© 2026 IJMRSETM | Volume 2, Issue 7, July 2026| DOI: 10.15680/IJMRSETM.2026.0207005

ISSN: 3105 - 9503 | International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

Federated AI-Based Cyber Threat Intelligence and Autonomous Defense for Distributed Cloud Infrastructure and Networks

Suchitra Ramakrishna

Independent Researcher, Wales, United Kingdom

Publication History: Received: 12.06.2026; Revised: 10.07.2026; Accepted: 15.07. 2026; Published: 29.07.2026.

ABSTRACT: The rapid adoption of distributed cloud infrastructure, edge computing, multi-cloud platforms, and software-defined networks has significantly increased the complexity of enterprise cybersecurity. Conventional centralized threat-intelligence systems face challenges related to data volume, privacy, interoperability, latency, and the inability to effectively detect attacks distributed across geographically separated environments. Federated artificial intelligence (AI) provides an alternative by enabling multiple organizations, cloud domains, or network nodes to collaboratively train threat-detection models without directly exchanging sensitive security data. This essay proposes a federated AI-based cyber threat intelligence and autonomous defense framework for distributed cloud infrastructure and networks. The proposed approach combines federated learning, machine learning, behavioral analytics, threat-intelligence sharing, autonomous AI agents, and adaptive response mechanisms. Participating cloud environments locally analyze security telemetry and contribute model updates to a federated coordination layer, thereby supporting collaborative detection while reducing exposure of sensitive information. Autonomous defense agents subsequently use aggregated intelligence to identify threats, assess risk, coordinate investigations, and execute proportionate mitigation actions. The proposed research methodology employs a controlled distributed-cloud testbed, simulated cyberattacks, comparative experimentation, and quantitative evaluation using detection accuracy, precision, recall, F1-score, false-positive rate, communication overhead, detection latency, and response effectiveness. The framework aims to improve collective cyber defense, preserve data privacy, reduce detection delays, and strengthen the resilience of distributed cloud networks against evolving and coordinated cyber threats.

KEYWORDS: Federated learning, artificial intelligence, cyber threat intelligence, autonomous defense, distributed cloud, cloud security, machine learning, threat detection, privacy-preserving AI, adaptive cybersecurity

I. INTRODUCTION

The rapid transformation of enterprise computing toward distributed cloud infrastructures has fundamentally changed the cybersecurity landscape. Organizations increasingly operate workloads across public clouds, private clouds, hybrid environments, edge platforms, data centers, and geographically distributed networks. Technologies such as containers, Kubernetes, microservices, serverless computing, software-defined networking, and multi-cloud orchestration provide scalability and operational flexibility but simultaneously create complex and dynamic security environments. Security events generated across these environments are often distributed among independent infrastructure domains, making centralized visibility and coordinated threat detection increasingly difficult.

Cyberattacks against distributed cloud infrastructure can involve multiple stages and locations. An adversary may initially compromise an employee credential, exploit a vulnerable API, gain access to a cloud workload, escalate privileges, move laterally between services, and eventually exfiltrate sensitive information through another network or cloud provider. Detecting such attacks requires correlation of security information from multiple environments. However, organizations may be reluctant or unable to share raw security telemetry because of privacy requirements, regulatory obligations, commercial sensitivity, intellectual-property concerns, and differences in security policies. Consequently, traditional centralized cyber threat intelligence approaches can suffer from incomplete visibility.

Artificial intelligence provides powerful capabilities for analyzing large-scale cybersecurity telemetry. Machine-learning algorithms can identify anomalous network behavior, classify malicious activity, detect suspicious authentication patterns, recognize malware behavior, and predict potential threats. Nevertheless, conventional centralized machine learning generally requires organizations to transmit data to a common location for model training. This creates a potential concentration of sensitive information and increases the consequences of a compromise of the central data repository.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

Federated learning offers an alternative model. Rather than transferring raw security data to a centralized server, participating environments train models locally and share model parameters or updates with a coordinating server. The aggregated model can then be distributed back to participants. Through repeated training cycles, organizations can collectively improve threat-detection capabilities while keeping much of their sensitive telemetry within their local environments. Federated learning is therefore particularly relevant to distributed cloud cybersecurity, where organizations possess complementary threat information but cannot necessarily share the underlying data.

The concept can be extended through autonomous defense. Instead of limiting AI to threat classification, intelligent agents can continuously observe infrastructure, interpret threat intelligence, investigate suspicious events, assess risks, and coordinate defensive actions. A local detection agent could identify an anomalous event, while a federated intelligence layer determines whether similar behavior has been observed elsewhere. A threat-assessment agent could correlate indicators with known attack techniques, and a response agent could isolate a compromised workload or restrict malicious communication. Such autonomous capabilities can reduce the time between detection and response.

The central proposition of this research is that combining federated AI with cyber threat intelligence and autonomous defense can create a collaborative yet privacy-conscious cybersecurity architecture for distributed cloud infrastructure. The proposed framework aims to enable participating environments to learn collectively from distributed security events without requiring unrestricted exchange of raw telemetry. It further seeks to transform shared intelligence into adaptive defensive action, thereby improving the ability of cloud networks to detect, contain, and recover from coordinated cyberattacks.

