International Journal on Science and Technology (IJSAT)
E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Cloud Network Anomaly Detection using Federated Learning and Explainable AI

Mr. Praveen Kumar[1], Reddy Idamakanti[2]

[1]CSE, [2]Btech
Dr.MGR Educational and Research Instittue

Abstract

Cloud computing environments handle vast amounts of data, making them prime targets for cyber threats such as Distributed Denial-of-Service (DDoS) attacks, ransomware, and insider threats. Traditional centralized anomaly detection methods pose significant privacy risks, scalability challenges, and high computational costs. To address these issues, we propose a privacy-preserving, federated learning (FL)-based anomaly detection model that enables decentralized threat detection without exposing raw data. Our approach integrates Explainable AI (XAI) techniques such as SHAP, LIME, and attention mechanisms to enhance interpretability and transparency, enabling security analysts to understand and validate AI-driven anomaly detections. We optimize model synchronization to reduce communication overhead. The proposed system ensures real-time threat detection, adaptability to evolving attack patterns. Experimental evaluations demonstrate improved accuracy, lower false positives, and enhanced explainability, making our approach a scalable and trustworthy solution for cloud network anomaly detection.

1. Introduction

Cloud environments handle large volumes of sensitive and dynamic data, making them highly vulnerable to a wide range of cyber threats, including data breaches, ransomware, and denial-of-service attacks. As organizations increasingly rely on cloud services, ensuring the confidentiality, integrity, and availability of data has become a critical challenge. To address these risks, effective and timely anomaly detection mechanisms are essential for identifying unusual activities that may indicate security breaches. However, traditional anomaly detection methods, which are typically centralized, pose significant challenges in cloud environments. Centralized approaches raise privacy concerns since sensitive data must be aggregated and processed in a central location. Moreover, they struggle to scale efficiently across distributed and heterogeneous cloud infrastructures and often face limitations in delivering real-time detection across geographically dispersed systems.

Federated Learning (FL) has emerged as a promising alternative by enabling decentralized model training across multiple clients without requiring the sharing of raw data. This decentralized approach enhances privacy preservation and supports scalability in distributed environments. Nevertheless, FL-based anomaly detection models suffer from a critical drawback: a lack of interpretability and transparency. The inability to explain or justify the model's predictions reduces trust and limits the practical adoption of such systems in security-sensitive contexts. To bridge this gap, Explainable AI (XAI) techniques have been introduced to provide insights into the inner workings and decision-making processes of complex AI models. Techniques such as SHAP (Shapley Additive explanations), LIME (Local Interpretable Model-

agnostic Explanations), and attention mechanisms can enhance transparency, foster trust, and support more informed decision-making by security analysts.

This research aims to develop a federated, explainable anomaly detection framework for cloud networks by integrating XAI techniques into the anomaly detection process. The primary goals include enhancing model explainability to improve user trust, evaluating the framework's accuracy, interpretability, and computational efficiency, and designing an FL-based anomaly detection model that achieves a balance between privacy preservation, scalability, and interpretability. Ultimately, this research contributes to the field by advancing privacy-preserving, transparent, and effective anomaly detection solutions for modern cloud infrastructures, addressing both security and explainability challenges in a cohesive framework.

1.1 Project Overview:

The project aims to enhance anomaly detection in cloud networks by leveraging Federated Learning (FL) and Explainable AI (XAI). Traditional centralized anomaly detection systems face challenges related to data privacy, scalability, and lack of transparency. To address these issues, the proposed system uses a decentralized FL approach, allowing individual cloud nodes to train anomaly detection models locally without sharing raw data, thereby preserving privacy and complying with regulations like GDPR. Model updates are securely aggregated on a central federated server, ensuring robustness against poisoning and adversarial attacks. To overcome the interpretability limitations of AI-based anomaly detection, XAI techniques such as SHAP, LIME, and attention mechanisms are integrated, providing transparent, actionable insights for security analysts. The system continuously adapts through analyst feedback, improving accuracy and resilience against evolving cyber threats. This approach not only strengthens data privacy and security but also improves scalability, transparency, and trust in AI-driven anomaly detection for modern cloud environments.

1.2 AIM OF THE PROJECT:

The aim of this project is to design and implement a robust, privacy-preserving, and explainable anomaly detection system tailored for cloud computing environments. By leveraging Federated Learning (FL), the system enables decentralized model training across multiple distributed data sources without the need to share raw data, thereby safeguarding sensitive information and mitigating privacy risks. Furthermore, the integration of Explainable AI (XAI) techniques—such as SHAP, LIME, and attention mechanisms—enhances the interpretability and transparency of the anomaly detection outcomes, allowing stakeholders to better understand, trust, and act upon the model’s predictions. This project aspires to achieve a balance between high detection accuracy, interpretability, scalability, and computational efficiency, ultimately contributing to more secure and trustworthy cloud network operations.

1.3 Scope of the Project:

The scope of this project extends to the comprehensive development of an advanced, privacy-preserving anomaly detection system specifically designed for cloud computing environments, leveraging the synergistic integration of Federated Learning (FL) and Explainable Artificial Intelligence (XAI) techniques. This system aims to address the critical need for secure, decentralized anomaly detection in the cloud, where large volumes of sensitive and distributed data are managed across multiple nodes and organizations. By employing Federated Learning, the project eliminates the requirement to centralize raw

data, thereby significantly enhancing data privacy and compliance with regulations such as GDPR and HIPAA, while enabling collaborative model training across diverse data sources.

At the core of the system lies the integration of XAI methodologies—such as SHAP (Shapley Additive explanations), LIME (Local Interpretable Model-agnostic Explanations), and attention-based mechanisms—which are incorporated to transform the traditionally opaque nature of machine learning models into transparent and interpretable tools. These explainability features empower cybersecurity analysts and decision-makers to understand the rationale behind anomaly detection outputs, facilitating more informed and confident responses to potential threats, reducing false positives, and building trust in AI-driven security systems.

The project also encompasses the implementation of mechanisms to ensure scalability and robustness in real-world cloud environments, addressing challenges such as data heterogeneity, non-IID (non-independent and identically distributed) data across nodes, communication efficiency, and adversarial resilience. By optimizing model aggregation, synchronization, and secure communication protocols, the system will maintain high detection accuracy and operational efficiency even as the number of participating nodes grows.

