iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** III **Month of publication:** March 2026

DOI: https://doi.org/10.22214/ijraset.2026.79052

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue III Mar 2026- Available at www.ijraset.com

Multi-Agent Generative AI Framework for Cloud Cost Prediction and Intelligent Resource Optimization

Vinoth kumar G[1], Narthana Hari S V[2], Samyucthan M[3], Angela Lincy N H[4]
[1, 2, 3]*Final Year,* [4]*Assistant Professor, Dept. of AI & DS, MNM Jain Engineering College, Chennai, India*

Abstract: ***Therapidadoptionofcloudcomputingservicessuch as Amazon Web Services, Microsoft Azure, and Google Cloud Platformhastransformedenterpriseinfrastructuremanagement by enabling elastic resource provisioning. However, dynamic re-source allocation mechanisms and complex pricing models often lead to unpredictable cloud expenditures and budget overruns. Conventionalcloudmonitoringdashboardsprimarilyprovidede- scriptiveanalyticsandlack predictiveinsights, anomalydetection, and automated optimization capabilities.***

This paper presents a Multi-Agent Generative Artificial In- telligence framework for proactive cloud cost prediction and intelligent resource optimization. The proposed system consistsof five cooperative agents, including a Data Ingestion Agent,Cost Prediction Agent, Anomaly Detection Agent, Optimization Recommendation Agent, and an LLM-Based Report Genera-tion Agent. A Long Short-Term Memory (LSTM) network is employed for multivariate time-series cost forecasting, while an Isolation Forest model is utilized to identify abnormal billing patterns. A risk-aware optimization mechanism translates an-alytical insights into actionable cost-saving recommendations, and a Generative AI-based Large Language Model produces explainable financial reports for decision-makers.

Additionally, ensemble learning models including Random Forest and XGBoost are integrated using a stacking-based meta model to enhance prediction robustness.

Experimental results indicate high forecasting accuracy based onthecoefficientofdetermination,robustanomalydetectionwith minimal false alerts, and efficient end-to-end system response, validating the framework's scalability, interpretability, and suit-ability for production deployment.

Index Terms: CloudComputing, Cloud CostPrediction, Gen- erative Artificial Intelligence, Multi-Agent Systems, LSTM, Iso-lation Forest, FinOps, Explainable AI, Resource Optimization

I. INTRODUCTION

Cloud computing has become a key part of changing how industries operate digitally. More and more companies are using public cloud services to run their apps, keep their data safe, and handle growing infrastructure needs, all without having to spend a lot of money upfront. The pay-as-you-go pricing model offers flexibility, letting businesses adjust their resourcesaccordingto howmuchworktheyneedtohandle atanytime.Thisflexibilityhasmadeiteasiertorespond quickly to changes and has simplified the management of the infrastructure.

Evenwiththesebenefits,managingfinancesinthecloudhas become a major issue. The way resources are provided, how systems automatically adjust to demand, and how services are used can all lead to costs that are hard to predict. When too manyresourcesaregiven,notused,orsetupwrong,costscan goupquicklyandit'shardtoseehowmuchisbeingspent. In big environments where many services and teams work at the same time, it's harder to find exactly where the costs are going up.

Many companies face unexpected jumps in their bills because of sudden increases in workloads, poorly set up automaticscalingsettings,orservicesthatwerestartedwithout beingnoticed.Costsoftengoupslowlybecausestoragekeeps growing without control and because the rules for managing how long data is kept are not working well. Cloud service providers provide cost monitoring dashboards, but these tools mainlyshowpastspendinginformation.Theydon'toftengive predictions or smart suggestions that help control costs before problems happen.

Another major issue is that the process for checking costsis done by hand. Financial and DevOps teams usually rely on spreadsheets and regular check-ins to look at billing informa- tion, which takes a lot of time and can lead to mistakes being missed.

©IJRASET: All Rights are Reserved

Without using predictive modeling and automated anomaly detection, companies have to make decisions after problems happen, like fixing budget issues only after they've gone over.

Newdevelopmentsinmachinelearningandartificialintelli- genceofferachancetogreatlyimprovehowfinancialmanage- ment works in the cloud. Predictive models can predict how much money will be spent by looking at past spending habits, and anomaly detection can spot unusual billing activity as it happens. When used with automated optimization strategies and clear reporting tools, these technologies can make things more transparent and help make better decisions faster.