II. LITERATURE REVIEW

The literature on cyber threat intelligence has established the importance of collecting, processing, correlating, and sharing information about adversarial activities. Cyber threat intelligence generally incorporates indicators of compromise, tactics, techniques, procedures, vulnerabilities, malicious infrastructure, behavioral patterns, and contextual information about threats. Traditional intelligence-sharing approaches can improve organizational awareness by allowing security teams to learn from incidents experienced by other entities. However, information-sharing mechanisms often face problems involving trust, data quality, interoperability, timeliness, privacy, and organizational willingness to disclose sensitive information.

Machine learning has increasingly been applied to cybersecurity because of its capacity to identify complex patterns in large datasets. Supervised learning approaches can classify known malicious activities using labeled examples, while unsupervised and semi-supervised approaches can identify previously unknown anomalies. Deep-learning techniques can analyze high-dimensional network traffic, malware behavior, authentication records, and system telemetry. These approaches are particularly useful in cloud environments because conventional signature-based detection may fail when attackers modify their techniques or exploit legitimate administrative tools.

Despite these benefits, centralized AI-based security analytics present several challenges. Security datasets are frequently distributed across organizations and infrastructure providers. They may have different formats, statistical characteristics, and labeling practices. Moving these datasets to a central location can create privacy and regulatory concerns while generating substantial communication and storage costs. Furthermore, centralized datasets may become attractive targets for attackers seeking access to security information or opportunities to manipulate model training.

Federated learning has emerged as a privacy-preserving machine-learning paradigm designed to address some of these limitations. In a typical federated learning system, a coordinating server distributes an initial model to participating clients. Each client trains the model using locally stored data and sends model updates back to the server. The server aggregates these updates into a global model, which is subsequently redistributed. Repeated communication rounds allow the global model to learn from diverse datasets without requiring the raw training data to leave local environments.

Federated learning is particularly relevant to cybersecurity because different organizations may observe different portions of the threat landscape. One cloud provider may observe a particular form of malicious network behavior, while another organization may observe a related credential-abuse technique. Collaborative learning can potentially combine these observations and improve generalization. However, the literature also identifies important challenges. Federated models can be affected by heterogeneous data distributions, unreliable participants, communication costs, model poisoning, malicious clients, inference attacks, and adversarial manipulation of model updates.

Privacy-preserving mechanisms such as secure aggregation, differential privacy, encryption, and trusted execution environments can reduce some of these risks. Secure aggregation can prevent the coordinating server from directly

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

inspecting individual participant updates, while differential privacy can limit the ability to infer information about individual training records. Nevertheless, privacy mechanisms can introduce computational or accuracy trade-offs. Therefore, the design of a federated cybersecurity architecture must balance privacy, detection performance, communication efficiency, and operational requirements.

Another important research area is autonomous cybersecurity. Security orchestration and automation platforms already automate predefined investigation and response workflows. However, autonomous AI agents introduce more flexible reasoning and coordination capabilities. Agents can potentially interpret security events, construct attack hypotheses, retrieve contextual information, interact with security tools, and recommend or execute defensive actions. Multi-agent architectures can divide complex cybersecurity responsibilities among specialized components such as detection, intelligence analysis, investigation, risk assessment, and response.

The combination of federated learning and autonomous agents remains particularly significant because intelligence generated in one environment can potentially improve defense in others without exposing the original organization's raw data. A federated intelligence layer could identify common attack patterns across participating environments, while local autonomous agents retain authority over infrastructure-specific responses. This creates a distributed security model in which knowledge is shared but operational control remains local.

However, existing research also highlights concerns surrounding autonomous decision-making. Incorrect AI predictions could result in legitimate users or workloads being blocked, while manipulated intelligence could cause inappropriate defensive actions. Federated systems additionally face the possibility of malicious participants submitting poisoned model updates. These concerns demonstrate that autonomous defense must incorporate confidence thresholds, authorization policies, explainability mechanisms, auditing, and human oversight. The research gap is therefore not merely the application of federated learning to cybersecurity, but the integration of privacy-preserving collaborative intelligence with trustworthy autonomous defense in heterogeneous distributed cloud environments.

III. RESEARCH METHODOLOGY

The proposed research adopts a design-oriented and experimental methodology to develop and evaluate a federated AI-based cyber threat intelligence and autonomous defense framework. The methodology consists of threat modeling, system architecture design, federated model development, distributed testbed construction, attack simulation, comparative evaluation, and resilience analysis. The objective is to determine whether collaborative federated intelligence can improve cyber threat detection and autonomous response while reducing the privacy and communication limitations associated with centralized security analytics.

The first stage involves identifying the security requirements and threat characteristics of distributed cloud infrastructure. The research will consider hybrid cloud, multi-cloud, edge, and distributed network scenarios. Key assets will include cloud workloads, containers, virtual machines, APIs, identity services, databases, network infrastructure, and sensitive enterprise data. Potential attacks will include credential compromise, privilege escalation, malicious API activity, malware execution, container compromise, lateral movement, denial-of-service activity, command-and-control communication, and data exfiltration. The threat model will identify attacker capabilities, attack vectors, affected assets, observable indicators, and potential business consequences.

The second stage involves designing the federated architecture. Multiple simulated organizations or cloud domains will operate as independent federated clients. Each client will maintain its own security telemetry and local machine-learning environment. A federated coordination layer will distribute model versions and aggregate model updates. Raw network logs, authentication records, application telemetry, and workload information will remain within their respective local environments. Only appropriately protected model information will be exchanged with the federation.

