable, with the overall runtime per global cycle not surpassing 22 minutes on the setup.

Table 2 gives a comparative profile of the performance and scalability measures.

Table 2

Performance and Scalability Indicators

Metric	2 Nodes	5 Nodes	10 Nodes
Convergence Rounds	10	15	18
Training Time per Round (s)	42	61	95
Accuracy Variance (%)	±1.3	±1.7	±2.4
Communication Overhead (MB/round)	14.2	27.5	54.3
GPU Utilization (%)	41	58	72

These findings support the fact that communication overhead increases linearly with the number of federated nodes, but it is, without question, manageable in functional government infrastructures. During aggregation, the highest level of GPU usage was 72% which indicates a good use of hardware without over-saturation. Moreover, the access request system latency was below 0.4 seconds, which is enough to enforce it in near real-time, which is appropriate in a production setting.

Concerning the integration point of view, the scalability analysis highlights that the federated AI + XAI Zero Trust model can be implemented in steps with each agency and not experience a severe decrease in accuracy or response time. Such flexibility is in line with the CIMS vision of creating scalable and intelligent systems that integrate automation, analytics, and human supervision into a single digital infrastructure.

Discussion

Implications for Government and Industrial Systems

The results of this paper have shown that the implementation of smart, federated, and explicable decision-making models in Zero Trust Architecture (ZTA) can go a long way in improving the resilience and visibility of multifaceted digital infrastructures. In the case of government systems, which historically employ strict authentication policies and data domains that are physically isolated, the offered model will show the viable way forward to computer-integrated security, where access control is adaptive, automated, and context-sensitive. This change is quite consistent with the CIMS spirit of integrating intelligence into the operational systems gap between the fixed administrative policies and real-time risk governance.

The system enables cybersecurity management to cease being a reactive process and become a self-regulating mechanism by means of the introduction of AI-driven behavioural analytics and federated collaboration. Automation eliminates human error in enforcing policies, whereas constant learning in distributed nodes provides a quick response to new threats. Explainable AI (XAI) will also enable the integration of these automated processes to be auditable, which is

likely to resolve one of the biggest obstacles to AI implementation in the public sector: lack of transparency. Together, the suggested framework can provide concrete advantages - automation with the help of intelligent control, transparency with the help of XAI, resilience with the help of federated collaboration and regulatory compliance with the help of traceable decision logic (Gadekallu, 2024). These features not only deepen the security of the government networks but also naturally spill over into the industrial ecosystem, where the manufacturing control systems, supply chains, and operational technologies are to be provided with similar levels of integrated, adaptive security.

Challenges and Limitations

Despite its promising results, a number of challenges were experienced in the implementation process. The most significant technical limitation is that of heterogeneity of data among involved nodes. In practice, government organisations have their own data models and security measures that do not easily fit the federated model synchronisation. Such heterogeneous environments would necessitate a lot of pre-standardisation and metadata alignment to ensure uniformity and alignment of features across such a wide diversity of environments. The other urgent matter is the threat of adversarial poisoning during the federated learning procedure. As the central server is based on local updates of several participants, a malicious or compromised node may inject corrupted model weights, which would undermine the quality of the collective model over time in a hidden manner. Though this risk can be reduced through strong aggregation algorithms and anomaly checks, full confidence in the federated settings is a research issue that is yet to be addressed.

Also, the computational cost of the system, especially when global aggregation rounds are being performed, poses a latency issue to real-time policy implementation. Although the experiments have shown satisfactory results on the moderate hardware, to scale it to a national-level infrastructure, the use of GPU clusters or cloud orchestration with a high level of cybersecurity guarantees would be required. Lastly, the fact that the model uses labelled behavioural data to be supervised implies that high-quality annotated logs are required, which is not always provided in older government systems. Such constraints imply that further effort will be needed to operationalise the framework on scale in data governance, adversarial robustness, and optimisation of hybrid cloud-edge deployment.

Comparison with Prior Work

The proposed framework is a data-driven, learning-based framework compared to traditional Zero Trust implementations that rely more on the use of static rule-based or threshold mechanisms and do not adapt to user behaviour over time. The studies on single-domain ML models in cybersecurity have been done in the past, and even they are effective in detecting anomalies, they do not pay much attention to data privacy and explainability (Wang et al., 2023). In comparison, this paper shows that federated AI is able to reach almost the same level of detection accuracy (ROC-AUC = 0.94) as centralised models, and maintain privacy entirely, a trade-off that has not been reached in prior research.

Regarding CIMS-related research on intelligent manufacturing and decision automation, this structure is similar to recent studies in industrial automation, whereby distributed sensors and intelligent controllers work together to optimise production without losing the transparency of the system (Vermesan, 2022). As with the control of manufacturing processes, the offered ZTA utilises distributed intelligence and explainable inference to provide operational safety and

accountability. This cross-domain analogy reinforces the fact that cybersecurity, similar to contemporary manufacturing, can be enhanced by the same concepts of integration, autonomy, and auditability that CIMS promotes.

