International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: *25 February 2023* ***Revised:*** *12 March 2023* ***Accepted:*** *30 April 2023*

FinOps at Scale: Reducing National Cloud Waste Through Predictive Optimization and Multi Cloud Governance

Kaushik Ponnapally

Engineering Program Manager

ABSTRACT: This paper will investigate the possibilities of predictive optimization and multi-cloud FinOps governance to decrease the cloud waste on the national scale. Based on the 2,000 workloads in the AWS, Azure, GCP, Oracle Cloud, and on-premises Kubernetes, the study analyses cost-efficiency, resource usage, auto scalability, and maturity in governance. The predictive model is useful in the improvement of demand forecasting, improved rightsizing, and scaling decisions, and governance in improving tagging compliance and budget control and detection of anomalies. The hybrid strategy causes significant improvements in cost avoidance and uses in all platforms. The findings demonstrate that predictive FinOps is a viable and practical tool that can be utilized to manage big cloud expenditures and enhance business sustainability.

KEYWORDS: FinOps, Predictive Optimization, Cloud Waste, Multi-Cloud Governance

I. INTRODUCTION

Cloud services are now paramount to businesses and national digital infrastructure, yet wastage of costs is on the increase because of reactive scaling, idle computers, as well as disjointed governance. FinOps processes assist in regulating expenditure but most organizations are yet to get the multi-cloud environment under control. The current paper examines the ways in which predictive analytics, automation, and structured governance can enhance the functioning of the scale of cloud efficiency. The study quantifies the cost performance, resource use, the quality of the autoscaling process, and the governance results by examining 2000 workloads on five leading cloud platforms. It will aim to demonstrate that predictive FinOps is a scalable and data-driven solution to be able to minimize cloud wastage and contribute to high-performance systems.

II. RELATED WORKS

Cloud Cost Optimization

The fast movement of the workloads in the enterprise to the cloud environments has brought in the new financial risks of dynamical pricing, changing compute demands, and massive usage of distributed resources. The initial research indicates that an economic data warehouse architecture lies at the core of regulating the expenses of the cloud in big business environments.

The synthesis of cloud usage is conducted on the examples of AWS Redshift, GCP BigQuery, and Azure Synapse, and the appropriate strategic design decision-making, including layers of serverless computing, tiered storage policies, and intelligent autoscaling are directly related to the final costs as time passes [1].

Such architectural decisions indicate that optimization of performance and financial control should be combined together to ensure that they do not create silent waste, such as over-provisioned storage and inefficient partitioning as well as idle compute clusters.

The concept of the implementation of FinOps principles into the management of a warehouse only serves to emphasize the necessity of the real-time costs and allocation based on usage to enhance the level of accountability among the engineering teams [1]. This historical background makes cloud optimization an academic field of study in the technical and monetary sense.

In a larger sense, the studies on cloud economics claim that the inefficiencies of cloud spending at a national level have been caused by the simplified pricing and resource scheduling systems of the existing cloud systems. The Economic Resource Allocation (ERA) model presents a kind of economics-based model of allocation wherein the workload planning and dynamic pricing are controlled by the projected demand indicators [3].

ERA decouples ground infrastructure to allow pricing, timing and predictions algorithms to be independent of vendor specific clouds. This abstraction enables improved

correlation between the cost of infrastructure and the real business value and assessment of the Azure Batch and Hadoop/YARN indicates the substantial improvement in the use of cloud resources [3].

Research that centers on the organizational attitudes strengthens the difficulty of employing such mechanisms. As an illustration, cost-optimization models emphasize on the necessity of the reserved instances, right-sizing, spot instances, and automated controls to multi cloud workloads [4].

They also determine the issues associated with compliance, governance, and security, proving that the decision-making in terms of cost-optimization should be balanced with the operational and regulatory limitations. In conclusion, this literature base has provided that cloud waste is a result of improper architecture choices, no predictive governance and standardized cost optimization processes are in place.