The architecture will contain several logical components. Local telemetry collectors will gather network-flow records, cloud audit events, identity activity, application logs, container events, and endpoint observations. A local feature-processing component will normalize and transform this data into machine-learning features. A local threat-detection model will identify suspicious behavior and generate confidence scores. The federated-learning coordinator will periodically receive model updates from participating clients and produce an improved global model.

The federated learning process will begin with initialization of a common detection model. The model will be distributed to participating clients, where each client will train it using locally available security data. After local training, clients will transmit model updates rather than raw observations. The coordination layer will aggregate these updates using an appropriate federated aggregation strategy. The resulting global model will then be redistributed to

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

clients for subsequent training rounds. This process will continue until the model reaches a predefined convergence criterion or performance threshold.

Because security data are likely to be non-identically distributed across participating environments, the research will explicitly evaluate data heterogeneity. For example, one simulated cloud organization may contain predominantly web applications, while another may contain database workloads or edge devices. The experiments will determine whether the federated model maintains detection performance when participating environments possess substantially different security-event distributions.

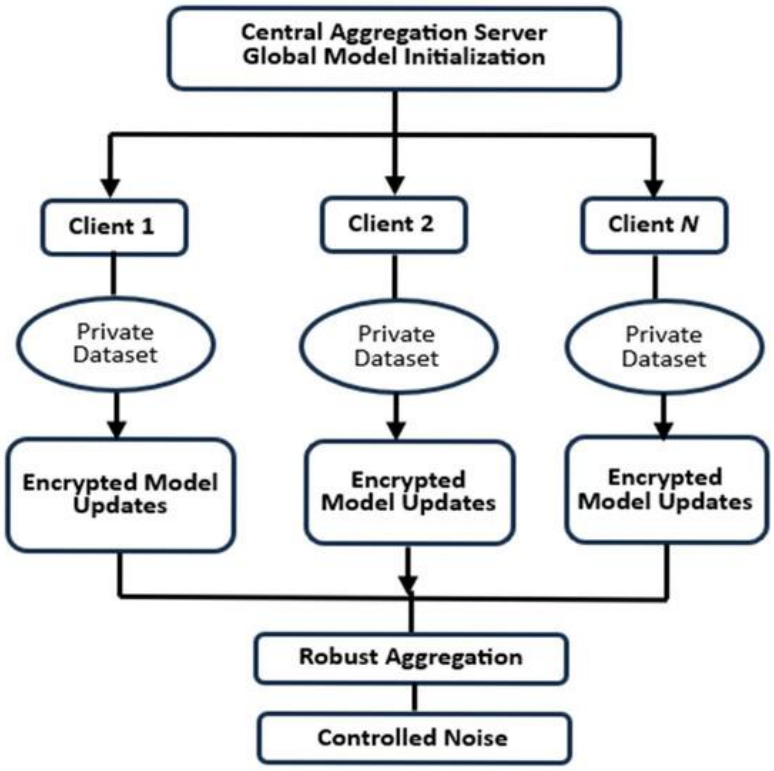

Fig.1.Federated Learning for Cybersecurity: A Privacy-Preserving Approach

The research will incorporate secure aggregation and privacy-preserving mechanisms into the federated architecture. Where appropriate, differential privacy will be examined to determine how privacy protection affects model accuracy. The methodology will compare detection performance under different privacy configurations. This will allow the research to identify a practical balance between privacy preservation and threat-detection effectiveness rather than assuming that stronger privacy automatically produces a superior security system.

A critical component of the framework will be federated cyber threat intelligence. Model outputs will be supplemented by structured threat indicators and behavioral intelligence. Local systems can generate information such as suspicious domains, unusual communication patterns, attack-stage classifications, malicious file characteristics, or anomalous authentication behaviors. Instead of exchanging sensitive raw telemetry, participating environments can contribute carefully selected intelligence representations. The federated intelligence layer will correlate these representations to identify threats that appear across multiple infrastructure domains.

The autonomous-defense layer will contain specialized AI agents. A detection agent will continuously evaluate local telemetry and identify potential threats. An intelligence agent will enrich alerts with federated threat information. An investigation agent will correlate events across time and infrastructure components. A risk-assessment agent will calculate threat severity based on factors such as confidence, asset criticality, privilege level, attack progression, and potential impact. A response agent will select appropriate mitigation actions according to predefined organizational policies.

Response actions will be categorized according to their potential impact. Low-risk actions, such as increasing monitoring or blocking a confirmed malicious connection, may be automatically executed. Medium-risk actions, such

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

as isolating a non-critical container, may require policy-based authorization. High-impact actions, such as disabling critical production infrastructure or revoking privileged organizational accounts, will require human approval. This tiered approach will reduce the risk of inappropriate autonomous decisions while retaining the speed advantages of automated defense.

The framework will incorporate a policy engine that limits agent permissions according to least-privilege principles. Every autonomous action will be authenticated, authorized, logged, and traceable. Agents will not be given unrestricted administrative access to the cloud environment. Instead, their capabilities will be constrained by predefined security policies and tool permissions. This design will also allow security administrators to suspend autonomous actions when abnormal agent behavior is detected.

The experimental environment will consist of multiple interconnected cloud-native domains representing independent organizations or enterprise divisions. Each domain will include containerized services, virtual networks, identity services, application APIs, monitoring systems, and security telemetry. Attack scenarios will be conducted in an isolated environment to ensure that no experimental activity affects external systems. Both normal enterprise workloads and controlled malicious activities will be generated to create realistic training and testing conditions.