Furthermore, the scope includes detecting a broad range of cyber threats—including but not limited to Distributed Denial of Service (DDoS) attacks, insider threats, ransomware, malware infections, and zero-day vulnerabilities—by continuously adapting to evolving attack patterns through incremental and continual learning strategies. This ensures that the anomaly detection system remains up-to-date and capable of recognizing novel attack vectors without compromising on performance or interpretability.

Beyond technical development, the project's scope also involves extensive evaluation of the system's performance through metrics such as detection accuracy, precision, recall, interpretability scores, training efficiency, communication overhead, and robustness against adversarial manipulations. It aims to produce a scalable, explainable, and privacy-compliant anomaly detection framework that can be effectively deployed in modern cloud infrastructures, ultimately contributing to more secure, transparent, and trustworthy cloud computing ecosystems.

2. LITERATURE REVIEW

[1] The research article titled **"*Cloud Network Anomaly Detection Using Machine and Deep Learning Algorithms*"** is authored by **Amira Abdallah, Aysha Alkaabi , Ghaya Alameri , Saida Hafsa Rafique** Their work focuses on addressing the growing need for robust security mechanisms in cloud networks, where Increasing data volumes and complex infrastructures make detecting anomalies a critical challenge.**aim of the project** is to design and develop an effective anomaly detection system tailored for cloud environments by leveraging both **machine learning (ML)** and **deep learning (DL)** techniques. The goal is to create a hybrid model capable of identifying malicious traffic patterns, thereby safeguarding cloud resources from cyberattacks. By integrating different algorithms and evaluating their performance, the study aims to enhance detection accuracy while reducing false alarms.

[2] The article "***Design of An Anomaly Detection Framework For Delay & Privacy-Aware Blockchain-Based Cloud Deployment*** "The authors A Venkata Nagarjun, Sujatha Rajkumar, presents a framework aimed at detecting anomalies in network delay patterns, focusing on applications within communication systems, IoT networks, and cloud infrastructures. Its scope involves improving real-time monitoring and predictive maintenance by identifying abnormal delays that could indicate underlying issues. The proposed system architecture includes several interconnected components: a data collection module that gathers real-time delay metrics from network nodes; a preprocessing unit for filtering and normalizing data; a feature extraction module that computes relevant statistical and temporal features; an anomaly detection engine leveraging machine learning algorithms to identify deviations from normal behavior; an alert system that triggers notifications or logs upon detection of anomalies; and a visualization dashboard providing real-time and historical anomaly reports. Despite its strengths, the framework faces drawbacks such as increased complexity and computational overhead, particularly when deployed in large-scale, real-time environments, and the need for frequent retraining or tuning to adapt to evolving network traffic patterns.

[3] The study titled *"**A Systematic Review on Anomaly Detection for Cloud Computing Environments**"* by Ali Dehghantanha, Kim-Kwang Raymond Choo, et al., published in 2021, provides a comprehensive evaluation of 215 research publications in the field. It offers valuable insights into the existing methods, challenges, and trends in anomaly detection within cloud environments. However, one limitation of this review is that it may not encompass the most recent advancements and developments that have emerged in the field after 2021, potentially leaving out newer techniques and approaches.

[4] The paper titled *"**Cloud-based Multiclass Anomaly Detection and Categorization Using Ensemble Learning**"* by Muhammad Asim, Muhammad Awais Azam, et al., published in 2022, introduces a Cloud-based Anomaly Detection (CAD) framework that leverages ensemble learning techniques to achieve high accuracy in detecting and categorizing anomalies. While the framework demonstrates strong performance, a key limitation lies in its potential requirement for extensive computational resources. This dependency may restrict its practical applicability in resource-constrained environments where computing power and infrastructure are limited.

[5] The paper titled ***"Anomaly Detection in Cloud Network: A Review"*** by Chukwuemeka Nwachukwu, Kehinde Durodola-Tunde, and Chukwuebuka Akwiwu-Uzoma, published in 2024, provides a thorough survey of existing techniques and challenges in detecting anomalies within cloud network environments. The study successfully identifies key research gaps and proposes potential solutions to address these limitations. However, a notable shortcoming of the survey is its lack of empirical validation, as the suggested approaches are not tested or evaluated through experimental results, leaving their practical effectiveness uncertain.

[6] The study titled *"**Machine Learning-Based Anomaly Detection in Cloud Virtual Machine Resource Usage**"* by Haili Wang, Jingda Guo, et al., published in 2023, explores the use of machine learning techniques to detect anomalies in the resource usage patterns of cloud virtual machines. The research primarily relies on the NSL-KDD dataset for training and evaluation. While the study provides valuable insights, a key limitation is the reliance on this dataset, which may not fully capture the complexity and

dynamics of modern cloud environments. This limitation raises concerns about the generalizability and applicability of the findings to current real-world cloud systems.

[7] The paper titled *"**Anomaly Detection in Cloud Computing Using Knowledge Graph and GraphSAGE**"* by Yuxin Su and Michael R. Lyu, published in 2023, presents a novel approach to anomaly detection by leveraging knowledge graphs combined with the GraphSAGE algorithm. This methodology effectively manages the complexity of diverse and dynamic cloud resources, enabling improved anomaly detection across interconnected systems. However, a notable limitation is its potential scalability challenge when applied to extremely large-scale cloud environments, where the computational and storage demands of processing massive graphs could hinder performance and efficiency.

[8] The paper titled *"AI-driven Anomaly Detection in Cloud Computing Environments"* by Ali Dehghantanha, published in 2024, explores a range of artificial intelligence techniques for detecting anomalies in cloud computing systems. The study provides a broad overview of AI-based methods and their potential benefits for improving security and performance in cloud environments. However, a key limitation is the absence of detailed implementation insights and real-world application results, making it difficult to assess the practical effectiveness and applicability of the proposed techniques in operational settings.

