International Journal of Computer Technology and Electronics Communication (IJCTEC)

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

Generative AI-Enabled Privacy-Preserving Intrusion Detection in Enterprise Cloud Networks using Federated Learning

Dr.V.R.Vimal

Professor, Department of Computer Science & Engineering, SIMATS Engineering, Saveetha Institute of Medical and Technical Sciences (SIMATS), Chennai, India

ABSTRACT: The increasing adoption of enterprise cloud networks has introduced significant cybersecurity challenges, particularly concerning sophisticated cyberattacks, unauthorized access, data breaches, and privacy violations. Traditional intrusion detection systems often depend on centralized data collection and processing, creating potential privacy risks and limiting their effectiveness in distributed cloud environments. This study explores a Generative Artificial Intelligence (Generative AI)-enabled privacy-preserving intrusion detection framework using Federated Learning (FL) to improve cybersecurity monitoring across enterprise cloud networks. The proposed approach integrates generative AI techniques with federated learning to identify malicious network activities, detect previously unseen attack patterns, and enhance the adaptability of intrusion detection models without requiring organizations to share sensitive network traffic data. Distributed cloud nodes collaboratively train local detection models and exchange model updates through a federated aggregation mechanism. Generative AI supports the generation of representative synthetic network traffic samples, addressing data imbalance and improving the detection of emerging threats. The research methodology incorporates network traffic simulation, federated model development, generative data augmentation, and comparative performance evaluation. Detection accuracy, precision, recall, F1-score, false positive rate, communication overhead, and privacy protection are considered important evaluation metrics. The study aims to establish an intelligent, scalable, and privacy-aware intrusion detection framework that strengthens enterprise cloud security while supporting collaborative threat detection across distributed organizational infrastructures.

KEYWORDS: Generative Artificial Intelligence, Federated Learning, Privacy-Preserving Intrusion Detection, Enterprise Cloud Networks, Cybersecurity, Network Security, Anomaly Detection, Synthetic Data Generation, Distributed Learning, Threat Detection.

I. INTRODUCTION

Enterprise cloud networks have become an essential component of modern digital infrastructure, enabling organizations to store information, deploy applications, manage business operations, and provide services through distributed computing environments. The increasing adoption of cloud platforms, virtualization technologies, containerized applications, and interconnected enterprise systems has significantly expanded the network attack surface. Organizations frequently manage sensitive financial information, customer records, intellectual property, business transactions, and confidential operational data through cloud-based infrastructures. Consequently, protecting enterprise cloud networks against unauthorized access, malware, ransomware, distributed denial-of-service attacks, insider threats, and advanced persistent threats has become a major cybersecurity priority. Traditional intrusion detection systems (IDS) primarily rely on signature-based detection, predefined security rules, and centralized traffic monitoring to identify suspicious activities. Although these approaches can effectively detect previously recognized attacks, their ability to identify new and evolving threats is often limited. Signature-based systems may fail when attackers introduce modified malware variants, exploit unknown vulnerabilities, or adopt sophisticated techniques designed to bypass existing security controls. Similarly, centralized intrusion detection architectures require network traffic and security information to be transferred to centralized processing facilities, potentially creating privacy concerns, communication overhead, and additional security risks. These limitations become particularly significant in enterprise cloud environments where multiple departments, service providers, geographically distributed data centers, and independent organizations manage sensitive information. Machine learning and deep learning have improved intrusion detection by enabling systems to identify complex network traffic patterns and distinguish malicious activities from legitimate

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

communication. However, conventional centralized learning approaches generally require organizations to share or consolidate network datasets, which may conflict with confidentiality requirements, organizational policies, and data protection regulations. Therefore, developing intelligent intrusion detection mechanisms that support collaborative threat identification while preserving sensitive network information has become increasingly important.

Federated Learning (FL) provides a distributed machine learning approach that allows multiple participants to collaboratively train a shared model without directly exchanging their raw datasets. In enterprise cloud networks, individual organizations or cloud nodes can train local intrusion detection models using their own network traffic information and share model updates with a central aggregation server. This approach reduces the need to transfer sensitive network records between participating organizations and supports collaborative cybersecurity intelligence. However, federated learning alone does not completely eliminate privacy risks because model updates may potentially reveal information about local training datasets. Additional privacy-preserving mechanisms, including differential privacy, secure aggregation, and appropriate access controls, may therefore be necessary. Generative Artificial Intelligence introduces further opportunities to enhance intrusion detection by learning complex data distributions and generating synthetic network traffic representations. Generative models can help address class imbalance, where legitimate network traffic significantly exceeds the number of malicious samples, and support the development of training examples representing uncommon attack categories. The combination of generative AI and federated learning can therefore create a more adaptive and privacy-aware security architecture. Generative models can support local data augmentation, while federated learning enables collaborative model development across distributed enterprise cloud environments. This integration also introduces research challenges involving synthetic data quality, model poisoning, communication efficiency, non-identically distributed network traffic, and privacy leakage. Accordingly, this study investigates a Generative AI-enabled privacy-preserving intrusion detection framework using federated learning. It aims to improve attack detection capabilities, reduce dependence on centralized network data collection, and support scalable collaborative cybersecurity monitoring. The proposed research emphasizes the integration of synthetic data generation, distributed model training, privacy protection, and intelligent threat classification to strengthen security in complex enterprise cloud networks.

