ijamrsd.com
JAMRSD
ISSN: 2583-6404
Journal of Advanced Multidisciplinary Research Studies and Development
Peer Reviewed Referred Journal Bi-monthly Open Access

CLOUDOPTIMA: A STRATEGIC VISION FOR A HOLISTIC CLOUD COST AND SUSTAINABILITY FRAMEWORK

Rushikesh Londhe[1], Dr. Akshata Nayak[2], Dr. Ajay Barve[3]
[1] *Masters Student (Computer Science), Computer Science Department, Kirti M. Doongursee College of Arts, Science and Commerce, Mumbai University, India, Maharashtra, Mumbai.*
Email: rushikeshlondhe942@gmail.com
[2] *HOD Computer Science, Computer Science Department, Kirti M. Doongursee College of Arts, Science and Commerce, Mumbai University, India, Maharashtra, Mumbai.*
[3] *HOD IT, Information Technology Department, Kirti M. Doongursee College of Arts, Science and Commerce, Mumbai University, India, Maharashtra, Mumbai.*

Abstract
CloudOptima is a conceptual, multi-dimensional framework that outlines how organizations can move from reactive cloud cost control to proactive, value-focused cloud financial and sustainability management. The framework combines four dimensions: technical efficiency, financial intelligence, governance and culture, and sustainability and risk, into one vision for future implementation. This paper does not describe any implemented model, prototype, or experiment. Instead, it presents a vendor-neutral architecture and a step-by-step roadmap to support the future development and adoption of CloudOptima in academic and industrial settings. The design is based on established cloud cost optimization practices, FinOps principles, and new ideas in carbon-aware computing. It outlines layers for data and telemetry, analytics and modeling, policy and governance, execution and automation, and user experience. A four-phase implementation path is suggested, beginning with basic visibility and tagging, moving through structured optimization routines, and ultimately including automation, analytics, and dual-objective cost and carbon decision-making. This contribution is a straightforward, stand-alone guide that shows how organizations, including teams with limited resources, could implement CloudOptima using open-source tools and free-tier cloud accounts. This conceptual paper lays the foundation for future prototypes and evaluation studies without exaggerating current technical progress.
Keywords: CloudOptima; Cloud Cost Optimization; Finops; Cloud Governance; Carbon-Aware Computing; Conceptual Framework.

► *Corresponding Author: Rushikesh Londhe*

1. Introduction

Cloud computing has become the go-to model for many modern applications. It provides quick setup, a global presence, and pay-as-you-go pricing (Babu & George, 2022). However, organizations often find it hard to manage cloud costs and relate spending to actual business value (Chaisiri et al., 2012).

Common problems include unused or forgotten resources, oversized instances, unmanaged test environments, and complicated bills that make it hard to see which teams or products are driving expenses. At the same time, boards and regulators increasingly want clear information about the environmental impact of digital infrastructure. They connect cloud strategies to broader sustainability and ESG goals.

Cloud Financial Operations (FinOps) has emerged to tackle these issues by encouraging shared responsibility for cloud costs among engineering, finance, and business teams. (Gupta & Amza, 2021; Jones & Smith, 2022) The FinOps Framework defines capabilities such as cost allocation, budgeting, forecasting, and optimization cycles, while cloud providers publish best-practice guidance, for example in cost optimization pillars of well-architected frameworks. However, many organizations adopt these ideas in a fragmented way: a dashboard for visibility, a tagging policy with limited enforcement, or occasional rightsizing exercises, without a unifying architecture or phased roadmap (Jones & Smith, 2022).

This paper proposes CloudOptima as a strategic, vision-only framework intended to integrate these scattered practices into a coherent, multi-dimensional model. The centre of gravity is conceptual: no code has been implemented and no experiments have been conducted. Instead, the paper defines the problem context, articulates high-level objectives, and describes a framework structured around four conceptual pillars and a five-layer architecture, together with a pragmatic four-phase implementation roadmap. The aim is to provide a blueprint that researchers or practitioners can adopt and later instantiate using open-source tools and cloud free-tier resources.

2. Problem Context and Vision Objectives

2.1 Multi-Dimensional Cloud Challenges

Many organizations struggle to see who owns which resources and why. This lack of visibility makes it difficult to allocate costs, justify budgets, or calculate return on investment.

On the operational side, issues with tagging, lifecycle management, and access control create shadow environments, duplicate services, and leftover test workloads that continue to rack up charges (Chaisiri et al., 2012; Herbst et al., 2013).