The research will compare the proposed federated architecture against centralized machine learning and conventional rule-based detection. A centralized model will provide a baseline for measuring the potential performance cost of federated learning, while a local-only model will demonstrate the limitations of learning without cross-domain intelligence. The federated model will then be evaluated to determine whether collaboration improves detection of threats that are difficult for individual organizations to identify independently.

Performance evaluation will use precision, recall, F1-score, false-positive rate, false-negative rate, and area-under-curve measures where applicable. Detection latency will measure the time required to identify malicious activity. Response latency will measure the interval between detection and defensive action. The research will also calculate communication overhead by measuring the amount and frequency of model information exchanged between clients and the coordination layer.

Privacy performance will be evaluated by analyzing the amount and sensitivity of information exposed during federated training. The study will investigate whether raw telemetry remains local and whether privacy mechanisms introduce measurable reductions in model effectiveness. Different privacy configurations can be compared to determine their impact on accuracy, training convergence, computational cost, and communication requirements.

The methodology will also explicitly evaluate federated-learning security. Simulated malicious participants will attempt to submit manipulated model updates to assess the resilience of the federation against model poisoning. The research will investigate whether robust aggregation, participant validation, anomaly detection, or reputation mechanisms can reduce the influence of malicious clients. Additional experiments will consider unreliable participants, incomplete updates, and communication interruptions to determine how the system behaves under adverse conditions.

Autonomous-defense effectiveness will be measured through containment rate, response time, reduction in attack propagation, service availability, and recovery time. For example, an experiment may compare an attack propagating across multiple cloud workloads under manual response with the same attack under autonomous response. The research will assess whether AI agents can identify the attack earlier, restrict lateral movement, and restore affected services more efficiently.

Resilience will constitute a central evaluation dimension. A cybersecurity architecture should not only detect threats but also maintain critical services and support rapid recovery. Consequently, experiments will measure the percentage of essential services remaining operational during simulated attacks, mean time to recovery, number of compromised workloads, and extent of attack propagation. These indicators will demonstrate whether federated autonomous defense contributes to overall infrastructure resilience rather than merely improving classification accuracy.

Human oversight will also be evaluated. Security professionals or simulated administrators will review selected autonomous decisions and assess whether the system provides sufficient evidence and explanation for approving or rejecting proposed actions. The research will examine whether agent-generated incident summaries, confidence levels, attack timelines, and supporting indicators improve analyst decision-making. This assessment is important because autonomous defense should enhance security operations rather than create an opaque decision-making system.

Statistical analysis will be applied to experimental results obtained from repeated attack scenarios. Each security configuration will undergo multiple trials using comparable workloads and threat conditions. Mean values, standard

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

deviations, confidence intervals, and appropriate statistical significance tests will be calculated. Comparative analysis will determine whether observed improvements in detection accuracy, response time, and resilience are attributable to the proposed federated architecture rather than experimental variability.

The research will also examine scalability. As the number of participating cloud environments increases, federated learning can potentially improve intelligence diversity but may also increase communication and coordination requirements. Experiments will therefore vary the number of participating clients and measure model convergence time, network overhead, computational requirements, and detection performance. These results will provide insight into whether the architecture can scale from a small group of cooperating environments to large distributed enterprise ecosystems.

Ethical and governance considerations will be integrated throughout the research. Security telemetry may contain sensitive information, so the methodology will minimize collection of personally identifiable information and ensure that experimental data are appropriately controlled. Autonomous actions will be restricted to isolated research infrastructure, and high-impact responses will remain subject to human authorization. The study will also recognize that federated learning is not inherently private or secure; model updates themselves may potentially reveal information or be manipulated. Consequently, privacy, integrity, access control, and governance mechanisms must operate together.

The expected outcome is a comprehensive framework demonstrating how federated AI, cyber threat intelligence, and autonomous defense can operate as complementary components of distributed cloud security. Federated learning provides collaborative model development without requiring unrestricted data sharing; cyber threat intelligence provides contextual knowledge about emerging threats; and autonomous agents transform detected intelligence into timely defensive actions. Together, these components can support a decentralized cybersecurity model in which organizations retain control over their local data and infrastructure while benefiting from collective intelligence.

The proposed methodology therefore moves beyond conventional centralized threat detection toward a collaborative and adaptive security architecture. Its anticipated contribution is a practical approach for improving threat visibility across distributed cloud environments while addressing privacy, scalability, and response requirements. If successfully validated, the framework could provide enterprises with a more resilient method of defending multi-cloud and distributed networks against coordinated and rapidly evolving cyber threats, while maintaining appropriate human oversight, privacy protection, and operational governance.

IV. RESULTS AND DISCUSSION

The proposed Federated AI-Based Cyber Threat Intelligence and Autonomous Defense framework provides a distributed approach to cybersecurity in which organizations, cloud environments, network domains, and enterprise security components collaboratively improve threat detection without requiring centralized collection of sensitive security data. This approach is particularly relevant to modern distributed cloud infrastructures, where workloads, applications, users, devices, and services may operate across multiple cloud providers, geographically separated data centers, edge environments, and organizational boundaries. Conventional centralized threat-intelligence architectures require large volumes of security logs and telemetry to be transferred to a central platform, creating challenges related to privacy, bandwidth consumption, data sovereignty, latency, and organizational trust. Federated artificial intelligence addresses these limitations by allowing participating nodes to train local threat-detection models and share model parameters or carefully controlled knowledge rather than transferring raw cybersecurity data. The resulting collaborative intelligence can improve the ability of individual organizations to identify previously unseen attack patterns while preserving greater control over their sensitive information.