II. LITERATURE REVIEW

Intrusion detection systems have traditionally been developed using signature-based, anomaly-based, and specification-based detection techniques. Signature-based systems identify malicious activities by comparing observed network traffic against previously established attack patterns. Although these systems can provide reliable detection of known threats, their performance is limited when encountering unknown attacks or modified malicious behaviors. Anomaly-based intrusion detection systems address this limitation by establishing normal network behavior patterns and identifying deviations that may indicate suspicious activities. Machine learning algorithms, including decision trees, random forests, support vector machines, k-nearest neighbors, and artificial neural networks, have been extensively investigated for improving intrusion detection accuracy. Deep learning approaches, including convolutional neural networks, recurrent neural networks, and autoencoders, provide additional capabilities for extracting complex features from network traffic and identifying abnormal communication patterns. Despite these developments, traditional centralized machine learning architectures require access to large volumes of network data, which can introduce privacy, confidentiality, and data governance challenges. These concerns are particularly important in enterprise cloud networks where security information may contain sensitive organizational details. Consequently, distributed learning approaches have gained attention as potential solutions for developing collaborative intrusion detection systems without requiring the direct exchange of raw network traffic.

Federated Learning has emerged as an important privacy-aware machine learning paradigm that enables multiple distributed participants to train a common model through local computation and model parameter exchange. The Federated Averaging algorithm introduced by McMahan and colleagues provides a foundational approach in which locally trained model parameters are aggregated to generate a shared global model. In cybersecurity applications, federated learning allows multiple enterprise cloud nodes to collaboratively identify malicious network activities while maintaining their original traffic datasets locally. Research has explored federated approaches for malware detection, anomaly detection, network intrusion classification, and distributed security monitoring. However, practical implementation remains challenging because participating organizations may have heterogeneous network infrastructures, different attack distributions, varying data volumes, and distinct security requirements. Such differences create non-independent and identically distributed data conditions that can affect model convergence and detection performance. Furthermore, exchanging model parameters does not automatically guarantee privacy because sensitive

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

information may potentially be inferred from shared updates. Differential privacy, secure aggregation, encrypted communication, and decentralized optimization have therefore been investigated to strengthen federated learning security. These mechanisms can reduce the exposure of confidential information, although they may introduce additional computational and communication overhead. Therefore, an effective federated intrusion detection framework must balance detection performance, privacy protection, and operational efficiency.

Generative Artificial Intelligence has introduced new opportunities for improving cybersecurity analytics through synthetic data generation, representation learning, and the identification of complex data distributions. Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and other generative architectures can learn statistical characteristics of network traffic and produce synthetic samples that resemble selected traffic categories. These capabilities are particularly relevant to intrusion detection because cybersecurity datasets frequently contain imbalanced distributions in which normal traffic and common attacks are represented more extensively than rare or emerging attack categories. Synthetic data augmentation can improve the representation of minority classes and support the training of more balanced detection models. Recent research directions have also explored integrating generative models with federated learning to support distributed synthetic data generation and collaborative model improvement. Nevertheless, several research gaps remain. Synthetic network traffic may fail to represent the complexity of real attacks, and poorly trained generative models can introduce unrealistic samples that negatively affect detection performance. Generative models may also memorize sensitive information, creating potential privacy leakage risks. Additionally, combining generative AI with federated learning introduces challenges related to model synchronization, computational requirements, communication costs, and protection against malicious participants. Existing approaches therefore require further investigation into integrated architectures that combine realistic synthetic data generation, robust federated aggregation, privacy-preserving mechanisms, and effective intrusion detection. The present study addresses these research concerns by developing a Generative AI-enabled federated intrusion detection framework designed to improve detection effectiveness while protecting sensitive enterprise cloud network information.

