International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

Privacy-Aware Enterprise Cloud Intrusion Detection Through Generative AI and Federated Learning Architectures

Dr. VUDASREENIVASARAO

School of Computing And Electrical Engineering, Bahir Dar University, Ethiopia

ABSTRACT: The increasing adoption of enterprise cloud computing has introduced significant cybersecurity challenges, particularly concerning sophisticated intrusion attempts, data privacy, and distributed threat detection. Traditional intrusion detection systems often depend on centralized data collection, which creates privacy risks, communication overhead, and limitations in identifying emerging cyber threats across heterogeneous cloud environments. This research proposes a privacy-aware enterprise cloud intrusion detection architecture integrating Generative Artificial Intelligence (GenAI) and Federated Learning (FL) to improve threat identification while maintaining data confidentiality. The proposed framework enables multiple enterprise cloud environments to collaboratively train intelligent intrusion detection models without exchanging sensitive raw network traffic or organizational information. Federated learning supports decentralized model development, while generative AI enhances threat analysis through synthetic security data generation, anomaly interpretation, and contextual threat classification. Privacy-enhancing mechanisms, including secure aggregation, differential privacy, and encrypted communication, are incorporated to strengthen information protection. The research methodology adopts a conceptual and experimental evaluation approach using representative network intrusion datasets and simulated distributed enterprise environments. Model performance is assessed through detection accuracy, precision, recall, F1-score, false-positive rate, communication overhead, and privacy considerations. The proposed architecture aims to establish a scalable, adaptive, and privacy-preserving cybersecurity framework capable of supporting intelligent intrusion detection across modern enterprise cloud infrastructures.

KEYWORDS: Enterprise Cloud Security, Intrusion Detection Systems (IDS), Generative Artificial Intelligence, Federated Learning, Privacy-Preserving Machine Learning, Network Anomaly Detection, Differential Privacy, Secure Aggregation, Cloud Cybersecurity, Distributed Threat Intelligence

I. INTRODUCTION

Enterprise cloud computing has become an essential component of modern organizational infrastructure because it provides scalable storage, flexible computing resources, application deployment capabilities, and cost-effective information management. Organizations increasingly depend on public, private, and hybrid cloud environments to manage business applications, customer information, financial transactions, and enterprise communication systems. However, the expansion of cloud infrastructure has significantly increased the cybersecurity attack surface, creating opportunities for unauthorized access, malicious activities, data breaches, ransomware, distributed denial-of-service attacks, and sophisticated network intrusions. Traditional security mechanisms, including signature-based intrusion detection systems and conventional firewall technologies, frequently experience difficulties identifying previously unknown attacks and rapidly evolving threat patterns. Consequently, intelligent intrusion detection approaches have become increasingly important for strengthening enterprise cloud security and improving organizational resilience against emerging cybersecurity threats.

Artificial Intelligence (AI) and machine learning have introduced advanced capabilities for identifying abnormal network activities and classifying potentially malicious behavior. Generative Artificial Intelligence further expands these capabilities by supporting synthetic security data generation, contextual threat analysis, and intelligent interpretation of complex security events. Nevertheless, conventional AI-based intrusion detection systems frequently require centralized training datasets containing sensitive organizational information, potentially creating privacy and regulatory compliance concerns. Federated Learning offers an alternative approach by enabling multiple cloud environments to collaboratively train machine learning models while maintaining local control over their datasets. Instead of transferring raw network traffic to a centralized repository, participating enterprise environments exchange

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

model updates through a coordinated learning process. Combining federated learning with generative AI creates opportunities for developing intelligent, adaptive, and privacy-aware intrusion detection architectures. Furthermore, integrating privacy-enhancing mechanisms such as differential privacy and secure aggregation can reduce the exposure of confidential information during collaborative model training.

This research focuses on developing a privacy-aware enterprise cloud intrusion detection framework that integrates generative AI and federated learning architectures. The proposed approach aims to address the limitations of centralized security intelligence by enabling collaborative threat detection across geographically distributed cloud infrastructures. It emphasizes the relationship between detection effectiveness, information confidentiality, computational efficiency, and system scalability. The research also investigates how generative models can support detection of previously unseen network behaviors without compromising organizational privacy. Through an integrated architecture, the study seeks to contribute to the development of reliable cloud cybersecurity solutions capable of addressing the security requirements of contemporary digital enterprises.