To address these challenges, this research introduces a multi-agentartificialintelligencesystemaimedatproviding predictive, diagnostic, and prescriptive insights into cloud costs. By combining time-series forecasting, anomaly detec- tion,risk-awareoptimization,andgenerativereportingintoone system, the framework is designed to help with proactive and sustainable management of cloud finances.

II. RELATED WORK

In recent years, there has been a lot of research on cloud cost management, choosing the right cloud provider, and opti-mizingservicessmartly,becausemulti-cloudenvironmentsare getting more complicated. Many research efforts have looked into ways to make things more standard, how to connect different parts of the system, better ways to make decisions, and using smart automation to improve how financial systems are managed and how they run day to day.

The approach handles the differences in how various cloud providers charge for their services by combining the FinOps Open Cost and Usage Specification (FOCUS) with Large Language Models (LLMs). The authors explain that cloud companies like AWS, Azure, and GCP create billing informa- tion in their own special formats, and these formats can have different levels of detail and ways of showing costs, which makes it hard to compare data across different providers. OPTIC offers an automated ETL process that takes cost data andputsitintoacommonSQLstructure,makingiteasier to work with. It also lets users ask questions in everyday language, which is then translated into SQL queries using a language model. The framework greatly helps with making FinOps environments more accessible and financially clear, butitsmainbenefitisintranslatingandhandlingbilling datainsteadofpredictingfuturecostsoroptimizingbased on unusual activity. Unlike other approaches, the framework introduced in this study does more than just translate. It also includes tools for predicting future trends, identifying unusual patterns,andmakingdecisionsthattakerisksintoaccount.The roleofcloudbrokersinautomaticallydistributingserviceshas been thoroughly studied in A Systematic Literature Review. The authors use a systematic literature review based on the PRISMAmethodtoexaminebrokeragesystemsthathelpwith choosingtherightservices,balancingworkloads,andlowering costs.Theirworkfocusesonautonomiccomputingideas like self-configuration, self-optimization, and self-healing to improveserviceefficiencyinmulti-cloudenvironments.Cloud brokers help make services easier to manage and resources used more efficiently. However, most current models forcloud brokering mainly focus on spreading out workloads and makingresponsesfaster,insteadofaccuratelypredictingcosts indetailorprovidingsmartfinancialinsightsbasedonunusual activity. This shows the gap in research that this work aims to fill, focusing on financial analytics and predictive intelligence as key parts.

Choosing a cloud provider has also been seen as a com- plicated decision that involves considering many factors, as explained in Cloud provier selection a comp. The authors consider cost, performance, and evaluation risk as different goalsthatmayconflictwitheachotherandusemethods from multi-objective optimization to model how to choose the best approach. Their method shows that choosing a provider requires balancing the need to save money with the need togetthebestperformance.However,thesemodelsusuallywork during the strategic selection phase and don't keep tracking or forecasting changes in billing once they are put into use. This study adds to those approaches by using real-time predictions and tools to spot unusual patterns, helping with managing costs as things happen. So far, earlier work has looked at standardization, ways to connect different parts, and choosing the best providers on their own. However, not much work brings together predictive modeling, anomaly detection, opti- mizationsuggestions,andexplainablereportingintoonesingle system.Thenewmulti-agentgenerativeAIsystemisdesigned to fill this gap by bringing together financial predictions,smart identification of unusual activities, and practical waysto improve performance, all within a system that can grow with a business and clearly explain how it works, making it idealforlargecompaniesmanagingtheirfinancialoperations.

©IJRASET: All Rights are Reserved

III. PROPOSED SYSTEM ARCHITECTURE AND METHODOLOGY

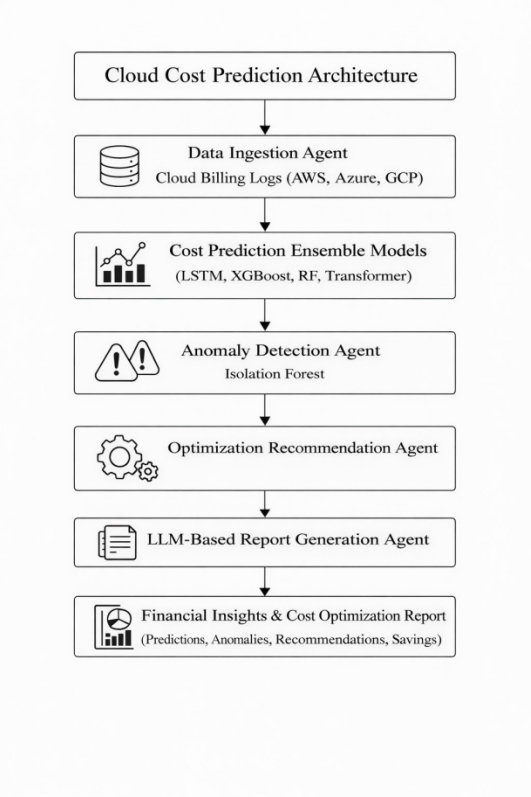


Fig. 1.System Architecture of the Proposed Multi-Agent Generative AIFramework for Cloud Cost Prediction