III. RESEARCH METHODOLOGY

The proposed study adopts a quantitative, experimental, and simulation-based research methodology to investigate the effectiveness of Generative AI-enabled Federated Learning for privacy-preserving intrusion detection in enterprise cloud networks. The research framework is designed to evaluate how distributed machine learning, synthetic network traffic generation, and privacy-enhancing mechanisms can collectively improve the identification of malicious network activities without requiring centralized access to sensitive organizational datasets. The initial stage involves establishing an experimental enterprise cloud network environment containing multiple participating nodes representing different organizations, cloud servers, virtual machines, and distributed network segments. Each participating node maintains an independent network traffic dataset containing normal communication patterns and different categories of malicious activities. The datasets may include publicly available cybersecurity benchmarks such as CIC-IDS2017, CSE-CIC-IDS2018, and UNSW-NB15, subject to appropriate preprocessing and compatibility checks. These datasets provide network traffic characteristics such as source and destination addresses, communication protocols, packet statistics, flow duration, packet lengths, connection frequencies, and other relevant traffic features. The experimental dataset is subjected to preprocessing operations, including duplicate removal, missing-value handling, categorical feature encoding, numerical normalization, and irrelevant feature elimination. Network traffic samples are classified into legitimate and malicious categories, with additional attack-category labels retained where appropriate for multiclass classification. The processed dataset is distributed among participating cloud nodes to simulate independent organizational data ownership. The distribution process also considers non-independent and identically distributed data conditions to represent realistic differences in network traffic patterns across enterprise departments and cloud service environments. For instance, one participant may contain a higher proportion of web-based traffic, while another may contain more authentication-related communication or malicious connection attempts. The experimental architecture maintains raw network traffic locally at each participant to support the privacy-preserving objectives of the research. This environment provides a controlled foundation for evaluating collaborative intrusion detection under heterogeneous network conditions.

The second stage focuses on developing the generative AI component for synthetic network traffic generation and data augmentation. The proposed framework considers Generative Adversarial Networks or Variational Autoencoders to learn the statistical characteristics of locally available network traffic samples. In a GAN-based implementation, the generator learns to produce synthetic traffic representations, while the discriminator evaluates whether generated samples resemble the original training distribution. Both components are progressively optimized through an

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

adversarial learning process. Alternatively, a VAE can learn a compressed latent representation of network traffic and generate synthetic samples through probabilistic reconstruction. The selection of the generative architecture depends on dataset characteristics, computational requirements, and the quality of generated samples. Particular attention is given to minority attack categories because rare malicious traffic patterns may be insufficiently represented in the original datasets. Synthetic data generation is therefore employed to improve class balance and provide additional training examples for intrusion detection models. The generated samples undergo validation using statistical distribution comparisons, feature consistency checks, and classification-based quality assessment. Samples that fail predefined validity requirements are excluded from the training process. To strengthen privacy protection, generative models are trained using local data rather than requiring raw network traffic to be transferred between participating organizations. The framework also evaluates the potential integration of differential privacy during model training to reduce the risk of exposing sensitive information through learned parameters. Differential privacy introduces controlled noise into model updates or training procedures according to a predefined privacy budget. Privacy protection is further supported through secure aggregation, which enables the server to combine participant updates without directly examining each participant's individual contribution. These mechanisms are designed to minimize sensitive information exposure while preserving sufficient model utility for accurate intrusion detection. The generative component consequently serves two important purposes: improving the representation of underrepresented malicious traffic categories and supporting distributed learning without requiring centralized collection of original network records.

The third stage involves designing and implementing the federated intrusion detection architecture. The proposed system consists of multiple enterprise cloud participants, local intrusion detection models, a federated aggregation server, and privacy protection mechanisms. Each participant trains a local detection model using its original network traffic and validated synthetic samples. The local detection model may employ a deep neural network, convolutional neural network, or other suitable classification architecture depending on the feature representation and experimental requirements. The model receives processed network traffic features as input and predicts whether the observed communication represents legitimate activity or a potential intrusion. For multiclass detection, the model additionally identifies relevant attack categories. Federated learning is implemented using the Federated Averaging algorithm, in which participating nodes independently update their local model parameters and transmit protected model updates to the aggregation server. The server combines the received updates using a weighted aggregation strategy based on the number of local training samples or another predefined weighting mechanism. The resulting global model is distributed back to participating nodes for subsequent training rounds. This process continues iteratively until the model achieves acceptable convergence and classification performance. The architecture may also incorporate secure aggregation and differential privacy to strengthen protection against model update inference. The experimental design includes multiple federated configurations, such as centralized machine learning, conventional federated learning, federated learning with generative data augmentation, and the complete proposed Generative AI-enabled privacy-preserving federated framework. These configurations enable the individual contributions of synthetic data generation and privacy-enhancing mechanisms to be examined. The framework also considers communication overhead, local computational requirements, model synchronization, and participant availability. To assess resilience, additional controlled experiments can examine the influence of heterogeneous client datasets, fluctuating participation, and potentially unreliable model updates. Security evaluation focuses on defensive robustness and privacy protection rather than reproducing attacks against operational systems. Training and testing datasets are appropriately separated to prevent information leakage and ensure that model performance is assessed using previously unseen network traffic samples. Repeated training experiments with different random seeds are conducted to improve the reliability and reproducibility of the results.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