[9] The paper titled *"**Maat: Performance Metric Anomaly Anticipation for Cloud Services with Conditional Diffusion**"* by Cheryl Lee, Tianyi Yang, et al., published in 2023, introduces an innovative two-stage paradigm designed to anticipate performance metric anomalies in cloud services using conditional diffusion techniques. This approach offers a proactive strategy for identifying potential issues before they impact service quality. However, a notable limitation is the need for further validation across diverse and varied cloud environments to ensure the generalizability and robustness of the proposed method in different operational contexts

[10] The paper titled *"**Cloud Shield: Real-time Anomaly Detection in the Cloud**"* by Zecheng He and Ruby B. Lee, published in 2021, presents a system designed for real-time anomaly detection in cloud computing environments. The research highlights the importance of timely identification and mitigation of anomalies to maintain cloud security and performance. However, a key limitation of this approach is its focus on real-time detection rather than faster-than-real-time anomaly anticipation, which may be necessary for proactively preventing issues in highly dynamic and mission-critical cloud systems.

[11] The paper titled *"**Online Self-Evolving Anomaly Detection in Cloud Computing Environments**"* by Haili Wang, Jingda Guo, et al., published in 2021, introduces a self-evolving anomaly detection framework designed to adapt and improve continuously over time within cloud environments. This approach shows significant promise in maintaining detection accuracy without requiring frequent manual updates. However, a notable limitation is the potential difficulty the framework may encounter in keeping pace with the rapid and unpredictable changes typical of modern cloud infrastructures, which could affect its adaptability and effectiveness in highly dynamic settings.

[12] The paper titled *"**Anomaly Detection in a Large-scale Cloud Platform**"* by Mohammad Saiful Islam, William Pourmajidi, et al., published in 2020, investigates anomaly detection techniques within the context of the IBM Cloud Platform. The study provides valuable insights into detecting anomalies in large-

scale, enterprise-level cloud infrastructures. However, a key limitation is that the findings are closely tied to the IBM Cloud's specific architecture and ecosystem, which may limit the generalizability and applicability of the results to other cloud platforms that operate under different architectures and configurations.

[13] Harshvardhan Chunawala and Pratikkumar Chunawala proposed "***a framework for advanced anomaly detection in cloud infrastructures using deep learning algorithms***", specifically Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs). The goal of the framework was to detect and categorize anomalies in real-time, leveraging a range of indicators, including CPU utilization, memory consumption, network traffic, and user behavior patterns. The framework demonstrated an ability to achieve higher accuracy and speed when compared to traditional anomaly detection methods.

[14] Sabbir M. Saleh, Ibrahim Mohammed Sayem, Nazim Madhavji, and John Steinbacher implemented a combination of Convolutional Neural Network (CNN) and Long Short-Term Memory (LSTM) models to detect unusual traffic patterns in cloud platforms. Their approach achieved impressive accuracies of 98.69% and 98.30% on the CSE-CIC-IDS2018 and CSE-CIC-IDS2017 datasets, respectively. The framework also focused on generating log files at different stages of the Continuous Integration/Continuous Deployment (CI/CD) pipeline, addressing security challenges in modern DevOps practices. Despite the successes, the approach relies heavily on existing datasets, which may not fully represent emerging threats in cloud environments. This limitation suggests that future work should explore newer datasets or develop methods to adapt to evolving security challenges in dynamic cloud platforms.

3. EXISTING SYSTEM

3.1 "Privacy-Preserving Anomaly Detection Using Federated Learning and Explainable AI" , presents a system that integrates federated learning (FL), Generative Adversarial Networks (GAN), and Explainable AI (XAI) for privacy-preserving anomaly detection in network traffic. The system architecture, depicted in

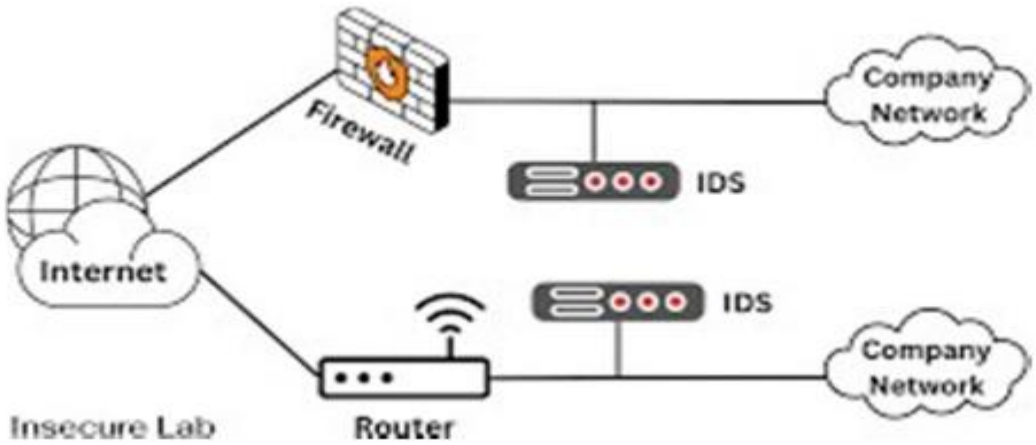

Figure 3.1

employs a federated learning framework where multiple cloud clients collaboratively train anomaly detection models (deep autoencoders, LSTM, Transformer-based) on their local network traffic data. A central server is responsible for aggregating the model updates received from the clients. To enhance the model's ability to detect novel, zero-day attacks, GANs are utilized to generate synthetic attack data, which is then used to augment the training process. For providing interpretability to the anomaly detection results, the system integrates explainable AI techniques, specifically SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations). These techniques help in assigning contribution values to individual network traffic features, allowing security analysts to identify the attributes that influenced the model's decision to flag a particular activity as anomalous. To ensure the privacy and security of the sensitive network traffic data during the federated learning process, the system employs homomorphic encryption and differential privacy for secure communication and model aggregation.

The key contributions of this research include the development of a privacy-preserving anomaly detection system that achieves high accuracy while maintaining the confidentiality of network data. The experimental results indicate that the FL-based approach reduces the risk of data leakage compared to centralized models, without compromising detection performance. The integration of SHAP and LIME enhances the transparency of the system, enabling security analysts to understand and validate the anomaly detection decisions. Furthermore, the FL-enabled model aggregation ensures robust generalization across distributed environments, making the system adaptable for real-world deployment in various cloud and IoT security infrastructures. However, this system also presents certain drawbacks and limitations. The paper does not provide a detailed analysis of the communication overhead inherent in the federated learning process, which can be a significant factor in cloud environments with a large number of participating clients. The computational resources required for training deep learning models and GANs on the client devices might be substantial, potentially limiting the scalability of the system to resource-constrained cloud instances. While the paper emphasizes the privacy benefits, a more in-depth discussion on how the system effectively handles highly heterogeneous network traffic data distributions across different cloud clients would be beneficial. Additionally, the complexity of integrating multiple advanced techniques (FL, GAN, XAI, encryption) might introduce challenges in terms of practical implementation and ongoing maintenance.

