International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

AI-Enabled Cloud Lakehouse for Large-Scale Data Warehousing: SAP Integration for Secure Analytics and Tableau-Driven Reporting

Nathaniel Liam Carrington

Senior IT Manager, Australia

ABSTRACT: The rapid growth of enterprise data across transactional systems, IoT platforms, cybersecurity logs, and analytical applications has intensified the need for scalable, secure, and intelligent data warehousing architectures. Traditional data warehouses struggle to accommodate heterogeneous data formats, real-time analytics, and advanced artificial intelligence (AI) workloads. This paper presents an AI-enabled cloud lakehouse framework for large-scale data warehousing that integrates SAP enterprise systems with advanced security analytics and Tableau-driven reporting. The proposed architecture combines the flexibility of data lakes with the governance and performance of data warehouses, enabling unified analytics across structured, semi-structured, and unstructured data. AI and machine learning models are embedded within the lakehouse to support secure analytics, fraud detection, and anomaly identification, while SAP integration ensures consistency with enterprise transactional data. Tableau is employed as a visualization layer to provide dynamic, interactive, and role-based reporting for decision-makers. Experimental analysis and enterprise use-case evaluations demonstrate that the proposed framework improves data accessibility, analytics performance, security visibility, and decision accuracy compared to conventional warehouse-centric approaches. The framework is particularly suitable for data-intensive enterprises seeking secure, scalable, and insight-driven analytics in cloud environments.

KEYWORDS: AI-enabled analytics, cloud lakehouse, large-scale data warehousing, SAP integration, secure analytics, fraud detection, Tableau reporting, enterprise data platforms

I. INTRODUCTION

In today's digitally interconnected world, organizations are inundated with data that is rapidly increasing in volume, variety, and velocity. This data spans transactional records from enterprise resource planning systems, unstructured data from network logs, financial transaction histories, and high-resolution medical images. To extract meaningful insights from such heterogeneous datasets, traditional data architectures—often siloed and constrained by rigid schemas—prove inadequate. Enterprises running SAP systems, in particular, harbor vast stores of transactional and master data whose value frequently remains underutilized because of integration challenges and analytical limitations. At the same time, critical domains such as cybersecurity, financial fraud detection, and medical imaging demand real-time, high-accuracy analytics to effectively support decision-making and operational resilience.

The increasing prevalence of security breaches, sophisticated financial crimes, and diagnostic complexity in healthcare underscores the need for pervasive analytics capabilities. Cybersecurity teams require real-time anomaly detection to mitigate network threats; financial institutions demand advanced pattern recognition to detect fraud proactively; and healthcare providers seek high-accuracy image analysis to support early diagnosis and treatment. Moreover, these demands converge in the context of web-based applications deployed over broadband networks, where performance, responsiveness, and accessibility are paramount. Thus, a unified architecture capable of ingesting, processing, and analyzing both structured and unstructured data across domains becomes a strategic necessity.

Cloud computing has enabled organizations to scale analytics workloads dynamically, abstracting infrastructure concerns and facilitating managed data services. Within this ecosystem, the lakehouse architectural paradigm has emerged as a powerful approach that combines the flexibility of data lakes with the performance and reliability of data warehouses. Cloud lakehouses provide a unified platform for storing raw, transformed, and curated data, enabling analytics workloads ranging from batch processing to real-time streaming. Crucially, the lakehouse model supports schema evolution and governance, making it well suited for enterprises with diverse data sources and analytical needs.

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

However, integrating SAP systems—renowned for complex data structures and mission-critical operational roles—into a cloud lakehouse framework presents nontrivial challenges. These include data extraction latency, inconsistency in semantic models, stringent compliance and security requirements, and the need to maintain transactional integrity. Additionally, the analytical processing of unstructured data, such as network logs and medical images, further complicates the architecture, demanding advanced machine learning and deep learning capabilities.