II. LITERATURE REVIEW

The development of intelligent intrusion detection systems has attracted considerable research attention because traditional cybersecurity approaches face increasing difficulties in identifying complex and previously unknown attacks. Conventional intrusion detection mechanisms primarily depend on predefined attack signatures, rule-based monitoring, and centralized network traffic analysis. Although these approaches provide effective protection against recognized threats, they frequently demonstrate limitations when dealing with zero-day attacks, evolving malicious behaviors, and heterogeneous enterprise cloud environments. Machine learning and deep learning techniques have consequently been introduced to improve network anomaly detection, malicious traffic classification, and automated security monitoring. Hernandez-Ramos et al. (2023) examined federated learning-based intrusion detection approaches and highlighted the potential of collaborative model training without transferring raw information between participating organizations. Similarly, Zhang et al. (2025) investigated federated intrusion detection techniques and identified challenges associated with communication overhead, heterogeneous datasets, distributed privacy requirements, and model aggregation. Makris et al. (2025) provided a comprehensive analysis of federated intrusion detection systems, emphasizing the importance of aggregation strategies and deployment considerations. These studies establish federated learning as a promising alternative to centralized security intelligence because enterprise organizations can collaboratively improve intrusion detection capabilities while maintaining local ownership of sensitive network information. However, the effectiveness of federated learning depends on appropriate model coordination, dataset distribution, communication efficiency, and protection against malicious model updates.

Recent developments in generative artificial intelligence have created additional opportunities for improving cybersecurity analytics and intrusion detection performance. Generative models, including Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), diffusion models, and Large Language Models (LLMs), can support synthetic security data generation, anomaly identification, data augmentation, and contextual interpretation of suspicious activities. These capabilities are particularly valuable when enterprise security datasets contain insufficient examples of rare or emerging attack categories. Nevertheless, generative AI introduces challenges concerning synthetic data reliability, privacy leakage, model hallucinations, and computational requirements. Bunko et al. (2026) highlighted that federated intrusion detection systems require additional privacy mechanisms because keeping raw data locally does not completely eliminate risks associated with gradient leakage, model inversion, and information inference. Their research emphasizes the importance of secure aggregation, differential privacy, and cryptographic protection. Liu et al. (2026) examined the integration of generative AI and federated learning for intrusion detection, identifying opportunities in synthetic traffic generation, anomaly detection, and intelligent alert explanation. The existing literature therefore indicates that federated learning and generative AI offer complementary advantages, but their combined application in privacy-aware enterprise cloud environments requires further investigation. The present research addresses this gap by proposing an integrated architecture that combines distributed learning, generative threat analysis, and privacy-enhancing mechanisms within a unified enterprise cloud intrusion detection framework.

III. RESEARCH METHODOLOGY

Research Design and Proposed Architectural Framework

The research adopts a design-oriented, quantitative, and experimental methodology to investigate the integration of generative artificial intelligence and federated learning for privacy-aware enterprise cloud intrusion detection. The

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

primary objective is to develop and evaluate an intelligent security architecture capable of identifying malicious network activities across distributed enterprise cloud infrastructures without requiring the centralized collection of sensitive network traffic. The proposed research framework consists of five interconnected architectural layers: enterprise cloud data collection, local preprocessing and feature extraction, federated learning coordination, generative AI-based security intelligence, and privacy-aware threat monitoring. The data collection layer represents multiple participating enterprise environments, including public cloud platforms, private cloud infrastructures, hybrid cloud systems, and virtualized enterprise networks. Each participating organization functions as an independent federated client that maintains ownership and control of its local security information. The preprocessing layer converts locally available network records into structured numerical representations suitable for machine learning analysis. The federated learning layer coordinates distributed model training through a central aggregation server without collecting the original training records. The generative AI layer supports synthetic security data generation, suspicious activity interpretation, and contextual threat classification. Finally, the monitoring layer provides security alerts, anomaly scores, and detection reports to enterprise security administrators. The research follows a comparative evaluation strategy involving conventional centralized machine learning, local-only learning, conventional federated learning, and the proposed privacy-aware generative federated learning architecture. These approaches will be compared under consistent experimental conditions to determine the contribution of each architectural component. The research design also considers practical enterprise requirements, including distributed resource availability, network communication limitations, confidentiality obligations, scalability, and operational reliability. The proposed architecture is conceptual at the design stage and will be assessed experimentally through controlled simulation rather than being presented as an already deployed production security system. This approach enables systematic investigation of the relationships between intrusion detection performance, privacy protection, and computational efficiency.

