World Journal of Advanced Engineering Technology and Sciences

eISSN: 2582-8266

Cross Ref DOI: 10.30574/wjaets

Journal homepage: https://wjaets.com/

(Review Article)

Check for updates

Role of AI in cloud cost optimization and FinOps (Financial Operations)

Pradeep Kurra *

Trace3, USA.

World Journal of Advanced Engineering Technology and Sciences, 2025, 15(01), 424-430

Publication history: Received on 26 February 2025; revised on 03 April 2025; accepted on 05 April 2025

Article DOI: https://doi.org/10.30574/wjaets.2025.15.1.0218

Abstract

Integrating Artificial Intelligence into cloud cost optimization and Financial Operations (FinOps) marks a transformative shift in enterprise cloud management strategies. Traditional approaches to cloud financial management have resulted in substantial resource wastage and cost inefficiencies across organizations of all sizes. By leveraging machine learning algorithms, predictive analytics, and automation, AI-enhanced FinOps solutions enable unprecedented cost forecasting accuracy, resource optimization, and financial governance. These technologies can analyze vast amounts of utilization data across multiple dimensions to identify complex patterns invisible to human analysts, anticipate future resource requirements, and automatically implement cost-saving measures without compromising performance. From LSTM-based forecasting models that capture temporal dependencies in cloud consumption to unsupervised learning techniques that detect spending anomalies in real-time, AI-powered tools are demonstrating remarkable efficacy in addressing the financial challenges of cloud computing. The practical applications of these technologies across financial services, e-commerce, and healthcare sectors provide compelling evidence of their capacity to deliver substantial ROI while enabling more precise capacity planning, dynamic resource allocation, and proactive cost management. As cloud environments grow in complexity, AI-driven FinOps represents a crucial evolution from reactive cost control to strategic financial management of cloud resources

Keywords: Cloud cost optimization; Artificial Intelligence; Financial Operations; Machine learning; Multi-cloud management

1. Introduction

The proliferation of cloud computing has fundamentally transformed how organizations build, deploy, and manage digital infrastructure. While cloud services offer unprecedented flexibility and scalability, they have introduced complex financial challenges. According to Rackspace's 2025 State of Cloud Report, organizations now waste approximately 37% of their cloud spend, with large enterprises reporting average annual cloud expenditures exceeding $18.3 million. The report reveals 82% of surveyed organizations struggle with accurate cloud cost forecasting, and 76% lack effective governance mechanisms to control spending across multiple cloud environments [1]. As organizations migrate more workloads to the cloud, they frequently encounter unexpected cost increases, billing complexity, and resource inefficiencies that significantly impact operational budgets. The dynamic pricing models of cloud services and the ease of provisioning resources create an environment where costs can quickly spiral out of control without proper governance mechanisms.

FinOps (Financial Operations) emerged as a discipline to address these challenges by bringing financial accountability to the cloud computing variable spend model. The State of FinOps 2025 Report indicates that organizations implementing mature FinOps practices achieve an average cost reduction of 31% within the first year of implementation, with the most advanced practitioners reporting savings of up to 46% through systematic optimization efforts. Furthermore, the report highlights that companies with established FinOps teams reduce their mean time to

* Corresponding author: Pradeep Kurra

detect cost anomalies from 18 days to just 4.3 hours, enabling significantly faster remediation of spending issues [2]. It represents a cultural shift that combines systems, best practices, and financial controls to optimize cloud spending while maximizing business value. However, traditional FinOps approaches often suffer from several limitations, including manual data analysis, reactive cost management, and siloed optimization efforts that cannot keep pace with the scale and complexity of modern cloud environments.

Artificial Intelligence (AI) is revolutionizing this landscape by introducing predictive capabilities, automation, and intelligent decision-making to cloud cost optimization. The Rackspace study found that AI-powered cloud optimization solutions can reduce cloud costs by an additional 23% beyond traditional optimization methods, with 67% of surveyed enterprises now employing some form of AI-driven cost management. These organizations report forecast accuracy improvements of 41% when using machine learning models compared to conventional forecasting techniques [1]. AI-powered FinOps solutions can process vast amounts of utilization data, identify complex patterns, and generate actionable insights at a scale beyond human capacity. These systems leverage machine learning algorithms to predict future spending patterns, detect anomalies in real-time, recommend optimal resource configurations, and automate cost-saving actions without human intervention.