A. System Overview

The proposed framework adopts a distributed multi-agent architecture in which specialized agents collaborate through API-basedcommunication.Initially,cloudbillingandresource utilization data are ingested and preprocessed to ensure con- sistency and analytical reliability. The processed data are then utilized by the Cost Prediction Agent to generate future cost estimates using an ensemble of prediction models including LSTM, XGBoost, and Random Forest. The outputs of these modelsarecombinedusingastacking-basedmetalearner to improve forecasting accuracy. Subsequently, the Anomaly DetectionAgentanalyzesbillingpatternstoidentifyabnormal cost behaviors. The predictive deviations and anomaly signals are aggregated into a risk score that guides the Optimization Recommendation Agent in generating cost-efficient resource management strategies. Finally, the LLM-Based Report Gen- eration Agent transforms numerical analytics into structured, explainable, and human-readable financial reports.

B. DataIngestionAgent

The Data Ingestion Agent aggregates heterogeneous billing andoperationalmetricsfrommultiplecloudserviceproviders, includingAWS,MicrosoftAzure,andGoogleCloudPlatform. The collected data consist of monthly billing statements, service-level cost breakdowns, and fine-grained resource uti- lization metrics such as CPU usage, memory allocation, stor- age consumption, and network bandwidth.

To ensure data reliability, preprocessing steps such as missing value imputation, feature normalization, and sliding window transformation are applied. Minor statistical noise is smoothed to maintain data stability without suppressing gen- uine anomalies. This preprocessing pipeline enables seamless integration of raw billing data into forecasting and anomaly detection modules.

C. CostPredictionAgent

Cloud cost prediction is formulated as a multivariate time- series regression problem in which historical resource utiliza- tionpatternsaremappedtofuturebillingcosts.Let$X_{t-k:t}$representhistoricalusagefeaturesoveraslidingwindowoflength$k$,and letC_{t+1}denote thepredictedcostforthenext billingcycle.

Theforecastingobjectiveistominimizethe Mean Squared Error loss between actual and predicted costs. Toeffectivelycapturenonlineartemporaldependencies,aLongShort-TermMemoryneuralnetworkisemployed. The LSTM architecture mitigates vanishing gradient issues and preserves long-term contextual information through gated memory mechanisms. This enables accurate modeling of sea- sonal billing cycles, workload growth trends, and long-term cost dependencies.

D. AnomalyDetectionAgent

The Anomaly Detection Agent identifies abnormal billing patterns using the Isolation Forest algorithm, an unsupervised learning technique well-suited for high-dimensional data. The modelassignsanomalyscoresbasedontheaveragepathlength required to isolate data points within randomly constructed trees. Shorter isolation paths indicate a higher likelihood of anomalous behavior.

This approach effectively detects sudden cost escalations, unexpectedserviceactivations,andconfiguration-inducedcost leakages. Unlike threshold-based methods, Isolation Forest dynamically adapts to evolving billing distributions, thereby reducing false positives while maintaining detection sensitiv- ity.

E. OptimizationRecommendationAgent

The Optimization Recommendation Agent translates pre- dictive deviations and anomaly signals into actionable cost-savingstrategiesusingaweightedriskaggregationmodel.The risk score incorporates prediction error magnitude, anomaly severity, resource usage growth rate, and cost variance. Based on the assessed risk level, the system recommends optimiza- tion actions such as instance rightsizing, termination of idle resources,adoptionofreservedpricingmodels,andimplemen- tation of storage lifecycle policies.

This risk-aware optimization mechanism ensures that cor- rective actions are financially prioritized and aligned with enterprise FinOps governance principles.

F. LLM-BasedReportGenerationAgent

The LLM-Based Report Generation Agent converts analyt- ical outputs into comprehensive financial intelligence reports. These reports include cost forecasts, explanations of detected anomalies, budget impact assessments, and clearly articu-lated optimization recommendations. The use of GenerativeAI enhances interpretability for non-technical stakeholders, thereby improving transparency and accelerating decision- making processes.