The architectural workflow begins when network monitoring components generate security-related observations within individual enterprise cloud environments. These observations may include connection metadata, traffic-flow statistics, protocol information, timestamps, packet counts, service requests, authentication events, and other relevant security indicators. Raw packet payloads and personally identifiable information should be excluded wherever they are unnecessary for the research objectives. Each enterprise client performs local preprocessing, including data cleaning, duplicate removal, missing-value treatment, categorical encoding, numerical normalization, and feature selection. The resulting dataset is divided into training, validation, and testing subsets using a consistent experimental protocol. Following preprocessing, local machine learning models are trained using the participating organization's own observations. Only authorized model updates or protected parameter representations are transmitted to the federated coordination server. The server aggregates eligible updates and distributes the resulting global model to participating clients. This process is repeated across multiple communication rounds until a predefined convergence criterion or training limit is reached. Generative AI operates as a complementary intelligence component rather than replacing the primary intrusion detection classifier. It supports the generation of additional representative training samples, identification of unusual behavioral patterns, and production of human-readable explanations for security alerts. To prevent generated information from being treated as confirmed evidence, synthetic observations are clearly distinguished from real network records and evaluated before inclusion in model training. The complete architecture is designed around the principle that local data ownership, collaborative learning, and controlled generative intelligence must operate together to achieve effective enterprise cloud protection.

Dataset Selection, Preprocessing, and Federated Learning Implementation

The experimental investigation will utilize publicly available cybersecurity datasets containing labeled network traffic and representative attack categories. Suitable benchmark datasets include CICIDS2017, CSE-CIC-IDS2018, and UNSW-NB15. These datasets provide different network traffic characteristics and security event categories, enabling comparative analysis across normal and malicious activities. CICIDS2017 includes representative benign traffic and multiple attack categories, while CSE-CIC-IDS2018 provides additional traffic observations for investigating intrusion detection under diverse conditions. UNSW-NB15 offers another benchmark for evaluating classification performance across normal and abnormal network activities. Dataset selection will consider availability, documentation, feature compatibility, licensing conditions, and suitability for controlled experimentation. The selected datasets will be divided into simulated enterprise clients to represent organizations operating separate cloud infrastructures. For example, different clients may receive distinct distributions of benign traffic, suspicious connections, and attack categories. Both independent and identically distributed (IID) and non-independent and identically distributed (non-IID) data arrangements will be investigated. The IID configuration provides a controlled baseline, whereas non-IID configurations represent realistic differences in network behavior, application usage, infrastructure configuration, and threat exposure across enterprises. Data preprocessing will involve removing irrelevant identifiers, handling incomplete

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

observations, transforming categorical features, scaling numerical variables, and addressing class imbalance. Feature selection techniques may include correlation analysis, mutual information, and model-based importance estimation. Class imbalance will be examined using appropriate weighting and controlled augmentation approaches rather than assuming that all attack categories occur equally frequently. The dataset partitioning process will prevent information leakage between training and testing environments. In particular, duplicate records and closely related observations will be managed carefully to reduce artificially inflated evaluation results. All preprocessing transformations will be fitted using training data and consistently applied to validation and testing data. The resulting dataset partitions will support reproducible experiments and meaningful comparisons between centralized and decentralized learning approaches.

Federated learning will be implemented using a distributed training architecture in which each simulated enterprise client independently trains a local intrusion detection model. A conventional aggregation strategy, such as Federated Averaging (FedAvg), will serve as the baseline coordination mechanism. During each communication round, participating clients receive the current global model, perform local training using their respective datasets, and submit authorized model updates to the coordination server. The server combines the received updates according to the selected aggregation strategy and distributes the updated global parameters. Additional aggregation approaches may be investigated to understand how different coordination mechanisms influence model stability and detection effectiveness under heterogeneous data distributions. The local detection model may employ a multilayer perceptron, convolutional neural network, or another suitable classification architecture depending on the characteristics of the selected dataset. The experimental design will maintain consistent model configurations wherever possible to ensure fair comparisons. The federated training process will evaluate different client participation levels, local training epochs, communication rounds, and model convergence conditions. Particular attention will be given to the effects of client heterogeneity, because enterprise networks may contain significantly different traffic patterns and attack distributions. Personalization may also be considered to allow individual organizations to adapt the shared model to their specific network conditions while retaining common threat knowledge. The model coordination process will incorporate participant authentication, controlled update submission, and basic integrity verification. Where computationally practical, robust aggregation mechanisms will be investigated to reduce the influence of abnormal or unreliable client updates. These mechanisms are important because federated learning does not automatically protect against malicious participants or compromised training processes. The methodology therefore evaluates federated learning not simply as a replacement for centralized training, but as a distributed collaborative security mechanism that requires independent privacy and integrity controls.