This paper examines how AI transforms cloud cost optimization across multiple dimensions, from predictive cost forecasting and dynamic resource allocation to anomaly detection and spending governance. By exploring theoretical frameworks and practical applications, we demonstrate how AI-driven FinOps represents a paradigm shift from reactive cost control to proactive financial management of cloud resources.

2. AI-Driven Cost Forecasting and Budgeting

Accurate forecasting and budgeting are the cornerstones of effective cloud financial management. Traditional approaches to this challenge often rely on simplistic trend analysis and manual calculations, resulting in forecasts that fail to account for the complexities of cloud usage patterns. AI has transformed this landscape by introducing sophisticated predictive modeling to analyze historical usage data across multiple dimensions and identify complex patterns invisible to human analysts.

Machine learning algorithms, particularly time series forecasting models like ARIMA (AutoRegressive Integrated Moving Average), Pand rophet, and deep learning approaches such as LSTM (Long Short-Term Memory) networks, excel at capturing the temporal dependencies and seasonal variations characteristic of cloud resource consumption. Research by Kollu et al. published in Wiley's Transactions on Emerging Telecommunications Technologies demonstrates that LSTM-based models achieve 31.8% higher accuracy in cloud cost prediction compared to traditional ARIMA models, with mean absolute percentage error (MAPE) rates reduced from 26.7% to just 9.8% when predicting VM workloads across multiple cloud platforms. Their study analyzing 29,495 VM instances across three cloud providers found that deep learning approaches outperformed statistical methods in 87.3% of test cases and were particularly effective at predicting irregular spikes in resource utilization that frequently lead to cost overruns [3]. These algorithms can ingest multi-dimensional data—encompassing compute usage, storage consumption, network traffic, and specialized services—to generate comprehensive spending forecasts with significantly higher accuracy than traditional methods.

AI forecasting models can identify usage patterns associated with business cycles, product launches, marketing campaigns, and other business activities influencing cloud consumption. According to Wipro's comprehensive analysis of 250 enterprise cloud implementations, AI-driven forecasting tools reduced prediction errors by 64.7% when accounting for seasonal business events. They detected spending pattern correlations that explained 83.2% of previously unexplained cost variations. Their study found that organizations implementing machine learning forecasting achieved an average reduction of $2.4 million in annual cloud spending through more accurate capacity planning, with the highest performers achieving ROI exceeding 430% on their AI forecasting investments within the first 12 months of implementation [4]. AI systems can precisely anticipate future resource requirements and associated costs by correlating historical cloud spending with these business events. This capability enables organizations to establish realistic budgets that align with business initiatives and growth projections.

Beyond simple point estimates, AI-driven forecasting provides probabilistic predictions that quantify uncertainty and risk. Wipro's research revealed that organizations implementing probabilistic AI forecasting reduced their budget contingency allocations from an average of 38.5% to just 14.2% while maintaining the same confidence level in budget adequacy. Furthermore, 72.8% of surveyed organizations cited improved financial planning confidence as a primary benefit of implementing AI-based cloud forecasting [4]. Instead of generating a single forecast value, these systems produce confidence intervals and probability distributions that help financial stakeholders understand the range of possible outcomes and make informed decisions about budget allocations and contingency planning. This approach is particularly valuable in volatile business environments where cloud usage may fluctuate unpredictably.

Furthermore, AI systems can continuously update their forecasting models as new data becomes available, learning from prediction errors and adapting to changing usage patterns. Kollu et al. demonstrated that self-adaptive LSTM models reduced forecasting errors by an additional 17.3% compared to static models when tested against rapidly evolving workloads [3]. This adaptive learning capability ensures that forecasts remain accurate even as organizations scale their cloud footprint, introduce new services, or change their application architectures. The result is a dynamic budgeting process that responds to changing conditions rather than rigid allocations based on outdated assumptions.

Table 1 AI-Powered Forecasting Capabilities and Performance Metrics [3, 4]

Capability	Traditional Methods	AI-Enhanced Approach
Pattern Recognition	Simple trend analysis	Multi-dimensional data analysis
Accuracy	High error rates	Significantly reduced MAPE
Business Event Correlation	Limited correlation capabilities	Strong pattern identification
Risk Quantification	Single-point estimates	Probabilistic predictions with confidence intervals
Model Adaptation	Static models	Self-adaptive learning models

3. Intelligent Resource Optimization and Rightsizing

Cloud environments often suffer significant inefficiencies due to overprovisioned resources, inappropriate instance types, and suboptimal configuration choices. According to a comprehensive analysis by Idowu and Olaoye, 78.6% of cloud instances are oversized, with the average CPU utilization across public cloud environments being just 12.3%, indicating massive resource waste. Their study examining 1,356 virtual machines across multiple organizations found that an average of 63.4% consistently overprovisioned memory, while storage allocations exceeded actual requirements by 57.2% in most cases [5]. AI-powered optimization systems address these challenges by analyzing workload characteristics, resource utilization patterns, and performance requirements to recommend optimal resource configurations that minimize costs without compromising performance.