3.2 "Federated Learning-Based Explainable Anomaly Detection for Industrial Control Systems" proposes the FedeX architecture, a federated learning-based system for explainable anomaly detection in Industrial Control Systems (ICSs). The system architecture, illustrated in Figure 3.2, is designed for smart factories with a distributed structure of IoT devices organized into multiple zones. Each zone is monitored by a local edge computing unit. In the first step, each edge device collects sensing data from the nodes within its zone, forming a local dataset. Subsequently, the edge device trains a local anomaly detection model using a lightweight Variational Autoencoder (VAE). After local training, each edge device uploads the weight matrix of its trained VAE model to a central cloud aggregator using the MQTT protocol, a standard for IoT environments. The cloud aggregator receives the weight matrices from all participating edge devices and aggregates them to create a new global model based on federated learning principles.

The updated global model is then distributed back to each edge device. For anomaly detection, the trained FedVAE model at each edge device is used in conjunction with Support Vector Data Description (SVDD)

to automatically determine the anomaly threshold. To provide explainability for the detected anomalies, the system periodically runs an XAI module using Shapley Additive explanations (SHAP) to interpret the predictions of the FedVAE-SVDD model and identify the contribution of each feature to the anomaly score.

The key contributions of the FedeX system include achieving high detection performance and real-time anomaly detection in a distributed ICS environment while preserving data privacy. The use of federated learning enables the global model to learn from the collective knowledge of all zones without requiring the

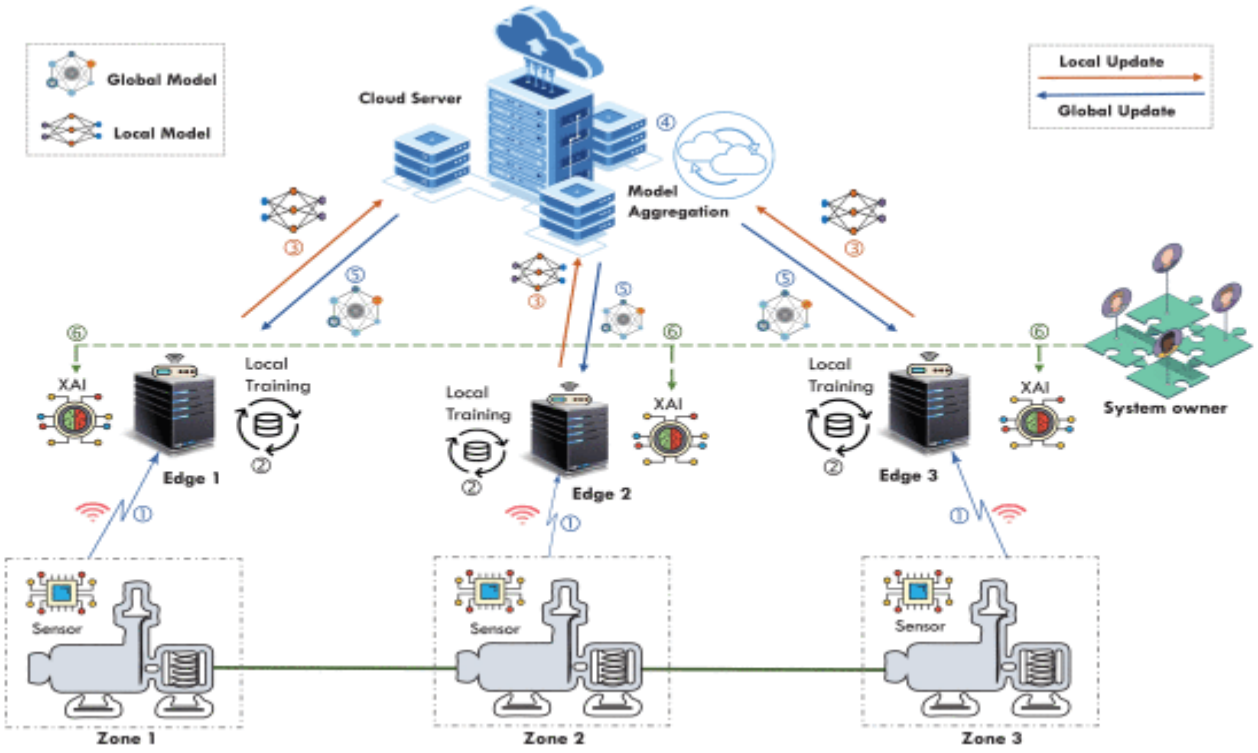

Figure 3.2

centralization of raw data. The integration of SHAP offers valuable explanations for detected anomalies, allowing domain engineers to understand the underlying reasons for the classifications. Experimental results demonstrate that the FedeX system outperforms several other existing anomaly detection solutions in terms of detection capability and response time. However, there are also certain drawbacks and limitations associated with this system. Implementing machine learning models, even lightweight ones like VAEs, on resource-constrained edge devices still incurs a computational cost, which might impact other critical services running on these devices. The paper notes that the SVDD phase of training can be particularly demanding in terms of CPU usage on edge devices like Raspberry Pi 4. There might be a trade-off between the simplicity of the lightweight models used and the potential detection performance achievable with more complex centralized models. Additionally, while federated learning reduces the need to send raw data, the communication overhead involved in exchanging model weights between the edge devices and the cloud server in each communication round could be a factor in large-scale deployments.