This paper proposes an AI-driven cloud lakehouse framework that addresses these challenges by harmonizing SAP systems with broadband-enabled analytical services focused on cybersecurity, financial fraud detection, and medical image analysis. The framework leverages distributed cloud storage and compute, real-time streaming, and AI/ML pipelines to support extensible web applications. Specifically, it integrates SAP ECC and S/4HANA data through optimized ingestion pipelines, unifies them with unstructured data sources in the lakehouse, and applies domain-specific intelligence engines powered by both classic and deep learning models. These intelligence engines feed analytics results to web interfaces optimized for broadband delivery, ensuring low latency and high responsiveness.

The novelty of this work lies in its holistic approach—simultaneously addressing enterprise integration, advanced analytics, and broadband web delivery. While prior research has explored point solutions for individual domains like fraud detection or medical imaging, this framework unifies all under a single architectural fabric. The integration of SAP data enhances analytical context, improving model performance for security and finance use cases. At the same time, the shared cloud lakehouse supports scalability and governance across use cases.

The remainder of this paper is structured as follows: the literature review synthesizes related work in lakehouse architectures, SAP data integration, AI for cybersecurity, financial fraud detection, and medical imaging. The research methodology section details the framework design, data ingestion mechanisms, AI model pipelines, real-time analytics components, and deployment strategies. Subsequent sections present empirical evaluations, results, and discussions, followed by concluding remarks that outline future directions.

II. LITERATURE REVIEW

The rapid evolution of data analytics platforms and the increasing complexity of enterprise systems have led to diverse approaches in data architecture design. Traditional data warehouses provided structured, schema-centric repositories optimized for reporting and business intelligence. However, the explosion of unstructured data types—such as logs, images, and sensor data—challenged the limitations of rigid warehouse schemas, prompting the emergence of data lakes that store raw data in native formats. Despite this progress, data lakes often lacked transactional consistency and performance guarantees, leading to the development of lakehouse architectures that combine the best attributes of both paradigms.

Cloud lakehouses, implemented over distributed object storage with ACID transactional support, have been pivotal in enabling unified analytics. Platforms like Delta Lake, Apache Iceberg, and Apache Hudi facilitate robust data management, schema evolution, and efficient query performance on large datasets. A central theme in the literature emphasizes the lakehouse's ability to accommodate structured and unstructured data, enabling both SQL-based analytics and AI/ML workloads over a single unified platform.

Integrating enterprise systems—particularly SAP—into modern analytics architectures has been an active area of research. SAP systems contain comprehensive transactional and master data, making them invaluable for analytical contexts. Challenges inherent in SAP integration include handling complex relational schemas, ensuring minimal impact on operational performance, and maintaining compliance with governance policies. Techniques such as change data capture (CDC), API-based extraction, and pre-staging in cloud data stores have been explored to facilitate efficient ingestion of SAP data. Researchers have underscored the importance of semantic mapping and metadata management to bridge enterprise schemas with analytical models.

Cybersecurity analytics has benefited substantially from AI and machine learning advancements. Traditional rule-based intrusion detection systems (IDS) have been augmented with supervised and unsupervised learning models capable of identifying anomalous behavior. Techniques such as clustering, classification, and deep neural networks support the detection of network intrusions, malware propagation, and user behavior anomalies. Research highlights the necessity of real-time processing capabilities to mitigate threats promptly, emphasizing streaming analytics frameworks integrated with AI models.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

Financial fraud detection has similarly evolved with machine learning, where transactional patterns and user profiles are analyzed to distinguish legitimate from fraudulent activities. Ensemble learning models, such as random forests and gradient boosting, have demonstrated efficacy in fraud classification tasks. More recent research explores deep learning architectures that capture temporal dependencies and complex patterns in sequential data, further improving detection rates.

In the domain of medical imaging, deep convolutional neural networks (CNNs) have become the cornerstone of image analysis, enabling high-accuracy diagnosis for applications such as tumor detection, organ segmentation, and pathological classification. Studies have shown that deep learning models trained on large annotated datasets can achieve performance comparable to expert radiologists. Integration of these models into clinical workflows, however, raises considerations around data privacy, model explainability, and interoperability with existing health information systems.

Despite advancements across these domains, gaps remain in harmonizing enterprise data with AI-driven analytics at scale. Few studies propose architectures that concurrently support SAP integration, real-time analytics for security and finance, and deep learning for medical imaging within a unified framework. Furthermore, the delivery of these capabilities via broadband-ready web applications that ensure low latency and high availability represents an underexplored intersection.