Machine learning algorithms can analyze historical resource utilization data—including CPU, memory, disk I/O, and network throughput—to identify usage patterns and peak demands across different periods. Research by Idowu and Olaoye demonstrates that AI-driven rightsizing systems can identify up to 91.2% of oversized instances with 98.7% accuracy, resulting in an average CPU utilization improvement from 12.3% to 38.7% while maintaining application performance SLAs. Their experiments with K-means clustering and reinforcement learning algorithms across 12 different cloud environments showed that machine learning models could detect and categorize resource usage patterns with 87.3% greater accuracy than manual analysis, leading to more precise capacity planning [5]. By understanding these patterns, AI systems can recommend appropriate instance types and sizes that match actual workload requirements rather than the often overestimated specifications provided by application teams. This rightsizing process alone can yield 30-50% cost savings in many organizations by eliminating wasted capacity.

Beyond simple rightsizing, AI enables intelligent workload placement across availability zones, regions, and cloud providers. A study by Teradata analyzing over 850 enterprise cloud deployments found that AI-based workload placement algorithms reduced overall costs by an average of 31.4% compared to manual placement strategies, with top performers achieving savings of up to 43.8%. The report highlights that cross-region optimization alone saved companies an average of $1.25 million annually, with one global retail organization reducing its multi-region deployment costs by $4.7 million through intelligent workload distribution algorithms [6]. AI systems can recommend optimal placement strategies that minimize total costs while meeting performance and reliability requirements by analyzing pricing variations, performance characteristics, and data transfer costs across different deployment options. This multi-dimensional optimization becomes increasingly valuable in multi-cloud environments where pricing models and service offerings vary significantly between providers.

AI-driven automation extends to dynamic resource scaling, enabling systems to adjust capacity in real-time based on current demand and predicted future requirements. The study by Idowu and Olaoye found that predictive auto-scaling reduced provisioning costs by 37.6% compared to reactive threshold-based policies while improving application response times by 23.2% during demand spikes. Their analysis of time-series forecasting models showed that LSTM networks outperformed traditional statistical methods by accurately predicting utilization spikes 18.4 minutes before they occurred, allowing proactive scaling actions [5]. Unlike traditional auto-scaling policies that rely on simple threshold-based rules, AI-powered scaling incorporates predictive modeling to anticipate demand changes before they occur. This proactive approach ensures resources are available when needed while avoiding the costs of maintaining excess capacity during periods of low demand.

Another critical area where AI delivers substantial value is optimizing storage resources. Teradata's report revealed that organizations implementing AI-driven storage optimization achieved average savings of 52.6% on storage costs, with one financial services company reducing its annual storage expenditure by $6.8 million through automated tiering and data lifecycle management. Their analysis found that AI-driven storage classification correctly identified access patterns for 94.7% of data objects, allowing automated migration policies to move 68.2% of rarely-accessed data to cold storage tiers without affecting performance [6]. Machine learning algorithms can analyze data access patterns to automatically implement tiered storage strategies, moving infrequently accessed data to lower-cost storage classes while keeping hot data on high-performance storage. These systems can also identify redundant data for deduplication and recommend appropriate data lifecycle policies, significantly reducing storage costs while maintaining accessibility and performance.

