



Adaptive AI Model Selection for Cloud Infrastructure Optimization: A Framework for Intelligent, Self-Regulating Computing Environments

Mallikarjuna Muchu*

Independent Researcher, USA

* **Corresponding Author Email:** mallikarjunamuchu1@gmail.com - **ORCID:** 0000-0002-5247-4450

Article Info:

DOI: 10.22399/ijcesen.4670
Received : 19 December 2025
Revised : 24 December 2025
Accepted : 04 January 2026

Keywords

Adaptive AI Model Selection,
Cloud Infrastructure
Optimization,
Event-Driven Automation,
Multi-Cloud Resource
Management,
Infrastructure as Code

Abstract:

Contemporary cloud computing ecosystems demand intelligent infrastructure management strategies that transcend traditional static provisioning models. The adaptive AI model selection framework presented herein addresses fundamental challenges in managing heterogeneous workloads across microservices architectures, data processing pipelines, IoT data streams, and AI inference engines through systematic integration of machine learning operations with event-driven automation mechanisms. The framework synthesizes cloud engineering principles, Infrastructure as Code methodologies, and intelligent model selection algorithms to enable real-time optimization based on telemetry analysis, historical workload patterns, and operational objectives. Through automated training pipelines, continuous evaluation protocols, and progressive deployment strategies with rollback capabilities, the framework facilitates a self-optimizing infrastructure that minimizes human intervention while maintaining service-level agreements and cost efficiency. The architecture comprises telemetry collection subsystems, versioned model repositories, intelligent selection engines implementing multi-criteria decision frameworks, and automated deployment orchestration utilizing canary patterns and circuit breakers. Event-driven automation enables real-time responsiveness through stream processing frameworks that evaluate optimization opportunities via windowed computations, complex event processing patterns, and stateful processing mechanisms. Enhancements to multi-cloud and hybrid environments support heterogeneous resource abstractions, cross-platform data movement limits, and vendor-specific operational behaviors by using cloud-agnostic abstraction layers and provider-specific adapters. The framework illustrates how smart and dynamic infrastructure operation can make organizations realize better resource utilization, better service reliability, and reduction of operational expenses in a distributed computing environment, as well as provision of regulatory compliance and data sovereignty requirements in a complex multi-cloud deployment.

1. Introduction

Modern cloud computing systems have transformed into a multiplexed and heterogeneous environment that concurrently runs microservice architectures, big-data processing pipelines, data streams of the Internet of Things (IoT), as well as artificial intelligence inference engines. The cloud adoption environment has been changing dramatically, and more companies are adopting multi-cloud environments to maximize their investment in infrastructure and operational capacity. Based on the extensive industry research that analyzes the trend of cloud computing, businesses are driving through complicated decisions regarding cloud

infrastructure management, and a major focus on cost optimization, workload placement models, and the incorporation of artificial intelligence functionalities into business processes [1]. The operational characteristics, performance needs, reliability limits, and cost optimization goals of each category of workloads impose unique management requirements that require highly complex management strategies and not just the conventional methods of workload provisioning. The choice and implementation of suitable artificial intelligence and machine learning models in the optimization of cloud infrastructure has become one of the core factors of operational efficiency, system resilience, and economic feasibility in

enterprise computing environments. The established methods of cloud infrastructure control are based largely on fixed systems and responsive intervention procedures, which are not sufficiently effective in the face of the dynamic and unforeseeable nature of the contemporary workload of computations. Studies conducted on energy-efficient resource arrangement in cloud data centers show that smart workload controls can ensure a substantial reduction in the operational cost without affecting quality of service guarantees, especially when dynamic consolidation strategies and predictive scaling systems are applied to respond to changes in workload [2]. The distributed systems and the growing demands of higher-quality service-level agreements and cost-containment pressures both dictate a paradigm change in how distributed infrastructure management is approached with greater intelligence and adaptability.

This scholarly article presents a comprehensive framework that synthesizes cloud engineering principles, Infrastructure as Code methodologies, and event-driven automation mechanisms to enable workload-aware AI model selection and deployment. The proposed framework addresses critical gaps in cloud operations research by establishing systematic protocols for real-time model recommendation based on telemetry data, historical workload patterns, and clearly defined operational objectives. Through the integration of automated training pipelines, continuous evaluation mechanisms, and intelligent rollback capabilities across hybrid and multi-cloud platforms, this approach facilitates the emergence of self-optimizing infrastructure systems that minimize human intervention while maximizing operational outcomes, ultimately contributing to organizational objectives of enhanced service reliability and optimized resource utilization.

2. Theoretical Foundations and Cloud Infrastructure Dynamics

The theoretical underpinnings of adaptive AI model selection for cloud infrastructure optimization reside at the intersection of distributed systems theory, machine learning operations, and control systems engineering. Cloud environments exhibit characteristics of complex adaptive systems, wherein multiple independent agents interact through network communications, resource sharing, and workload distribution mechanisms. The fundamental challenge in modern cloud infrastructure management centers on achieving efficient resource utilization while maintaining quality of service guarantees across diverse workload types. Studies on energy-sensitive

resource allocation have revealed that there are natural trade-offs in cloud data centers between resource utilization and energy usage, and evidence has shown that data center servers take a considerable amount of power even when idle, and that can be seized in the opportunity to develop intelligent consolidation strategies that move workloads to fewer physical hosts during low demand times [2]. Such dynamic consolidation strategies entail complex prediction systems in order to predict the future needs of the resources and preemptively modify the infrastructure configurations before performance breakdown has taken place.

