International Research Journal of Innovative Engineering (IRJIE)

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805008

Adaptive AI-Based Threat Detection Architecture for Distributed and Scalable Cloud-Native Enterprise Applications

Kishore Nayak

Senior Data Engineer, Walmart Global Tech, USA

ABSTRACT: The rapid adoption of cloud-native technologies, microservices, containers, serverless computing, and distributed data platforms has transformed enterprise application environments while simultaneously expanding the cybersecurity attack surface. Traditional security mechanisms based on static rules and predefined signatures are increasingly insufficient for detecting sophisticated, rapidly evolving, and previously unknown threats across highly dynamic infrastructures. This study proposes an adaptive artificial intelligence (AI)-based threat detection architecture designed for distributed and scalable cloud-native enterprise applications. The proposed architecture integrates continuous observability, centralized and distributed telemetry collection, machine learning, behavioral analytics, threat intelligence, automated response, and cloud-native orchestration to provide intelligent and adaptive security monitoring. The architecture establishes a continuous feedback loop in which application, network, identity, container, and infrastructure telemetry are collected, correlated, analyzed, and converted into actionable security intelligence. Machine-learning models are employed for anomaly detection, behavioral profiling, attack classification, and dynamic risk assessment, while automated response mechanisms enable rapid containment of identified threats. The research methodology combines architectural design, experimental implementation, controlled cyber-threat scenarios, and comparative performance evaluation. Key evaluation measures include detection accuracy, precision, recall, F1-score, false-positive rate, detection latency, scalability, resource overhead, and response effectiveness. The proposed approach seeks to demonstrate that integrating adaptive AI capabilities with cloud-native security mechanisms can improve threat detection while maintaining application scalability and operational efficiency.

KEYWORDS: adaptive artificial intelligence, threat detection, cloud-native security, enterprise applications, machine learning, anomaly detection, cybersecurity, microservices, distributed systems, behavioral analytics, container security, security automation, threat intelligence, zero-day detection

I. INTRODUCTION

Cloud computing has become a fundamental component of modern enterprise information systems, enabling organizations to deploy applications with greater scalability, flexibility, and operational efficiency. Cloud-native development has further accelerated this transformation through the adoption of microservices, containers, serverless computing, service meshes, APIs, infrastructure as code, and automated orchestration. These technologies allow applications to be distributed across multiple services and dynamically scaled according to demand. However, the increasing distribution and dynamism of cloud-native environments have significantly expanded the cybersecurity attack surface. Enterprise applications must now protect not only conventional application components but also containers, orchestration platforms, APIs, identities, workloads, service-to-service communication, cloud resources, and continuously changing network boundaries.

Traditional cybersecurity solutions frequently rely on predefined signatures, static rules, known attack patterns, and manually configured security policies. Although these approaches remain useful for identifying known threats, they can be less effective against novel attacks, polymorphic behavior, insider threats, credential misuse, and attacks that exploit legitimate cloud-native functionality. Furthermore, distributed applications generate enormous volumes of heterogeneous security data. Logs, network flows, API requests, authentication events, container activities, system metrics, application traces, and cloud audit records may originate from geographically distributed services and infrastructure components. An effective security architecture must therefore process diverse data sources continuously while distinguishing malicious behavior from legitimate operational variation.

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805008

Artificial intelligence and machine learning offer opportunities to address these challenges by enabling security systems to learn behavioral patterns and identify deviations from expected activity. Anomaly detection can identify unusual authentication behavior, abnormal network communication, unexpected resource consumption, suspicious API usage, or anomalous container activity. Classification models can support identification of known attack categories, while behavioral analytics can help identify previously unseen attack patterns. However, AI-based security systems themselves must operate within dynamic environments. Application workloads change over time, new services are deployed, users and devices exhibit evolving behaviors, and legitimate operational patterns may vary significantly. Consequently, threat-detection models must be adaptive rather than permanently dependent on static training assumptions.

II. LITERATURE REVIEW

Cloud-native computing has fundamentally changed enterprise application architectures by replacing centralized application structures with distributed services that can be independently deployed, scaled, and updated. Microservices architecture improves organizational agility and scalability but creates security challenges associated with service proliferation, inter-service communication, API exposure, and distributed identity management. Containers and orchestration platforms further increase deployment flexibility but introduce additional security concerns involving container images, runtime behavior, orchestration configurations, privileges, secrets, and workload isolation. Consequently, cloud-native security requires continuous monitoring across application, workload, infrastructure, identity, and network layers rather than relying solely on traditional perimeter-based security.