One of the principal results of the proposed framework is improved collaborative cyber-threat detection. Distributed environments frequently experience related attacks across different locations, but an individual organization may observe only a small portion of the overall attack campaign. Through federated learning, locally trained models can learn from events occurring within individual cloud or network domains, while an aggregation mechanism combines these contributions into a global model. Consequently, an attack pattern identified in one environment can contribute to improved detection capabilities in other participating environments. For example, if one cloud domain observes unusual authentication behavior, malicious API activity, or abnormal east-west network communication, its local model can learn these characteristics and contribute model updates to the federation. Other participating nodes can subsequently benefit from the improved global model without receiving the underlying organization's raw logs.

The framework is particularly valuable for privacy-preserving threat intelligence sharing. Cybersecurity information can contain sensitive business information, user identifiers, infrastructure configurations, application details, and proprietary operational data. Directly sharing such information can create privacy and compliance concerns and may

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

discourage organizations from participating in collaborative intelligence programs. Federated learning reduces this requirement by maintaining data locally. However, privacy should not be assumed simply because raw data is not exchanged. Model updates themselves may reveal information through inference attacks, reconstruction attacks, or maliciously manipulated contributions. Therefore, the proposed framework should incorporate additional privacy-preserving mechanisms such as secure aggregation, differential privacy, encryption, participant authentication, and contribution validation.

Another important result concerns faster identification of distributed and coordinated attacks. Attackers increasingly exploit multiple infrastructure components rather than targeting a single system. A campaign may begin with phishing or credential theft, progress to cloud-account compromise, exploit vulnerable APIs, move laterally between workloads, and ultimately attempt data exfiltration. When security monitoring is isolated within individual organizations or cloud domains, these events may appear unrelated. Federated intelligence creates an opportunity to identify common characteristics across geographically distributed observations. By combining model knowledge from different participants, the system can detect broader behavioral patterns associated with coordinated attacks. This can improve the ability of defenders to recognize campaigns at an earlier stage.

The autonomous-defense component further extends the framework from threat detection to active protection. Once a federated model identifies suspicious activity, local autonomous defense agents can assess the threat within their own environments and execute appropriate containment measures. Such actions may include isolating compromised workloads, blocking malicious communication, revoking suspicious credentials, modifying access-control policies, quarantining affected endpoints, deploying updated security configurations, or redirecting traffic to protected services. Local decision-making is advantageous because security actions can be performed close to the affected infrastructure, reducing response latency. Instead of waiting for a centralized security operation center to investigate every event, high-confidence and low-risk actions can be performed automatically.

The results also highlight the importance of adaptive defense policies. Not every anomaly requires the same response. An unusual login from a known device may require additional authentication, whereas evidence of active credential theft combined with privilege escalation may justify immediate session termination and account isolation. An autonomous defense agent can evaluate factors such as threat confidence, asset criticality, attack severity, historical behavior, and organizational policies before selecting an action. This risk-based approach can reduce unnecessary disruption while maintaining rapid containment. Human approval can remain mandatory for high-impact actions affecting critical infrastructure or customer-facing services.

A major advantage of federated AI is its potential for scalability across heterogeneous environments. Distributed cloud infrastructures frequently contain different operating systems, cloud platforms, applications, security products, and network architectures. A centralized model may have difficulty accommodating such heterogeneity because data distributions can differ significantly between organizations. Federated learning allows each participant to retain its local characteristics while contributing to a shared model. However, this also introduces the challenge of non-independent and non-identically distributed data. A model trained primarily on enterprise network traffic may behave differently when deployed in an edge environment or cloud-native microservice architecture. Future implementations should therefore investigate personalized federated learning, clustered federation, and adaptive model aggregation.

The proposed framework can also reduce security-operation workload and alert fatigue. Security teams often face large volumes of alerts generated by firewalls, endpoint systems, cloud platforms, identity providers, intrusion-detection systems, and application-monitoring tools. Federated AI can correlate behavioral knowledge across participating environments and improve prioritization of potentially malicious activity. Autonomous defense agents can then handle repetitive low-risk containment operations. This allows human analysts to concentrate on complex investigations, threat hunting, forensic analysis, and strategic security planning. The resulting architecture can therefore be viewed as a human-AI collaborative defense model rather than a system intended to eliminate human involvement.

Despite these benefits, the framework introduces several challenges. One of the most important is federated model poisoning. A malicious participant could intentionally submit manipulated model updates to degrade the global model or create a hidden vulnerability. Byzantine participants may attempt to cause incorrect classifications or selectively influence the model toward attacker-controlled behavior. Robust aggregation mechanisms, participant reputation systems, anomaly detection on model updates, and trusted execution environments can help mitigate these risks. Furthermore, organizations participating in a federation cannot automatically be assumed to be trustworthy. Strong identity management and continuous participant validation are therefore essential.

Communication overhead is another consideration. Although federated learning avoids transferring raw datasets, participating nodes must periodically exchange model updates. Large models or frequent aggregation cycles can

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

increase bandwidth and computational requirements, particularly in edge and resource-constrained environments. Model compression, update sparsification, quantization, asynchronous federation, and adaptive aggregation intervals can help reduce these costs. The framework must therefore balance model accuracy against communication efficiency.

