International Journal of Research Publications in Engineering, Technology and Management (IJRPETM)

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

Intelligent Cloud Cost Governance Driven by FinOps Automation and Machine Learning-Based Resource Optimization

Dr.K. Anbazhagan

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

Publication History: Received: 02.10.2026; Revised: 27.10.2026; Accepted: 01.11. 2026; Published: 15.11.2026

ABSTRACT: The rapid adoption of cloud computing has transformed enterprise IT infrastructure by providing scalable resources, flexible deployment models, and on-demand computing capabilities. However, increasing cloud consumption has introduced significant challenges related to infrastructure expenditure, resource utilization, budget management, and financial accountability. Traditional cloud cost management approaches often depend on manual monitoring, static budgeting, and periodic financial reporting, which are insufficient for managing dynamic enterprise workloads. This research proposes an intelligent cloud cost governance framework driven by FinOps automation and machine learning-based resource optimization. The proposed framework integrates cloud usage monitoring, cost allocation, predictive analytics, automated policy enforcement, and intelligent resource provisioning to improve financial efficiency. Machine learning algorithms analyze historical consumption patterns, workload behavior, utilization metrics, and pricing information to predict future expenditure and identify optimization opportunities. Automated FinOps mechanisms support budget alerts, anomaly detection, unused resource identification, rightsizing, and workload scheduling. The framework also incorporates governance policies to establish accountability across departments, projects, and cloud environments. By combining financial intelligence with automated infrastructure management, organizations can improve cloud cost visibility, reduce unnecessary expenditure, and maintain operational performance. The proposed approach provides a structured foundation for implementing continuous cloud financial management across hybrid and multi-cloud environments while supporting scalable, transparent, and data-driven enterprise cloud governance.

KEYWORDS: Cloud Cost Governance, FinOps Automation, Machine Learning, Cloud Resource Optimization, Predictive Cost Analytics, Cloud Financial Management, Intelligent Resource Allocation.

I. INTRODUCTION

Cloud computing has become an essential foundation for modern enterprise information technology because it provides flexible infrastructure, scalable computing resources, storage services, and application deployment capabilities. Organizations increasingly adopt public, private, hybrid, and multi-cloud environments to support business applications, digital transformation, data analytics, artificial intelligence, and enterprise operations. Although cloud computing reduces the requirement for maintaining physical infrastructure, it introduces complex financial management challenges associated with consumption-based pricing and continuously changing resource requirements. Cloud expenditure can increase unexpectedly when computing instances remain underutilized, storage resources accumulate unnecessary data, network traffic increases, or applications consume resources beyond their actual requirements. Organizations operating large cloud environments must therefore establish effective governance mechanisms to monitor spending, allocate financial responsibilities, and optimize resource consumption. Traditional infrastructure budgeting practices are not sufficient for these environments because cloud expenditure changes according to workload behavior, service configuration, resource availability, and operational demand. Consequently, intelligent cloud cost governance has become an important requirement for organizations seeking to achieve financial transparency and operational efficiency. FinOps provides a collaborative operational model that connects engineering, finance, and business teams through continuous cost visibility, accountability, and optimization. However, implementing FinOps at enterprise scale requires automated monitoring, accurate forecasting, intelligent recommendations, and governance policies capable of responding to changing cloud conditions. Machine learning

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

offers significant opportunities to improve these processes by analyzing historical usage information, recognizing consumption patterns, predicting future costs, and identifying abnormal resource behavior. Integrating machine learning with automated FinOps practices can transform cloud financial management from a reactive reporting activity into a proactive and continuously optimized governance process.