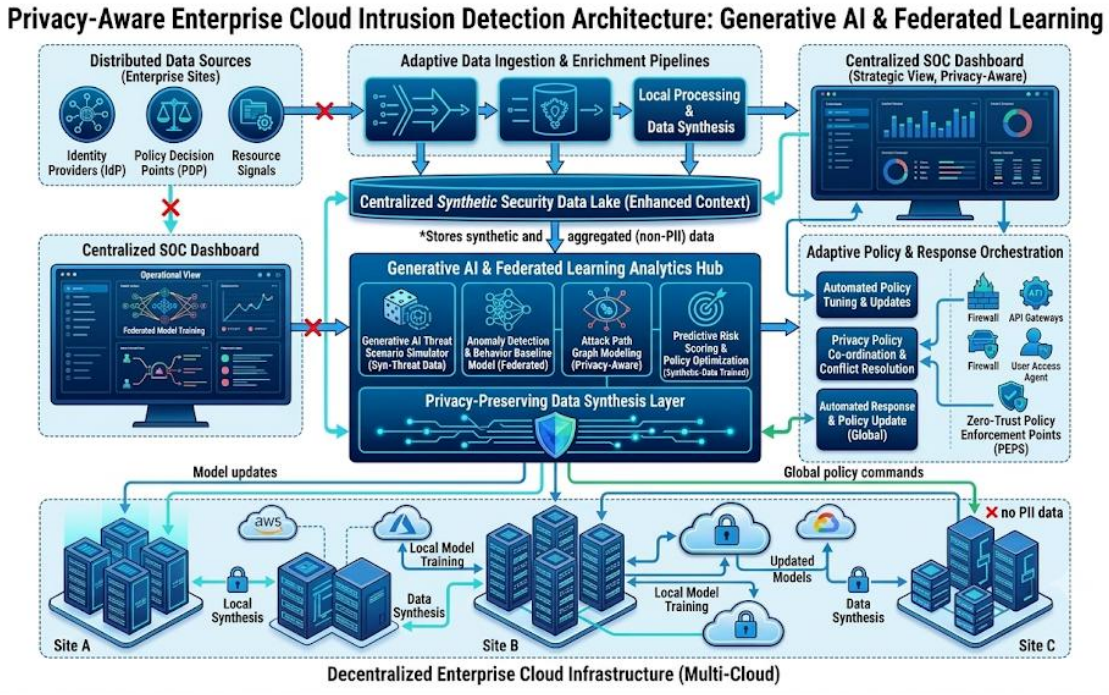


Figure 1: High-level architectural diagram for a Privacy-Aware Enterprise Cloud Intrusion Detection system leveraging Generative AI and Federated Learning.

International Journal of Advanced Research in Computer Science & Technology (IJARCST)

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

Generative AI Integration and Privacy-Preserving Mechanisms

Generative artificial intelligence will be incorporated into the proposed framework to improve the representation, interpretation, and analysis of security-related information. The research will investigate suitable generative architectures, including GANs, VAEs, and transformer-based language models, according to their computational requirements and intended functions. GANs or VAEs may be used to generate representative synthetic network-flow observations for underrepresented attack categories. Their purpose is to improve training-data diversity and reduce the effects of class imbalance. Generated samples will undergo statistical and distributional validation to determine whether they sufficiently resemble the relevant training characteristics without reproducing confidential records. The synthetic data generation process will be restricted to approved local datasets or privacy-controlled environments. Transformer-based language models may support security-event summarization, alert explanation, and contextual interpretation of model-generated anomaly indicators. For instance, when the intrusion detection classifier identifies unusual network behavior, the generative component can produce an understandable explanation based on approved event metadata and predefined security knowledge. The generative component will not independently authorize defensive actions or treat generated explanations as verified evidence. Instead, security analysts will review high-priority alerts and use the generated explanations as decision-support information. The study will compare detection performance with and without synthetic data augmentation to identify its actual contribution. It will also examine the potential risks of low-quality synthetic observations, distribution mismatch, and generated-data memorization. To evaluate the usefulness of generative AI, relevant indicators will include synthetic-data utility, downstream classification improvement, diversity of generated samples, and the consistency of security explanations. A carefully controlled experimental procedure will distinguish improvements resulting from additional training data from improvements caused by the federated learning process itself. This separation is necessary to understand whether generative AI contributes measurable benefits to enterprise cloud intrusion detection rather than merely increasing architectural complexity.