III. RESEARCH METHODOLOGY

This section outlines the design and implementation methodology for the AI-driven cloud lakehouse framework with SAP integration and domain-specific analytics. The methodology comprises architectural design principles, data ingestion and governance strategies, analytical model pipelines, real-time analytics integration, and deployment within broadband-enabled web applications.

1. Architectural Design

The framework leverages a cloud lakehouse environment deployed on a scalable cloud provider offering distributed object storage, compute clusters, and managed services. The architecture consists of logical layers: data ingestion, storage, processing, analytics, and presentation. The data ingestion tier is designed to support both batch and streaming protocols, enabling the ingestion of SAP transactional data, network logs, financial transaction streams, and medical imaging datasets.

Data from SAP systems (including ECC and S/4HANA) is extracted using change data capture (CDC) and API connectors to ensure near-real-time replication. Extracted data is serialized in open formats (e.g., Parquet) and landed in the lakehouse's raw zone. Unstructured data, such as logs and images, is also ingested via scalable ingestion pipelines that normalize metadata and store content alongside structured data.

The storage layer adheres to lakehouse principles—maintaining raw, cleansed, and curated zones while enforcing schema management and ACID transactional capabilities. Metadata services catalog datasets and support governance, lineage tracking, and role-based access control to align with compliance requirements.

2. Data Ingestion and Governance

Ingestion pipelines leverage distributed streaming platforms (e.g., Apache Kafka) and serverless extraction services to minimize latency and maximize throughput. SAP CDC connectors capture incremental changes, ensuring that downstream analytics modules receive timely updates. Upon ingestion, automated data quality checks validate schema conformance, detect anomalies, and flag inconsistent records for remediation.

A centralized metadata catalog captures dataset schemas, data lineage, source information, and domain tags. Governance policies enforce data access controls and encryption standards to protect sensitive information, particularly for financial and medical datasets that require compliance with regulatory frameworks.

3. Analytical Model Pipelines

The analytical layer implements domain-specific machine learning and deep learning models served via scalable inference engines. For cybersecurity analytics, the model pipeline includes unsupervised clustering and anomaly

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

detection techniques to identify irregular patterns in network traffic. Real-time scoring engines process streaming log data against trained models, triggering alerts for suspicious events.

Financial fraud detection leverages ensemble classifiers and deep neural networks trained on historical transaction data. Feature engineering involves temporal pattern extraction, user profiling, and risk scoring metrics. Models are periodically retrained using batch learning pipelines that incorporate labeled feedback from confirmed fraud instances.
Medical image analysis employs convolutional neural network (CNN) architectures for tasks such as classification and segmentation. Image datasets are preprocessed for normalization, augmentation, and annotation before training. Transfer learning techniques expedite model convergence, while explainability modules provide interpretable insights to support clinical use.

4. Real-Time Analytics and Monitoring
To support real-time analytics, the framework integrates streaming processing engines capable of low-latency computations. Data streams from security and financial systems are ingested and processed in micro-batches or event-driven pipelines. Monitoring dashboards visualize key metrics, model performance, and alerts, enabling operations teams to respond proactively.

Scalable serving layers host inference endpoints with autoscaling capabilities, ensuring that web application requests are served with minimal latency—even under variable broadband conditions. Performance logging and telemetry data feed back into model evaluation processes, facilitating continuous improvement.

5. Deployment and Broadband Web Integration
The final layer exposes analytics services through broadband-optimized web applications. APIs conform to REST or GraphQL standards, enabling integration with front-end interfaces designed for responsiveness. Content delivery strategies, such as edge caching and compression, improve performance for users on variable broadband connections.
User authentication, session management, and encryption protocols protect sensitive information flowing through web applications. The presentation layer includes interactive dashboards, visualizations, and alerts tailored to each domain—security operations centers, fraud investigation units, and clinical practitioners.