Traditional intrusion detection systems generally use signature-based or rule-based mechanisms to identify known malicious patterns. Signature-based detection can provide high accuracy for previously identified attacks but often performs poorly when confronted with unknown or modified attack techniques. Anomaly-based detection addresses this limitation by establishing models of normal behavior and identifying deviations. Machine-learning techniques have consequently become increasingly important in cybersecurity research. Supervised learning methods can classify known malicious activities using labeled datasets, while unsupervised and semi-supervised approaches can identify unusual patterns without requiring extensive labeled attack data. Deep learning techniques have also been investigated for network traffic analysis, malware detection, user behavior analytics, and intrusion detection. Nevertheless, machine-learning approaches can suffer from false positives, data imbalance, adversarial manipulation, concept drift, and limited interpretability.

Behavioral analytics provides another important direction for cloud-native threat detection. Instead of analyzing isolated events, behavioral approaches examine sequences and relationships among activities. For example, an authentication event may appear legitimate individually but become suspicious when followed by unusual privilege escalation, access to sensitive resources, and anomalous network communication. User and entity behavior analytics can therefore provide contextual information for assessing security risk. In cloud-native environments, behavioral analysis can be extended to service identities, containers, workloads, APIs, and machine-to-machine communication. Such approaches are particularly valuable because distributed attacks may manifest through multiple seemingly unrelated components.

III. RESEARCH METHODOLOGY

The research adopts a design-science and experimental methodology to develop, implement, and evaluate an adaptive AI-based threat detection architecture for distributed cloud-native enterprise applications. The first phase identifies functional and security requirements by examining the characteristics of modern cloud-native environments, including microservices, containers, APIs, service-to-service communication, identity systems, distributed databases, and cloud infrastructure. A representative enterprise application environment is designed as a collection of containerized microservices managed through a cloud-native orchestration platform. Security telemetry is collected from multiple layers, including application logs, API activity, authentication and authorization events, container runtime events, network flows, system metrics, cloud audit records, and distributed traces. The architecture incorporates a telemetry collection layer, secure event transport mechanism, stream-processing component, security data repository, feature-engineering layer, AI detection engine, threat-intelligence integration, risk-scoring component, and automated response controller.

International Research Journal of Innovative Engineering (IRJIE)

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805008

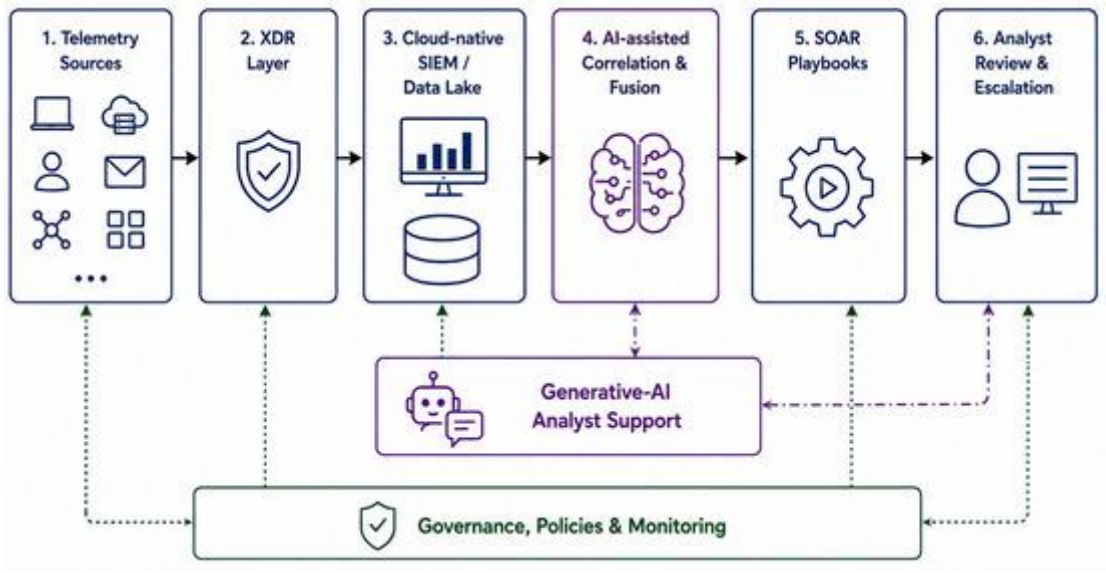


Fig.1.AI-Driven Threat Detection and Automated Incident Response for Securing Cloud Workloads