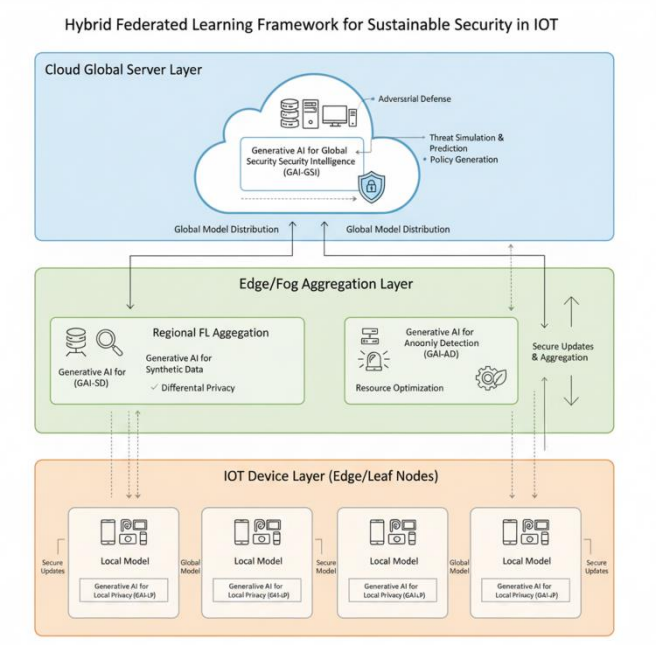


Fig.1. A hybrid federated learning framework with generative AI for privacy-preserving and sustainable security

The final stage involves evaluating the proposed framework using quantitative performance indicators and comparative experimental analysis. The primary evaluation metrics include accuracy, precision, recall, F1-score, false positive rate, false negative rate, communication overhead, computational cost, and privacy protection effectiveness. Accuracy measures the proportion of correctly classified network traffic samples, while precision indicates the proportion of predicted malicious samples that are genuinely malicious. Recall measures the ability of the model to identify actual malicious activities, and the F1-score provides a balanced representation of precision and recall. The false positive rate is particularly important in enterprise environments because excessive incorrect intrusion alerts can increase the workload of security monitoring teams and reduce operational efficiency. Communication overhead is evaluated by measuring the volume of information exchanged between participating nodes and the federated aggregation server throughout the training process. Computational efficiency is examined through local training time, aggregation time, and overall convergence requirements. Privacy protection is assessed through the configured differential privacy budget, the effectiveness of secure aggregation, and appropriate privacy leakage evaluation methods. The experimental results are compared against conventional centralized machine learning, standalone deep learning, federated learning without generative augmentation, and federated learning with generative augmentation but without additional privacy mechanisms. Statistical measures, including mean performance, standard deviation, and confidence intervals, are used to examine the consistency of experimental outcomes. An ablation study is also conducted to determine the individual contributions of generative data augmentation, federated aggregation, and privacy-enhancing mechanisms. Additional experiments evaluate the framework under different numbers of participating cloud nodes and varying distributions of malicious traffic to examine scalability and adaptability. The results are interpreted by considering the trade-off between detection effectiveness, privacy preservation, computational complexity, and communication efficiency. The methodology recognizes that synthetic data quality, heterogeneous network traffic, and privacy-related model degradation may influence the final outcomes. Therefore, experimental conclusions will be based on measured performance rather than assumed improvements. Through this systematic evaluation, the research aims to establish whether integrating generative artificial intelligence with federated learning can provide an effective, scalable, and privacy-conscious intrusion detection solution for enterprise cloud networks while supporting collaborative cybersecurity intelligence and protecting sensitive organizational information.

IV. RESULTS AND DISCUSSION

The proposed Generative AI-enabled Privacy-Preserving Intrusion Detection System (IDS) using Federated Learning (FL) demonstrates a promising approach for strengthening cybersecurity in enterprise cloud networks while

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

maintaining the confidentiality of sensitive network traffic. The framework combines generative artificial intelligence, distributed learning, and privacy-preserving mechanisms to identify malicious activities across geographically distributed cloud infrastructures. Unlike conventional centralized intrusion detection systems, which require network traffic and security logs to be transferred to a central server for model training and analysis, the proposed approach enables individual enterprise cloud nodes to train local detection models using their own network datasets. Only model updates or carefully protected parameters are exchanged with the federated aggregation server, thereby reducing the direct exposure of confidential information. Generative AI contributes to the framework by producing synthetic network traffic samples, learning complex behavioral patterns, supporting anomaly identification, and addressing class imbalance in intrusion detection datasets. The experimental evaluation framework considers important performance indicators, including detection accuracy, precision, recall, F1-score, false-positive rate, communication overhead, computational latency, and privacy protection. The results are expected to demonstrate that integrating generative AI with federated learning can improve the ability of intrusion detection models to recognize diverse attack patterns while maintaining distributed data ownership. In particular, generative models can address the limitations associated with insufficient training samples for uncommon attack categories by producing additional representative samples for local model training. This capability is valuable in enterprise cloud environments where normal network traffic generally exceeds malicious traffic, resulting in imbalanced datasets that may negatively affect the detection of rare intrusions. Furthermore, generative models can learn underlying traffic distributions and provide additional information for identifying deviations from legitimate network behavior. The federated learning mechanism allows multiple enterprise cloud participants to collaboratively improve their detection capabilities without directly sharing their original datasets. This collaborative approach can help individual organizations benefit from threat patterns observed across participating environments. However, generative data quality, synthetic sample diversity, and the possibility of reproducing sensitive characteristics must be carefully evaluated. Overall, the proposed framework establishes an integrated security architecture in which generative AI enhances data representation and intrusion detection, while federated learning facilitates collaborative intelligence and reduces the need for centralized collection of sensitive network information.