From a technical perspective, architectures often focus on performance and delivery speed. As a result, systematic rightsizing, autoscaling, storage tiering, and placement optimization are rarely used to their full potential (Gong et al., 2010; Herbst et al., 2013).

Culturally, engineering, finance, and leadership teams often use different terms and have different goals. This communication gap hinders discussions about costs and sustainability, which slows down decision-making.

Environmentally, decisions on region, instance type, and workload timing can greatly affect energy use and carbon emissions. However, these factors seldom come up in daily cloud design and operations (Jindal & Di, 2022).

Current frameworks tackle many of these issues, but they often appear as separate features or best practices instead of as parts of a cohesive, ready-to-use design..

2.2 Vision-Level Objectives of CloudOptima

Given this context, CloudOptima is defined through clear vision-level objectives rather than empirical targets:

- **Holistic scope:** The framework aims to unify technical optimization, financial strategy, governance and culture, and sustainability and risk into one coherent model, avoiding the common trap of separate, disconnected initiatives.
- **Progressive adoption**: CloudOptima is designed to support step-by-step adoption. Organizations can start with basic visibility and manual practices, then move towards automation and analytics as they improve**.**

- **Vendor-agnostic design**: Concepts are presented in a provider-neutral way. This allows implementation across major public clouds and hybrid environments without relying on a single commercial platform.
- **Democratized implementation**: Early phases are intentionally limited so they can be carried out using open-source tools and free-tier cloud accounts. This makes the framework available to students, small teams, and organizations with limited resources.
- **Feedback-driven operation:** CloudOptima focuses on regular review loops. This includes routine cost and performance reviews, updating policies, and learning from issues, so practices get better over time.
- **Sustainability awareness:** After establishing financial and reliability baselines, environmental impact becomes an important factor in choosing regions, scheduling, and making architectural decisions.
- The rest of the paper translates these goals into a conceptual framework and an implementation roadmap, clearly stating what has not yet been implemented**.**

3. Conceptual CloudOptima Framework

3.1 Four Conceptual Pillars

CloudOptima is organized around four conceptual pillars that group related practices and capabilities. These pillars are derived from common themes in existing financial operations frameworks and research on sustainable computing.

1. Technical Efficiency

This pillar encompasses resource-level practices such as rightsizing instances, tuning autoscaling policies, selecting appropriate storage tiers, optimizing data transfer patterns, and adopting efficient architectural patterns (for example, serverless for spiky workloads) (Ali-Eldin et al., 2012; Gong et al., 2010). At the conceptual level, CloudOptima defines these as reusable patterns and decision rules that future implementations could codify into recommendation engines or guardrails.

2. Financial Intelligence

This pillar covers cost allocation, showback, chargeback, budgeting, forecasting, and commitment planning, such as how much to commit to reserved capacity versus on-demand usage. The framework describes which analyses and views are required, such as per-team cost reports and simple forecast ranges, without suggesting specific algorithmic implementations (Chaisiri et al., 2012; Calheiros et al., 2015).

3. Governance and Culture

This pillar focuses on roles, policies, and behaviors. Key elements include standardized tagging policies, ownership models for accounts and applications, defined cost review schedules, and ways to build cost awareness, such as internal dashboards or simple gamification (Gupta & Amza, 2021). CloudOptima treats these as processes and policies that can be adopted even without advanced tools.

4 Sustainability and Risk

This pillar brings environmental and risk factors into cloud decisions. It encourages using carbon intensity and energy source data to compare regions and evaluating regulatory and data residency constraints. It also highlights the importance of resilience and vendor lock-in when designing cost strategies (Jindal & Di, 2022). In the CloudOptima vision, these aspects are included in policies and design guidelines instead of existing as separate sustainability projects.

5. These four pillars provide a structured approach and help ensure that future CloudOptima implementations remain balanced across technical, financial, organizational, and environmental concerns.These four pillars provide thematic structure and ensure that future CloudOptima implementations remain balanced across technical, financial, organizational, and environmental concerns.

3.2 Layered Architecture (Vision Only)

Complementing the pillars, CloudOptima is expressed as a five-layer architecture that indicates how capabilities might be realized technically in future work.

1. Data and Telemetry Layer

This layer aggregates utilization metrics (CPU, memory, I/O), logs, resource metadata (including tags), and billing records across cloud accounts. It also anticipates ingestion of external data such as regional carbon intensity and energy mix from public or provider-supplied APIs (Islam et al., 2012). In this paper, it is defined logically, without specific tools or schemas.