Privacy preservation will be treated as a fundamental methodological requirement throughout the federated training and evaluation process. The framework will combine data locality with additional privacy-enhancing mechanisms because keeping raw data within individual organizations does not provide complete protection against information leakage. Differential privacy will be investigated as a method for limiting the amount of information that model updates reveal about individual training records. Privacy parameters, including the privacy budget (ε) and, where applicable, the failure probability (δ), will be documented for experiments using formal differential privacy mechanisms. The relationship between privacy strength and intrusion detection performance will be examined by comparing different privacy configurations. Secure aggregation will also be incorporated to prevent the coordination server from directly inspecting individual client updates when supported by the experimental environment. Cryptographic protection may be considered for securing model-update communication and limiting unauthorized access to distributed training information. The research will additionally investigate threats such as membership inference, model inversion, and malicious client updates through controlled privacy and robustness assessments. These evaluations will use approved test procedures and synthetic or benchmark information rather than exposing operational enterprise systems. Privacy evaluation will examine whether the implemented safeguards provide measurable protection and what performance costs they introduce. Communication channels will employ authenticated and encrypted transport, while access to training coordination components will follow role-based authorization principles. Generated samples will also be evaluated for possible memorization of sensitive training information. The research will document the privacy assumptions, threat model, protection mechanisms, and known limitations associated with each experimental configuration. Importantly, the study will distinguish between data confidentiality, model-update confidentiality, and model integrity because these represent separate security objectives. The integration of differential privacy, secure aggregation, and appropriate cryptographic controls is expected to provide a more comprehensive privacy framework than relying exclusively on decentralized data storage.

Performance Evaluation, Comparative Analysis, and Validation

The final methodological stage will evaluate the proposed architecture through quantitative performance measurement and comparative experimentation. The primary detection metrics will include accuracy, precision, recall, F1-score, false-positive rate, false-negative rate, and confusion matrix analysis. Accuracy will measure the proportion of correctly classified observations, while precision will indicate the reliability of predicted intrusion alerts. Recall will measure the proportion of actual malicious observations successfully identified by the detection system. The F1-score will provide a balanced measure of precision and recall, particularly when the dataset contains imbalanced attack categories. False-positive and false-negative rates will be examined separately because excessive false alarms increase security analyst workloads, whereas undetected intrusions can expose enterprise infrastructure to substantial risks. The proposed privacy-aware generative federated learning architecture will be compared with centralized machine learning,

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

conventional federated learning without additional privacy mechanisms, federated learning with privacy controls, and federated learning incorporating generative augmentation. This comparison will help identify the independent and combined contributions of the proposed components. System efficiency will be assessed through local training duration, model convergence time, communication volume, memory consumption, computational requirements, and detection latency. Scalability will be evaluated by increasing the number of simulated enterprise clients and examining how system performance changes as additional participants join the federated training process. Privacy-related evaluation will document differential privacy parameters, aggregation protection, and the effectiveness of selected leakage assessments. The experimental framework will also consider model robustness under changing traffic distributions and different client participation conditions. All experiments will be repeated under consistent configurations where possible, and results will be reported using appropriate descriptive statistics. Where sufficient experimental repetitions are available, confidence intervals and statistical significance tests may be used to determine whether observed differences are meaningful. Ablation experiments will separately remove the generative component, differential privacy, and secure aggregation to examine their individual contributions. This systematic evaluation will provide a clearer understanding of the trade-offs between detection effectiveness, privacy protection, and operational overhead.