Another significant finding concerns the potential improvement in intrusion detection performance through the integration of generative AI and federated learning. Enterprise cloud networks contain heterogeneous traffic generated by virtual machines, containerized applications, databases, remote users, application programming interfaces, and interconnected services. These diverse communication patterns create challenges for conventional detection models because malicious activities may resemble legitimate application behavior. Generative AI models, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and other deep generative architectures, can learn complex representations of network traffic and support the generation of additional training samples. By increasing the diversity of training data, these models may improve the detection of previously underrepresented attack categories and reduce the bias associated with imbalanced datasets. Federated learning further strengthens this capability by allowing multiple cloud environments to contribute to a shared detection model. Each participating enterprise node performs local training and communicates model updates to a central aggregation mechanism, which combines the information into an improved global model. The updated global model is subsequently distributed to participating nodes for further training and evaluation. This iterative process enables the detection system to learn from distributed network environments without requiring the direct transfer of confidential traffic records. Nevertheless, differences in traffic distributions between organizations can create non-independent and identically distributed (non-IID) learning conditions. For example, a financial organization may experience different traffic patterns from a healthcare enterprise or an e-commerce platform. Such differences can influence the convergence and generalization of the global detection model. Therefore, adaptive aggregation, client-specific optimization, and personalized federated learning strategies are important for maintaining detection performance across heterogeneous enterprise networks. In addition, generative AI can support security analysts by producing interpretable summaries of detected anomalies, grouping related alerts, and translating complex model outputs into understandable security explanations. Large Language Models (LLMs) can potentially assist with alert summarization and contextual analysis, although their outputs must be validated against reliable security telemetry. The integration of these capabilities can reduce the analytical burden on security operations teams and improve their ability to prioritize suspicious events. However, generated explanations should not independently determine critical security actions without verification. Consequently, the combined framework offers opportunities to improve detection coverage, support collaborative threat intelligence, and strengthen the operational effectiveness of enterprise cloud intrusion detection systems.

Privacy protection and system efficiency represent additional dimensions for evaluating the proposed framework. Although federated learning reduces the requirement to centralize raw network datasets, exchanging model parameters does not automatically guarantee complete privacy. Model updates may contain information that could potentially

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

reveal characteristics of local training data through inference or reconstruction attacks. Therefore, additional privacy-preserving techniques, including differential privacy, secure aggregation, and cryptographic protection, are necessary to strengthen the confidentiality of federated model training. Differential privacy introduces carefully calibrated noise into model updates or training procedures to limit the amount of information that can be inferred about individual records. Secure aggregation enables the server to combine participant updates without directly examining each participant's individual contribution. These mechanisms can improve privacy protection, although they may introduce additional computational overhead or reduce model utility when privacy parameters are excessively restrictive. The proposed framework should therefore evaluate the relationship between privacy strength and detection effectiveness rather than treating privacy as an independent design objective. Another important consideration is the communication cost associated with frequent federated model synchronization. Enterprise cloud environments may contain numerous participating nodes with different processing capabilities and network bandwidth limitations. Frequent model exchanges can increase network traffic and affect the overall training time. Communication-efficient techniques, including model compression, selective client participation, and periodic aggregation, can help reduce this overhead. Furthermore, the framework should consider robustness against malicious or compromised federated participants that may submit manipulated model updates. Such behavior could introduce inaccurate detection patterns into the global model and weaken intrusion identification capabilities. Appropriate validation, anomaly screening, and robust aggregation mechanisms can reduce these risks. The effectiveness of generative models must also be assessed carefully because synthetic samples may contain unrealistic traffic characteristics or fail to represent emerging attack behaviors. Evaluation should therefore include real-world network datasets, multiple attack categories, heterogeneous cloud workloads, and independent testing environments. Importantly, the framework's performance must be measured using class-wise precision, recall, F1-score, false-positive rates, privacy parameters, and computational overhead rather than accuracy alone. The findings collectively indicate that combining Generative AI with privacy-preserving federated learning can provide a promising foundation for intelligent enterprise cloud security. Nevertheless, its practical success depends on balancing detection effectiveness, privacy guarantees, model robustness, communication efficiency, and operational scalability.