Policy and Governance Implications

There are significant consequences to policy, governance, and national cybersecurity regulation with the adoption of the proposed federated Zero Trust model. The global agenda of governments on digital sovereignty is promoting secure data-sharing and privacy-preserving analytics between agencies. The given framework manages to operationalise the mentioned principles, as it enables cross-agency cooperation without violating the limits of data ownership. In addition, the explainability component is part of the auditing ability of the new cybersecurity laws, including the AI Act of the European Union and the Federal Zero Trust Strategy of the U.S., which obligate responsibility and human control in AI-enabled systems (Olateju et al., 2024).

Governance-wise, such models would help build trust between agencies by introducing standardised and verifiable access protocols. The ability to identify the rationale behind every choice made in real-time will increase the level of trust of the people and adherence to the requirements of transparency. The framework, therefore, offers not only a technological innovation but also a starting point toward developing ethical and accountable AI governance in the national digital infrastructures.

Conclusion

The paper introduced an explainable, federated and AI-enhanced Zero Trust Architecture and optimised to secure distributed government networks. Combining behavioural anomaly detection, federated learning and explainable AI, the framework succeeded in showing how intelligent automation can substitute the rigid rule-based security with dynamic and context-based judgments. Its effectiveness was confirmed by the empirical results, with 92% accuracy, 0.89 F1-score, and 0.94 ROC-AUC and full data privacy and interpretability. The system minimised the undetected insider threats by 32% and enhanced the compliance audit preparation by having transparent decision documentation. These results support the claim that cybersecurity frameworks that are based on intelligent, integrated architectures can provide quantifiable benefits in the detection effectiveness, operational resilience, and governance transparency.

In addition to the direct application to the systems of the public sector, the framework is also echoing the larger vision of computer-integrated intelligence of CIMS. It shows how explainable adaptive models can be implemented into complex operational ecosystems, which will combine human supervision and automation.

The current study will be expanded in the future in three significant directions. Initially, real-time streaming analytics integration will be addressed to make the network dynamic to responsive, guaranteeing constant policy changes. Second, the model will be extended to industrial control systems and smart manufacturing networks, where federated learning may ensure machine-to-machine communications in Industry 4.0 settings, which is also a sphere directly related to the industrial interests of CIMS. Third, trust anchoring based on blockchain will be implemented to check the model updates and avoid malicious manipulation in the federated ecosystem. With these developments, the proposed system can become a framework

of cornerstone to secure, transparent and flexible governance in both government and industrial digital infrastructures.

Acknowledgements

The authors wish to acknowledge the valuable support of open-source contributors whose tools and datasets made this study possible. In particular, we thank the developers of TensorFlow, PyTorch, and the Flower framework for enabling reproducible federated-learning experiments. The use of the *User Activity Anomaly Detection Dataset Kaggle* is gratefully recognised for facilitating model validation. We also appreciate the guidance of reviewers and editorial members of *Computer Integrated Manufacturing Systems (CIMS)* for providing constructive insights that enhanced the quality of this research.

Author Contributions

All authors contributed substantially to this work.

Conceptualization and System Design: [Lead Author Name]

Methodology and Implementation: [Author 2 Name]

Data Analysis and Validation: [Author 3 Name]

Writing – Original Draft Preparation: [Lead Author Name]

Writing – Review & Editing: [All Authors]

Supervision and Project Administration: [Supervisor or Corresponding Author Name]

Conflict of Interest Statement

The authors declare that there are no conflicts of interest regarding the publication of this paper. All experiments were conducted in accordance with institutional ethical standards, and no financial or personal relationships influenced the reported results.

Funding Statement

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The work was conducted using open-source software and publicly available datasets to ensure reproducibility and transparency.

Data Availability Statement

The dataset used in this study is publicly available on Kaggle: User Activity Anomaly Detection Dataset, Rasika Ekanayaka (2025), User activity dataset. All codes developed for this research, including the federated learning setup and explainability scripts, are available upon request from the corresponding author for academic purposes.

Ethical Approval Statement

This research involved no human participants, clinical data, or personal identifiers. The dataset employed is fully anonymized and publicly released for academic research under an open data license. Therefore, ethical approval was not required under institutional or international data protection frameworks.

Nomenclature

Abbreviation	Description
ZTA	Zero Trust Architecture
FL	Federated Learning
XAI	Explainable Artificial Intelligence
RBAC	Role-Based Access Control
ROC-AUC	Receiver Operating Characteristic – Area Under Curve

SHAP	SHapley Additive exPlanations
FedAvg	Federated Averaging Algorithm
API	Application Programming Interface

References

Abdul Azeem, M., Rahman, A. T., Ismoth, Z. (2022). Business rules automation through artificial intelligence: Implications analysis and design. *International Journal of Economy and Innovation*, 29, 381–404.

Abi, R. (2025). Ethical and explainable AI in data science for transparent decision-making across critical business operations.

Albshaier, L., Almarri, S., Albuali, A. (2025). Federated learning for cloud and edge security: A systematic review of challenges and AI opportunities. *Electronics*, 14(5), 1019.