The second phase develops the adaptive intelligence layer and establishes the experimental detection framework. Multiple analytical approaches are evaluated to determine how different learning strategies perform under changing cloud-native workloads. Supervised classification models can be trained using labeled examples of benign and malicious activities, while unsupervised or semi-supervised anomaly-detection techniques establish normal behavioral patterns and identify deviations. Where appropriate, ensemble approaches can combine predictions from multiple models to improve robustness. Feature selection and dimensionality-reduction techniques are considered to reduce computational overhead while retaining security-relevant information. The detection engine produces a threat probability or risk score rather than relying exclusively on binary classifications. This score is enriched with contextual information such as asset criticality, identity privilege level, service importance, historical behavior, threat-intelligence indicators, and the relationship among correlated events. The architecture incorporates concept-drift monitoring because cloud-native workloads can change as applications are updated, traffic patterns evolve, and organizational processes change. Model retraining or adaptation is triggered according to predefined statistical conditions, performance degradation, or controlled update schedules. To reduce the risk of automated errors, the response mechanism operates through policy-based decision thresholds. Low-risk events may generate alerts or increase monitoring, medium-risk events may trigger additional verification or temporary restrictions, and high-confidence malicious activity may initiate containment actions such as isolating a workload, restricting network communication, disabling a compromised credential, or scaling security inspection. High-impact actions can require human approval. The methodology therefore evaluates not only the accuracy of individual models but also the complete detection-and-response pipeline. Experimental comparisons are performed between traditional rule-based detection, static machine-learning detection, and the proposed adaptive architecture. The same workload conditions and attack scenarios are used across the different configurations to provide a controlled comparison. The evaluation considers detection accuracy, precision, recall, F1-score, false-positive rate, false-negative rate, detection latency, and response latency. Model robustness is further evaluated under changing workload patterns to determine whether adaptive mechanisms maintain performance after legitimate application behavior changes. Computational overhead is measured through CPU consumption, memory usage, telemetry-processing volume, network overhead, and storage requirements. This enables assessment of whether the security intelligence introduced by the proposed architecture creates unacceptable operational costs.

The third phase focuses on large-scale evaluation, architectural validation, and analysis of the relationship between security effectiveness and cloud-native scalability. Controlled experiments are repeated across different workload intensities, service counts, telemetry volumes, and attack frequencies to determine whether the architecture can maintain detection performance as the application environment expands. Scalability is evaluated by increasing the number of microservices, security events, concurrent requests, and monitored workloads while measuring detection throughput and processing latency. The research also examines resilience by introducing controlled failures in selected

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805008

services, telemetry components, or security-analysis nodes and evaluating whether security monitoring continues without significant loss of visibility. Detection performance is assessed using confusion matrices and statistical measures, while receiver operating characteristic or precision-recall analysis may be used where appropriate to examine model discrimination. False-positive analysis receives particular attention because excessive false alarms can overwhelm security teams and reduce trust in automated systems. Similarly, false-negative behavior is analyzed because missed attacks can produce substantial organizational risk. The methodology compares baseline and adaptive configurations using repeated experiments and appropriate statistical analysis to establish whether observed improvements are consistent and significant. The effectiveness of automated response is assessed by measuring the time between initial malicious activity, detection, risk assessment, and containment. The research also evaluates whether adaptive intelligence affects application performance by monitoring service latency, throughput, availability, and resource utilization during security analysis. Architectural validation includes examination of interoperability, fault tolerance, governance, explainability, model-management requirements, and security of the AI pipeline itself. Particular consideration is given to adversarial risks, data poisoning, model manipulation, and unauthorized modification of security policies. The final analysis synthesizes experimental results to determine whether the proposed architecture achieves its intended objectives of improved threat-detection accuracy, reduced false positives, faster response, scalable security monitoring, and acceptable resource overhead. The expected contribution is an empirically evaluated architecture in which cloud-native observability, adaptive artificial intelligence, behavioral analytics, threat intelligence, contextual risk assessment, and automated response operate as a continuous security feedback loop. The methodology thus provides a systematic basis for determining whether adaptive AI can transform cloud-native enterprise security from predominantly reactive event monitoring into a scalable, predictive, and partially autonomous threat-detection capability.