The effectiveness of the global model might also be affected by significant heterogeneity in the data distributions across different ICS zones

3.3 "Comprehensive Analysis over Centralized and Federated Learning-based Anomaly Detection in Networks with Explainable AI (XAI)" presents a comparative analysis of anomaly detection systems based on both centralized learning (CL) and federated learning (FL) architectures, both enhanced with Explainable AI (SHAP). The system architecture for the centralized learning model involves collecting network data from various sources within the network and storing it in a central, trustable data collector. A Deep Neural Network (DNN) is then trained using this entire centralized dataset for anomaly detection. After training, SHAP is applied to the CL model to analyze its predictions and identify anomalies in network flows. In contrast, the federated learning model divides the network into multiple security domains, each acting as a client. Each client has its own local dataset and trains a local DNN for anomaly detection. Model updates are exchanged with a central parameter server for aggregation, and SHAP is applied locally at each client to generate explanations based on the local model and data.

The key contributions of this study include a comprehensive comparison of centralized and federated learning approaches for network anomaly detection, both integrated with the SHAP explainable AI technique. The experimental results on network datasets (UNSW-NB15 and NSL-KDD) reveal that while federated learning offers advantages in terms of data privacy, it generally achieves lower accuracy in anomaly detection compared to centralized learning models. Furthermore, the performance of the federated learning model is found to be sensitive to the percentage of anomalies present in the local data of each client. The study utilizes SHAP to provide global explanations of feature importance in both centralized and federated learning settings. However, the research also identifies several drawbacks and limitations. The lower accuracy of the federated learning models compared to centralized models suggests a potential trade-off between privacy and detection performance. The significant impact of the anomaly distribution across clients on the federated learning model's performance could be a limitation in real-world scenarios with varying data characteristics. The computational intensity of SHAP, especially for large datasets, might pose challenges for scalability. Additionally, the observed linear relationship between feature importance changes in centralized and federated learning models was primarily noticeable in the top features, indicating that the explainability might not be consistent across all features. The underlying anomaly detection models

in both architectures are black-box DNNs, and while SHAP provides post-hoc explanations, the

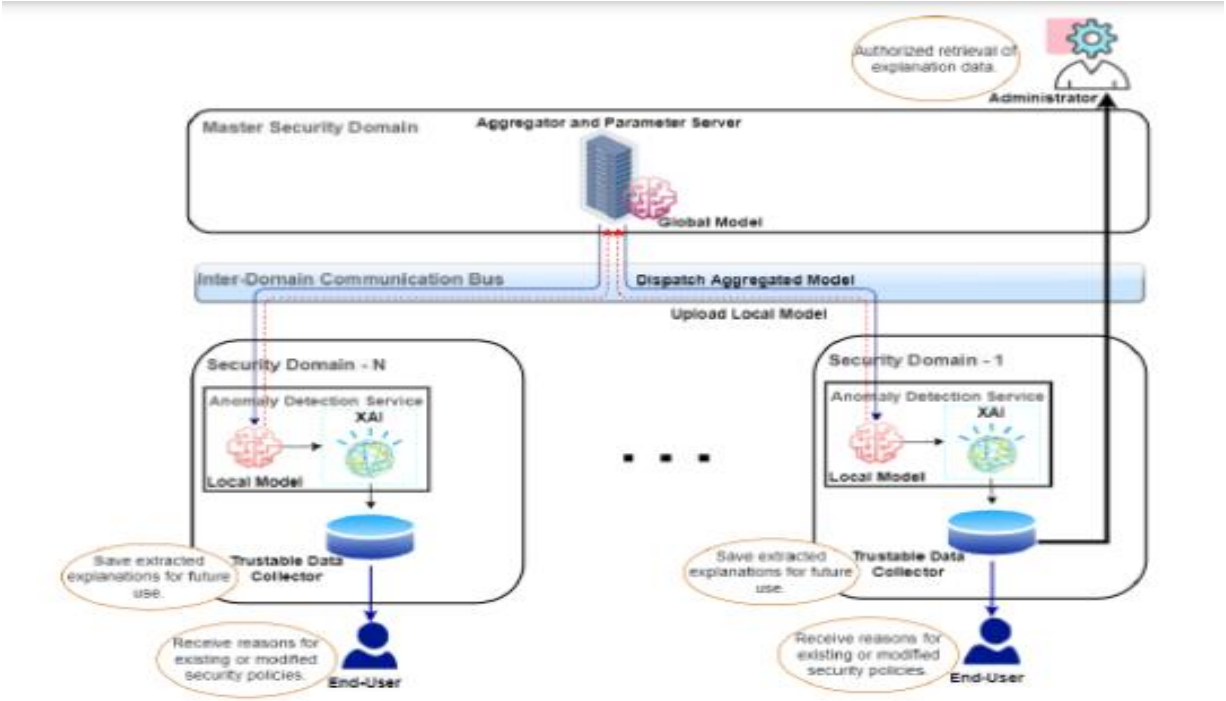

Figure – 3.3

Based on the analysis of these three existing systems, our research paper aims to address the following limitations: the accuracy gap observed in federated learning-based cloud anomaly detection compared to centralized approaches, the need for more efficient and scalable explainable AI methods in federated settings, and the enhancement of the robustness of federated learning against data heterogeneity and potential attacks in cloud environments.

3.4 PROBLEM STATEMENT

The increasing adoption of federated learning (FL) in cloud and IoT environments has led to advancements in anomaly detection systems, particularly in maintaining data privacy while achieving accurate detection of network anomalies. However, challenges persist, particularly regarding the accuracy gap observed in FL-based systems compared to centralized learning approaches, as well as the computational cost and communication overhead involved in federated settings. Additionally, existing systems often face limitations in handling heterogeneous data distributions across different cloud or IoT clients, and the need for efficient and scalable explainable AI (XAI) methods remains critical. Although methods like SHAP and LIME have been integrated to enhance interpretability, they are computationally intensive and may not offer consistent explanations across all features. Furthermore, existing models struggle to address emerging threats, such as zero-day attacks, and fail to fully exploit the benefits of federated learning, including its ability to adapt to distributed environments. This research aims to bridge the accuracy gap between federated and centralized anomaly detection systems, optimize XAI methods for federated environments, and improve the robustness of federated learning models against data heterogeneity and potential security threats in cloud infrastructures.