The proposed research focuses on developing an intelligent cloud cost governance framework that integrates FinOps automation with machine learning-based resource optimization. The framework is designed to collect infrastructure consumption metrics, financial information, workload characteristics, and resource utilization data from cloud environments. These datasets are processed through analytical components to identify inefficient resource allocation, predict expenditure trends, and detect unexpected cost variations. Machine learning models evaluate CPU utilization, memory consumption, storage activity, network usage, application workload patterns, and historical billing information to recommend appropriate infrastructure configurations. Automated governance mechanisms translate these recommendations into operational actions such as resource rightsizing, unused resource identification, scheduled shutdowns, storage optimization, and budget-based alerts. The architecture also introduces cost allocation mechanisms using organizational tags, project identifiers, departments, and application ownership information to improve financial accountability. Predictive cost analysis enables financial teams to estimate future cloud expenditure and identify potential budget overruns before they affect organizational financial planning. Furthermore, policy-driven automation ensures that optimization activities follow predefined security, availability, and performance requirements. The major objective is to establish a continuous feedback mechanism connecting cloud resource monitoring, financial analytics, intelligent decision-making, and automated optimization. The research aims to improve cost visibility, reduce unnecessary expenditure, enhance resource utilization, and support sustainable enterprise cloud operations. By combining predictive intelligence with FinOps governance principles, the proposed framework establishes a foundation for more efficient and accountable cloud financial management across complex enterprise environments.

II. LITERATURE REVIEW

Cloud financial management has developed significantly with the increasing adoption of consumption-based infrastructure services. Early cloud cost management approaches primarily concentrated on billing reports, manual resource monitoring, and periodic expenditure analysis. These methods provided organizations with basic visibility into infrastructure expenses but offered limited capabilities for predicting future consumption or automatically correcting inefficient resource allocation. Cloud service providers introduced cost management dashboards, usage reports, budgeting tools, and resource recommendation services to support financial transparency. However, managing expenditure across multiple cloud providers remains challenging because different platforms use different pricing structures, billing units, discount mechanisms, and resource measurement techniques. FinOps emerged as an operational framework connecting finance, engineering, and business teams to establish shared responsibility for cloud expenditure. Its principles emphasize visibility, accountability, optimization, and continuous improvement. Research and industry practices have increasingly highlighted the importance of cost allocation, organizational tagging, unit economics, budget monitoring, and financial forecasting. Effective cloud governance requires organizations to understand not only the total infrastructure expenditure but also the financial contribution of individual applications, departments, workloads, and business services. Nevertheless, traditional FinOps implementations frequently depend on predefined thresholds and manually reviewed recommendations, limiting their ability to respond quickly to changing workload conditions. Therefore, intelligent automation and predictive analytics are increasingly relevant to modern cloud financial governance.

Machine learning and predictive analytics have introduced new possibilities for improving cloud infrastructure utilization and expenditure forecasting. Regression algorithms, decision trees, random forests, gradient boosting, and time-series forecasting techniques can analyze historical billing information and infrastructure metrics to estimate future resource requirements. These models can identify relationships between application demand, computing utilization, storage consumption, and financial expenditure. Anomaly detection algorithms are particularly useful for identifying unexpected increases in cloud costs caused by configuration errors, uncontrolled workload execution, abnormal data transfers, or unauthorized resource provisioning. Clustering techniques can group similar workloads according to utilization behavior and identify opportunities for consolidation and resource optimization. Reinforcement learning has also attracted attention for dynamic resource allocation because it can learn infrastructure provisioning strategies through repeated interactions with changing workload environments. Such approaches can support decisions involving virtual machine sizing, container allocation, autoscaling, and workload scheduling. However, machine learning-based optimization introduces challenges related to training data quality, model accuracy, prediction

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

uncertainty, and changing application requirements. Incorrect optimization recommendations can negatively affect application performance, service availability, and user experience. Therefore, intelligent cloud cost management requires coordination between predictive models, operational constraints, and governance policies rather than relying exclusively on cost minimization.

Recent enterprise cloud management approaches increasingly combine automation, observability, and policy-driven governance to support large-scale infrastructure operations. Cloud monitoring platforms collect utilization information from virtual machines, containers, databases, storage services, and application environments. These measurements provide the operational information required for identifying inefficient resources and evaluating optimization opportunities. Infrastructure-as-Code technologies and automated orchestration platforms enable organizations to implement consistent resource configurations and enforce governance policies across cloud environments. Automated workflows can perform scheduled shutdowns, identify unattached storage volumes, recommend instance rightsizing, and generate financial alerts. Despite these developments, several research gaps remain. Many existing solutions address cost monitoring, forecasting, anomaly detection, and resource provisioning as separate activities instead of integrating them into a unified governance architecture. Cross-cloud cost normalization, accurate workload-level allocation, explainable optimization recommendations, and policy-aware automated actions remain important challenges. Additionally, organizations must balance financial savings with application performance, security, compliance, and availability requirements. The proposed research addresses these limitations by combining FinOps principles, machine learning-based predictive analytics, automated resource optimization, and centralized governance into an integrated cloud cost management framework.