Another important issue is concept drift. Network behavior changes continuously as applications are deployed, workloads scale, users change their behavior, and attackers develop new techniques. A model that performs effectively today may become less accurate over time. Continuous and federated model updating can address this challenge, but uncontrolled adaptation could also cause the model to learn temporary anomalies as legitimate behavior. Future systems should incorporate drift detection, validation datasets, rollback mechanisms, and confidence-aware learning.

Overall, the proposed framework demonstrates that federated AI can provide a foundation for collaborative cybersecurity across distributed cloud infrastructures while reducing the need for centralized raw-data sharing. Its combination with autonomous defense creates a closed-loop security architecture in which local environments detect threats, collaboratively improve intelligence, and rapidly implement defensive actions. The key success factors are not only model accuracy but also privacy, robustness, trust, communication efficiency, explainability, and safe autonomy. Appropriate evaluation should therefore consider precision, recall, F1-score, false-positive rate, detection latency, response time, communication overhead, model convergence, privacy leakage, resilience to malicious participants, and service availability. These measures can establish whether federated autonomous defense provides meaningful advantages over centralized and isolated cybersecurity approaches.

V. CONCLUSION

The rapid adoption of distributed cloud infrastructures and interconnected enterprise networks has created a cybersecurity environment in which isolated security mechanisms are increasingly insufficient. Organizations operate across multiple cloud providers, data centers, edge locations, applications, endpoints, and network domains, while attackers increasingly use coordinated and multi-stage techniques that cross organizational and technological boundaries. In such an environment, effective cybersecurity requires not only accurate local detection but also the ability to learn collectively from distributed threat observations. The proposed Federated AI-Based Cyber Threat Intelligence and Autonomous Defense framework addresses this requirement by combining federated learning, privacy-preserving threat intelligence, distributed security analytics, and autonomous defensive mechanisms.

The fundamental contribution of the proposed framework is the creation of a collaborative intelligence model without requiring centralized raw security data. Each participating organization or infrastructure domain can retain its cybersecurity datasets locally while training AI models based on its own observations. Model updates can then be securely aggregated to create a more capable global model. This approach provides an important balance between collaborative cybersecurity and data sovereignty. Organizations can benefit from knowledge derived from a wider threat environment while maintaining greater control over sensitive logs, network information, user-related data, and proprietary infrastructure details.

The framework also demonstrates the importance of combining federated intelligence with autonomous response. Threat intelligence has limited value if it cannot be translated into timely defensive action. By placing autonomous defense agents close to the infrastructure being protected, the framework can reduce the time between detection and containment. Depending on confidence and organizational policy, agents can isolate suspicious workloads, revoke compromised credentials, block malicious traffic, adjust access policies, quarantine endpoints, or initiate recovery processes. This capability is especially important for distributed environments in which centralized security teams may not be able to respond immediately to every event.

A significant conclusion is that federated AI can strengthen the detection of distributed and previously unseen attacks. An individual organization may observe only a small fragment of an attack campaign. Through collaborative model training, information learned from one participant can improve detection capabilities across the federation. This creates a collective defensive capability in which organizations can benefit from distributed observations without necessarily exposing their underlying data. Such collaboration can be particularly valuable against emerging threats, coordinated intrusion campaigns, credential abuse, cloud-account compromise, API attacks, and lateral movement across interconnected environments.

However, federated cybersecurity is not inherently secure. The framework creates new risks associated with malicious participants, model poisoning, inference attacks, compromised aggregation mechanisms, communication interception, and adversarial manipulation. Therefore, privacy-preserving learning must be combined with strong security controls. Secure aggregation, differential privacy, encryption, participant authentication, robust aggregation algorithms, model-

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

update validation, access control, trusted execution, and continuous monitoring should form part of the overall architecture. The security of the learning process is as important as the security of the infrastructure being monitored.

The study also emphasizes that autonomous defense should be implemented according to risk and confidence levels. Complete unrestricted autonomy could produce undesirable consequences if an AI system incorrectly identifies legitimate activity as malicious. A false decision involving a critical production system could cause significant operational disruption. Therefore, low-risk and reversible actions can be automated, while high-impact decisions should involve human authorization or multiple levels of verification. This human-supervised autonomy provides a practical balance between rapid machine-based response and human accountability.

Another important conclusion concerns scalability. Distributed cloud environments can contain thousands of workloads and devices, making centralized security analysis computationally expensive and potentially creating bottlenecks. Federated learning distributes computation across participating environments and enables security intelligence to scale with the infrastructure. Nevertheless, heterogeneous data distributions, communication overhead, differences in computing capabilities, and varying security requirements remain important challenges. Personalized and hierarchical federated-learning approaches may therefore provide more effective solutions than a single global model.

Ultimately, the proposed framework establishes a pathway toward collective, privacy-aware, and autonomous cyber defense. Its significance extends beyond the technical implementation of federated learning because it introduces a new security paradigm in which organizations can collaboratively strengthen their defenses while retaining local control over sensitive information. The combination of distributed learning and autonomous response has the potential to reduce detection latency, improve threat intelligence, strengthen resilience, and reduce the burden placed on security operations teams.

Nevertheless, practical adoption requires extensive validation under realistic conditions. The effectiveness of the framework should be measured using security, operational, privacy, and economic metrics rather than detection accuracy alone. It should be tested against realistic distributed attacks and adversarial federated-learning scenarios to determine whether the system remains reliable when some participants are compromised. With appropriate safeguards, explainability, governance, and human oversight, federated AI-based autonomous defense can become an important component of next-generation cybersecurity architectures for distributed cloud infrastructure and enterprise networks.