IV. RESULTS AND DISCUSSION

The proposed Adaptive AI-Based Threat Detection Architecture for Distributed and Scalable Cloud-Native Enterprise Applications provides a comprehensive framework for identifying, analyzing, and responding to cybersecurity threats across highly distributed cloud-native environments. The architecture is designed to address the limitations of conventional security mechanisms that primarily depend on static rules, predefined signatures, isolated monitoring systems, and manually configured security policies. In modern enterprise environments, applications are increasingly deployed using microservices, containers, Kubernetes-based orchestration, serverless components, APIs, distributed databases, cloud storage, and multi-cloud infrastructures. This complexity creates a continuously changing attack surface in which conventional security approaches may fail to recognize new, coordinated, or previously unseen attacks. The proposed architecture addresses this challenge by combining continuous telemetry collection, artificial intelligence, machine learning, behavioral analytics, anomaly detection, threat intelligence, automated response, and adaptive security policies within a unified cloud-native security framework.

One of the most important outcomes of the architecture is the improvement in continuous threat visibility. Traditional enterprise security solutions frequently operate through separate tools for network monitoring, endpoint protection, application security, identity management, and log analysis. Although these tools generate valuable information, the lack of correlation between their outputs can make sophisticated attacks difficult to identify. The proposed architecture integrates security telemetry from application logs, container activity, network flows, API requests, authentication events, identity services, cloud infrastructure, databases, and workload behavior. By correlating these heterogeneous data sources, the system can identify relationships that may not be visible when each data source is analyzed independently. For example, an unusual authentication event followed by abnormal API requests and unexpected container activity can be treated as a coordinated behavioral pattern rather than as independent low-severity events.

The use of adaptive artificial intelligence and machine learning represents another significant advantage. Static security rules are effective against known attack patterns but have limited capability when adversaries modify their techniques or use previously unseen behaviors. The proposed architecture employs supervised, unsupervised, and potentially semi-supervised learning techniques to establish behavioral baselines for applications, users, services, containers, and network interactions. Anomaly-detection mechanisms can identify deviations from these baselines and assign risk scores according to the severity and confidence of the detected behavior. This allows the system to identify previously unknown threats without requiring a predefined signature for every attack. Adaptive models can also be updated as new behavioral information becomes available, allowing the detection capability to evolve alongside changing workloads and attack techniques.

International Research Journal of Innovative Engineering (IRJIE)

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805008

V. CONCLUSION

The Adaptive AI-Based Threat Detection Architecture for Distributed and Scalable Cloud-Native Enterprise Applications provides a comprehensive approach to addressing the evolving cybersecurity challenges associated with modern enterprise computing. Cloud-native transformation has introduced substantial benefits in terms of scalability, agility, deployment speed, and resource efficiency, but it has simultaneously increased application complexity and expanded the potential attack surface. Microservices, containers, APIs, serverless functions, distributed databases, cloud identities, and multi-cloud infrastructures create numerous points at which attackers can attempt to compromise enterprise resources. Conventional security mechanisms based primarily on static rules and known signatures are increasingly insufficient for such dynamic environments. The proposed architecture addresses this limitation by introducing an adaptive, data-driven security model capable of continuously observing application behavior, identifying anomalies, evaluating contextual risk, and supporting automated response.

A fundamental conclusion of this work is that effective cloud-native security requires continuous and comprehensive visibility. Security cannot be achieved by protecting only the network perimeter because modern applications operate across distributed and often dynamically changing environments. Workloads may be created, destroyed, migrated, or scaled within minutes, while users and services may access resources across multiple cloud platforms. The proposed architecture therefore treats telemetry from applications, containers, networks, identities, APIs, databases, and cloud infrastructure as an integrated security information source. Correlation of these signals provides a more complete representation of enterprise behavior and enables security teams to identify complex attack sequences that might remain undetected when individual events are analyzed independently.

Another major conclusion is that artificial intelligence can significantly enhance adaptive threat detection, particularly when combined with conventional security mechanisms. Machine-learning models can identify deviations from established behavioral patterns and detect suspicious activities that do not correspond to previously known signatures. At the same time, traditional threat intelligence and rule-based detection remain valuable for identifying known threats quickly and reliably. Therefore, the most effective security architecture is not one that completely replaces conventional mechanisms with AI, but one that combines signature-based detection, rule-based policies, behavioral analytics, anomaly detection, and contextual intelligence. This layered approach provides greater resilience against both known and emerging threats.