4. PROPOSED SYSTEM

4.1 OBJECTIVES

The primary objective of this research is to develop a novel federated learning framework for cloud anomaly detection that not only maintains data privacy but also achieves comparable accuracy to centralized approaches. This will ensure that cloud tenants can benefit from the collaborative power of federated learning without compromising the confidentiality of their data. A key challenge in federated learning is handling heterogeneous data distributions across different cloud tenants or regions. To address this, the framework will incorporate adaptive client weighting and personalized model adjustments, enabling the system to dynamically adapt to the unique characteristics of each tenant's data. This will improve the model's generalization and accuracy across diverse environments.

Another key objective is to integrate an efficient and scalable explainable AI (XAI) module that utilizes techniques like SHAP (Shapley Additive Explanations) and LIME (Local Interpretable Model-Agnostic Explanations). These techniques will be tailored for federated learning, ensuring that high-quality and consistent explanations for detected anomalies can be provided at both global and local levels. This will help security analysts understand the reasoning behind the model’s anomaly detection decisions, enhancing transparency and trust in the system.

To further improve the efficiency of the federated learning process, the framework will focus on minimizing the communication overhead that is often a limiting factor in large-scale cloud deployments. This will be achieved through techniques such as model compression and selective parameter updates, which will reduce the amount of data exchanged between cloud clients and the central server, thereby optimizing the overall performance.

Lastly, security is a critical concern in federated learning, especially in cloud environments where the risk of model poisoning and information leakage is high. To mitigate these risks, the system will incorporate secure aggregation protocols and differential privacy techniques to safeguard the integrity of the model training process and ensure that sensitive data is not exposed during communication or aggregation. By addressing these challenges, the proposed system will provide a robust, privacy-preserving, and scalable solution for anomaly detection in cloud environments.

4.2 ARCHITECTRE

The Anomaly Detection Engine plays a critical role in identifying deviations or suspicious patterns in data streams. It employs a variety of machine learning algorithms, including supervised, unsupervised, and deep learning techniques, to detect anomalies in real-time or batch mode. For instance, it can identify DDoS attacks by analyzing abnormal spikes in network traffic. Complementing this engine is the Explainable AI (XAI) module, which provides interpretable explanations for the system's decisions using techniques like SHAP and LIME. This transparency is crucial for validating alerts, ensuring regulatory compliance, and building trust in automated systems.

A key feature of this architecture is the Federated Learning Server (Central Aggregator), which enables privacy-preserving model training. Instead of sharing raw data, local nodes train models on their own data

and send only model updates to the server. The server aggregates these updates to improve a global model, ensuring data privacy—a vital requirement for industries like healthcare and finance. This approach also reduces bandwidth usage and supports edge intelligence, allowing real-time local inference for IoT devices.

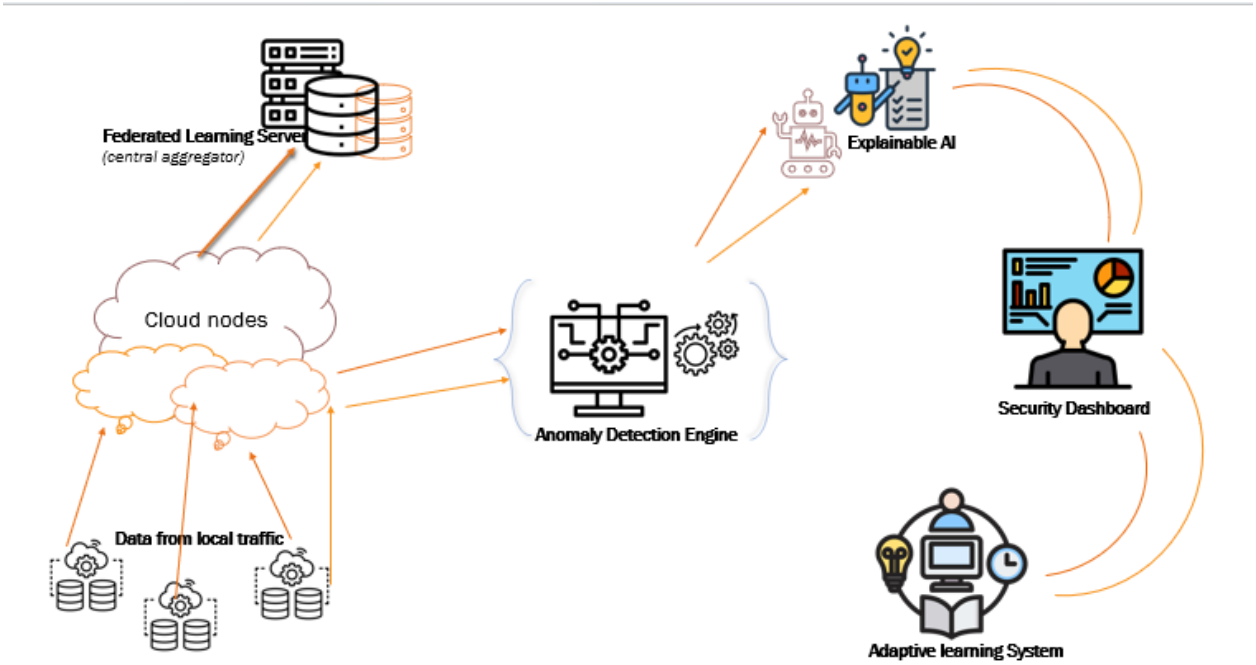

Figure – 4.2

Data for the system is sourced from Local Traffic, which includes network logs, IoT sensor readings, and user activity records. This data remains decentralized, processed on-premises or at the edge, minimizing latency and avoiding cross-border data transfer issues. The insights derived from this data are presented through a Security Dashboard, which offers real-time alerts, visual analytics like heatmaps and risk scoring, and drill-down capabilities for deeper investigation. This dashboard caters to different user roles, from SOC analysts handling triage to executives monitoring high-level threats.

To ensure continuous improvement, the Adaptive Learning System updates models incrementally based on new data and feedback. It addresses challenges like concept drift and reduces false positives by incorporating analyst corrections. Finally, the system includes Conditional Logic ("If anomaly detect") to automate responses. Detected anomalies can trigger notifications, mitigation actions (e.g., blocking malicious IPs), or escalation protocols, enabling rapid response and reducing manual workload.