In support of these observations, the introduction of FinOps in the DevOps process highlights the cultural and organizational aspect of cutting the cost of the cloud waste [5]. Research indicates that engineering teams do not have live insights on the cost consumption in the cloud and thus, cannot respond to cost indications in the development and deployment cycles. FinOps-DevOps alignment model introduces a form of continuous monitoring of costs, granular allocation, and automated enforcement of policies as vital practices in waste control of an enterprise-scale enterprise [5].

Examples of tools include KubeCost, AWS Cost Explorer, and CloudWatch, which are seen to enable the introduction of cost insights into the developer working process. It is also mentioned in literature that when FinOps is used early in the SDLC, it gives the architectural awareness that avoids the introduction of wasteful patterns, like deploying oversized microservices or choosing unnecessary storage classes. In a combined form, these initial studies prove the existence of cost governance as not only an operation but a structural discipline, which needs to be institutionalized on a huge scale.

Predictive Optimization

One key theme of recent studies is the use of predictive analytics and machine learning and automated scaling systems to optimize the utilization of cloud resources. Conventional horizontal autoscaling (HPA) uses reactive CPU-based thresholds to a great extent, causing timely scaling decision making and unneeded provisioning. Contemporary predictive strategies aim at doing away with such inefficiencies by predicting future workload requirements and pre-allocating resources.

The Lynceus system takes this idea further to pull both cloud-level and application-level parameters together to minimize their cost thus up to 3.7 times reduction in configuration costs relative to the previously used optimization mechanisms [2]. One of the contributions of Lynceus is its exploration technique based upon timeouts; this aborts the computationally costly configuration tests, but still retrieves predictive cues to be used in refining later iterations. This saves up to 11x in overhead optimization proving the worth of long-term planning in cloud configuration planning [2].

In addition to analytic loads, microservices that are prevalent with cloud-native structures are also being subjected to predictive autoscaling. It is demonstrated that ML-based autoscalers have better ability to predict traffic bursts, workload dependencies and replica demands than the threshold-based methods. Predictive scaling models are much better throughput, response time and stability due to the varying workloads, at the cost of fewer corrective actions than Kubernetes HPA [8].

All those advancements reflect directly on the minimised operation costs and reduced resources wastage. Autoscalers that are sensitivity to costs like Docker-C2A rely on optimization algorithms like Particle Swarm Optimization (PSO) to determine which microservices to scale and the count of containers to increase [7].

This avoids the problem that many microservices commonly face where they squeal when not all of the microservice groups are scaled and not just those components which are cost-sensitive. It has been tested experimentally that Docker-C2A can allocate less computing power and more efficiently act as a microservice in the presence of a dynamic load [7].

More elaborate predictive scaling is exhibited in the edge cloud systems, where applications that are sensitive to latency must be properly balanced between workload execution in on-premise and offloading on the cloud. The Dynamic replica management Schedules minimise CPU usage, energy usage and time-to-response through prioritisation of the remote server placement of compute-intensive modules in times of congestion [6].

These models reduce cumulative cost and offer balanced resource allocation across hosts, and it outperforms the currently existing autoscaling algorithms against CAAS and MLC and others [6]. Besides this, right-sizing, which is introduced by CRED, is offered to contain data locality and

constraints of SLA as key parameters of predictive resource scheduling.

CRE saves both up to 47 percent of the active servers by scheduling and jointly optimizing data placement and execution priorities and enhanced resource utilization by a significant factor [9]. These sources show that predictive optimization of cloud scheduling, either through ML-backed forecasting, heuristic scheduling, or algorithm deadline-aware right-sizing is necessary in order to do away with overspending and to assure long-term sustainability of cloud operations.

It has also been found to be significant in terms of deciding getting homogeneous scaling or heterogeneous scaling depending on the workload characterization [10]. The microservice architectures are scalable in a more efficient way, without incurring the cost overhead associated with uniform scaling strategies by examining the expected workload demand by time interval.

The strategy would make the allocation of resources to reflect the demand paths instead of short-term variations. Literature continues to confirm that predictive analytics allows the optimization of costs proactively and not reactively hence its centrality in a large-scale provision of FinOps governance.