IV. EXPERIMENTAL SETUP AND EVALUATION

A. DatasetDescription

Theexperimentaltestingoftheproposedmulti-agentframe- work was done using a detailed dataset that includes 24 monthsofcloudbillinginformationfromvariouscloudservice types. The dataset has detailed information about the costs of different services and includes precise data on how resources are used, like how much CPU, memory, storage, and network bandwidth each virtual machine is using. These features were chosentoshowthemainreasonswhycloudcostsriseinactual business settings.

To make the data more like real-world situations, synthetic seasonal patterns were introduced into the dataset. These patterns help show how workloads change over time in real systems, such as monthly reports, busy times when lots of people use the system, and heavy workloads at the end ofeachquarter.Additionally,controlledspikesinanomalieswere added to mimic abnormal billing issues that happen becauseof wrong setup in auto-scaling, unexpected service starts,poorstoragemanagement, andsuddenincreasesinworkloads. This approach allowed for a thorough check of the anomaly detection system while keeping the real statistical features of actual cloud usage data intact.

B. DataPreprocessingandPreparation

Before training the model, the billing data was carefully prepared through a thorough process to make sure it was reliable and consistent over time. Missing values in the billing and usage records were filled in using statistical methods to avoidbreaksinthetime-seriesdata.Numberswereadjustedto the same scale to help the model work smoothly and prevent some numbers from having too much influence.

The historical billing data were turned into overlapping sliding time windows to create supervised learning sequences thataregoodforforecastingmultivariatetimeseries.Thedata wassplitinorderovertimetoseparatethetrainingandtesting sets, making sure that information from the future wasn't included in the training data by accident. This evaluation method is very similar to how things work in the real world, where predictions have to be made only based on past data.

C. EvaluationMetrics

Theperformanceofthecostforecastingpartwaschecked using common regression measurements that are usually used intime-seriespredictionwork.MeanAbsoluteError(MAE) wasusedtomeasuretheaveragedifferencebetweenthe predicted cloud costs and the real costs, giving a clear idea of how accurate the predictions were. Root Mean Squared Error (RMSE)wasusedtohighlightbiggermistakesandcheckhow wellthemodelworkswhenbillingconditionsareunstable. TheR-squaredvaluewasusedtoshowhowmuchofthe variation in cloud costs is explained by the forecasting model, which helps to assess how well the model predicts outcomes. The system checked how well it found anomalies by using classificationmeasures likeprecision,recall,andF1score.Precisionlookedathowmanyoftheflaggedbillingissueswereactuallycorrect,whilerecallmeasuredhoww ellthe modelfoundalltherealabnormalcosts.TheF1-scoregave afairwaytocheckhowwelldetectionworksbybalancing precision and recall. Also, the false positive rate was checked tounderstandhowoftenwronganomalyalertshappen.This is especially important in big companies to prevent too many alertsthatcanmakepeopleignorerealissuesandkeeptrustintheautomaticmonitoringtools.

D. SystemPerformanceMetrics

Besides checking how well the system can predict and detectthings,wealsolookedatotherperformancemeasuresat the system level to see if the proposed framework is practical to use in real situations. The time it took to create cost predictions and calculate anomaly scores for new billing data was measured as the detection latency. This measure shows how quickly the system can respond in situations where real- time monitoring is needed.

The time it took the LLM-based reporting agent to gen- erate the report was also checked to see how efficiently it could turn analytical results into clear and organized financial summaries that people can easily read. These performance metricstogethershowhowwellthemulti-agentframework canhandlelarge-scaleworkandworkinrealtimewithincloud environments used for actual production.

E. EvaluationObjective

Theexperimentwassetuptocheckhowaccuratethe cost prediction and anomaly detection parts are, as well ashow strong, able to handle more work, and quick the whole multi-agentsystemperforms.Byusingpredictiveperformance metrics along with system-level efficiency measurements, the evaluation gives a complete picture of how well the frame- work is suited for intelligent cloud financial governance in enterprise-scale environments.

V. RESULTS AND DISCUSSION

This part takes a close look at the proposed Multi-Agent Generative AI framework by checking how well it predicts, findsunusualpatterns,usescomputerresourcesefficiently,and how well the whole system works. The results are looked at not just by numbers but also by how useful they are in real- world cloud setups used by companies. The comparison with standardmodelsshowsthatthenewdesignoffersbetterresults and more dependable performance in real-world situations involving cloud cost management.