Table 2 AI-Driven Resource Optimization Techniques and Their Impact [5, 6]

Technique	Application Area	Benefits
Rightsizing	Instance type selection	Elimination of wasted capacity
Workload Placement	Multi-cloud distribution	Cost-optimized deployment
Predictive Auto-scaling	Dynamic resource allocation	Reduced provisioning costs and improved response times
Storage Optimization	Data lifecycle management	Significant cost reduction through tiered storage
Usage Pattern Analysis	Resource utilization assessment	Enhanced capacity planning

4. Automated Anomaly Detection and Cost Governance

As cloud environments grow in complexity, detecting cost anomalies manually becomes increasingly challenging. Unexpected spending spikes can result from various factors, including misconfigurations, resource leaks, unauthorized usage, or changes in cloud provider pricing. According to research by Nwachukwu et al., organizations experience an average of 9.7 significant cloud cost anomalies quarterly, with each incident resulting in a mean excess expenditure of $31,850 before detection and remediation. Their study analyzing 143 enterprise cloud environments found that unexpected cost spikes took an average of 8.3 days to identify through manual processes, with remediation efforts adding another 3.6 days, resulting in considerable financial waste [7]. AI-powered anomaly detection systems address this challenge by continuously monitoring spending patterns across all cloud resources and services to identify unusual activity that may indicate waste or potential issues.

Machine learning algorithms, particularly unsupervised learning techniques such as clustering, isolation forests, and autoencoder neural networks, excel at identifying outliers in multi-dimensional data without requiring predefined thresholds. Research by Nwachukwu et al. demonstrated that ensemble models combining isolation forests and LSTM networks achieved 96.3% accuracy in detecting cloud spending anomalies, compared to just 63.7% for traditional threshold-based approaches. Their experimental deployment across 38 distinct cloud environments revealed that deep learning models detected 89.4% of cost anomalies within 47 minutes of occurrence, while threshold-based systems identified only 42.8% of anomalies and took an average of 31.2 hours [7]. These systems establish normal behavior baselines for different resource types, services, and organizational units, enabling them to detect deviations that may indicate problems. Unlike rule-based systems that rely on static thresholds, AI-powered anomaly detection can adapt to evolving usage patterns and identify subtle anomalies that would escape human attention.

When anomalies are detected, AI systems can automatically categorize them based on their characteristics, estimate their financial impact, and generate contextual alerts with actionable information. According to Kummarapurugu's comprehensive research on AI-powered cloud governance, advanced anomaly detection systems reduced mean time to detection (MTTD) from 6.2 days to just 1.7 hours and decreased mean time to resolution (MTTR) from 28.5 hours to 6.8 hours across a sample of 215 enterprise cloud environments. His analysis of 1,876 cost anomaly incidents found that AI systems correctly identified the root cause in 83.7% of cases without human intervention, enabling automated remediation workflows that resolved 62.4% of anomalies without manual intervention [8]. For example, the system might identify a sudden increase in storage costs due to an exponential growth in log files for a specific application, enabling teams to address the root cause quickly before costs escalate further. This context-rich notification represents a significant improvement over generic alerts, indicating a spending threshold has been exceeded.

Beyond reactive anomaly detection, AI-driven governance systems implement proactive controls to prevent cost overruns. Nwachukwu et al. found that organizations implementing AI-powered governance tools reduced their unexpected cloud expenses by an average of 76.8% and decreased policy violations by 87.3% within eight months of deployment. Their case study of a financial services organization revealed that predictive governance systems prevented 94.6% of potential cost anomalies by identifying and addressing risky configurations before they resulted in significant expenditures [7]. These systems can automatically enforce tagging policies, implement resource quotas, and temporarily restrict access to expensive services when spending approaches predefined limits. By combining these preventive measures with continuous monitoring and anomaly detection, organizations can maintain strict financial governance while still allowing teams the flexibility to provision resources as needed.

AI governance platforms also facilitate cost allocation and chargeback processes by automatically analyzing resource usage patterns and attributing costs to appropriate business units, projects, or applications. Kummarapurugu reported that machine learning-based allocation models increased cost attribution accuracy by 41.6% compared to traditional tagging-based approaches, resulting in an average of $1.78 million in previously unallocated costs being correctly assigned to responsible departments. His survey of 87 large enterprises found that organizations implementing AI-based chargeback systems increased cost recovery from shared cloud resources by 36.8% and improved budgetary compliance by 42.3% compared to organizations using manual allocation methods [8]. Machine learning algorithms can identify shared resource usage patterns and implement sophisticated allocation models that distribute costs fairly based on actual consumption rather than simplistic allocation rules. This capability enables organizations to implement accountability frameworks incentivizing cost-efficient behavior across all cloud consumers.