The validation process will focus on reproducibility, experimental reliability, and the practical relevance of the proposed architecture to enterprise cloud environments. The experimental implementation may use Python-based machine learning frameworks, federated learning simulation libraries, and standard data analysis tools. Model configurations, dataset partitions, preprocessing procedures, training parameters, privacy settings, and evaluation metrics will be documented to support repeatability. The evaluation will include multiple experimental scenarios representing different enterprise network sizes, traffic distributions, attack frequencies, and communication conditions. A baseline scenario will establish conventional detection performance, while subsequent scenarios will introduce federated coordination, generative augmentation, and additional privacy controls. The results will be interpreted according to predefined research questions: whether federated learning can achieve useful detection performance without centralizing sensitive training data; whether generative AI can improve the detection of underrepresented attack categories; whether differential privacy and secure aggregation can reduce exposure risks; and whether the combined architecture can operate efficiently across multiple enterprise cloud clients. The research will also consider implementation limitations, including computational resource requirements, communication delays, generative model reliability, privacy-utility trade-offs, and differences between benchmark datasets and real-world enterprise traffic. Ethical and governance considerations will include dataset licensing, responsible handling of security information, exclusion of unnecessary personal information, and controlled access to experimental records. The proposed framework will be considered successful only if its experimental results demonstrate an acceptable balance between intrusion detection effectiveness, privacy protection, and computational efficiency. No specific improvement percentage or performance advantage will be assumed before experimentation. The final analysis will identify architectural strengths, limitations, and opportunities for future research, including real-time deployment, adaptive federated aggregation, explainable security intelligence, and cross-cloud interoperability. Through this validation process, the study aims to establish a technically grounded framework for privacy-aware enterprise cloud intrusion detection that combines collaborative learning and generative intelligence while respecting organizational confidentiality requirements.

IV. RESULTS AND DISCUSSION

The proposed privacy-aware enterprise cloud intrusion detection framework integrates Generative Artificial Intelligence (GenAI) and Federated Learning (FL) to improve cybersecurity monitoring, threat identification, and data privacy across distributed enterprise cloud environments. The framework is designed to detect malicious activities, including unauthorized access, abnormal network traffic, distributed denial-of-service (DDoS) attacks, malware communication, privilege escalation, and suspicious user behaviour. The expected results indicate that combining federated learning with generative AI can improve intrusion detection capabilities while reducing the need to transfer sensitive enterprise information to centralized servers. Federated learning enables multiple cloud nodes or organizational departments to collaboratively train a shared detection model without directly exchanging their raw security logs and network traffic. Generative AI further supports the detection process by learning complex traffic patterns, generating representative synthetic attack samples, and identifying previously unseen or evolving threats. The integration of these technologies provides an opportunity to improve detection accuracy, reduce false-positive alerts, and strengthen the adaptability of enterprise security systems. Furthermore, privacy-preserving mechanisms, including secure model aggregation, encrypted communication, and differential privacy, can help minimize the exposure of

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

confidential information during collaborative model training. The framework also supports continuous learning, allowing intrusion detection models to adapt to changing network conditions and emerging cyberattack patterns.

The discussion highlights that the combination of generative AI and federated learning offers several advantages over conventional centralized intrusion detection systems. Traditional security architectures often require collecting large volumes of network traffic and system logs in a centralized location, creating potential privacy risks and increasing communication overhead. In contrast, the proposed architecture distributes model training across participating cloud environments, enabling organizations to retain control over their sensitive information. Generative AI enhances the analytical capabilities of the framework by identifying complex relationships between network events and generating synthetic datasets to address class imbalance in cybersecurity datasets. This is particularly useful because real-world intrusion datasets frequently contain significantly fewer examples of sophisticated attacks than normal network activities. However, the framework also presents challenges, including computational complexity, communication latency, non-independent and identically distributed (Non-IID) data across participating organizations, and the possibility of adversarial manipulation of federated model updates. Additionally, generative models may produce inaccurate or misleading outputs if they are trained using incomplete or biased datasets. Therefore, model validation, explainable AI techniques, secure aggregation, and continuous performance monitoring are essential to ensure reliable intrusion detection. Overall, the proposed architecture provides a promising direction for developing scalable, intelligent, and privacy-aware enterprise cloud security systems.