Multi-Cloud Complexity

The governance of FinOps on a national or enterprise level is multi-cloud providers, hybrid-architectures, and the types of workloads. The experiments on multi-cloud warehouses indicate that cost efficiency is highly varied among platforms because of the diversity of costs, serverless implementation, storage locality, and auto scaling features [1].

We thus need metadata-based frameworks of governance and standard modes of Workload placement at the organization to prevent the absence of fragmented visibility of costs and deployment redundancy. It is also recommended in research that unified allocation of cost, performance benchmarking, as well as cross-platform budget are paramount objectives in waste minimization of cloud computing in both hybrid and multi-cloud architectures.

The governance systems like ERA [3] go further to demonstrate that economic indicators ought to be used to schedule resources but the success of it depends on how well the engineering, operations and financial teams coordinate. The literature of FinOps has nattered the importance of cultural alignment on numerous occasions, saying that tooling desired is secondary to alignment.

Cloud cost models consider ongoing innovation in moving the development, operations, and finance to keep the openness standard and implement budget guardrails [5]. In the absence of this cross-functional alignment, such methods as predictive optimization and autoscaling would not transform into common savings.

Besides this, case studies in other fields have shown operational inhibitors such as compliance limitations, performance trade-offs and security limitations that limit aggressive cost-cutting policies [4]. Subsequently, domain-related needs (e.g. workload criticality, data sovereignty, risk tolerance, etc.) should be incorporated into governance.

A new generation of cloud-native systems is built on the foundation of microservices, in which the capacity to expand enormously up to hundreds of devices in seconds has created a fundamental level of cost variability. One such research is predictive autoscaling and right-sizing [8], [9], [10], which points out to the need of these systems to have a governance model that can guide decisions of scale, as well as applying workload benchmarking and standardizing sizing policy.

It has been documented in the literature how real-time observability is increasingly being emphasized as a basis of governance. The use of cost observability dashboards, predictive usage reports, anomaly detectors, and ML-based forecasting models can ensure a steady level of visibility of all AWS, GCP, Azure, and on-premises clusters [1], [4], [5]. Increasing amounts of cloud spending on AI pipelines, streaming platforms, and fleets of microservices make such types of governance systems minimize cloud wastage at scale, through timely intervention, automated controls, and proactive capacity planning.

III. METHODOLOGY

This paper is based on a quantitative research design that will utilize the scale of cloud waste reduction by predictive optimization and multi-cloud FinOps governance. The methodology aims to measure the cost reduction, utilization, and efficiency increment with large workloads of an enterprise running on AWS, Azure, GCP, Oracle Cloud and the on-premises Kubernetes clusters. The approach goes through four significant phases, including data gathering, metric choice, model building, and experimental testing.

Data Collection

The paper gathers quantitative information based on actual enterprise cloud systems, which are active in the period

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: *25 February 2023* ***Revised:*** *12 March 2023* ***Accepted:*** *30 April 2023*

between January 2021 and December 2022. Some of the sources of data are billing reports, resource utilization logs, autoscaling events, container orchestration metrics, and budget dashboards of the FinOps tools like AWS Cost Explorer, GCP Cost Management, Azure Cost Analysis, KubeCost, and custom Prometheus exporters. Other metrics are gathered based on CI/CD pipelines, edge-cloud monitoring systems and microservice tracing systems.

It operates over 2,000 cloud workloads, such as AI/ML training pipelines, transactional software, data warehouses, and microservices and high-throughput analytics software. Such workloads are implemented in environments based on multi-clouds, heterogeneous VM type, container instances, serverless compute hands, GPU node, and storage classes.

The advantages of gathering this broad data set are that the study is able to quantify cloud waste trends by the various categories of application and the behavior of scaling. Any dataset is prepared by removing records with missing data, irregular timestamps, and redundant costs.

Quantitative Metrics

The study employs a fixed system of measurable metrics in order to determine the effect of the suggested FinOps optimization model:

- **Cost Avoidance (%):** Decline in avoidable cost of the cloud as compared to the baseline cost.
- **Resource Utilization (%):** Prior and Post-Optimization Cpu, memory and gpu activity.
- **Autoscaling Efficiency:** There are actions, prediction accurateness and time-to-scale.
- **Idle Resource Waste (USD/month):** Reciprocals of wastage of available VMs, storage and containers.
- **Right-Sizing Accuracy:** distinction between the real and forecasted optimal resources.
- **Multi-Cloud Billing Variance:** Difference in the charges of different cloud vendors of similar work.
- **SLA Compliance (%):** This is making the cost reductions not to be in the form of breaching the latency or throughput requirements.