A. CostPredictionPerformance

TheresultsfromtheexperimentshowthattheLSTM-based model for predicting costs works much better than traditional methods like regression and ensemble learning, based on all the different ways we measured its performance. The model we created has an R squared value of 0.96, which means it accounts for about 96

Although ensemble learning models were incorporated to improve robustness, the LSTM component contributed most significantlytocapturingtemporaldependenciesandachieving higher forecasting accuracy.

The lower MAE and RMSE values also show that the forecast errors stay small, even when there are big changes in workload or unexpected scaling events. Lower RMSE shows that the model is more reliable when there are big changes in costs, meaning it doesn’t get too affected by short-term noisy data in billing.

LinearRegressionperformslesseffectivelybecauseitrelies on the idea that there is a straight-line connection betweenhow resources are used and the costs involved. In real cloud setups, costs change in complicated ways because of things like non-linear scaling, automatic resource adjustments, and how different services rely on each other, which can’t all be properly shown with simple linear models.

Random Forest and XGBoost help make predictions more accuratebyconsideringhowdifferentfeaturesinteractinnon-linearwaysandbyunderstandingthecomplicatedconnections between various usage factors. However, these tree-based ensemble models handle each observation separately and do not have a clear way to remember past events. Because ofthis,theydon'tworkaswellwhenitcomestoshowinghow long-term billing is affected or how changes in settings can cause costs to come up later.

Unlike the other model, the LSTM uses a special structure thatletsitrememberimportantinformationovertimebyusing controlledmemoryunits.Thissetupletsthemodelunderstand overall patterns, regular changes, and the way billing happens overtime,whichmakesitespeciallygoodforpredictingcloud spending.

TABLEI
COSTPREDICTIONPERFORMANCECOMPARISON

Model	MAE	RMSE	R^2
LinearRegressio	8.7%	11.2%	0.88
n	6.3%	8.1%	0.92
RandomForest			
XGBoost	5.4%	7.3%	0.94
LSTM	4.8%	6.5%	0.96

TheLSTMmodelachievedthehighestforecastingaccuracy, demonstratingitseffectivenessincapturingcomplextemporal billing patterns.

B. AnomalyDetectionandSystemEfficiency

Thepartthatfindsunusualpatterns,whichusestheIsolation Forestmethod,workswellaccordingtoallthetestingmea- surements. A precision rate of 95.2% shows that most of the flaggedissuesarerealbillingproblemsandnotjustregular changesinhowthingsnormallyoperate.Thisisespecially importantinlargebusinesssettingswheretoomanyfalse warnings cancausepeopletoignorerealissuesandlose confidence in the automatic tools used to watch over systems. A recall value of 94.1% shows the model is good at finding mostunusualcostevents,whichhelpsreducethechanceof missing out on financial risks. The F1-score of 94.6% showsa good balance between how accurately anomalies are found andhowmanyarecaught, makingsurethatanomaliesareidentifiedreliablywithoutmissingtoomany.

Thesystemhasafalsepositiverateof3.7%,whichis quite low. This shows that it works reliably without sending too many false warnings. Unlike static systems that use fixed limits which might not keep up with increasing workloads,the Isolation Forest automatically learns and adapts to the actual billing data patterns. This flexibility allows the modelto tell the difference between real growth in workloads and unexpected increases in costs due to errors in setup, too much automatic scaling, or services being turned on by mistake.

Combining anomaly detection signals with predictions of future issues helps improve the understanding of the situation. When forecast errors match up with anomaly scores, the system increases the risk severity, which helps with better prioritization and assessing financial risks.

C. SystemPerformanceMetrics

Besides checking how well the system predicts and detects issues, they also looked at other performance measures to see if the system can work smoothly in real time. The LSTM- basedforecastingmoduletookanaverageof1.2seconds foreachinferencecycle,showingitcanprocesstime-series data efficiently in sequence. The anomaly detection was done in about 0.8 seconds, showing how fast the Isolation Forest algorithm works.

TABLEII
SYSTEM-LEVELPERFORMANCEMETRICS

Metric	Average Time (seconds)
Cost Forecasting Latency	1.2
AnomalyDetectionLatency	0.8
ReportGenerationTime	1.9
End-to-EndResponseTime	3.9

TheAI-poweredreportingagentcreatedorganizedfinancial summaries in about 1.9 seconds on average. This module changes complex numbers into easy-to-understand financial information,helpingpeoplefromtechnicalandfinancialback- grounds talk to each other more clearly.