2. Analytics and Modelling Layer

This layer hosts the analytical logic needed for CloudOptima: cost allocation rules, simple budget variance calculations, anomaly detection on spending patterns, and basic forecasting of cost or demand. It also includes conceptual models for scoring different placement options based on cost and carbon (Calheiros et al., 2015; CostAnom Project Team, 2021). The paper describes the purpose of these models but does not implement or evaluate them.

3. Policy and Governance Layer

This layer encodes rules such as required tags, allowed regions, spending thresholds, and optimization strategies (for example, "cost first" versus "balanced cost–carbon"). It may be expressed as textual policies, decision matrices, or future policy-as-code rules that are enforced by automation.

4. Execution and Automation Layer

This layer represents scripts, scheduled tasks, or services that will eventually carry out policy-driven actions: adjusting instance sizes, enforcing start/stop schedules, applying lifecycle rules to storage, or generating recommendations (Chen et al., 2022; DB-Tune Research Consortium, 2022). At present, it is defined only by expected behaviors and interfaces. Implementation details are left for future work.

5. Experience and Collaboration Layer

This layer offers customized views and workflows for different stakeholders. It includes dashboards for engineers, cost reports for finance, and summaries for leadership. It follows best practices for sharing understanding and promoting collaboration among roles. In the CloudOptima vision, this layer makes sure insights and decisions are easy to access and act on.

6. Defining these layers early gives later prototypes a clear spot to integrate, which lowers the risk of creating random, one-off scripts that are hard to merge or expand.