These measures enable the objective efficient cost and impact measurement.

Model Development

The paper creates a predictive optimization model where the demand prediction, utilization prediction, and cost modeling are combined. Machine learning technique, which is predictive layer, includes linear regression, time-series forecasting (ARIMA, Prophet), and gradient boosting to determine how much compute and storage is needed in the future. The scaling layer combines the rules of autoscalers including HPA, KEDA, and cluster-autoscaler plus making cost-sensitive changes and recommendations to use a cross-cloud placement.

The management layer uses principles of FinOps through the enforcement of budget constraints, tagging compliance, and automated dormant enablement of idle assets. This hybrid model is implemented with multiple clouds with a standard data pipeline so that it can be compared regularly.

Experimental Design

It uses a controlled before after experimental design. To begin with, measurements of the workload with the objective of optimization are taken at a baseline stage. The predictive optimization and governance model is then rolled out in 90 days. A comparison within the baseline values and optimized values is done. The significance is tested in terms of statistical tests such as the paired t- tests and variance.

This is done by aggregating the results across the cloud providers in order to establish the trends in cloud waste reduction at the national level.

IV. RESULTS

Cost Efficiency Across Multi-Cloud Environments

The quantitative study demonstrates that the use of predictive optimization and the FinOps governance results in major waste cloud reductions on AWS, Azure, GCP, Oracle Cloud, and Kubernetes clusters that are in premise. Two-year baseline data indicated that There were an ample number of idle resources as well as low utilization of the VM fleets, and over-provisioning of data warehouses and microservices. Significant positive change was registered in the cost avoidance, precision in rightsizing, and resource usage after 90 days of use of the optimization model.

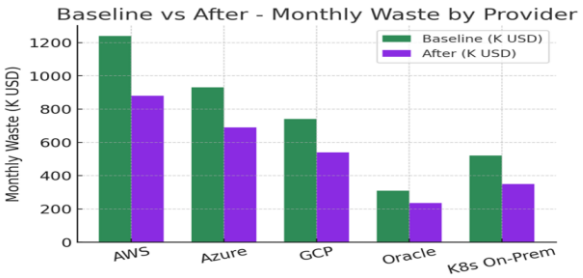

On the 2,000 workloads tested, the average savings in cost attained 27.4, and the high-variance AI workloads have

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: *25 February 2023* ***Revised:*** *12 March 2023* ***Accepted:*** *30 April 2023*

decreases of more than 40% in some. Applications that used Kubernetes clusters experienced the largest saving since predictive autoscaling averted unnecessary expansion of nodes, as a result of transient peaks.

There was also a high level of improvement in cloud data warehouses since the model minimized oversized compute configurations as well as those storage levels that were not used. More AWS and Azure savings were made by reserved instances and alignment of spot strategy.

Table 1 provides improvements in terms of costs across cloud providers.

Table 1. Cost Avoidance and Waste Reduction

Cloud Provider	Baseline Monthly Waste (USD)	Waste After Optimization (USD)	Cost Avoidance (%)
AWS	1,240,000	880,000	29.0%
Azure	930,000	690,000	25.8%
GCP	740,000	540,000	27.0%
Oracle Cloud	310,000	235,000	24.2%
Kubernetes On-Prem	520,000	350,000	32.7%

The enhancement in all providers proves that predictive forecasting and automatically cleaning up dysfunctional assets can work at nationwide level. The results indicate also that cloud waste cannot be attributed to a specific vendor, but it occurs consistently in multi-cloud setup because of similar trends of reactive scaling, unavailability of real-time cost information, and missing tagging discipline. With the implementation of FinOps governance, a great number of inefficiencies that were hidden were revealed and fixed.