Thetotaltimeittakesforthesystemtorespondfromstartto finish is about 3.9 seconds, which shows that the coordinated multi-agent system works quickly enough to meet real-time requirements. This ability to respond quickly lets companies keep a close eye on their cloud costs all the time and helps them fix any small changes in the bill before they turn intobigspendingproblems.Themodulardesignalsomakesiteasy to scale, letting each part work best on its own as more data comes in.

D. OverallDiscussionandInsights

Theresultsfromtheexperimentsshowthatcombiningdeep learning, unsupervised anomaly detection, and Generative AI for reporting in a single multi-agent system works well. Each part does something different in terms of analysis: the LSTM model gives reliable and precise cost predictions, the Isolation Forest helps find unusual billing patterns, and the LLM-based reporting agent makes the results easier to understand and explain.

Unlike regular cloud cost dashboards that mainly showpast data, the new framework offers smarter financial insights that can predict, diagnose, and suggest solutions for costs. Whenagentsworktogether,thesystemcanspotpotentialcost problems early, suggest ways to save money, and share useful information in a way that's easy to understand.

The results show that the proposed framework works well, giving accurate predictions, effectively finding unusual ac- tivity, having few false alarms, and running quickly enoughfor real-time use. These features show that it is useful, can grow with bigger systems, and works well for smart financial management in large cloud setups.

VI. CONCLUSION

ThispaperpresentedaMulti-AgentGenerativeArtificialIn- telligence framework for intelligent cloud cost prediction and resource optimization. The proposed system integrates pre- dictive modeling, anomaly detection, risk-aware optimization, and automated report generation into a unified architecture. Unliketraditionalclouddashboardsthatprovideonlyhistorical insights, theproposedframework enablesproactivefinancial decision-makingusingpredictiveanalyticsandintelligentrec- ommendations.

TheLSTM-basedmultivariatetime-seriesmodeleffectively capturescomplexcostpatterns,workloadvariations,andlong- termbilling dependencies. TheIsolationForest-basedanomaly detection module successfully identifies abnormal cost behav- iorswithminimalfalsealerts.Furthermore,therisk-awareop- timization mechanism converts analytical insights into action- able cost-saving strategies. The LLM-based reporting agent improves interpretability by transforming analytical outputs into structured and human-readable financial summaries.

Experimental evaluation demonstrates strong performance, achieving a high prediction accuracy with an R^2 score of 0.96 and a low false positive rate of 3.7

VII. FUTURE WORK

Although the proposed framework demonstrates effective performance in cloud cost prediction and anomaly detection, several enhancements can be explored in future work. Rein- forcement learning techniques can be incorporated to enable autonomous decision-making for dynamic resource allocation and real-time cost optimization. Such an approach can allow the system to continuously learn from cloud usage patternsand automatically adapt to changing workloads.

Transformer-based time-series models, such as attention- driven architectures, can be investigated to improve long-term forecasting accuracy in highly dynamic cloud environments. These models can better capture temporal dependencies and complex usage patterns compared to traditional deep learning approaches.

Additionally, real-time streaming analytics frameworks can beintegratedtosupportcontinuousmonitoringofcloudbilling data and enable faster anomaly detection.

This would help organizations identify unusual cost spikes instantly and take proactive corrective actions.

Federated learning approaches may also be explored to en- able privacy-preserving multi-cloud intelligence without shar- ing sensitive billing or infrastructure data. This would allow multiple organizations or cloud platforms to collaboratively improve model performance while maintaining data confiden- tiality.

Furthermore, sustainability-aware optimization strategies can be introduced to balance cost reduction with energy- efficient cloud resource utilization. Incorporating green com- puting metrics and carbon-aware scheduling can make the proposed framework more environmentally sustainable.

©IJRASET: All Rights are Reserved

These enhancements can extend the proposed framework towardafullyautonomous,adaptive,andintelligentcloudcost management system capable of operating efficiently in real- world multi-cloud environments.

VIII. ACKNOWLEDGMENT

Theauthorswouldliketoexpresstheirsinceregratitude to the Department of Artificial Intelligence and Data Science, MisrimalNavajeeMunothJainEngineeringCollege, Chennai, for providing the necessary academic support, infrastructure, and research facilities to carry out this work successfully.