III. RESEARCH METHODOLOGY

1. Research Design and Proposed Framework Architecture: The research adopts a design-oriented experimental methodology to develop an intelligent cloud cost governance framework integrating FinOps automation and machine learning-based resource optimization. The proposed architecture consists of five interconnected layers: cloud data collection, data preprocessing and integration, machine learning analytics, FinOps governance, and automated resource optimization. The cloud data collection layer gathers information from infrastructure services, including virtual machines, container clusters, databases, storage platforms, network services, and application workloads. Relevant parameters include CPU utilization, memory consumption, storage capacity, input/output operations, network traffic, instance configuration, operating duration, and service-level expenditure. Financial information is collected from cloud billing systems, usage reports, pricing databases, and resource consumption records. The data integration layer consolidates these heterogeneous datasets into a centralized analytical repository. Data normalization techniques standardize resource identifiers, timestamps, billing units, and service categories to support consistent financial analysis. Organizational metadata, including department identifiers, application names, project codes, and ownership tags, is incorporated to establish accurate cost allocation. The machine learning layer analyzes historical usage patterns and predicts future expenditure, while the FinOps governance layer evaluates financial policies, budget limitations, and organizational optimization requirements. Finally, the automation layer executes approved optimization actions through cloud management interfaces and orchestration services. A feedback mechanism continuously evaluates the financial and operational impact of implemented decisions. This architecture is intended to support continuous monitoring and optimization across public, private, hybrid, and multi-cloud infrastructure environments. The methodology emphasizes modularity, scalability, transparency, and controlled automation to ensure that each architectural component can operate independently while exchanging information through standardized interfaces.

2. Data Collection, Preprocessing, and Feature Engineering: The second methodological component focuses on collecting and preparing infrastructure consumption and financial datasets for machine learning analysis. The experimental dataset is designed to include historical cloud billing records, infrastructure utilization measurements, workload execution information, resource configuration details, and pricing information. Data can be obtained from cloud monitoring services, billing exports, infrastructure telemetry platforms, and controlled cloud workload experiments. The dataset includes time-based observations representing hourly or daily resource consumption across different infrastructure categories. Important input features include CPU utilization percentage, memory utilization percentage, storage usage, network transfer volume, number of active instances, instance type, workload execution duration, service category, regional pricing, and historical expenditure. Additional features are created to identify temporal and operational characteristics, such as weekday, weekend, peak usage period, monthly consumption trend, and application workload classification. Data preprocessing involves removing duplicate observations, correcting inconsistent timestamps, handling missing values, and identifying abnormal measurements. Numerical attributes are

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

standardized or normalized according to the requirements of the selected machine learning algorithms. Categorical features such as cloud provider, resource type, service category, and workload classification are transformed using suitable encoding techniques. Feature engineering creates additional indicators, including resource utilization efficiency, estimated cost per workload, average operating duration, idle resource ratio, and cost variation percentage. For expenditure forecasting, historical observations are organized into chronological training, validation, and testing datasets to prevent future information from leaking into model training. Where multiple cloud providers are involved, financial measurements are converted into a common reporting currency and normalized using consistent billing definitions. Data quality checks are applied to verify the completeness and reliability of the prepared dataset. These procedures establish a consistent analytical foundation for identifying resource inefficiencies, forecasting expenditure, and generating optimization recommendations.