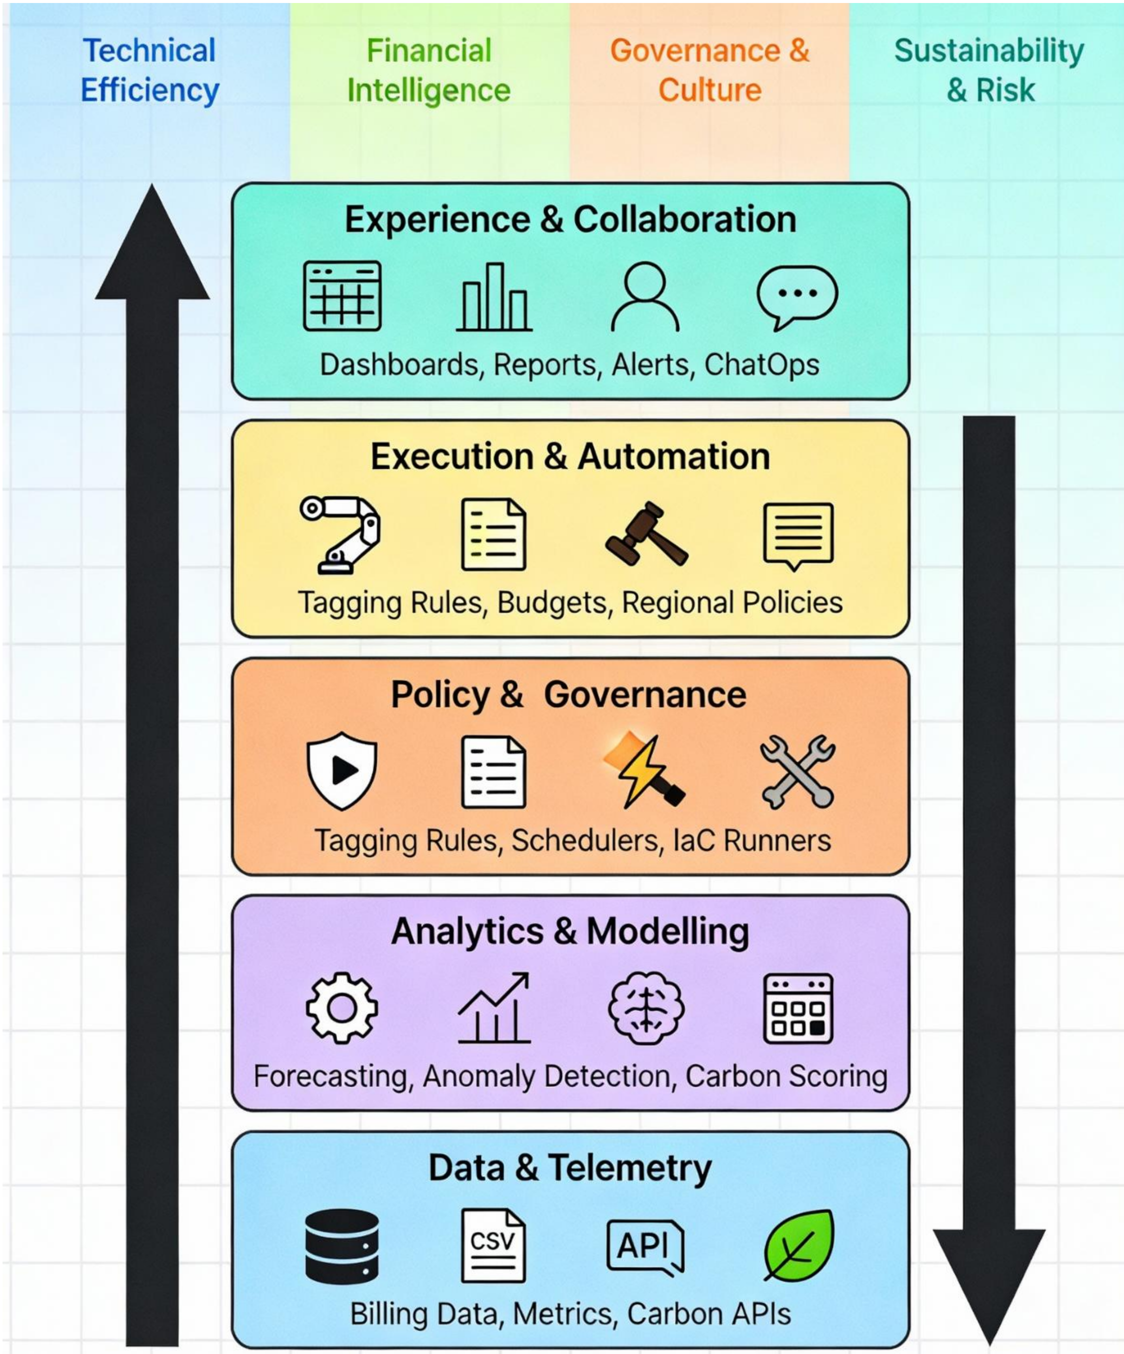

4. Equations, Figures and Tables

Equations, Figures, and Tables In this conceptual paper, equations, figures, and tables act as design tools and illustrate potential CloudOptima logic rather than showing outputs from applied models.

4.1 Example Dual-Objective Equation

To illustrate how CloudOptima could integrate cost and sustainability in future policies, a simple dual-objective score for choosing a cloud region r can be defined as:

$$\text{Score}_r = \alpha \cdot C_r^* + (1 - \alpha) \cdot E_r^*$$

In this expression, C_r^* represents a normalized measure of cost for region r, E_r^* represents a normalized measure of carbon intensity, and α is a weight between 0 and 1 that expresses how strongly the organization prioritizes cost compared with emissions. A lower score indicates a more desirable region under the chosen weighting. This equation is not implemented or evaluated here; it simply demonstrates the type of policy logic that CloudOptima might use in later stages. (Jindal & Di, 2022)