4.3 MODULE DESCRIPTION

The Federated Anomaly Detection System comprises eight key modules that work together to provide a comprehensive privacy-preserving solution for identifying and responding to anomalies. The Cloud Computing Module forms the foundation, offering scalable infrastructure for model training and data processing through virtual machines and auto-scaling resources. At the core, the Anomaly Detection Module utilizes machine learning models like Random Forest and LSTM to identify abnormal patterns in data streams, while the Explainable AI (XAI) Module provides interpretable explanations of detections

using techniques like SHAP and LIME to ensure transparency. The Federated Learning Module enables privacy-preserving distributed training by securely aggregating model updates from edge nodes without sharing raw data. Local data collection and preprocessing is handled by the Edge Data Processing Module, which sanitizes data and extracts features before secure upload. Detected anomalies are visualized and managed through the Security Operations Module's dashboard, which integrates with existing SIEM systems. Continuous improvement is achieved via the Adaptive Learning Module, which monitors performance and automatically retrains models, while the Response Automation Module executes predefined actions like notifications and mitigations when threats are detected. Together, these modules create an end-to-end system capable of processing data from thousands of edge nodes with sub-100ms latency for critical detections while maintaining compliance with regulations like GDPR and HIPAA through its federated learning approach and robust security controls. The system's modular design allows for flexible deployment across various use cases including cybersecurity, IoT monitoring, and healthcare applications.

5. Conclusion and Future Scope

5.1 Conclusion:

The sentence is describing a part of a research paper in which the main findings and contributions are discussed. This part will identify what the research has accomplished and why it is significant. The system proposed in the paper was developed to overcome the shortcomings of existing methods for cloud anomaly detection. That is, although current methods may have some flaws—like being less accurate, having poor data privacy protection, or systems difficult for humans to understand—the proposed system tries to overcome them .One of the key concerns of the system is enhancing accuracy in anomaly detection in cloud environments. Anomalies may include suspicious activities such as cyberattacks, system crashes, or unauthorized entry. With enhanced accuracy of detection, the system minimizes false alarms and ensures real threats are detected reliably. This enhances the effectiveness and reliability of the system for organizations using cloud services. Another significant contribution is the system's focus on preserving privacy. Conventional centralized machine learning methods tend to involve exposing huge quantities of sensitive information to a master server, which poses security as well as privacy issues. To address this, the system adopts federated learning—a process by which data remains localized on various devices or nodes, and model updates (not original data) are exchanged. This strategy greatly lowers the threat of data breaches and assures data protection laws compliance .Lastly, the system adds explainable AI methods to improve interpretability. Most AI models are "black boxes" where it is hard to know why they make a particular prediction. With explainability, the system enables users, analysts, and stakeholders to view why an anomaly was detected. This fosters trust in the system and enables security teams to make more informed decisions based on the AI output.

In short, the proposed system has several contributions: it provides greater accuracy in anomaly detection, ensures data privacy via federated learning, and enhances interpretability by employing explainable AI techniques. These improvements collectively make the system a more robust and practical solution than current methods in cloud anomaly detection.

5.2 Future Scope :

Advanced Federated Learning for Non-IID and Skewed Data

A key direction for future research is exploring more sophisticated federated learning techniques to handle **non-independent and identically distributed (non-IID) and highly skewed data** across cloud environments. Traditional federated learning methods struggle with data heterogeneity among clients, leading to degraded model performance. Investigating approaches such as **personalized federated learning, federated meta-learning, or hierarchical aggregation** could improve learning in scenarios where client data distributions vary significantly. Research should also focus on optimizing communication and computation efficiency in such complex federated setups.

Integrating Diverse Explainable AI Techniques

Another promising research avenue is the integration of additional **explainable AI (XAI) techniques** beyond SHAP and LIME. Techniques like **attention mechanisms** could offer insights into which parts of the input data influence the model's decisions, particularly for sequential or structured data. Similarly, **counterfactual explanations**—which show how input changes could alter outcomes—could enhance interpretability for end-users. Combining multiple XAI methods may provide **more comprehensive, multi-faceted explanations**, helping different stakeholders (e.g., security analysts, developers, users) understand detected anomalies from various perspectives.

Evaluation Across Diverse Cloud Environments and Datasets

Future work should also prioritize evaluating the proposed anomaly detection system on **diverse, real-world cloud datasets** and across different cloud deployment models—**public, private, and hybrid clouds**. Such evaluations would test the model's robustness, scalability, and adaptability in practical settings. Performance metrics should consider not only detection accuracy but also factors like latency, scalability, and cost-efficiency across deployment scenarios. Validation in heterogeneous environments would strengthen the system's generalizability and reliability for real-world applications.

Addressing Security Threats in Federated Learning

An essential research direction involves tackling **security challenges unique to federated learning in anomaly detection**. Specifically, developing **robust defenses against model poisoning attacks, inference attacks, and information leakage** is critical to preserving the integrity and privacy of the system. Techniques such as secure aggregation, differential privacy, and adversarial training could be explored to mitigate these threats. Future work should aim to balance **privacy, security, and model performance**, ensuring that federated learning remains trustworthy in sensitive cloud environments.

Edge Computing Integration for Real-Time Anomaly Detection

Combining **edge computing with federated learning** presents another promising research frontier. By offloading anomaly detection and explanation generation closer to data sources (i.e., edge devices), systems could achieve **lower latency, reduced bandwidth usage, and improved real-time responsiveness**. Research should explore how to distribute model components effectively across cloud and edge layers, addressing challenges such as model partitioning, incremental updates, and resource constraints at the edge.

Continuous Adaptation to Evolving Cloud Threats

Finally, developing methods for **continuous monitoring and adaptation** of anomaly detection models and explanation mechanisms is crucial given the **dynamic and evolving nature of cloud infrastructures and cyber threats**. This involves creating systems capable of **online learning, concept drift detection, and automatic retraining** to maintain high detection performance over time. Future work could also explore integrating feedback loops from human analysts to iteratively refine both the detection and explanation modules.

References

[1] Privacy-Preserving Anomaly Detection Using Federated Learning and Explainable AI

Author(s) (if available). (Year). *Privacy-Preserving Anomaly Detection Using Federated Learning and Explainable AI*. [Conference/Journal Name], vol. 14, pp. [pages]. DOI (if available).