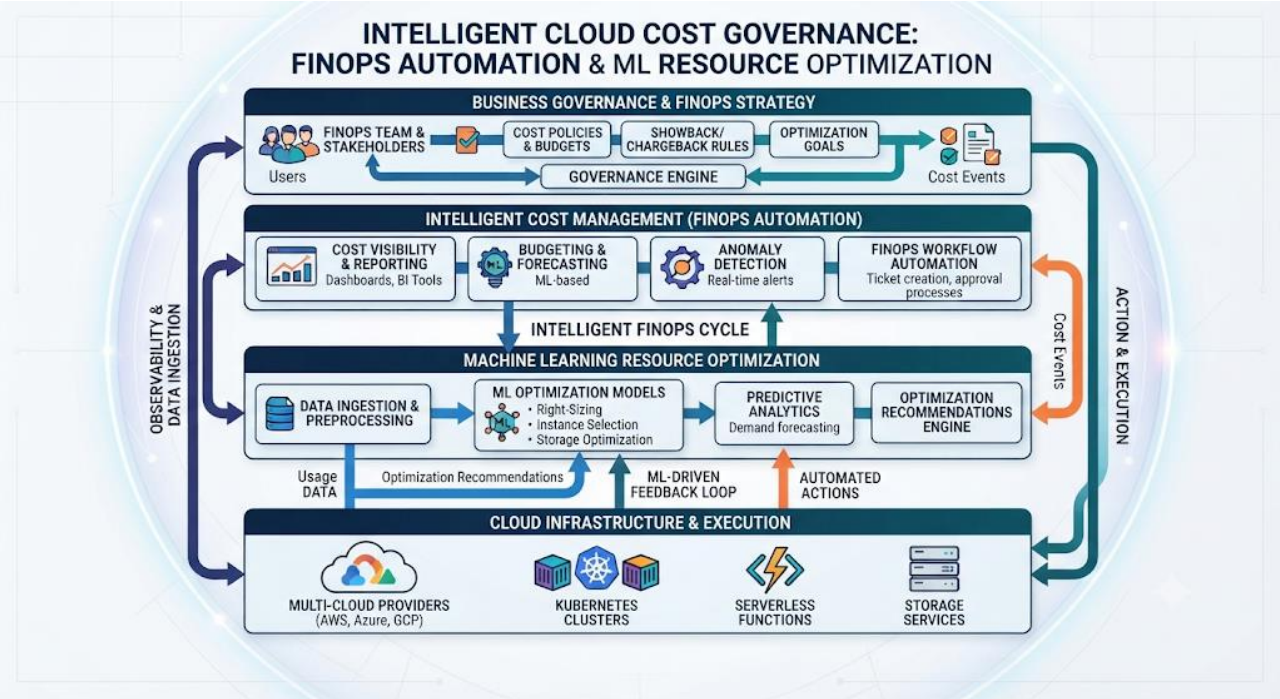

Figure 1: Intelligent Cloud Cost Governance and FinOps Automation Architecture.

3. Machine Learning Model Development and Intelligent Cost Prediction: The third component develops predictive models for cloud expenditure forecasting, resource utilization analysis, and abnormal cost detection. Multiple machine learning algorithms are considered to evaluate their suitability for different analytical tasks. Linear regression provides a baseline for estimating financial expenditure based on infrastructure consumption variables. Random forest and gradient boosting models are used to identify nonlinear relationships between resource utilization, workload behavior, and cloud expenditure. Time-series forecasting approaches analyze sequential billing information to estimate future spending and identify seasonal consumption patterns. Models such as ARIMA or suitable recurrent neural network architectures can be evaluated when sufficient chronological data is available. Anomaly detection is performed using statistical thresholds, isolation forest, or clustering-based techniques to identify unusual expenditure patterns. The model training process uses historical consumption data, while validation datasets support hyperparameter selection and model comparison. Important evaluation metrics include mean absolute error, root mean squared error, mean absolute percentage error, precision, recall, and F1-score, depending on the prediction task. The forecasting models generate estimated cloud expenditure for upcoming operational periods and compare predictions with predefined organizational budgets. Resource optimization models analyze utilization patterns to identify oversized instances, underutilized containers, idle storage, and workloads suitable for scheduled execution. A recommendation engine combines predicted consumption with resource pricing information and application performance constraints. It generates proposed actions such as reducing instance capacity, consolidating workloads, adjusting autoscaling policies, or scheduling non-production resources during low-demand periods. Reinforcement learning can be investigated as an advanced optimization technique for environments requiring adaptive resource allocation. Its reward function should

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

consider financial savings alongside response time, resource availability, service-level objectives, and operational stability. Model explainability techniques, including feature importance and SHAP-based analysis, can provide insights into the factors influencing individual predictions and optimization recommendations. This enables financial and engineering teams to understand why a particular resource has been identified as inefficient and evaluate whether the proposed action is operationally appropriate.