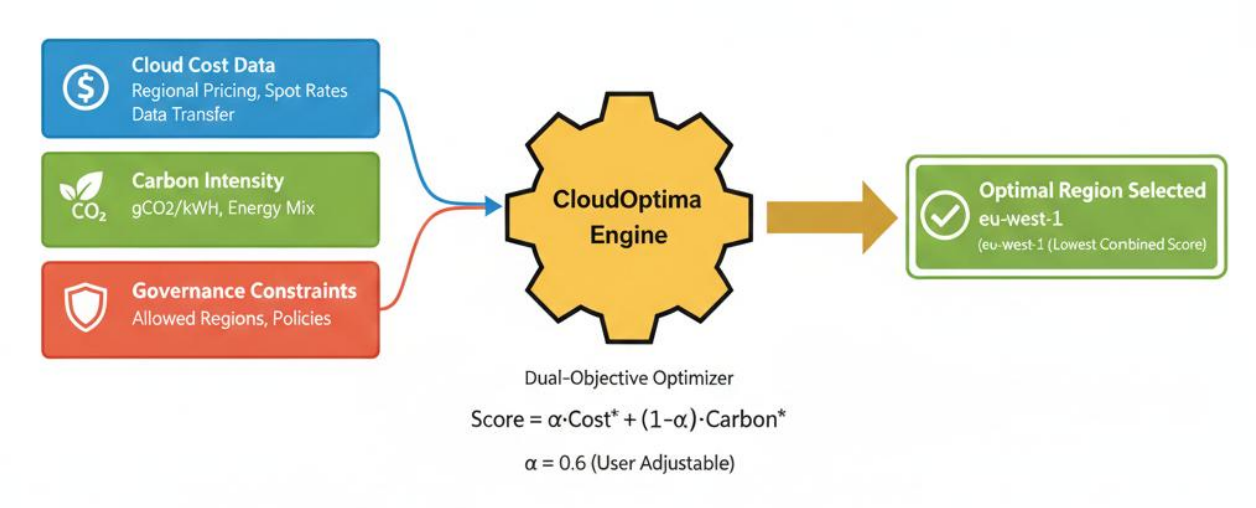

4.2 Conceptual Table: Phased Capability Rollout

The following table summarises a four-phase roadmap for implementing CloudOptima capabilities over time. It is intended as a planning aid rather than a report of completed work. (Jones & Smith, 2022)

Phase	Primary Focus	Example Capabilities (Vision)
1	Visibility and Tagging	Basic tagging policy, resource inventory, spend dashboards
2	Structured Optimization	Rightsizing reviews, non-production schedules, initial budgets
3	Automation and Guardrails	Scripted tag enforcement, idle resource alerts, simple guardrails
4	Analytics and Sustainability	Forecasting, scenario simulation, dual cost–carbon decision policies

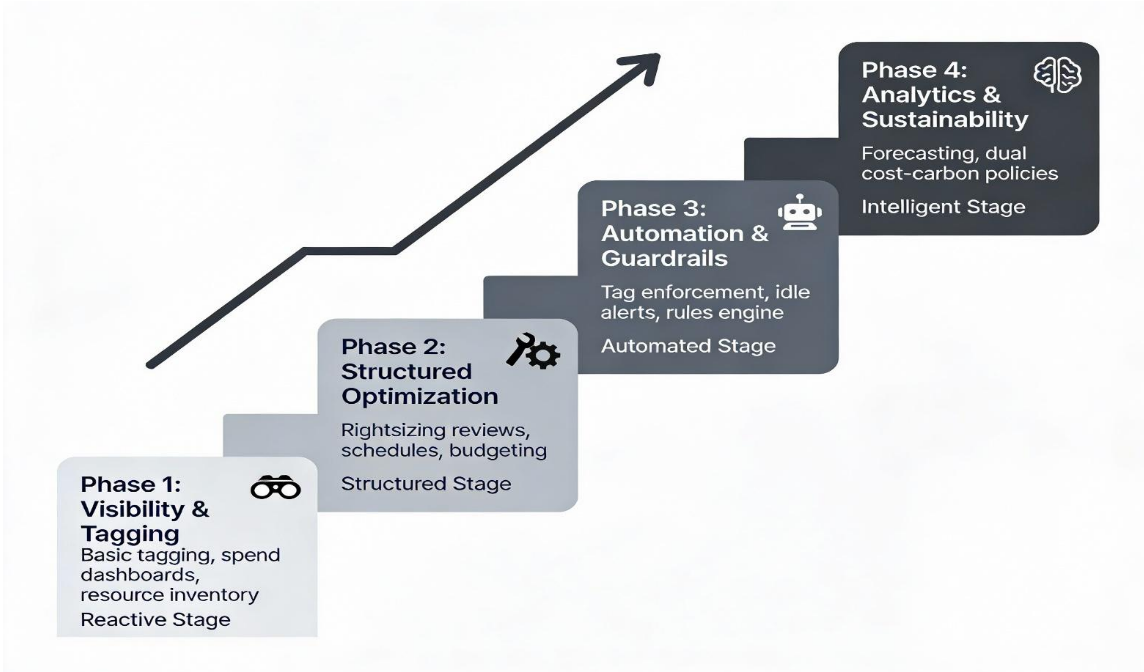

This phased view helps readers see how CloudOptima can be adopted incrementally, with early benefits available even before advanced analytics or carbon-aware logic are in place.

5. Conclusion

This paper introduces CloudOptima as a strategic framework for managing cloud costs and sustainability. It starts with the observation that organizations often treat cloud financial operations, technical optimization, and sustainability as separate tasks. This leads to fragmented tools and inconsistent practices. To address this issue, CloudOptima combines four main pillars: technical efficiency, financial intelligence, governance and culture, and sustainability and risk. It also features a five-layer architecture that organizes future capabilities clearly.

The paper's contribution is mainly conceptual. It does not present any working prototype, algorithm, or experimental results. Instead, it provides a clear blueprint and a four-phase implementation roadmap. This roadmap can guide students, researchers, and practitioners when they are ready to build and test CloudOptima components using open-source tools and free-tier cloud accounts. The early phases emphasize visibility and basic governance, while the later phases focus on automation, analytics, and dual-objective cost and carbon policies.