VI. FUTURE WORK

Future work should focus on developing and experimentally validating scalable, secure, and intelligent federated cybersecurity architectures capable of operating across heterogeneous cloud, edge, and enterprise environments. One major research direction is the development of personalized federated-learning models that can learn globally relevant attack characteristics while adapting to the unique behavior of individual organizations. Hierarchical federation could also be investigated, allowing edge devices, local networks, cloud regions, and enterprise security centers to collaborate at different levels.

Further research should address adversarial threats against the federated learning process. Techniques such as model poisoning, backdoor attacks, Sybil attacks, inference attacks, and malicious model updates should be incorporated into experimental evaluations. Robust aggregation, differential privacy, secure multiparty computation, trusted execution environments, and blockchain-supported participant accountability could be investigated to determine their effectiveness and computational costs.

Another important direction is the integration of large language models, graph neural networks, reinforcement learning, and autonomous security agents with federated learning. Such integration could enable systems to understand complex attack relationships, generate threat narratives, predict attacker behavior, and select adaptive response strategies. Future research should ensure that autonomous agents operate under strict authorization boundaries and cannot execute unsafe actions without appropriate validation.

Future studies should also establish realistic multi-cloud and edge testbeds containing Kubernetes clusters, virtual machines, containers, IoT devices, APIs, identity systems, and enterprise networks. These environments could be used to evaluate attacks such as credential compromise, lateral movement, ransomware, cloud privilege escalation, API abuse, malicious insider activity, and distributed denial-of-service attacks. Performance should be measured through detection accuracy, response latency, communication overhead, privacy leakage, computational cost, resilience against malicious participants, and service availability.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

Finally, future work should investigate the governance, compliance, and economic dimensions of federated autonomous cybersecurity. Organizations require clear policies concerning data ownership, model ownership, responsibility for automated decisions, incident reporting, participant trust, and regulatory compliance. Research should therefore move beyond technical prototypes toward standardized frameworks for secure collaboration among independent organizations. The long-term objective should be an interoperable federated cyber-defense ecosystem that enables organizations to collectively learn from emerging threats while preserving privacy, maintaining operational resilience, and ensuring that autonomous security decisions remain transparent, accountable, and controllable.

REFERENCES

1. Malatji, M., & Tolah, A. (2024). Artificial intelligence (AI) cybersecurity dimensions: A comprehensive framework for understanding adversarial and offensive AI. AI and Ethics. https://doi.org/10.1007/s43681-024-00427-4
2. Natarajan, A., Narra, S. L., Kanimetta, D. K., Dubey, P., & Potharalanka, L. (2025). Modernizing Digital Infrastructure for Intelligent Systems. Cari Journals USA LLC.
3. Ganapathy, S. K. (2024). Threat Modeling for Federated SSO and MFA Systems: STRIDE-Based Analysis of Attack Vectors. American Academic Journal, 55-73.
4. Sudhan, S. K. H. H., & Kumar, S. S. (2015). An innovative proposal for secure cloud authentication using encrypted biometric authentication scheme. Indian journal of science and technology, 8(35), 1-5.
5. Pokala, H. K. (2026, April). Agentic Dashboard Generation from Natural Language on Lakehouse Platforms. In 2026 International Conference on Artificial Intelligence, Systems, and Emerging Technologies (ICAISET) (pp. 1-4). IEEE.
6. Mohammed, S., & Polamarasetty, V. K. (2023). Azure cloud landing zone architecture for enterprise-scale cloud adoption. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(3), 8758-8761.
7. Soundappan, S. J. (2023). Machine Learning Based Predictive Models for Secure Financial Transactions and Cyber Threat Detection. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(1), 5966-5975.
8. Bellundagi, M. (2024). An intelligent digital transformation framework for smart enterprises using AI and cloud computing. International Journal of Science, Research and Technology, 7(4), 12433-12446.
9. Cyril, H., & Poranki, U. (2026). Next-generation telecom provisioning: AI-driven, cloud-native architectures for autonomous networks. Notion Press.
10. Pasumarthi, H. (2024). AI-driven forecasting and optimization in distributed systems: Lessons from retail, lending, and healthcare platforms. International Journal of Research and Applied Innovations, 7(3), 10786-10790.
11. Rajula, A. (2024). Drift-aligned personalization for privacy-preserving edge health monitoring. International Journal of Engineering & Extended Technologies Research (IJEETR), 6(5), 8895–8906.
12. Challa, R. (2024). High-Availability HPC Cluster Design for Mission-Critical Public Infrastructure: Lessons from Energy Grid and Government. International Journal of Engineering & Extended Technologies Research (IJEETR), 6(6), 9305-9309.
13. Bandaru, P. K. (2023). Validation methodologies for over-the-air software updates in modern automotive platforms. International Journal of Computer Technology and Electronics Communication (IJCTEC), 6(6), 8159–8165.
14. Akiri, C. K., Jayabalan, K., Lopes, J., Kareem, S. A., & Tabbassum, A. (2025, March). Generative AI for real-time cloud security: Advanced anomaly detection using GPT models. In 2025 IEEE Conference on Computer Applications (ICCA) (pp. 1-6). IEEE.
15. Gopakumar, S. (2026, April). Tenancy-Aware AI Automation for B2B SaaS Admin Workflows. In 2026 IEEE International Conference on Smart Sustainable Systems for Computer and Engineering Applications (3SCEA) (pp. 77-83). IEEE.
16. Vemireddy, S. (2024). Secure and scalable intelligent service architectures for next-generation enterprise applications. International Journal of Engineering & Extended Technologies Research (IJEETR), 6(4), 8177–8182.
17. Tyagi, N. (2025). Privacy Preserving AI in Financial Sector-Balancing Utility, Security and Compliance. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 8(5), 12795-12802.
18. Padmanabham, S. (2022). Secure enterprise architecture for pharmacy benefit management platforms. International Journal of Engineering & Extended Technologies Research, 4(5), 5381–5386.
19. Naraparedy, V. S. R., & Yerramilli, S. K. (2024). Devops Compliance-as-Code. Universal Library of Engineering Technology., 01 (02), 47–54.
20. Anand, L. (2024). AI-Powered Cloud Cybersecurity Architecture for Risk Prediction and Threat Mitigation in Healthcare and Finance. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 7(Special Issue 1), 5-12.