4. FinOps Automation, Governance Policies, and Experimental Evaluation: The final component integrates machine learning outputs with automated FinOps workflows and evaluates the effectiveness of the complete framework. The governance layer establishes financial policies for resource ownership, budget allocation, tagging compliance, expenditure thresholds, and optimization approval. Every cloud resource is associated with an application, project, department, or responsible business unit wherever possible. This enables financial reports to represent consumption at both infrastructure and business-service levels. Budget monitoring mechanisms continuously compare actual expenditure with predicted spending and generate notifications when financial thresholds are approached or exceeded. The anomaly detection component identifies unusual cost increases and provides relevant resource information for investigation. The recommendation engine proposes optimization actions, while the automation layer executes approved changes through cloud provider APIs, infrastructure management tools, and orchestration workflows. High-impact actions, including resource deletion, database resizing, and production workload modifications, require human approval. Reversible and low-risk operations, such as non-production shutdown scheduling, may be automated under predefined policies. A monitoring mechanism evaluates the impact of each action by comparing resource utilization, financial expenditure, and application performance before and after implementation. The experimental evaluation compares the proposed framework with conventional cloud cost management approaches using indicators such as total cloud expenditure, resource utilization efficiency, forecasting error, anomaly detection performance, optimization execution time, and budget compliance. Controlled experiments can simulate varying workload conditions, including normal traffic, peak demand, idle periods, and unexpected resource consumption. Comparative analysis determines whether predictive and automated governance provides measurable improvements over static threshold-based management. Security controls, access permissions, audit logs, and policy validation are incorporated to maintain operational accountability. The research methodology therefore evaluates the framework not only through its ability to reduce expenditure but also through its capacity to preserve application performance, maintain governance consistency, and support responsible enterprise cloud operations.

IV. RESULTS AND DISCUSSION

The proposed intelligent cloud cost governance framework establishes an integrated approach for managing cloud expenditure through predictive analytics, automated FinOps operations, and resource optimization. The expected experimental outcomes indicate that centralized collection of billing information and infrastructure utilization metrics can improve financial visibility across enterprise cloud environments. By combining resource-level consumption information with application and departmental ownership metadata, organizations can identify the financial contribution of individual workloads and business services. Machine learning models provide additional analytical capabilities by recognizing relationships between infrastructure utilization and expenditure patterns. Predictive cost estimation enables organizations to identify potential budget overruns earlier than conventional periodic reporting methods. Similarly, anomaly detection supports the identification of unusual expenditure variations associated with unexpected workload execution, idle resources, or inefficient configurations. These capabilities can improve financial accountability and provide engineering teams with actionable information for controlling infrastructure consumption. However, actual improvements must be established through controlled experiments using real or clearly documented simulated cloud consumption datasets.

Resource optimization represents another important outcome of the proposed framework. Machine learning-based analysis can identify oversized virtual machines, underutilized containers, unnecessary storage allocations, and workloads that can be scheduled during lower-demand periods. Automated rightsizing recommendations can improve the relationship between allocated infrastructure capacity and actual workload requirements. Integration with autoscaling and orchestration mechanisms allows resource provisioning to respond more effectively to changing application demand. These activities can reduce unnecessary infrastructure consumption while maintaining required operational capacity. The effectiveness of optimization depends on the quality of historical data, the accuracy of predictive models, and the constraints associated with individual applications. For example, reducing computing resources without considering peak workload requirements may increase application response time. Therefore, the proposed framework incorporates performance indicators and service-level objectives into its optimization decisions.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

The combination of financial and operational measurements provides a more balanced approach than cost reduction alone. Explainable recommendations and human approval mechanisms further improve the reliability of automated resource management.

The integration of FinOps governance and machine learning also strengthens enterprise financial management practices. Automated budget alerts, cost allocation reports, anomaly notifications, and optimization workflows establish a continuous governance process. Financial teams can use predictive expenditure information for budget planning, while engineering teams can investigate resource inefficiencies and implement corrective actions. Business departments can obtain clearer visibility into the financial impact of their applications and operational activities. Nevertheless, implementation challenges include inconsistent resource tagging, differences in cloud provider billing structures, incomplete telemetry, and the computational overhead of maintaining predictive models. Organizations must also establish appropriate permissions and audit mechanisms before allowing automated infrastructure modifications. Overall, the proposed framework provides a foundation for improving financial transparency, resource utilization, and cloud governance. Quantitative claims regarding cost savings, forecasting accuracy, and performance improvements should be reported only after completing the proposed experimental evaluation.