V. CONCLUSION

The proposed study presents a privacy-aware enterprise cloud intrusion detection architecture by integrating Generative Artificial Intelligence and Federated Learning to address the growing cybersecurity challenges associated with distributed cloud computing environments. The architecture combines collaborative model training, intelligent threat analysis, and privacy-preserving mechanisms to provide an effective approach to identifying malicious network activities without requiring organizations to share their raw security data. Federated learning enables multiple enterprise cloud environments to contribute to a common intrusion detection model while maintaining local data ownership and confidentiality. At the same time, generative AI improves the system's ability to recognize complex attack behaviours, identify unusual network patterns, and generate synthetic samples for underrepresented attack categories. This combination offers a foundation for improving intrusion detection accuracy, reducing false-positive rates, and enhancing the ability of security systems to respond to evolving cyber threats. The architecture also demonstrates the importance of integrating privacy protection into intelligent cybersecurity solutions rather than treating data confidentiality as a separate consideration. By incorporating secure aggregation, encrypted communication, and differential privacy, the framework supports collaborative cybersecurity intelligence while minimizing unnecessary exposure of sensitive enterprise information. Consequently, the proposed approach represents a potential advancement over conventional centralized intrusion detection architectures.

Furthermore, the study emphasizes that effective enterprise cloud security requires a balance between detection performance, computational efficiency, scalability, and privacy preservation. Although generative AI and federated learning provide significant opportunities, their practical implementation requires careful consideration of communication costs, model convergence, data heterogeneity, and adversarial threats. The reliability of the proposed architecture depends on the quality of locally available security datasets, the robustness of generative models, and the effectiveness of privacy-preserving federated optimization techniques. Explainable AI can additionally improve security analysts' understanding of detected anomalies by providing interpretable information about suspicious activities and potential attack characteristics. The integration of automated threat prioritization and continuous model updates can further strengthen enterprise incident response capabilities. Overall, the proposed framework establishes a conceptual foundation for intelligent, distributed, and privacy-preserving cloud intrusion detection. It demonstrates how generative AI and federated learning can complement each other to support collaborative threat intelligence while protecting organizational data confidentiality. The architecture can contribute to the development of next-generation enterprise cybersecurity systems capable of addressing increasingly sophisticated attacks in dynamic cloud environments.

VI. FUTURE WORK

Future research will focus on implementing and experimentally evaluating the proposed privacy-aware enterprise cloud intrusion detection architecture using realistic and publicly available cybersecurity datasets, including CICIDS2017,

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

CSE-CIC-IDS2018, and UNSW-NB15. The experimental implementation will investigate the effectiveness of different federated learning algorithms, such as Federated Averaging (FedAvg), Federated Proximal (FedProx), and personalized federated learning approaches, under different cloud network conditions. Generative AI techniques, including Variational Autoencoders (VAEs), Generative Adversarial Networks (GANs), and transformer-based models, can be explored for synthetic attack traffic generation, anomaly identification, and minority-class augmentation. Future experiments should measure important performance indicators, including accuracy, precision, recall, F1-score, false-positive rate, detection latency, communication overhead, and computational resource consumption. Comparative evaluations between centralized learning, conventional machine learning, standalone generative AI, and federated learning-based intrusion detection will help establish the effectiveness of the integrated architecture. Additionally, experiments involving different numbers of participating cloud nodes and heterogeneous enterprise datasets can provide insights into scalability, convergence behaviour, and model generalization. Such investigations will help determine whether the proposed architecture can maintain consistent detection performance when deployed across geographically distributed cloud infrastructures and organizations with different security requirements.

Another important direction for future work involves strengthening the security, privacy, explainability, and operational reliability of the proposed framework. Advanced privacy-preserving methods, including adaptive differential privacy, homomorphic encryption, secure multiparty computation, and trusted execution environments, can be investigated to reduce the risk of information leakage during federated model training. Future studies should also examine the framework's resilience against poisoning attacks, model inversion, membership inference, and adversarial manipulation of generative models. Blockchain-based mechanisms may be explored for maintaining transparent records of federated model updates and supporting accountability among participating organizations. In addition, explainable artificial intelligence techniques such as SHAP and LIME can be incorporated to provide interpretable explanations of intrusion detection decisions and improve the confidence of security analysts. Future implementations may also integrate large language models to summarize security alerts, correlate multiple attack indicators, and generate understandable incident reports. However, such capabilities should operate under human supervision to prevent inaccurate automated security decisions. Finally, real-time deployment in hybrid cloud, multi-cloud, edge computing, and Internet of Things environments will be investigated to evaluate the framework under practical operational conditions. These future developments can transform the proposed architecture into a more robust, scalable, adaptive, and trustworthy enterprise cloud intrusion detection solution capable of protecting sensitive organizational resources against emerging cybersecurity threats.