ISSN: 3105 - 9503 International Journal of Multidisciplinary Research in Science, Engineering, Technology and Management (IJMRSETM)

21. Bajarang Bhagwat, V. (2023). Optimizing payroll to general ledger reconciliation: Identifying discrepancies and enhancing financial accuracy. Journal of Advance and Future Research, 1(4).
22. Wadhwa, R. (2023). Ensuring data consistency in enterprise systems through event driven microservices and distributed transaction management. International Journal of Computer Science and Engineering Research and Development, 6(1), 24-51.
23. Jeyaraj, S. A. L., Kumar, S., Revathy, S., Yenigalla, G., Krishna, K. B., & Jayabalan, K. (2024). Machine Learning Algorithms for E-Commerce Security: A Practical Approach. In Strategies for E-Commerce Data Security: Cloud, Blockchain, AI, and Machine Learning (pp. 361-385). IGI Global Scientific Publishing.
24. Narra, S. L. (2025). The Future of Endpoint Security: Autonomous Agents and Self-Healing Systems. Journal Of Multidisciplinary, 5(7), 109-117.
25. Jayaraman, S., Rajendran, S., & P, S. P. (2019). Fuzzy c-means clustering and elliptic curve cryptography using privacy preserving in cloud. International Journal of Business Intelligence and Data Mining, 15(3), 273-287.
26. Raja, G. V. (2020). Metadata gets a makeover: The machine learning approach. International Journal of Computer Technology and Electronics Communication, 3(6), 2900-2903.
27. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. Indian Journal of Science and Technology, 9, 44.
28. Mohammad Ali, M. A., Md Shahadat Hossain, M. S. H., Md Whahidur Rahman, M. W. R., & Md Shahdat Hossain, M. S. H. (2025). AI-Driven Predictive Modeling to Detect and Prevent Financial Fraud in US Digital Payment Systems. AI-Driven Predictive Modeling to Detect and Prevent Financial Fraud in US Digital Payment Systems, 5(12), 228-255.
29. Suddala, V. R. A. K. (2024). Machine learning for operational excellence: Real-world applications. International Journal of Future Innovative Science and Technology (IJFIST), 7(6), 13917.
30. Hossain, M. D., Akter, F., Rahaman, M., Biswas, B., & Hasan, H. M. (2026). Hierarchical Federated Learning with Heavy-Hitter Detection for Real-Time Cyber Threat Mitigation in Large-Scale Networks. Available at SSRN 6340898.
31. Hoque, M. J., Hasan, M. M., Khatun, M. M., Akter, F., & Mohammad, A. R. (2021). Impact of COVID-19 on Consumer Buying Behavior During COVID-19 Pandemic Using Data Analytics. Journal of Business and Management Studies, 3(2), 296-307.
32. Sivakumer, D. (2024). The impact of technology industry layoffs on project managers: An empirical study in the USA. International Journal of Research and Applied Innovations (IJRAI), 7(4), 11166–11177.
33. Chaba, A. (2025). From chatbot to agent: Designing agentic AI for autonomous customer journeys in digital commerce. International Journal of Computer Technology and Electronics Communication, 8(3), 10781–10786.
34. Teja, T. V., Bharadwaj, D., Surabhi, K. R., Pokala, H. K., & Kavithamani, B. (2026, April). AI-Driven Quality of Service in Wireless Sensor Network Using Dueling Double Deep Q-Network with Lyapunov Drift-Plus-Penalty Method for Mobile Networks. In 2026 2nd International Conference on Intelligent Systems and Computational Networks (ICISCN) (pp. 1-6). IEEE.
35. Kundurthy, O. H., Ghadiyaram, R., & Vanam, L. (2025, September). Global AI Regulation Review: Comparative Insights from EU, US and APAC. In International Conference on Intelligent Computing and Communication (pp. 422-434). Cham: Springer Nature Switzerland.
36. Alex Mathew. (2023). Threat defense through cyber fusion. International Journal of Computer Science and Mobile Computing, 12(1), 24–27. https://doi.org/10.47760/ijcsmc.2022.v12i01.003
37. Papachappa, H. M., Kumar, A., & Badam, N. (2026, June). Riemannian Intent Manifolds for Session-Based Search: A Joint Embedding Predictive Architecture for Intent Forecasting. In 2026 15th Mediterranean Conference on Embedded Computing (MECO) (pp. 1-8). IEEE.
38. Cagle, A., & Ahmed, A. M. C. (2024). Architecting enterprise AI applications: A guide to designing reliable, scalable, and secure enterprise-grade AI solutions. Springer.