V. CONCLUSION

Intelligent cloud cost governance has become an important requirement for enterprises adopting large-scale cloud infrastructure and consumption-based computing services. Traditional financial management approaches based on manual monitoring, periodic billing analysis, and static budgeting provide limited capabilities for managing dynamic workloads and continuously changing infrastructure requirements. This research proposes an integrated cloud cost governance framework driven by FinOps automation and machine learning-based resource optimization. The proposed architecture combines cloud usage monitoring, financial data integration, predictive cost analytics, anomaly detection, automated governance policies, and intelligent resource provisioning. Machine learning algorithms analyze historical expenditure and resource utilization patterns to estimate future financial requirements and identify infrastructure inefficiencies. These analytical capabilities support improved decision-making regarding resource rightsizing, workload scheduling, storage optimization, and infrastructure allocation. The framework also introduces cost allocation mechanisms that associate cloud expenditure with organizational departments, applications, and business projects. This approach strengthens financial accountability and provides a clearer understanding of infrastructure consumption across enterprise environments. The integration of automated FinOps processes further supports budget monitoring, expenditure alerts, optimization recommendations, and continuous governance. By connecting financial information with operational intelligence, the proposed framework transforms cloud cost management into a proactive and data-driven process. Its modular architecture can support deployment across different cloud platforms while maintaining flexibility for changing enterprise requirements.

The proposed research highlights the importance of balancing financial optimization with operational performance, security, availability, and organizational governance. Machine learning-based recommendations should be evaluated against workload characteristics and service-level requirements before automated actions are implemented. Human approval, explainability, audit logging, and policy validation are particularly important for high-impact infrastructure modifications. The experimental methodology provides a structured approach for assessing the framework through forecasting accuracy, resource utilization efficiency, anomaly detection performance, expenditure reduction, and budget compliance. Although the proposed architecture establishes a comprehensive foundation, its practical effectiveness depends on reliable datasets, appropriate model selection, consistent tagging practices, and secure integration with cloud management interfaces. Organizations can use the framework to improve collaboration between finance, engineering, and business teams while establishing continuous financial accountability. The integration of predictive analytics and automated resource management offers opportunities to improve infrastructure efficiency and reduce avoidable cloud expenditure without compromising application reliability. In conclusion, intelligent cloud cost governance driven by FinOps automation and machine learning represents a promising approach for developing transparent, scalable, and financially sustainable enterprise cloud environments.

VI. FUTURE WORK

Future research can extend the proposed framework through advanced machine learning techniques, reinforcement learning, and multi-agent intelligent automation. More sophisticated predictive models can be developed to improve expenditure forecasting across complex enterprise workloads with seasonal demand, irregular usage patterns, and

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

rapidly changing resource requirements. Deep learning architectures, including long short-term memory networks, temporal convolutional networks, and transformer-based time-series models, can be evaluated for predicting infrastructure consumption and financial expenditure. Reinforcement learning can further support adaptive resource provisioning by continuously evaluating the relationship between infrastructure allocation, application performance, and financial cost. Multi-agent architectures may enable specialized intelligent agents to manage different operational responsibilities, including billing analysis, anomaly investigation, workload classification, budget monitoring, and resource optimization. These agents could coordinate their decisions through a shared governance engine to provide more comprehensive cloud financial management. Future implementations can also incorporate generative AI-based interfaces that allow financial administrators and cloud engineers to interact with expenditure information using natural language. Such interfaces could generate financial summaries, explain cost variations, identify optimization opportunities, and provide recommendations based on organizational policies. However, generated recommendations should be validated against trusted billing records and infrastructure telemetry before implementation. Additional research can investigate automated model retraining and drift detection to ensure that predictive capabilities remain effective as workloads and pricing structures change over time. Another important research direction involves expanding the framework to support large-scale hybrid and multi-cloud environments with different pricing models, billing mechanisms, and infrastructure management interfaces. Future implementations can investigate cloud cost normalization methods that provide consistent expenditure comparisons across multiple service providers and geographical regions. Integration with Kubernetes, serverless computing, edge infrastructure, and enterprise data platforms can improve the framework's applicability to modern distributed applications. Advanced governance mechanisms may incorporate policy-as-code, automated compliance validation, and risk-aware optimization to ensure that financial decisions remain consistent with organizational security and regulatory requirements. Sustainability can also become an additional optimization objective by considering energy consumption, carbon emissions, and workload placement alongside financial expenditure. Future experiments should evaluate the proposed framework using real enterprise billing datasets and production-inspired workload environments over extended periods. Comparative studies can examine different forecasting algorithms, optimization strategies, and levels of automation under varying workload conditions. Research should also investigate explainable AI, privacy-preserving analytics, and federated learning to enable collaborative financial intelligence without unnecessarily centralizing sensitive organizational information. Finally, the development of standardized FinOps evaluation metrics could support consistent assessment of cost efficiency, forecasting reliability, governance maturity, and operational performance. These developments can transform the proposed framework into a comprehensive intelligent financial operations platform capable of supporting autonomous, sustainable, and accountable cloud infrastructure management.