REFERENCES

1. Marston, S., Li, Z., Bandyopadhyay, S., Zhang, J., & Ghalsasi, A. (2011). Cloud computing—The business perspective. Decision Support Systems, 51(1), 176–189. https://doi.org/10.1016/j.dss.2010.12.006
2. 2. Schmidhuber, J. (2015). Deep learning in neural networks: An overview. Neural Networks, 61, 85–117. https://doi.org/10.1016/j.neunet.2014.09.003
3. Ponugoti, M. (2023). Frameworks for ensuring compliance in digital platform governance. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(6), 7575-7586.
4. Reddy, K. V. (2022). Power-aware verification closure using AI-enhanced UVM and dynamic voltage and frequency scaling models. International Journal of Research and Applied Innovations, 5(4), 7424–7436.
5. Gurram, A. (2023). A Framework for Achieving Zero-Downtime Data Migration Using Resilient Patterns for Production Database Schema Evolution. International Journal of Enhanced Research in Management & Computer Applications.
6. Kalakoti, M. K. R. (2023). AI-driven infrastructure orchestration for large-scale e-commerce platforms during global product launches. International Journal of Computer Technology and Electronics Communication, 6(4), 7388–7399.
7. Shewale, V. (2022). 3rd party and supply chain risk in oil & gas. International Journal of Emerging Trends in Computer Science and Information Technology, 3(4), 189–194. https://doi.org/10.63282/3050-9246.IJETCSIT-V3I4P118
8. 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.
9. Bitragunta, S. L. V. (2024). A novel AI-blockchain-edge framework for fast and secure transient stability assessment in smart grids. International Journal For Multidisciplinary Research, 6(6).
10. Karuturi, Y. (2023). Energy-aware capacity governance for large-scale SAP HANA deployments. International Journal of Research and Applied Innovations, 6(3), 6000–6009.

| ISSN: 2347-8446 | www.ijarcst.org | editor@ijarcst.org |A Bimonthly, Peer Reviewed & Scholarly Journal|

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

DOI:10.15662/IJARCST.2024.0706043

11. Basavala, S. R. (2022). Security assessment framework for microservices-based applications using automated penetration testing. International Journal of Humanities and Information Technology, 4(1–3), 327–341.
12. Namdeo, A. (2023). Generative synthetic data pipelines for balanced BI training. International Journal of Emerging Research in Engineering and Technology, 4(1), 228–235. https://doi.org/10.63282/3050-922X.IJERET-V4I1P124
13. Jinna, K. R., & Thilagaraj, S. V. (2022). From application server to pod: A readiness-scoring framework for migrating legacy J2EE middleware workloads to enterprise Kubernetes platforms. International Journal of Engineering & Extended Technologies Research (IJEETR), 4(3), 4953-4963.
14. Beeram, S. (2023). Energy efficiency and carbon-aware workload scheduling in Azure. International Journal of Research and Applied Innovations, 6(5), 9626–9628. https://doi.org/10.15662/IJRAI.2023.0605019
15. 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 (IJRPETM), 5(1), 6230-6240.
16. 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.
17. Gazi, M. S., Tasnim, M., Enayet, M. A., Chowdhury, J. A., Ullah, S., Akter, A., Rahman, M. R., & Khan, M. A. N. (2024). A novel framework for memory-resident malware investigation and forensic evidence reconstruction. International Journal of Research and Applied Innovations, 7(3), 10910–10921.
18. Potdar, A. (2022). Hybrid Sovereign Cloud Framework for Artificial Intelligence Powered Enterprise Analytics and Secure Data Integration. International Journal of Future Innovative Science and Technology (IJFIST), 5(5), 9265.
19. Randhi, K. S. (2024). A scalable integration framework for automated utility revenue leakage detection and remediation. International Journal of Research and Applied Innovations, 7(2), 10536–10555.
20. 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.
21. Karrothu, A. (2024). Automated risk-based security controls for vulnerability detection and prevention across cloud-native software development lifecycles. International Journal of Humanities and Information Technology, 6(03), 147–166.
22. Bellundagi, M. (2023). Blockchain based secure data sharing framework for smart applications. International Journal of Humanities and Information Technology, 5(2), 95–109.
23. Shaik, N. (2022). Engineering global resilience: Multi-cluster Istio for zero-downtime traffic management. International Journal of Computer Technology and Electronics Communication, 5(5), 5813–5825.