The authors extend their heartfelt thanks to Ms. Angela Lincy N H, Assistant Professor, Department of Artificial Intelligence and Data Science, for her continuous guidance, valuable suggestions, and technical expertise throughout the development of this research. Her encouragement and con- structive feedback played a vital role in improving the quality of this work.

Theauthorsalsoacknowledgethevaluablecontributions of the co-authors for their collaborative efforts, technical support,andactiveinvolvementinmodeldevelopment,exper- imentation, and validation of the proposed framework. Their cooperation and dedication significantly contributed to the successful completion of this research work.

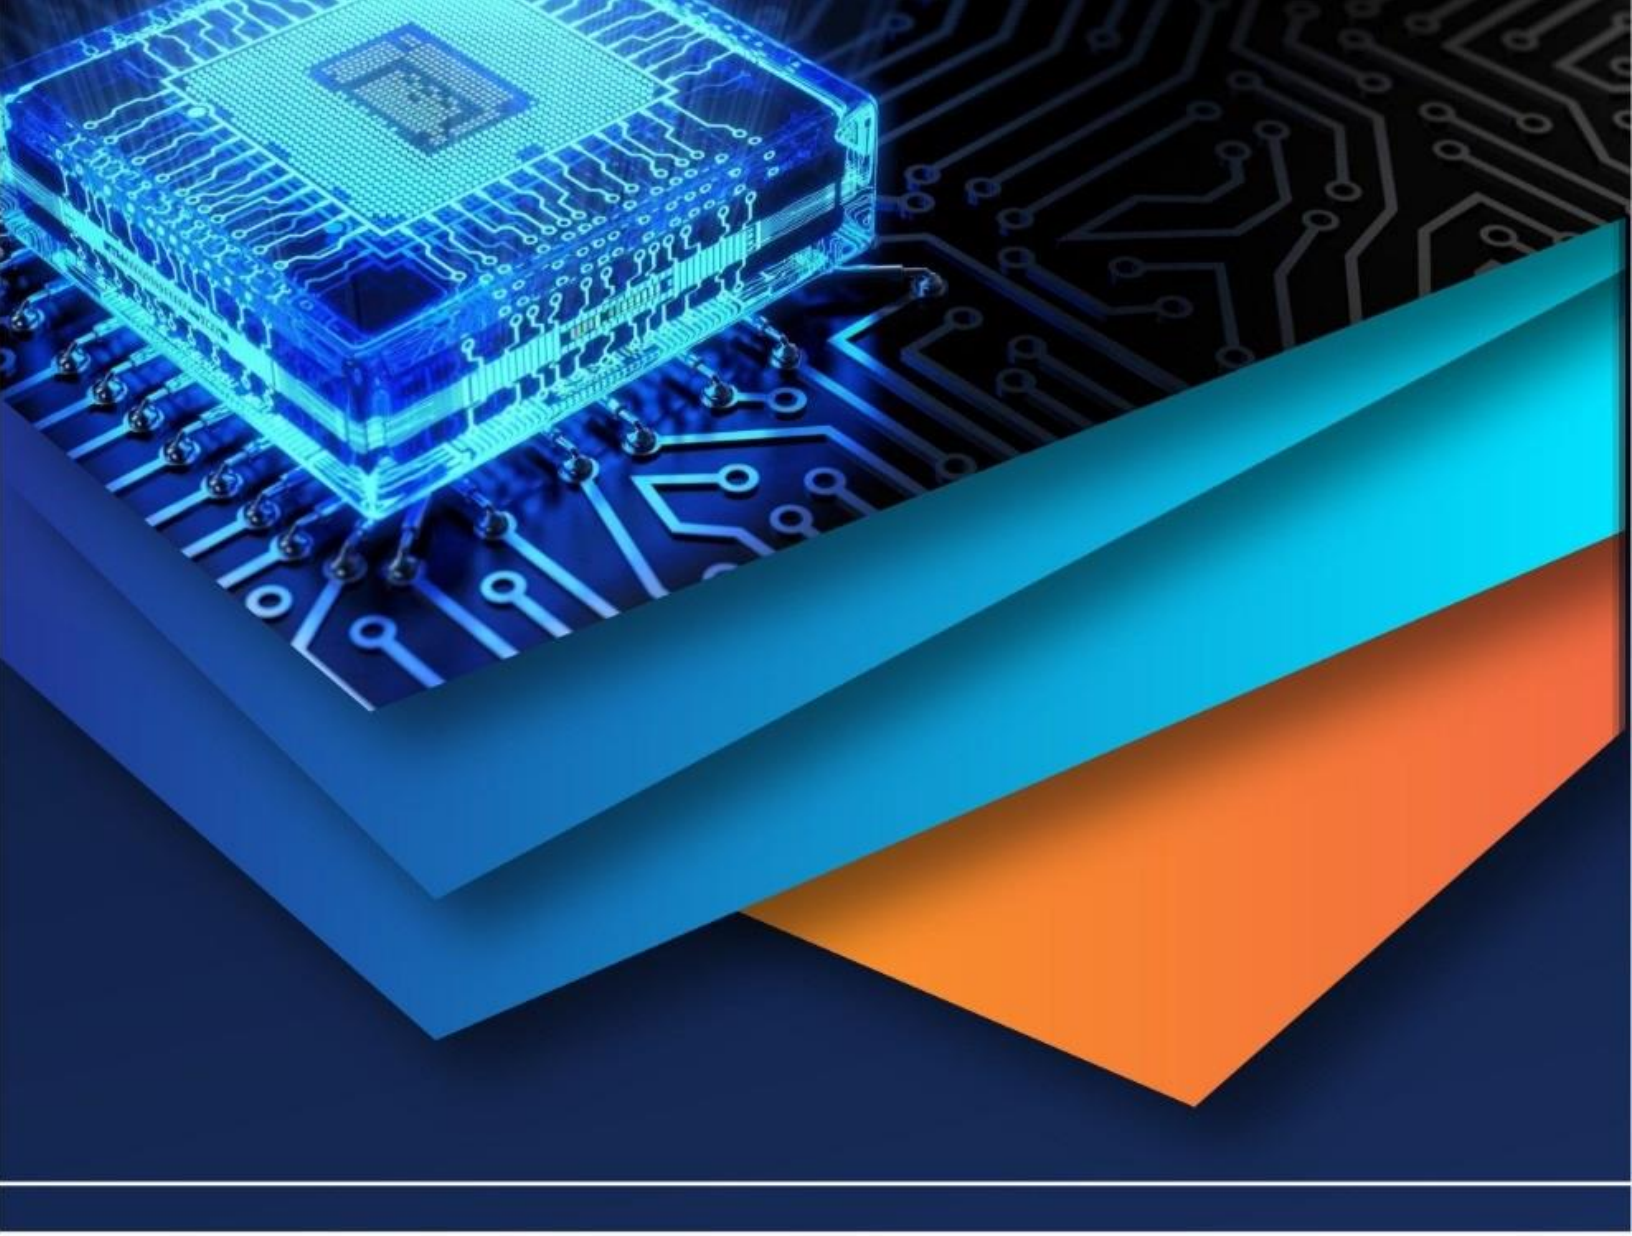
The authors also express their appreciation to the faculty members of the AI & DS department for their support and insightfuldiscussionsduringtheresearchprocess.Theirinputs helped refine the methodology and strengthen the proposed framework.