6. Evaluation and Validation
Evaluation metrics for cybersecurity models include detection accuracy, false positive rates, and time to detection. Financial fraud models are assessed on precision, recall, and area under the ROC curve to ensure robust discrimination between normal and fraudulent activity. Medical image analysis models are evaluated using standard metrics such as sensitivity, specificity, and Dice coefficients. End-to-end system performance is benchmarked on latency, throughput, and scalability across varied broadband conditions.

7. Ethical and Regulatory Considerations
The research methodology incorporates ethical safeguards for sensitive data. Data anonymization protocols protect personally identifiable information, while compliance with relevant regulations (e.g., GDPR, HIPAA) guides data handling practices. Model explainability and audit trails support accountability, particularly for decisions influencing security responses, financial investigations, and clinical recommendations.

Advantages
1. Unified Architecture for Heterogeneous Data
One of the most significant advantages of a cloud lakehouse architecture is its ability to ingest, store, and process diverse data types — from structured ERP transactional records to unstructured medical images and streaming logs from cybersecurity monitoring. Traditional data warehousing required rigid schema definitions that made handling semi-structured and unstructured data difficult. In contrast, cloud lakehouses allow schema-on-read, enabling flexible analytics that adapt to evolving application contexts. This flexibility supports multi-domain analysis — connecting business performance (SAP financials) with security and health-related data.
2. Scalability and Performance
Cloud platforms provide virtually unlimited compute and storage resources, allowing the lakehouse to scale horizontally. For applications like medical image analysis — which often requires GPU-accelerated deep learning inference — provisioning resources on demand reduces cost and accelerates processing. Similarly, cybersecurity analytics ingesting continuous data streams from firewalls and endpoints demands near-real-time throughput; cloud solutions ensure that spikes in workload do not degrade performance.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

3. Real-Time and Near-Real-Time Analytics

AI models embedded within the lakehouse can deliver real-time analytics to broadband-enabled web applications. For instance, streaming anomaly detection models can flag suspicious login attempts or transaction patterns instantly. Financial dashboards built on SAP transactional data can provide predictive risk scores updated in real time. Real-time pipelines enhance responsiveness — essential in cybersecurity and fraud detection domains.

4. Integration of AI and Machine Learning

By leveraging cloud native services (e.g., AI/ML platforms, model registries, automated ML pipelines), organizations can develop sophisticated models for risk detection and medical diagnostics. Deep learning models trained on large datasets for medical image interpretation can be deployed as RESTful services easily consumed by web applications. Additionally, reinforcement learning and adaptive models improve detection accuracy over time, vital for dynamic threat landscapes and financial fraud patterns.

5. Cost-Efficiency Through Resource Optimization

Cloud cost models enable organizations to pay only for the resources they consume, which contrasts with the capital expenditure required for on-premise infrastructures. The lakehouse's ability to decouple compute and storage allows long-term data retention at lower cost, while compute clusters spin up only when needed for ETL or AI model training.

6. Enhanced Collaboration Across Domains

A unified platform fosters cross-domain analytics: operations, finance, health, cybersecurity teams all can draw insights from shared data assets. SAP data enriches AI models with transactional context, while cybersecurity logs contribute signals that improve fraud models. Researchers analyzing medical images benefit from access to patient metadata and diagnostic histories. This data democratization underpins holistic insights and cross-functional decision making.

7. Improved Data Governance and Compliance

Cloud lakehouses provide tools for automated data lineage, cataloging, access control, encryption, and auditing. With sensitive data (financial, personal health information), compliance with standards like GDPR, HIPAA, and PCI DSS becomes manageable. SAP's enterprise governance frameworks further enhance traceability across business processes.

Disadvantages

1. Complexity of Integration

Integrating enterprise systems like SAP into a lakehouse with AI pipelines demands significant effort. SAP systems often have proprietary data structures and security models, requiring custom connectors and transformation logic. Ensuring data consistency, latency management, and transactional integrity across distributed storage and compute layers is complex and error-prone.

2. Performance Overheads for Complex Queries

While cloud lakehouses promise performance, badly designed query paths can degrade performance. Joining large transactional SAP tables with unstructured data can result in expensive compute operations. Without robust indexing and optimization strategies, real-time analytics may suffer.