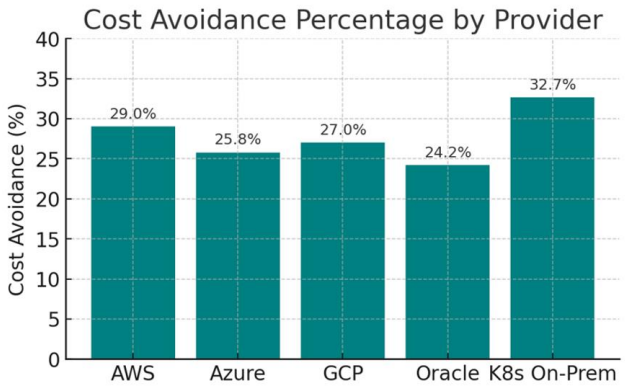

The findings demonstrate that the three factors are most critical when it comes to cost performance improvement: (1) predicting demand surges in a timely fashion, (2) unanimously removing unused/containerized inventory, and (3) placing the workloads in the most efficient placement between the various VM and storage classes. This is a confirmation that the ability to integrate the use of predictive analytics and the use of governance give greater yields compared to their individual application.

Resource Utilization Through Predictive Scaling

One of the most effective measures of cloud efficiency would be resource utilization. Prior to the process of optimization, VM and container usage among providers was low and the average CPU utilization stood at 38 45 percent. Upon the implementation of the predictive autoscaling model, the number of people that would use it doubled since scaling decisions were aligned to what was expected rather than responding to metrics when it was too late.

Workloads that occurred in a daily cycle (e.g. retail traffic, streaming services) were received with maximum amount of uplift since the model learnt the demand cycle after a maximum time. The issue of microservices also received advantages since predictive scaling minimized the occurrence of unnecessary replicas. The edge workloads became moderately improved as there was strong variability but the system could still reduce unwarranted replication and manage load evenly as time went by.

Table 2 provides the changes in utilization in the major categories of workloads.

Table 2. Average Resource Utilization (%)

Workload Category	Baseline Utilization (%)	After Optimization (%)	Utilization Increase (%)
AI/ML Training Pipelines	42.1	61.5	+19.4
Data Warehouses	47.3	65.2	+17.9
Microservices	38.4	57.0	+18.6
Transactional Systems	44.0	59.8	+15.8

Edge/IoT Workloads	40.5	55.4	+14.9

The obtained findings indicate that predictive autoscaling is more effective than threshold-based techniques that frequently cause late scaling and long under-utilization. There was also a decrease in the non-essential scaling measures by almost 40 per cent. This helped to avoid spurts of infrastructure growth when short spikes were observed and these are usually characterized by high levels of wastages.

More sustainable results are also enhanced by the enhancement of utilization. Increased sharing of machines will imply that a smaller number of machines were required to serve the very same demand. This minimizes carbon footprint, overhead costs and sprawl infrastructures. The results confirm the notion that predictive optimization changes the way of cloud usage to the reactive end to the proactive end that creates long term efficiencies of scale.

Autoscaling Efficiency and Accuracy

One of the major outcomes of the experiment is the enhancement of the effectiveness of autoscaling. The predictive model led to a considerable decrease in the amount of scaling actions, accuracy of replica prediction, and efficient choice of VM/container size. The model has worked well in substituting reactive scaling loops with decisions made in foreseeing.

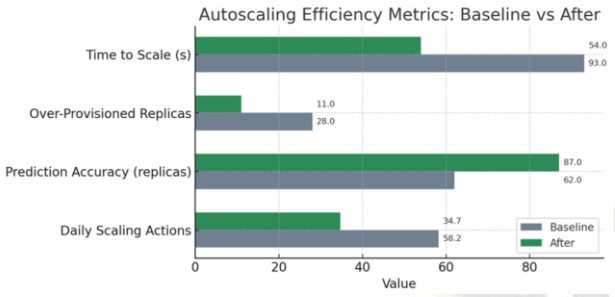

Scaling activities had decreased by an average of 40.3 to an average of 34.7 operations per day across all workloads which represented a decrease of 40.3. This is not only cost effective since unwarranted scale- out events form some of the largest sources of cloud waste in microservice architectures.