REFERENCES

1. Tong, Z., Chen, H., Deng, X., Li, K., & Li, K. (2020). A scheduling scheme in the cloud computing environment using deep Q-learning. *Information Sciences, 502*, 1170–1191. https://doi.org/10.1016/j.ins.2019.10.035
2. Jauro, F., Chiroma, H., Gital, A. Y., Almutairi, M., Abdulhamid, S. M., & Abawajy, J. H. (2020). Deep learning architectures in emerging cloud computing architectures: Recent development, challenges and next research trend. *Applied Soft Computing, 96*, 106582. https://doi.org/10.1016/j.asoc.2020.106582
3. Uddandarao, D. P. (2026). The reality of an AI implementation. *Communications of the ACM, 69*(8), 8–9. https://doi.org/10.1145/3821443
4. Yelagam, C. K. (2022). Interactive vector annotation layers and distributed payload buffer serialization for cloud-native document review interfaces. *International Journal of Computer Technology and Electronics Communication, 5*(5), 5826–5833. https://doi.org/10.15680/IJCTECE.2022.0505011
5. Garg, D. K., & Bansal, D. K. (2026). Garg's principle: Horizon-based safety stock optimization through dynamic net availability. *International Journal of Emerging Trends in Computer Science and Information Technology, 7*(1), 391–398. https://doi.org/10.63282/3050-9246.IJETCSIT-V7I1P155
6. Bandhela, R. R. (2022). Advancing banking systems with federated learning and a fuzzy-based blockchain framework for secure and efficient transactions. *Journal of Informatics Education and Research, 2*(2), 24–32. https://jier.org/index.php/journal/article/view/3127/2539
7. Chinnam, N. B. (2024). Observability-driven performance testing and reliability validation for high-volume distribution center system migrations. *International Journal of Engineering & Extended Technologies Research, 6*(2), 7983–7995.
8. Prakash, N. (2026, July). A scalable LLM-as-judge framework for automated evaluation of generative AI systems with domain-adaptive calibration. In *2026 International Conference on Intelligent and Sustainable AI Systems (ICOSAAS)* (pp. 1366–1373). IEEE.

|www.ijrpetm.com | ISSN: 2454-7875 | editor@ijrpetm.com |A Bimonthly, Peer Reviewed & Scholarly Journal|