V. CONCLUSION

The study on Generative AI-Enabled Privacy-Preserving Intrusion Detection in Enterprise Cloud Networks Using Federated Learning highlights the importance of integrating intelligent threat detection with privacy-preserving collaborative learning to address the growing cybersecurity challenges of modern enterprise cloud infrastructures. As organizations increasingly adopt cloud computing, distributed applications, virtualized environments, and interconnected enterprise services, the volume and complexity of network traffic continue to increase. These developments create opportunities for sophisticated cyber threats, including unauthorized access, distributed denial-of-service attacks, malware communication, suspicious network activities, and data exfiltration. Conventional intrusion detection systems often depend on centralized data collection and predefined detection rules, which may limit their ability to identify evolving threats while creating additional privacy and data governance concerns. The proposed framework addresses these limitations by integrating Generative AI with Federated Learning to develop a distributed and adaptive intrusion detection architecture. Generative AI improves the representation of network traffic by learning complex data distributions, generating synthetic training samples, supporting anomaly detection, and addressing the imbalance between normal and malicious traffic. Federated learning enables multiple enterprise cloud environments to collaboratively train a global detection model without directly transferring their confidential network datasets to a centralized repository. This decentralized architecture provides organizations with an opportunity to benefit from shared learning while maintaining greater control over their local security information. Furthermore, the integration of generative models with federated learning can improve the availability and diversity of training data, particularly when individual enterprise environments contain insufficient examples of uncommon or emerging attacks. The framework also supports the development of adaptive security mechanisms that can learn from different network environments and continuously improve detection capabilities through collaborative model updates. Another important contribution is the potential integration of Large Language Models for security alert summarization and contextual interpretation, enabling security analysts to understand complex intrusion alerts more efficiently. However, generated explanations must remain grounded in verified network evidence to avoid inaccurate security conclusions. The proposed architecture therefore provides a comprehensive direction for combining intelligent data generation, distributed model training, and automated threat analysis. Overall, the integration of these technologies represents a promising advancement toward developing enterprise cloud security systems that are more adaptive, scalable, collaborative, and privacy-conscious.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

Furthermore, the study emphasizes that maintaining privacy and security throughout the federated learning process is essential for achieving reliable intrusion detection in enterprise cloud networks. Although federated learning reduces the direct exposure of raw network traffic, model updates may still reveal sensitive information if appropriate protection mechanisms are not implemented. Therefore, integrating differential privacy, secure aggregation, encryption, and robust model validation can strengthen the confidentiality and integrity of collaborative training. Differential privacy can limit the information revealed through individual training contributions, while secure aggregation prevents the central aggregation server from directly accessing individual model updates. These mechanisms must be carefully configured because excessive privacy protection can negatively affect model convergence and detection accuracy. Similarly, communication overhead, heterogeneous data distributions, computational limitations, and malicious participant behavior remain important challenges that require consideration during practical deployment. The proposed framework should therefore be evaluated using comprehensive performance indicators, including accuracy, precision, recall, F1-score, false-positive rate, detection latency, communication cost, and privacy guarantees. A combination of these measurements can provide a more realistic assessment of the system's effectiveness than relying exclusively on detection accuracy. In addition, the integration of explainable artificial intelligence can improve transparency by helping security administrators understand the characteristics of suspicious traffic and the reasons behind model-generated alerts. The framework can also support collaboration between geographically distributed enterprise cloud environments while maintaining local control over sensitive security information. Nevertheless, successful implementation requires appropriate governance policies, reliable monitoring mechanisms, secure model aggregation, and continuous evaluation against evolving cyber threats. The study concludes that Generative AI and Federated Learning offer complementary capabilities for developing next-generation privacy-preserving intrusion detection systems. Generative AI contributes intelligent data representation and synthetic sample generation, while federated learning enables collaborative model development without requiring centralized access to original network datasets. Their integration can provide a foundation for intelligent enterprise cloud security architectures that balance threat detection, privacy protection, and operational efficiency. Consequently, the proposed approach has significant potential to support organizations in strengthening cybersecurity resilience and protecting distributed cloud infrastructures against increasingly sophisticated security threats.