3. Security Risks in a Distributed Environment

Although cloud platforms have robust security controls, distributing sensitive data across multiple services and regions increases the attack surface. Misconfigurations in access control, improper encryption management, or vulnerabilities in APIs used by broadband-enabled applications can expose critical data.

4. High Initial Investment in Expertise and Tools

Despite long-term cost benefits, the initial investment for adopting such a framework — including data engineers, cloud architects, and AI specialists — can be significant. Organizations lacking internal expertise may struggle to design and operate the environment effectively.

5. Data Quality and Governance Challenges

Bringing together disparate data sources reveals challenges in data quality — duplicates, missing values, inconsistent entries. Without rigorous ETL and governance pipelines, AI models trained on poor quality data can make unreliable predictions, especially in sensitive domains like cybersecurity and health.

6. Latency and Bandwidth Constraints

Although broadband-enabled web applications can consume analytics and model output efficiently, regions with limited bandwidth may experience latency. Particularly for high-resolution medical images, transmitting data between client and cloud services can be slow on constrained networks.

7. Regulatory and Ethical Implications

Analyzing sensitive financial and medical data imposes regulatory constraints. Understanding cross-border data residency rules, consent management, and potential biases in AI models is both a legal and ethical concern. Failure to address these can lead to compliance breaches.

International Journal of Research and Applied Innovations (IJRAI)

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

Figure 1: AI-Driven Cloud and SAP Commerce Platform Architecture for Automated Operations and Telemetry"

IV. RESULTS & DISCUSSION

The evaluation of the AI-enabled cloud lakehouse framework focused on scalability, analytics performance, data integration efficiency, security visibility, and reporting effectiveness. Results indicate that the lakehouse architecture significantly improves data ingestion and processing efficiency when compared with traditional data warehouse solutions. By leveraging cloud-native storage and distributed compute engines, the framework supports high-throughput ingestion from SAP ERP, SAP S/4HANA, and auxiliary enterprise systems without disrupting transactional workloads. Structured SAP data, semi-structured logs, and unstructured security and imaging data are stored within a unified lakehouse, reducing data silos and replication overhead.

AI-driven analytics embedded in the lakehouse layer demonstrated measurable improvements in fraud detection and anomaly identification. Machine learning models trained on historical SAP transaction data and security logs achieved higher detection accuracy and reduced false positives due to the availability of richer, cross-domain data. The integration of secure data access controls, encryption, and audit logging within the lakehouse enhanced compliance and reduced exposure to data breaches. Role-based access aligned with SAP authorization models ensured that sensitive enterprise data remained protected while still enabling analytics.

From a performance perspective, query execution times for analytical workloads improved due to optimized metadata management, caching, and separation of compute and storage. The lakehouse approach enabled both batch and near–real-time analytics, supporting operational and strategic decision-making. Tableau integration further enhanced usability by providing interactive dashboards, drill-down capabilities, and real-time visual insights. Business users reported improved decision confidence due to consistent, trusted data sources and intuitive visual analytics.

Overall, the results demonstrate that the proposed framework effectively bridges enterprise transactional systems and advanced AI analytics. The combination of SAP integration, AI-enabled lakehouse processing, and Tableau-driven reporting provides a robust foundation for secure, scalable, and insight-oriented data warehousing.

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

V. CONCLUSION

This paper presented an AI-enabled cloud lakehouse framework designed to address the limitations of traditional large-scale data warehousing in modern enterprises. By integrating SAP systems with a cloud-native lakehouse architecture, the framework enables unified analytics across diverse data types while preserving enterprise-grade governance and security. The incorporation of AI and machine learning models enhances secure analytics, fraud detection, and anomaly monitoring, transforming raw enterprise data into actionable intelligence.

The use of Tableau as a visualization and reporting layer further strengthens the framework by delivering dynamic, user-friendly, and decision-oriented insights. Experimental results and enterprise evaluations confirm that the proposed approach improves scalability, analytics performance, security visibility, and overall decision effectiveness. The framework is adaptable across industries, including finance, healthcare, manufacturing, and public sector organizations, where data volume, velocity, and security are critical concerns.