Finally, the authors acknowledge the contributions of re- searchersandscholarswhoseworkhasbeencitedinthispaper. Their studies in cloud computing, cost optimization, anomaly detection, and deep learning provided important insights that significantly influenced the design and implementation of the proposed system.

REFERENCES

[1] G. N. Junior and C. Marcon, “OPTIC – Optimized Translation andInteraction for Cloud-Costs: Use of FOCUS Standardization and LLMIntegration for Cloud Billing Analysis,” IEEE International Conferenceon Cloud Computing, 2025.
[2] S. Santhosh, “Cloud-Based Software Development Lifecycle,” Interna-tional Journal of Software Engineering and Applications, June 2023.
[3] Y. Zhang, T. Zhang, and Q. Li, “Cloud Resource Cost Prediction UsingDeep Learning Approaches,” IEEE Transactions on Cloud Computing,2022.
[4] A.Choudhary,“ComparativeStudyofVariousCloudServiceProviders,” InternationalJournalofComputerApplications,2022.
[5] N.Bhatte,“ComparisonofDifferentCloudProviders,”InternationalJournal of Advanced Research in Computer Science, June 2022.
[6] A. M. Mohamed, “Multicriteria Optimization Model for Cloud ProviderSelection,” Journal of Cloud Computing, May 2019.
[7] F.T.Liu,K.M.Ting,andZ.H.Zhou,“IsolationForest,”IEEEInternational Conference on Data Mining, 2008.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)