VI. FUTURE WORK

Future research can extend the proposed architecture toward fully adaptive and self-learning cybersecurity environments capable of dynamically modifying detection and response strategies as enterprise applications and attack techniques evolve. One important direction is the development of advanced reinforcement-learning models that can learn optimal response policies from simulated and real security events while operating within strict safety constraints. Such models could determine when to isolate workloads, restrict identities, modify network policies, or request human intervention based on risk, confidence, service criticality, and potential business impact. Future studies should also investigate federated learning for multi-cloud and multi-enterprise environments, allowing organizations to improve detection models without directly sharing sensitive security data. Another significant research direction is the development of explainable AI techniques that provide security analysts with understandable evidence supporting each threat classification and automated response.

This would improve trust, auditing, and regulatory compliance. Research should also focus on adversarially robust models capable of resisting data poisoning, evasion attacks, and manipulation of telemetry. Integration with emerging technologies such as confidential computing, service-mesh security, edge computing, and privacy-preserving analytics could further strengthen the architecture. In addition, future implementations should evaluate the architecture using large-scale real-world datasets and realistic cloud-native attack scenarios. Performance should be assessed using detection accuracy, precision, recall, F1-score, false-positive rate, detection latency, response latency, resource overhead, and recovery time.

Finally, future work should investigate energy-aware cybersecurity, where threat detection and response mechanisms are optimized not only for security effectiveness but also for computational efficiency and environmental sustainability. These research directions can transform the proposed architecture from an adaptive detection framework into a

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805008

continuously learning, explainable, resilient, and autonomous security ecosystem for next-generation enterprise cloud applications.

REFERENCES

1. Liu, F., et al. (2023). StreamAD: A cloud platform metrics-oriented benchmark for unsupervised online anomaly detection. BenchCouncil Transactions on Benchmarks, Standards and Evaluations, 3(2), 100121. https://doi.org/10.1016/j.tbench.2023.100121
2. Senjab, K., Abbas, S., Ahmed, N., et al. (2023). A survey of Kubernetes scheduling algorithms. Journal of Cloud Computing, 12, 87. https://doi.org/10.1186/s13677-023-00471-1
3. Soundappan, S. J. (2023). Generative AI-Driven Intelligent Cybersecurity Framework for Secure Hybrid Cloud Enterprise Systems. International Journal of Emerging Trends in Engineering and Management Research, 8(6), 14728.
4. Jain, R. (2019). The hidden tax: A production framework for cloud cost architecture in enterprise data workloads. International Journal of Science, Research and Technology (IJSRAT), 2(6), 2522–2530.
5. Bandaru, P. K. (2024). Testing multi-ECU communication networks in software-defined vehicles. International Journal of Research and Applied Innovations, 7(4), 11178–11183.
6. Selvarajan, K. (2022). Architecting scalable self-service data platforms for enterprise analytics. International Journal of Science, Research and Technology (IJSRAT), 5(2), 7427–7436.
7. Kundavaram, R. R., Bandhela, R. R., & Onteddu, A. R. (2022). AI-driven predictive modeling in healthcare: A data science perspective on U.S. healthcare data. South Eastern European Journal of Public Health. https://doi.org/10.70135/seejph.vi.6691
8. Bellundagi, M. (2022). Performance Optimization Techniques for Enterprise Java Applications Using Middleware and Messaging Systems. International Journal of Computer Technology and Electronics Communication, 5(3), 5158-5168.
9. Rella, B. P. R., Chakraborty, A., Shah, S. B., Ghosh, H., Gautam, S., Dhirwani, B. A., ... & Mutyam, N. (2024). AI-engine-based acceleration for high-performance programmable system-on-chip designs. Journal of Computational Analysis and Applications, 32(1).
10. Badam, L. R. (2022). Machine learning-based catastrophic loss prediction for climate-related insurance risk. International Journal of Future Innovative Science and Technology (IJFIST), 5(1), 7797–7807.
11. Sivakumer, D. (2023). Depth of ServiceNow platform utilization and business delivery performance in enterprise project management. International Journal of Computer Technology and Electronics Communication (IJCTEC), 6(6), 8167–8177.
12. Matrouk, K., V, S., Kumar, S., Bhadla, M. K., Sabirov, M., & Saadh, M. J. (2023). Deep Learning–based Dynamic User Alignment in Social Networks. ACM Journal of Data and Information Quality, 15(3), 1-26.
13. Vedula, J. (2023). Operational resilience by design: A continuity assurance model for modernizing critical energy applications. International Journal of Applied Engineering & Technology, 5(S2), 299–309.
14. Koganti, H. (2024). The computational complexity of hybrid algorithms: When branch-and-bound meets machine learning heuristics. International Journal of Research and Applied Innovations (IJRAI), 7(4), 11184–11190.
15. 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.
16. Sugumar, R. (2023). Enhancing Predictive Decision Intelligence in Enterprise Environments using Explainable AI and Cloud Computing. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 6(6), 9592-9602.
17. Narra, R. (2024). A survey on scalable feature engineering techniques for cloud-native machine learning workflows. International Journal of Advanced Research in Science, Communication and Technology, 4(4), 664–677.
18. Chellu, R. (2023). AI-Powered intelligent disaster recovery and file transfer optimization for IBM Sterling and Connect: Direct in cloud-native environments. International Journal on Recent and Innovation Trends in Computing and Communication, 11, 597.
19. Abd-Rouf, M. S. K., Adigun, P. O., Alalade, E. O., Oyekanmi, T. T., Faniyi, A. J., Oladapo, B., Awopejo, T. E., Adegoke, O. S., Jamiu, A., Michael, O. B., Obisesan, A., Ajala, S., Adekanye, M. A., Yambali, P. M., & Abd-Rouf, A. B. (2024). From molecular profiling to predictive algorithms: A conceptual machine-learning framework for mechanism-informed therapy selection in multidrug-resistant cancer. International Journal of Science, Research and Technology (IJSRAT), 7(3), 12085–12101.