Future work will include selecting specific capabilities from this framework, implementing them on one or more public cloud platforms, and designing studies to measure their impacts on cost, performance, and environmental footprint. Until this work is done, CloudOptima should be seen as a ready-to-implement vision. It provides a structured, vendor-neutral foundation that combines insights from cloud optimization and carbon-aware computing into a single design.

References

1. Ali-Eldin, A., Tordsson, J., & Elmroth, E. (2012). An adaptive hybrid elasticity controller for cloud infrastructures. In *Proceedings of the IEEE/IFIP Network Operations and Management Symposium (NOMS 2012)* (pp. 204-212). IEEE.

2. Babu, S., & George, T. (2022). Comparative study on cloud cost optimization strategies in multi-cloud environments. *International Journal of Cloud Applications and Computing, 12*(4), 57-73.
3. Bobroff, N., Kochut, A., & Beaty, K. (2007). Dynamic placement of virtual machines for managing SLA violations. In *Proceedings of the 10th IFIP/IEEE International Symposium on Integrated Network Management (IM 2007)* (pp. 119-128). IEEE.
4. Calheiros, R. N., Masoumi, E., Ranjan, R., & Buyya, R. (2015). Workload prediction using ARIMA model and its impact on cloud applications' QoS. *IEEE Transactions on Cloud Computing, 3*(1), 449-458.
5. Chaisiri, S., Lee, B. S., & Niyato, D. (2012). Optimization of resource provisioning cost in cloud computing. *IEEE Transactions on Services Computing, 5*(2), 164-177.
6. Chen, X., Wang, H., Ma, Y., Zheng, X., & Guo, L. (2022). MicroScaler: Fine-grained scaling for microservices using request-level metrics. In *Proceedings of the 2022 USENIX Annual Technical Conference (USENIX ATC '22)* (pp. 489-504). USENIX Association.
7. CloudGain Research Group. (2023). A gamified platform for promoting cost-awareness among developers. *Journal of Cloud Economics, 4*(2), 88-101.
8. CostAnom Project Team. (2021). Using cloud billing data for security anomaly detection. In *Proceedings of the ACM Workshop on Cloud Security* (pp. 45-56). ACM Press.
9. Dashti, S. E., & Rahman, M. S. (2020). A genetic algorithm-based approach for cost-effective and reliable resource allocation in cloud. *Cluster Computing, 23*(4), 3017-3035.
10. DB-Tune Research Consortium. (2022). An automated cost and performance tuner for cloud-native databases. *Proceedings of the VLDB Endowment, 15*(11), 2702-2715.
11. FuncFlow Development Team. (2021). A cost-aware scheduler for serverless workflows. *IEEE Transactions on Serverless Computing, 2*(3), 112-125.
12. Gong, Z., Gu, X., & Wilkes, J. (2010). PRESS: Predictive elastic resource scaling for cloud systems. In *Proceedings of the 6th International Conference on Network and Service Management (CNSM 2010)* (pp. 9-16). IEEE.
13. Gupta, A., & Amza, C. (2021). FinOps in practice: A framework for cloud cost optimization using tagging and accountability. In *Proceedings of the 2021 IEEE International Conference on Cloud Computing (CLOUD)* (pp. 156-165). IEEE.
14. Herbst, N. R., Kounev, S., & Reussner, R. H. (2013). Elasticity in cloud computing: What it is, and what it is not. In *Proceedings of the 10th International Conference on Autonomic Computing (ICAC 2013)* (pp. 23-27). USENIX Association.
15. Islam, S., Keung, J., Lee, K., & Liu, A. (2012). Empirical prediction models for adaptive resource provisioning in the cloud. *Future Generation Computer Systems, 28*(6), 155-162.
16. Jindal, A., & Di, S. (2022). Carbon-aware and cost-efficient scheduling in geo-distributed cloud data centers. *IEEE Transactions on Green Communications and Networking, 6*(2), 1021-1035.
17. Jones, T., & Smith, L. (2022). The FinOps journey: A maturity model for cloud financial management. *ACM Transactions on Management Information Systems, 13*(4), Article 35.
18. Khazaei, H., Misic, J., & Misic, V. B. (2012). Performance analysis of cloud computing centers using M/G/m/m+r queuing systems. *IEEE Transactions on Parallel and Distributed Systems, 23*(5), 936-943.