Table 3 Comparison of Traditional and AI-Enhanced Anomaly Detection Methods [7, 8]

Feature	Traditional Approach	AI-Enhanced Capability
Detection Timeframe	Manual detection processes	Near real-time identification
Root Cause Analysis	Generic alerts	Contextual problem identification
Preventive Controls	Static thresholds	Predictive governance systems
Cost Allocation	Tag-based attribution	Machine learning-based allocation models
Policy Enforcement	Manual intervention	Automated compliance measures

5. Case Studies: AI-Driven FinOps in Practice

Numerous real-world implementations across industries substantiate the theoretical benefits of AI-powered FinOps. This section examines several case studies that demonstrate the tangible impact of AI on cloud cost optimization and financial governance.

5.1. Financial Services: Global Investment Bank

A global investment bank with over $5 billion in annual cloud spending implemented an AI-driven FinOps platform to address escalating costs across their multi-cloud environment. According to a detailed case study published by Sanka, this financial institution was experiencing annual cloud cost growth of 38.7% despite efficiency initiatives, with 31.4% of cloud resources identified as underutilized. Their study highlighted that the bank's infrastructure sprawled across 428 distinct AWS accounts, 213 Azure subscriptions, and 97 Google Cloud projects, creating significant governance challenges and optimization opportunities [9]. The system analyzed three years of historical usage data across these platforms to identify optimization opportunities and predict future spending patterns. Within six months of deployment, the AI system identified $78 million in annual savings opportunities through rightsizing, reserved instance purchases, and workload redistribution. The predictive forecasting capabilities reduced variance between budgeted and actual cloud spending from 23% to 4%, significantly enhancing financial planning accuracy.

The bank's most valuable finding came from the AI's analysis of storage usage patterns, which revealed that 67% of their data was accessed less than once per quarter. The organization reduced storage costs by automatically implementing lifecycle policies that moved this cold data to archival storage tiers by $34 million annually while maintaining necessary access speeds for frequently used data. The anomaly detection system also prevented a potential $12 million unexpected cost by identifying a misconfigured data transfer process within hours of its introduction. Sanka's analysis indicates that the bank achieved a 427% ROI on their AI-powered FinOps implementation within the first year, with total cost avoidance reaching $97.3 million over 18 months. Their machine-learning models achieved 94.8% accuracy in predicting monthly cloud expenditures across all three major cloud providers, enabling much more precise financial planning [9].

5.2. E-Commerce: Online Retail Platform

A rapidly growing e-commerce platform struggled with unpredictable cloud costs that fluctuated dramatically during promotional events and seasonal peaks. According to research by Venkatesh, this retailer was experiencing cost variations of up to 312% between normal operations and peak seasonal events, making financial planning nearly impossible. The study details how their cloud costs for Black Friday events spiked from an average daily spend of $187,500 to over $584,000 within 72 hours, creating significant budget challenges [10]. Their implementation of an AI-powered cost optimization platform created a predictive scaling model that analyzed historical traffic patterns, promotional calendars, and external factors like weather and social media sentiment to anticipate resource requirements with unprecedented accuracy.

The AI system's dynamic resource allocation capabilities enabled the platform to maintain 99.99% availability during their highest-traffic sales event while reducing compute costs by 42% compared to their previous approach of static over-provisioning. The anomaly detection component identified several unused development environments and zombie resources that collectively cost $1.3 million annually. At the same time, the intelligent instance selection model reduced database costs by 28% by recommending appropriate instance families based on actual query patterns rather than generic specifications. Venkatesh's case study notes that the platform reduced its overall cloud spend by $7.2 million annually while supporting 143% year-over-year growth in transaction volume. The retailer's implementation of reinforcement learning algorithms for automatic scaling decisions reduced peak event provisioning costs by 37.6% compared to the previous year while maintaining 147ms average API response times across 2.3 million concurrent user sessions [10].

5.3. Healthcare: Medical Imaging Provider

A healthcare organization specializing in medical imaging processing deployed an AI-driven FinOps solution to optimize their compute-intensive workloads. Sanka's analysis reveals that this healthcare organization was processing over 11.3 petabytes of medical imaging data annually, with cloud computing representing 47.8% of its operational expenditure. Before implementing AI-driven optimization, the organization maintained 83 dedicated GPU instances running at just 23.7% average utilization, representing significant wasted capacity [9]. The system analyzed GPU utilization patterns across thousands of imaging processing jobs to identify optimal scheduling strategies and instance selections. By intelligently distributing workloads across spot instances, reserved instances, and on-demand resources based on urgency, reliability requirements, and current market pricing, the organization reduced processing costs by 61% while meeting all performance SLAs.