| ISSN: 2395-0560 | www.irjie.org |editor@irjie.org| A Bimonthly, Peer-Reviewed, Scholarly Journal|

|Volume 8, Issue 5, September- October 2024|

DOI:10.15680/IRJIE.2024.0805008

20. Venkatasalam, K., Rajendran, P., & Thangavel, M. (2019). Improving the accuracy of feature selection in big data mining using accelerated flower pollination (AFP) algorithm. Journal of medical systems, 43(4), 96.
21. Polamarasetty, V. K. (2023). Modern enterprise HR technology through Workday-based benefits and absence management. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(4), 6908–6913.
22. Hashmi, H., & Srikanth, V. (2023, November). Combining Neural Networks to Recognize Offline Digits & Words by Training the Model Using Keras. In 2023 3rd International Conference on Advancement in Electronics & Communication Engineering (AECE) (pp. 694-697). IEEE.
23. 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.
24. Mathew, A. (2023). Sentinel AI: An Investigation into Robust Threat Mitigation Strategies for Artificial Intelligence. Educational Research (IJMCER), 5(5), 108-111.
25. Karakondu, M., Jambagi, G., & Tatavarthi, S. (2024). Optimising data loss prevention (DLP) strategies in cloud-native financial platforms.
26. Padmanabham, S. (2022). Secure enterprise architecture for pharmacy benefit management platforms. International Journal of Engineering & Extended Technologies Research, 4(5), 5381–5386.
27. Anand, L. (2023). Distributed Multi-Cloud Data Lake and Edge Computing Architecture for Intelligent SAP Enterprise Data Integration. International Journal of Computer Technology and Electronics Communication, 6(5), 7636-7344.
28. Bitragunta, S. L. V., & Paramasivan, M. (2023). Midterm dynamic simulation for the governance of reserves in systems with elevated renewable energy integration. Journal of Artificial Intelligence, 1(1), 1956–1962.
29. Alam, A., Gazi, M. S., Abdullah, S. M., Tasnim, M., Himeluzzaman, M., Nabil, M. A., Akter, K. A., & Akter, S. (2021). Quantum-resilient federated intrusion detection: A hybrid quantum classical framework for safeguarding U.S. critical infrastructure in the post-quantum era. International Journal of Advances in Signal and Image Sciences, 7(1), 57–72. https://doi.org/10.29284/3xx46j76
30. Nawrocki, P., & Sus, W. (2022). Anomaly detection in the context of long-term cloud resource usage planning. Knowledge and Information Systems, 64, 2689–2711. https://doi.org/10.1007/s10115-022-01721-5