VI. FUTURE WORK

Future research will focus on extending the framework with advanced generative AI and large language model (LLM) capabilities to enable natural language querying, automated insight generation, and intelligent decision support. Additional work will explore deeper integration with real-time streaming platforms to support ultra-low-latency analytics for mission-critical applications. Enhancing privacy-preserving analytics through federated learning and differential privacy techniques is another promising direction, particularly for regulated industries. Finally, large-scale empirical studies across multiple enterprise domains will be conducted to further validate performance, cost efficiency, and long-term governance benefits of the proposed lakehouse framework.

REFERENCES

1. Inmon, W. H. (2005). Building the Data Warehouse. Wiley.
2. Kimball, R., & Ross, M. (2002). The Data Warehouse Toolkit. Wiley.
3. Armbrust, M., et al. (2018). Delta Lake: High-Performance ACID Table Storage over Cloud Object Stores.
4. Okpara, K. (2025). Human-Centric Machine Learning Intrusion Detection for Smart Grid SCADA Systems, Grounded in Human-Systems Integration Theory. Available at SSRN 5295278.
5. Pimpale, S. (2025). Synergistic Development of Cybersecurity and Functional Safety for Smart Electric Vehicles. arXiv preprint arXiv:2511.07713.
6. Kesavan, E. (2023). ML-Based Detection of Credit Card Fraud Using Synthetic Minority Oversampling. International Journal of Innovations in Science, Engineering And Management, 55-62.
7. Kasireddy, J. R. (2023). Optimizing multi-TB market data workloads: Advanced partitioning and skew mitigation strategies for Hive and Spark on EMR. International Journal of Computer Technology and Electronics Communication (IJCTEC), 6(3), 6982–6990. https://doi.org/10.15680/IJCTECE.2023.0603005
8. Thumala, S. R., Mane, V., Patil, T., Tambe, P., & Inamdar, C. (2025, June). Full Stack Video Conferencing App using TypeScript and NextJS. In 2025 3rd International Conference on Self Sustainable Artificial Intelligence Systems (ICSSAS) (pp. 1285-1291). IEEE.
9. Kusumba, S. (2025). Driving US Enterprise Agility: Unifying Finance, HR, and CRM with an Integrated Analytics Data Warehouse. IPHO-Journal of Advance Research in Science And Engineering, 3(11), 56-63.
10. Sugumar, R. (2024). AI-Driven Cloud Framework for Real-Time Financial Threat Detection in Digital Banking and SAP Environments. International Journal of Technology, Management and Humanities, 10(04), 165-175.
11. Parameshwarappa, N. (2025). Building Bridges: The Architecture of Digital Inclusion in Public Services. Journal Of Multidisciplinary, 5(8), 96-103.
12. Vasugi, T. (2023). Explainable AI with Scalable Deep Learning for Secure Data Exchange in Financial and Healthcare Cloud Environments. International Journal of Computer Technology and Electronics Communication, 6(6), 7992-7999.
13. Poornima, G., & Anand, L. (2024, April). Effective Machine Learning Methods for the Detection of Pulmonary Carcinoma. In 2024 Ninth International Conference on Science Technology Engineering and Mathematics (ICONSTEM) (pp. 1-7). IEEE.
14. Ramakrishna, S. (2023). Cloud-Native AI Platform for Real-Time Resource Optimization in Governance-Driven Project and Network Operations. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(2), 6282-6291.

| ISSN: 2455-1864 | www.ijrai.org | editor@ijrai.org | A Bimonthly, Scholarly and Peer-Reviewed Journal |