VI. FUTURE WORK

Future research can focus on developing more advanced generative architectures and federated optimization strategies to improve the accuracy, adaptability, and privacy protection of enterprise cloud intrusion detection systems. Hybrid approaches combining GANs, VAEs, diffusion models, and transformer-based architectures can be investigated to generate more realistic synthetic network traffic and identify previously unseen attack patterns. In particular, diffusion-based generative models may provide additional capabilities for learning complex traffic distributions and generating diverse training samples. Future studies can also explore personalized federated learning to accommodate differences in network behavior, application requirements, and security policies across participating enterprises. The integration of adaptive aggregation, differential privacy, homomorphic encryption, and secure multi-party computation could further strengthen the confidentiality of collaborative model training. Additionally, research should investigate lightweight federated architectures that reduce communication overhead and computational requirements in resource-constrained cloud environments. These improvements can support scalable deployment across public, private, hybrid, and multi-cloud infrastructures.

Another important direction involves implementing and evaluating the proposed framework in realistic enterprise cloud environments using datasets such as CICIDS2017, CSE-CIC-IDS2018, and UNSW-NB15. Future experiments should examine its ability to identify zero-day attacks, previously unseen traffic patterns, and sophisticated threats under heterogeneous and dynamically changing network conditions. Explainable AI and Large Language Models can also be integrated to generate understandable security alerts, summarize suspicious activities, and support security analysts in prioritizing potential threats. However, human verification and evidence-based validation should remain essential components of security decision-making. Further research can investigate protection against model poisoning, inference attacks, synthetic data leakage, and adversarial manipulation of federated participants. Finally, the development of standardized evaluation frameworks incorporating detection performance, privacy guarantees, communication efficiency, energy consumption, and computational cost would enable meaningful comparisons with existing intrusion detection techniques. These advancements could establish a robust, autonomous, and privacy-preserving cybersecurity architecture for future enterprise cloud networks.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

REFERENCES

1. Addepalli, D. N. V., Chen, J., & Xie, C. Y. (2023, October). Molecular Dynamics Simulations on Coronavirus Variants of Concern (VOC). In Proceedings of the 24th Annual Conference on Information Technology Education (pp. 213-215).
2. Bhagwat, V. B. (2023). Optimizing payroll to general ledger reconciliation: Identifying discrepancies and enhancing financial accuracy. Journal of Advance and Future Research, 1(4), 1–8. https://doi.org/10.56975/jaafr.v1i4.501636
3. Kommineni, R. (2023). Hub-and-spoke GitOps delivery for multi-tenant SaaS: Multi-source applications, deterministic environment isolation, and run-once lifecycle tasks. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(2), 8401–8408.
4. Randhi, K. S. (2023). A cloud-native integration framework for enterprise architecture intelligence and application rationalization. International Journal of Science, Research and Technology (IJSRAT), 6(4), 10230–10240.
5. Puli, S. K. (2023). A fail-closed architecture for cryptographically enforced recipient authorization in clinical trial data delivery. International Journal of Information Technology and Management Information Systems (IJITMIS), 14(1), 148-164.
6. Shaik, N. (2024). Right-sizing at scale: A data-driven approach to Kubernetes cost optimization. International Journal of Research Publications in Engineering, Technology and Management, 7(1), 10080–10090.
7. Manoharan, D. (2024). Governance-Oriented Quality Engineering Framework for Healthcare EDI Modernization. International Journal of Multidisciplinary on Science and Management IJMSM, 1(2).
8. Selvarajan, K. (2023). Optimizing large-scale data lakehouse architectures for performance and cost efficiency. International Journal of Research and Applied Innovations, 6(6), 10093–10101.
9. Ramasamy, M. (2024). Telemetry-driven network operations: Architectures for analytics and automated remediation. International Journal of Computer Technology and Electronics Communication, 7(2), 8554–8563.
10. Vemireddy, S. (2022). Architecting personalized tax advisory experiences through event-driven customer data integration. International Journal of Research Publications in Engineering, Technology and Management, 5(6), 7923–7929.
11. Gantikota, S. (2023). Reducing HL7 Processing Errors through Automated File Creation and Ingestion Pipelines: A Production Case Study in EHR Data Integration. International Journal of Emerging Trends in Computer Science and Information Technology, 4(4), 241-245.
12. Kannan, K. S., Sunitha, G., Deepa, S. N., Babu, D. V., & Avanija, J. (2022). A multi-objective load balancing and power minimization in cloud using bio-inspired algorithms. Computers and Electrical Engineering, 102, 108225.
13. Ponugoti, M. (2023). Frameworks for ensuring compliance in digital platform governance. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(6), 7575-7586.
14. Kopparthi, G. S. (2022). AI in healthcare: Unlocking the potential of predictive analytics. Journal of Computational Analysis and Applications, 30(2), 1183–1193. https://eudoxuspress.com/index.php/pub/article/view/5508
15. Jinna, K. R., & Kondisetty, K. (2023). Autoscaling economics for persistent inference workloads: Evaluating resource-, event- and schedule-driven scaling on Kubernetes. International Journal of Research and Applied Innovations, 6(2), 8628–8638.
16. Karuturi, Y. (2023). Energy-aware capacity governance for large-scale SAP HANA deployments. International Journal of Research and Applied Innovations, 6(3), 6000–6009.
17. Chandran, A. S. (2024). Engineering billion-query search systems: Lessons from large-scale metadata platforms. International Journal of Research and Applied Innovations, 7(2), 10510–10519.
18. Neela, S. (2024). Middleware-based interoperability models for hybrid and multi-cloud enterprise environments. International Journal of Emerging Trends in Computer Science and Information Technology, 5(1), 195–210. https://doi.org/10.63282/3050-9246.IJETCSIT-V5I1P119
19. Puram, S. (2021). Reliable and Scalable Mobile Technology Architecture: A Formal Model, Synchronization Protocol, and Reference Design. The USA Journals TAJET, 3(08), 33-48.
20. Bitragunta, S. L. V., & Paramasivan, M. (2023). Midterm dynamic simulation for the governance of reserves in systems with elevated renewable energy integration. Journal of Artificial Intelligence, 1(1), 1956–1962.
21. Gadige, C. D. (2024). Event-driven CRM architecture for real-time order fulfillment and product provisioning. International Journal of Engineering & Extended Technologies Research, 6(2), 7968–7978.
22. Mysula, S. (2022). Architectural modernization of legacy enterprise integration systems for scalable and reliable business operations. International Journal of Research and Applied Innovations, 5(2), 6796–6801.