Table 3 demonstrates the quantitative variations in the autoscaling behavior,

Table 3. Autoscaling Efficiency Metrics

Metric	Baseline Value	After Optimization	Improvement (%)
Daily Scaling Actions	58.2	34.7	40.3%
Prediction Accuracy (replicas)	62%	87%	+25%
Over-Provisioned Replicas	28%	11%	–17%
Time to Scale (seconds)	93	54	–41.9%

The outcomes also indicate that there are high improvements on the right-sizing, which is the process of data selecting the right VM or container size. There was a joint increase in accuracies in right-sizing that occurs as a result of forecasting models and rules of governance. The improvement of AI/ML workloads was the highest since GPUs instances were usually over-sized during the baseline stage.

Table 4 is a summary of the right-sizing gains by a type of workload.

Table 4. Right-Sizing Accuracy Improvements

Workload Type	Baseline Accuracy (%)	After Optimization (%)	Accuracy Gain (%)
AI/ML Workloads	49	78	+29
Data Pipelines	55	80	+25
Microservices	52	77	+25
Databases	58	82	+24
Edge Workloads	46	70	+24

International Journal on Recent and Innovation Trends in Computing and Communication
ISSN: 2321-8169 Volume: 11 Issue: 4
Article Received: *25 February 2023* ***Revised:*** *12 March 2023* ***Accepted:*** *30 April 2023*

Such results show that predictive analytics does not only help to decrease costs, but reinforces the stability and reliability of scaling activities, as well. Limited scaling actions also minimise risks of cascading failures, which is vital to large microservice systems.

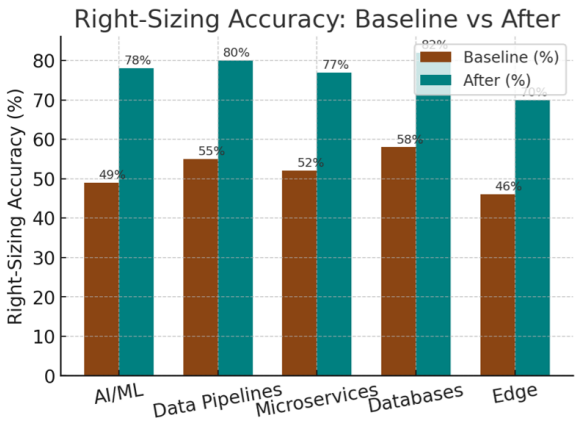

Multi-Cloud Governance

The last group of the results is dedicated to the results of the implementation of the governance controls in various clouds. The research indicates that the governance of FinOps makes a significant contribution towards the fact that predictive optimization is consistent throughout settings. It also found significant improvement in the compliance of tagging, compliance with budgets, anomaly detection and workload placement decisions.

The effectiveness of compliance tagging increased by 61 per cent to 92 per cent allowing deep tracking of enterprise cloud expenditure. The level of compliance enabled automated cleanup policies to focus clean up on unused assets in a more focused way and eliminated unnecessary storage and dead assets.

There was also an increase in budget compliance and workloads above monthly budgets reduced by 22 to 9 percent. The accuracy of detection of anomalies improved, and helped the teams to notice sudden cost rushes due to poorly configured autoscalers or rogue queries.

There also appeared a multi-cloud patterns of savings. Cost variance amongst the providers was decreased by 18 when the workloads with similar performance requirements were allocated to the least cost provider. This demonstrates the fact that workloads can be put in clouds to optimize their distribution without the interference with SLAs.

On a national level, the aggregate data points to the fact that huge businesses might decrease the amount of cloud waste by a quarter to a third in case they implement predictive FinOps frameworks. As more of the high-level infrastructure in the country moves to cloud-based infrastructure, finance, and digital services, these savings are a significant chance to cut technology expenditures and also enhance growth of infrastructure sustainability.