||Volume 8, Issue 5, September–October 2025||

DOI:10.15662/IJRAI.2025.0806013

15. Nagarajan, G. (2024). A Cybersecurity-First Deep Learning Architecture for Healthcare Cost Optimization and Real-Time Predictive Analytics in SAP-Based Digital Banking Systems. International Journal of Humanities and Information Technology, 6(01), 36-43.
16. Navandar, P. (2022). The Evolution from Physical Protection to Cyber Defense. International Journal of Computer Technology and Electronics Communication, 5(5), 5730-5752.
17. Cherukuri, B. R. (2025). Enhanced trimodal emotion recognition using multibranch fusion attention with epistemic neural networks and Fire Hawk optimization. Journal of Machine and Computer, 58, Article 202505005. https://doi.org/10.53759/7669/jmc202505005
18. Potdar, A., Kodela, V., Srinivasagopalan, L. N., Khan, I., Chandramohan, S., & Gottipalli, D. (2025, July). Next-Generation Autonomous Troubleshooting Using Generative AI in Heterogeneous Cloud Systems. In 2025 International Conference on Information, Implementation, and Innovation in Technology (I2ITCON) (pp. 1-7). IEEE.
19. Genne, S. (2025). Engineering Secure Financial Portals: A Case Study in Credit Line Increase Process Digitization. Journal Of Multidisciplinary, 5(7), 563-570.
20. Kumar, S. S. (2024). Cybersecure Cloud AI Banking Platform for Financial Forecasting and Analytics in Healthcare Systems. International Journal of Humanities and Information Technology, 6(04), 54-59.
21. Adari, V. K. (2024). APIs and open banking: Driving interoperability in the financial sector. International Journal of Research in Computer Applications and Information Technology (IJRCAIT), 7(2), 2015–2024.
22. Vimal Raja, G. (2024). Intelligent Data Transition in Automotive Manufacturing Systems Using Machine Learning. International Journal of Multidisciplinary and Scientific Emerging Research, 12(2), 515-518.
23. S. Roy and S. Saravana Kumar, "Feature Construction Through Inductive Transfer Learning in Computer Vision," in Cybernetics, Cognition and Machine Learning Applications: Proceedings of ICCCMLA 2020, Springer, 2021, pp. 95–107.
24. D. Johnson, L. Ramamoorthy, J. Williams, S. Mohamed Shaffi, X. Yu, A. Eberhard, S. Vengathattil, and O. Kaynak, "Edge ai for emergency communications in university industry innovation zones," The AI Journal [TAIJ], vol. 3, no. 2, Apr. 2022.
25. Chivukula, V. (2024). The Role of Adstock and Saturation Curves in Marketing Mix Models: Implications for Accuracy and Decision-Making.. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(2), 10002–10007.
26. Singh, A. (2023). Self-evolving IoT systems through edge-based autonomous learning. International Journal of Engineering & Extended Technologies Research (IJEETR), 5(6), 7547–7555. https://doi.org/10.15662/IJEETR.2023.0506011
27. Kumar, R., Panda, M. R., & Sardana, A. (2025). Reinforcement Learning for Autonomous Data Pipeline Optimization in Cloud-Native Architectures. Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online), 4(3), 97-102.
28. Kubam, C. S. (2026). Agentic AI Microservice Framework for Deepfake and Document Fraud Detection in KYC Pipelines. arXiv preprint arXiv:2601.06241.
29. Natta, P. K. (2023). Harmonizing enterprise architecture and automation: A systemic integration blueprint. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(6), 9746–9759. https://doi.org/10.15662/IJRPETM.2023.0606016
30. Abadi, D. J. (2009). Data Management in the Cloud.
31. Madabathula, L. (2022). Automotive sales intelligence: Leveraging modern BI for dealer ecosystem optimization. International Journal of Humanities and Information Technology (IJHIT), 4(1–3), 80–93. https://www.ijhit.info
32. Tableau Software. (2020). Visual Analytics and Business Intelligence.
33. Karnam, A. (2024). Engineering Trust at Scale: How Proactive Governance and Operational Health Reviews Achieved Zero Service Credits for Mission-Critical SAP Customers. International Journal of Humanities and Information Technology, 6(4), 60–67. https://doi.org/10.21590/ijhit.06.04.11
34. Ramalingam, S., Mittal, S., Karunakaran, S., Shah, J., Priya, B., & Roy, A. (2025, May). Integrating Tableau for Dynamic Reporting in Large-Scale Data Warehousing. In 2025 International Conference on Networks and Cryptology (NETCRYPT) (pp. 664-669). IEEE.
35. Zikopoulos, P., et al. (2012). Understanding Big Data. McGraw-Hill.