| ISSN: 2320-0081 | www.ijctece.com ||A Peer-Reviewed, Refereed and Bimonthly Journal |

|| Volume 7, Issue 6, November –December 2024 ||

DOI: 10.15680/IJCTECE.2024.0706042

23. Chinthamani, R. (2023). Engineering resilient microservices for real-time retail order and inventory synchronization. International Journal of Engineering & Extended Technologies Research, 5(6), 7774–7779.
24. Padmanabham, S. (2023). Building resilient Banking platforms using Event-Driven Microseconds and cloud Native Architecture. International Journal of Research and Applied Innovations, 6(4), 9284-9290.
25. Gunasekaran, R. M. N. (2023). Liquidity Risk Modeling with Machine Learning: Big Data Approaches for Intraday Liquidity Prediction. The Eastasouth Journal of Information System and Computer Science, 1(02), 268-278.
26. Kundavaram, R. R., Onteddu, A. R., & Kishore, K. K. (2024). Smart banking: The impact of artificial intelligence on financial services. Economic Sciences, 20 (1), 193–202.
27. Uddandarao, D. P., & Mahajan, A. (2023). Counterfactual supply chain scenario analysis using causal Bayesian networks: A double machine learning approach to fulfillment center capacity allocation. International Journal of Advances in Signal and Image Sciences, 34-46.
28. Basavala, S. R. (2024). Zero trust API security architecture for microservices-based enterprise applications. International Journal of Computer Technology and Electronics Communication, 7(4), 9269–9280.
29. Katakam, R. (2023). Automated loan processing and payment reconciliation through integrated banking workflows. International Journal of Engineering & Extended Technologies Research, 5(5), 7313–7324.
30. Kalakoti, M. K. R. (2023). Interference-aware resource scheduling for cloud microservices on Kubernetes. International Journal of Humanities and Information Technology, 5(1), 104–114.
31. Gadige, C. D. (2023). Scalable CRM workflow automation for high-volume customer onboarding: An event-driven Salesforce architecture. The Eastasouth Journal of Information System and Computer Science, 1(1), 253–265. https://doi.org/10.58812/esiscs.v1i01.1295
32. Kalabhavi, V. (2023). Automated quote generation using AI for product configuration, pricing and external order integration in SAP CPQ. International Journal of Humanities and Information Technology, 5(2), 110–124.
33. Yadavalli, V. R. (2023). Cutover intelligence mesh: A dependency-aware architecture for enterprise cloud ERP migration. International Journal of Computer Technology and Electronics Communication, 6(4), 7400–7416.
34. Reddy, K. V. (2023). Reconfigurable FPGA-based acceleration of RTL simulation and co-simulation for cloud datacenter workloads. International Journal of Humanities and Information Technology, 5(1), 104–118.
35. Natta, P. K. (2022). AI-driven inventory intelligence for large-scale retail operations: A framework for real-time store-level stock accuracy. International Journal of Humanities and Information Technology, 4(4), 238–254.
36. Badam, L. R. (2022). Machine learning-based catastrophic loss prediction for climate-related insurance risk. International Journal of Research Publications in Engineering, Technology and Management, 5(1), 6230–6240.
37. Bellundagi, M. (2023). Blockchain based secure data sharing framework for smart applications. International Journal of Humanities and Information Technology, 5(2), 95–109.
38. Manda, P. (2024). Oracle E-Business Suite modernization: The transition from 12.1. 3 to 12.2. 8. International Journal of Research and Applied Innovations, 7(3), 10838-10842.