[2] Federated Learning-Based Explainable Anomaly Detection for Industrial Control Systems

Author(s) (if available). (Year). *Federated Learning-Based Explainable Anomaly Detection for Industrial Control Systems*. [Conference/Journal Name], vol. 120, pp. [pages]. DOI (if available).

[3] Comprehensive Analysis over Centralized and Federated Learning-based Anomaly Detection in Networks with Explainable AI (XAI)

Author(s) (if available). (Year). *Comprehensive Analysis over Centralized and Federated Learning-based Anomaly Detection in Networks with Explainable AI (XAI)*. [Conference/Journal Name], vol. 139, pp. [pages]. DOI (if available).

Worked sites :

Works cited

1. Anomaly Detection in Cloud Network: A Review - BIO Web of Conferences, accessed on May 5, 2025, https://www.bio-conferences.org/articles/bioconf/pdf/2024/16/bioconf_iscku2024_00019.pdf
2. AI-driven anomaly detection in cloud computing environments - International Journal of Science and Research Archive, accessed on May 5, 2025, https://ijsra.net/sites/default/files/IJSRA-2024-2184.pdf
3. A Review of Anomaly Detection Systems in Cloud Networks and Survey of Cloud Security Measures in Cloud Storage Applications - Scientific Research Publishing, accessed on May 5, 2025, https://www.scirp.org/journal/paperinformation?paperid=55903
4. What is Anomaly Detection? - AWS, accessed on May 5, 2025, https://aws.amazon.com/what-is/anomaly-detection/
5. What Is Anomaly Detection? - CrowdStrike, accessed on May 5, 2025, https://www.crowdstrike.com/en-us/cybersecurity-101/next-gen-siem/anomaly-detection/
6. What is Anomaly Detection? - Wiz, accessed on May 5, 2025, https://www.wiz.io/academy/anomaly-detection
7. What Is Anomaly Detection? | IBM, accessed on May 5, 2025, https://www.ibm.com/think/topics/anomaly-detection
8. What is Cloud Anomaly Detection? Best Practices & How It Works, accessed on May 5, 2025, https://www.velotix.ai/glossary/cloud-anomaly-detection/
9. Cloud Anomaly Detection: Guide for 2024 - Eyer.ai, accessed on May 5, 2025,

https://www.eyer.ai/blog/cloud-anomaly-detection-guide-for-2024/
10. How Anomaly Detection Helps in Cloud Threat Detection - Uptycs, accessed on May 5, 2025, https://www.uptycs.com/blog/threat-research-report-team/proactive-threat-hunting-with-anomaly-detection-in-the-cloud
11. Cloud Anomaly Detection with AI: Supply Chain Attack Example - Orca Security, accessed on May 5, 2025, https://orca.security/resources/blog/cloud-anomaly-detection-with-ai/
12. What Is Anomaly Detection? Examples, Techniques & Solutions - Splunk, accessed on May 5, 2025, https://www.splunk.com/en_us/blog/learn/anomaly-detection.html
13. An Introduction to the Federated Learning Standard - Auburn University, accessed on May 5, 2025, https://www.eng.auburn.edu/~szm0001/papers/GetMobile_FL21.pdf
14. Privacy - Preserving Anomaly Detection Using Federated Learning and Explainable AI, accessed on May 5, 2025, https://www.ijraset.com/research-paper/privacy-preserving-anomaly-detection-using-federated-learning-and-explainable-ai
15. Privacy - Preserving Anomaly Detection Using Federated Learning and Explainable AI, accessed on May 5, 2025, https://www.researchgate.net/publication/389431418_Privacy_-_Preserving_Anomaly_Detection_Using_Federated_Learning_and_Explainable_AI
16. What is Explainable AI (XAI)? - IBM, accessed on May 5, 2025, https://www.ibm.com/think/topics/explainable-ai
17. Interplay between Federated Learning and Explainable Artificial Intelligence: a Scoping Review The work in this paper was supported by the VALIDATE project grant 101057263 from the EU HORIZON-RIA. - arXiv, accessed on May 5, 2025, https://arxiv.org/html/2411.05874v2
18. Interplay between Federated Learning and Explainable Artificial Intelligence: a Scoping Review The work in this paper was supported by the VALIDATE project grant 101057263 from the EU HORIZON-RIA. - arXiv, accessed on May 5, 2025, https://arxiv.org/html/2411.05874v1
19. Anomaly Detection with Machine Learning: Techniques and Applications | DoiT, accessed on May 5, 2025, https://www.doit.com/anomaly-detection-with-machine-learning-techniques-and-applications/
20. Anomaly Detection - Infracost, accessed on May 5, 2025, https://www.infracost.io/glossary/anomaly-detection/
21. What are cloud cost anomalies? | FinOps Glossary - Zesty.co, accessed on May 5, 2025, https://zesty.co/finops-glossary/cloud-cost-anomalies/
22. Understanding Anomaly Detection - Middleware, accessed on May 5, 2025, https://middleware.io/blog/anomaly-detection/
23. Principles and Components of Federated Learning Architectures - arXiv, accessed on May 5, 2025, https://arxiv.org/html/2502.05273v2
24. Federated Learning and Data Mesh: how it enhances data architecture - www.apheris.com, accessed on May 5, 2025, https://www.apheris.com/resources/blog/federated-learning-and-data-mesh
25. Federated Learning: A Privacy-Preserving Approach to Collaborative AI Model Training, accessed on May 5, 2025, https://www.netguru.com/blog/federated-learning
26. Privacy-Preserving Federated Learning: Understanding the Costs and Benefits, accessed on May 5, 2025, https://rtau.blog.gov.uk/2024/02/22/privacy-preserving-federated-learning-understanding-the-costs-and-benefits/
27. Federated learning: Decentralised training for privacy-preserving AI - STL Partners, accessed on May 5, 2025, https://stlpartners.com/articles/edge-computing/federated-learning/

28. Applications of Federated Learning in Mobile Health: Scoping Review - PMC, accessed on May 5, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC10186185/
29. What is federated learning? - IBM Research, accessed on May 5, 2025, https://research.ibm.com/blog/what-is-federated-learning
30. Explainable AI, LIME & SHAP for Model Interpretability | Unlocking AI's Decision-Making, accessed on May 5, 2025, https://www.datacamp.com/tutorial/explainable-ai-understanding-and-trusting-machine-learning-models