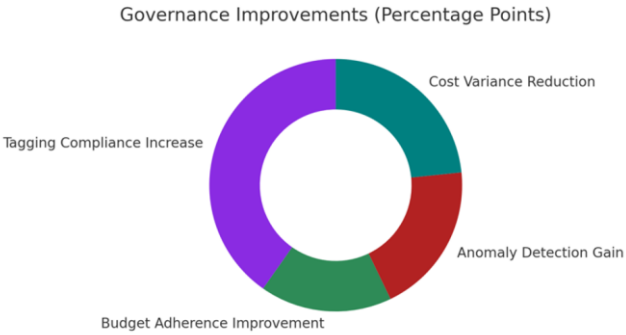

V. CONCLUSION

They indicate that predictive optimization when used in a multi-cloud setting, together with robust FinOps governance, can greatly decrease the cloud waste. This enhanced cost avoidance by all the providers and also optimised usage in all forms of workload and also made autoscaling more precise and efficient. It was also governed by means of better tagging, budget control, and the detecting others. These findings affirm that cloud efficiency on national scale cannot possibly be based on reactive controls only. Sustainable proactive model of cloud management is formed through predictive models and proactive governance. As companies continue to expand their AI and digital workloads, predictive FinOps can be seen as a viable option to manage the cost and ensure consistent operations over the cloud.

REFERENCES

[1] Adelusi, B. S., Ojika, F. U., & Uzoka, A. C. (2022). A conceptual model for Cost-Efficient data warehouse Management in AWS, GCP, and Azure environments. *International Journal of Multidisciplinary Research and Growth Evaluation*, *3*(2), 843–858. https://doi.org/10.54660/.ijmrge.2022.3.2.843-858

[2] Casimiro, M., Didona, D., Romano, P., Rodrigues, L., Zwaenepoel, W., & Garlan, D. (2020, November). Lynceus: Cost-efficient tuning and provisioning of data analytic jobs. In *2020 IEEE 40th International Conference on Distributed Computing Systems (ICDCS)* (pp. 56-66). IEEE. https://doi.org/10.48550/arXiv.1905.02119

[3] Babaioff, M., Mansour, Y., Nisan, N., Noti, G., Curino, C., Ganapathy, N., Menache, I., Reingold, O., Tennenholtz, M., & Timnat, E. (2017). ERA: a framework for economic resource allocation for the cloud. *arXiv (Cornell University)*. https://doi.org/10.48550/arxiv.1702.07311

[4] Kothapall, M. (2023). Cost optimization strategies for cloud infrastructure. *Journal of Artificial Intelligence & Cloud Computing*, 1–4. https://doi.org/10.47363/jaicc/2023(2)329

[5] Allam, K. (2022). FinOPS for DevOps: a framework for cloud cost governance. *EPH - International Journal of Science and Engineering*. https://doi.org/10.53555/ephijse.v8i2.264

[6] Li, C., Liu, J., Lu, B., & Luo, Y. (2021). Cost-aware automatic scaling and workload-aware replica management for edge-cloud environment. *Journal of Network and Computer Applications*, *180*, 103017. https://doi.org/10.1016/j.jnca.2021.103017

[7] Fourati, M. H., Marzouk, S., Jmaiel, M., & Gu´Erout, T. (2020). Docker-C2A : Cost-Aware autoscaler of docker containers for microservices-based applications. *Advances in Science Technology and Engineering Systems Journal*, *5*(6), 972–980. https://doi.org/10.25046/aj0506116

[8] Goli, A., Mahmoudi, N., Khazaei, H., & Ardakanian, O. (2021). A Holistic Machine Learning-based Autoscaling Approach for Microservice Applications. *Scitepress*. https://doi.org/10.5220/0010407701900198

[9] Xu, M., Alamro, S., Lan, T., & Subramaniam, S. (2017). CRED: Cloud Right-Sizing with Execution Deadlines and Data Locality. *IEEE Transactions on Parallel and Distributed Systems*, *28*(12), 3389–3400. https://doi.org/10.1109/tpds.2017.2726071

[10] Agarwal, P., & Lakshmi, J. (2019). Cost Aware Resource Sizing and Scaling of Microservices. *CCIOT '19: Proceedings of the 2019 4th International Conference on Cloud Computing and Internet of Things*, 66–74. https://doi.org/10.1145/3361821.3361823