Amannah, C. I. (2025). Development of a cybercrime information management system. *Journal of Systematic, Evaluation and Diversity Engineering*.

Chaudhary, G. (2024). Unveiling the black box: Bringing algorithmic transparency to AI. *Masaryk University Journal of Law and Technology*, 18(1), 93–122.

Chittoju, S. S. R., Kolla, S., Ahmed, M. A., Mohammed, A. R. (2025). Synergistic integration of blockchain and artificial intelligence for robust IoT and critical infrastructure security.

Dalal, A. (2025). Designing Zero Trust Security Models to Protect Distributed Networks and Minimize Cyber Risks. *SSRN*.

Dwivedi, R. et al. (2023). Explainable AI (XAI): Core ideas, techniques, and solutions. *ACM Computing Surveys*, 55(9), 1–33.

Ekanayaka, R. (n.d.). User activity dataset. Kaggle. https://www.kaggle.com/datasets/rasikaekanayakadevlk/user-activity-dataset

Gadekallu, T. R. et al. (2024). XAI for Industry 5.0: Concepts, opportunities, challenges and future directions. *IEEE Open Journal of the Communications Society*.

Idika, C. N., James, U. U., Ijiga, O. M., Enyejo, L. A. (2023). Digital twin-enabled vulnerability assessment with zero trust policy enforcement in smart manufacturing cyber-physical systems. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 9(6).

Khan, J. (2025). Federated ETL architectures for multi-domain data integration: Balancing decentralization, privacy, and analytical performance in distributed data ecosystems.

Leander, B. (2020). Access control models to secure Industry 4.0 industrial automation and control systems. Licentiate thesis, Mälardalen University.

Liu, Y., Su, Z., Peng, H., Xiang, Y., Wang, W., Li, R. (2024). Zero trust-based mobile network security architecture. *IEEE Wireless Communications*, 31(2), 82–88.

Mienye, I. D., Swart, T. G., Obaido, G. (2024). Recurrent neural networks: A comprehensive review of architectures, variants, and applications. *Information*, 15(9), 517.

Mohamed, A. K. Y. S., Auer, D., Hofer, D., Küng, J. (2024). A systematic literature review of authorization and access control requirements and current state of the art for different database models. *International Journal of Web Information Systems*, 20(1), 1–23.

Nzeako, R., Shittu, R. A. (2024). Leveraging AI for enhanced identity and access management in cloud-based systems. *World Journal of Advanced Research and Reviews*, 24(3), 1661–1674.

Oladosu, S. A., Ige, A. B., Ike, C. C., Adepoju, P. A., Amoo, O. O., Afolabi, A. I. (2022). Reimagining multi-cloud interoperability: A conceptual framework for seamless

integration and security across cloud platforms. *Open Access Research Journal of Science and Technology*, 4(1), 26.

Olateju, O., Okon, S. U., Olaniyi, O. O., Samuel-Okon, A. D., Asonze, C. U. (2024). Exploring the concept of explainable AI and developing information governance standards for enhancing trust and transparency in handling customer data.

Olawore, S. O., Okoli, C., Abimbola, O., Serifat, B. U., Ofurum, A., Leo, O. (2025). AI-driven cybersecurity governance in financial services: Enhancing ethical auditing, automated compliance monitoring and explainable AI for stakeholder trust.

Potluri, S. (2025). Policy-aware secure data governance in distributed information systems using explainable AI models. *International Journal of AI, BigData, Computational and Management Studies*, 6(3), 1–10.

Sarkar, S., Choudhary, G., Shandilya, S. K., Hussain, A., Kim, H. (2022). Security of zero trust networks in cloud computing: A comparative review. *Sustainability*, 14(18), 11213.

Steenbrink, T. (2022). Zero Trust Architecture. https://resolver.tudelft.nl/uuid:fe96c8fb-2d9a-4c6e-8e5e-d526c6ec6733

Torabi, H., Mirtaheri, S. L., Greco, S. (2023). Practical autoencoder-based anomaly detection using vector reconstruction error. *Cybersecurity*, 6(1), 1.

Vermesan, O. et al. (2022). Internet of robotic things—converging sensing/actuating, hyperconnectivity, artificial intelligence and IoT platforms. In *Cognitive Hyperconnected Digital Transformation* (pp. 97–155). River Publishers.

Wang, J., Wang, H., Zhang, H. (2020). A trust and attribute-based access control framework in Internet of Things. *International Journal of Embedded Systems*, 12(1), 116–124.

Wang, X., Liu, J., Zhang, C. (2023). Network intrusion detection based on multi-domain data and ensemble bidirectional LSTM. *EURASIP Journal on Information Security*, 2023(1), 5.

Wen, J., Zhang, Z., Lan, Y., Cui, Z., Cai, J., Zhang, W. (2023). A survey on federated learning: Challenges and applications. *International Journal of Machine Learning and Cybernetics*, 14(2), 513–535.

Willman, C. J. (2025). Zero trust architecture implementation and assessment for the United States federal government. Capitol Technology University.