||Volume 9, Issue 6, November-December 2026||

DOI:10.15662/IJRPETM.2026.0905003

9. Gudi, S. (2026). A data quality-aware real-time healthcare data infrastructure for early patient deterioration detection. *Acta Innovations*, 1738–1750.
10. Sethupathy, A. (2024). AI-powered cybersecurity mesh for financial networks: A real-world deployment case study. *International Research Journal of Modernization in Engineering Technology and Science, 6*(11), 2456–2471.
11. Kalakoti, M. K. R. (2025). Automated disaster recovery infrastructure for HIPAA-regulated healthcare systems: A cloud-native implementation using infrastructure as code. *International Journal of Computational and Experimental Science and Engineering, 11*(3). https://doi.org/10.22399/ijcesen.3928
12. Javvadi, V. P., Ganapathy, R. S., Sri, M. S., Nekkalapudi, K., Pochincharla, S., & Akula, T. V. (2026, July). Detecting and mitigating flaky tests in continuous integration A decision catalogue for quality engineering teams. In *2026 International Conference on Intelligent and Sustainable AI Systems (ICOSAAS)* (pp. 1313–1320). IEEE.
13. Gaddam, N. (2022). Unified state machine normalization across heterogeneous clinical systems: An architectural pattern for enterprise status federation. *American International Journal of Computer Science and Technology, 4*(3), 60–86.
14. Manoharan, D. (2026). Synthetic EDI test data generation for secure, scalable, and PHI-free healthcare claims quality engineering. *Journal of International Crisis and Risk Communication Research, 9*(1), 358–371.
15. Shewale, V. (2022). IT/OT convergence: A zero trust architecture for energy sector. *American International Journal of Computer Science and Technology, 4*(5), 52–60. https://doi.org/10.63282/3117-5481/AIJCST-V4I5P106
16. Vemireddy, S. (2024). Machine learning-driven expert matching and personalized recommendation frameworks for intelligent tax advisory platforms. *International Journal of Research and Applied Innovations, 7*(1), 10207–10220.
17. Billa, P. K. (2023). Architectural strategies for modernizing legacy SAP ERP landscapes through S/4HANA transformation. *International Journal of Research Publications in Engineering, Technology and Management, 6*(2), 8409–8422. https://doi.org/10.15662/IJRPETM.2023.0602006
18. Narra, S. L. (2025). Demystifying zero trust architecture: Why it's not just a buzzword. *International Journal of Computing and Engineering, 7*(10), 1–16.
19. Shekh, J., Shahin, M. B. N., Islam, M. J., Suvo, M. H., & Nandy, D. (2024). ML framework for predictive hazard monitoring in smart manufacturing environments. *International Journal of Scientific Research in Science, Engineering and Technology, 11*(6), 583–592. https://doi.org/10.32628/IJSRSET2512553
20. Uddandarao, D. P. (2026). The reality of an AI implementation. *Communications of the ACM, 69*(8), 8–9. https://doi.org/10.1145/3821443
21. Yelagam, S. K. (2023). A qualitative decision procedure for rollback or retain-and-retry after partial cloud provisioning failures. *International Journal of Computer Technology and Electronics Communication, 6*(3), 7025–7032.
22. Dhaduk, H. (2025). Hybrid deep learning framework for predicting CRISPR-Cas9 off-target activities using experimentally validated mismatch and indel datasets. *American International Journal of Computer Science and Technology, 7*(2), 177–186.
23. Meghani, Viren and Singh Siddhu, Harpreet and Sethi, Mayank, An Integrated Framework for Compliance-Aware Engineering in Regulated Enterprise (July 29, 2026). Available at SSRN: https://ssrn.com/abstract=7400518 or http://dx.doi.org/10.2139/ssrn.7400518
24. Chinthamani, R. (2023). A dependency-aware orchestration framework for failure-safe change automation across distributed enterprise systems. *International Journal of Research Publications in Engineering, Technology and Management, 6*(3), 8785–8802.
25. Badam, L. R. (2025). AI-based risk/compliance decision framework. *American International Journal of Computer Science and Technology, 7*(2), 167–176. https://doi.org/10.63282/3117-5481/AIJCST-V7I2P113
26. Kande, J. (2025). Real-time anomaly detection. In *Recent trends in artificial intelligence and IoT: Second International Conference, ICAII 2024, Jamshedpur, India, October 25–26, 2024, revised selected papers*. Springer Nature.
27. Nendrambaka, S. K. (2024). Recent advances and innovations in SAP S/4HANA Cloud and SAP BTP and SAP AI: Integration strategies and latest developments. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 10*(6), 1743–1751.
28. Gudimetla, S. (2025). No-code automation platforms for incident management: Transforming engineering operations. *Journal of Information Systems Engineering and Management, 10*, 441–460.
29. Devineni, A., Prabagaran, A., Mali, R. K., Mishra, M. K., Nekkalapudi, K., & Velu, R. (2026, June). Intelligent incident prediction in banking cloud operations: A machine learning framework for proactive SRE response. In *2026 5th International Conference on Computer Networks, Big Data and IoT (ICCBI)* (pp. 1072–1078). IEEE.