The AI platform's workload prediction capabilities enabled the organization to purchase reserved instances and savings plans confidently, generating an additional 43% cost reduction for their baseline capacity. Perhaps most importantly, the system's continuous optimization approach adapted to changing cloud provider pricing models and new instance types, automatically shifting workloads to more cost-effective options as they become available without requiring manual analysis or intervention. Venkatesh noted that the healthcare provider maintained its 98.2% SLA while reducing per-image processing costs from $3.27 to $0.89, enabling it to expand services to underserved communities. Implementing intelligent queueing and batch processing algorithms increased throughput by 37.4% while reducing costs by 18.3%, ultimately enabling the organization to process 41.7% more studies within the same budget envelope [10].

Table 4 Industry-Specific AI-FinOps Implementation Outcomes [9, 10]

Organization Type	Implementation Focus	Results
Financial Services	Multi-cloud optimization	Dramatic cost reduction and improved forecasting accuracy
E-Commerce	Dynamic resource allocation	Maintained availability while reducing compute costs
Healthcare	Compute-intensive workload optimization	Significant per-image processing cost reduction

6. Conclusion

Integrating AI technologies into cloud financial management represents a definitive evolution in how organizations approach cost optimization and governance in increasingly complex multi-cloud environments. Through sophisticated machine learning algorithms and predictive analytics capabilities, these solutions transcend the limitations of traditional FinOps practices by enabling proactive rather than reactive cost management. The dramatic improvements in forecasting accuracy achieved through techniques like LSTM networks and ensemble models allow organizations to establish realistic budgets that align with business initiatives while reducing contingency allocations. Similarly, AI-driven resource optimization delivers substantial efficiency gains through intelligent rightsizing, workload placement, and dynamic scaling that would be impossible to achieve through manual analysis alone. Most significantly, automated anomaly detection and governance systems provide the vigilance and responsiveness needed to prevent costly misconfigurations and policy violations in rapidly evolving cloud environments. The case studies from diverse industry sectors demonstrate that AI-powered FinOps delivers measurable financial benefits while supporting business growth and innovation. As organizations continue to accelerate their cloud adoption, the strategic implementation of these technologies will increasingly differentiate those that achieve financial sustainability in their cloud operations from those that struggle with unpredictable costs and operational inefficiencies. The evidence suggests that AI-driven cloud cost optimization has moved beyond theoretical potential to become an essential component of mature cloud financial management practices.

References

[1] Rackspace, "The 2025 State of Cloud Report," January 14, 2025. [Online]. Available: https://www.rackspace.com/resources/2025-state-cloud-report-infographic

[2] FinOps Foundation, "State of FinOps 2025 Report." [Online]. Available: https://data.finops.org/

[3] Praveen Kumar Kollu et al., "Comparative analysis of cloud resources forecasting using deep learning techniques based on VM workload traces," Wiley, 15 January 2024. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/ett.4933

[4] Wipro Ltd., "AI-Powered Strategies for Cloud Cost Optimization." [Online]. Available: https://www.wipro.com/infrastructure/ai-powered-strategies-for-cloud-cost-optimization/

[5] Marvel Idowu, Godwin Olaoye, "Optimizing Cloud Resource Allocation with Machine Learning: A Smarter Approach," ResearchGate, February 2025. [Online]. Available: https://www.researchgate.net/publication/389262778_Optimizing_Cloud_Resource_Allocation_with_Machine_Learning_A_Smarter_Approach

[6] Pat Alvarado, "The Ultimate Guide to AI Driven Cloud Cost Optimization," Teradata Insights, August 27, 2024. [Online]. Available: https://www.teradata.com/insights/ai-and-machine-learning/guide-to-ai-driven-cloud-cost-optimization

[7] Chukwuemeka Nwachukwu et al., "AI-driven anomaly detection in cloud computing environments," International Journal of Science and Research Archive, 2024, 13(02), 692–710, 11 November 2024. [Online]. Available: https://ijsra.net/sites/default/files/IJSRA-2024-2184.pdf

[8] Charan Shankar Kummarapurugu, "AI-Powered Cloud Cost Governance Systems: Automating Budget Management, Resource Allocation, and Policy Compliance for Optimized Cloud Expenditure," ResearchGate, December 2023. [Online]. Available: https://www.researchgate.net/publication/389814332_AI-Powered_Cloud_Cost_Governance_Systems_Automating_Budget_Management_Resource_Allocation_and_Policy_Compliance_for_Optimized_Cloud_Expenditure