The heterogeneity of workloads is the core issue that requires delicate optimization plans depending on the peculiarities of applications.

Microservices designs have bursty short-period computational profiles with a focus on low-latency interaction and horizontal scalability, necessitating infrastructure designs that place a high value on network response and fast container coordination. Data processing pipelines exhibit a long-run resource consumption profile with a known periodicity and batch-based execution models that, in many cases, enjoy the advantage of reserved capacity allocations that spread cost across long periods of execution. Internet of Things deployments have created new demands on the edge computing capabilities, where processing needs to be carried out in close physical proximity to data to meet the latency requirements of real-time analytics and control applications. Studies examining IoT architectural frameworks identify vision elements including ubiquitous sensing capabilities, heterogeneous device integration, dynamic service composition, and intelligent data processing at multiple hierarchical levels from edge devices through fog computing layers to centralized cloud resources [3]. These architectural considerations necessitate intelligent model selection mechanisms that account for deployment topology, network latency characteristics, and the computational capabilities of resource-constrained edge devices.

The concept of Infrastructure as Code has fundamentally transformed cloud resource management by enabling declarative specification of infrastructure configurations, version control integration, and reproducible deployment processes. IaC frameworks facilitate the codification of infrastructure policies, resource dependencies, and operational constraints as executable artifacts that can be subjected to automated testing, validation, and continuous integration workflows. This programmatic approach creates the foundational substrate upon

which adaptive AI model selection mechanisms can operate, enabling dynamic reconfiguration of cloud resources in response to model recommendations without manual intervention. The integration of machine learning into resource management systems requires careful consideration of prediction accuracy, model training overhead, and the temporal dynamics of workload patterns. Energy-aware allocation research demonstrates that combining historical workload data with current system state enables more accurate predictions of future resource demands, with empirical evaluations showing that intelligent allocation heuristics can achieve substantial energy savings compared to baseline approaches while maintaining service level agreement compliance through careful management of performance degradation risks during consolidation operations [2].

Event-driven architectures can be used to achieve the level of real-time responsiveness required by adaptive cloud optimization systems by enabling the decoupling of infrastructure monitoring, model inference, and resource provisioning via asynchronous message passing. The cloud workloads presented by this architectural pattern naturally follow the time behavior of the irregular arrival, variable realistic execution patterns, and unpredictable resource demands that demand the capability of continuous monitoring and respond with much expediency. This development in event-driven cloud management is similar to larger trends in the design of distributed systems, in which loosely coupled components interact via a narrow interface of well-defined messages instead of a call-and-response interface. The importance of event-driven processing in processing high-velocity streams of data generated by distributed sensors is explored in research studies on the workings of IoT architectures, with architectural models that include message brokers, stream processing engines, and support more complex event processing operations than merely extracting actionable insights out of raw telemetry data [3]. These same architectural principles apply directly to cloud infrastructure optimization, where telemetry streams from thousands of compute instances, storage systems, and network devices must be processed in real-time to identify optimization opportunities and trigger automated remediation actions.

3. Framework Architecture and Intelligent Model Selection Mechanisms

The proposed framework architecture comprises four primary subsystems that work in concert to enable continuous, automated optimization of cloud infrastructure configurations. The telemetry

collection subsystem aggregates multi-dimensional operational data from distributed infrastructure components, including compute utilization metrics, network traffic patterns, storage input/output characteristics, application-level performance indicators, and cost attribution data. Modern cloud environments generate tremendous volumes of operational telemetry, creating both opportunities and challenges for intelligent management systems. Industry surveys examining cloud computing trends reveal that organizations struggle with visibility and control across multi-cloud environments, with significant portions of cloud spending going to waste through overprovisioned resources, unused reserved instances, and suboptimal workload placement decisions [1]. Feature engineering transformations extract temporally relevant patterns through time-series decomposition techniques that isolate trend components, seasonal variations, and residual fluctuations, enabling models to distinguish between systematic patterns that can be exploited for prediction and stochastic variations that represent fundamental uncertainty.

Dimensionality reduction approaches compress high-dimensional telemetry streams into compact feature vectors suitable for efficient model inference while preserving the information content necessary for accurate optimization decisions. The model repository maintains versioned collections of AI and ML models specifically designed for cloud infrastructure optimization tasks, spanning multiple paradigm categories including time-series forecasting models for resource demand prediction, anomaly detection models for reliability monitoring, reinforcement learning agents for dynamic resource allocation, and classification models for workload type identification. Each model maintains associated metadata describing its training data characteristics, performance benchmarks, computational requirements, and operational constraints, enabling the intelligent selection engine to evaluate candidate models against current infrastructure state and workload requirements. Research into intelligent workload management for hybrid cloud environments demonstrates the complexity of model selection decisions, showing that optimal placement strategies must consider not only computational requirements but also data locality constraints, network bandwidth limitations, security policy requirements, and economic factors, including the comparative costs of on-premises versus cloud execution [4].

The clever selection engine will introduce the multi-criteria decision system that compares candidate models with workload-specific goals, working limits, and integrates past performance

statistics, real-time telemetry functionalities, capacity limitations of the infrastructure, and cost-efficiency studies to find the best model-workload combinations. Meta-learning methods make use of the cross-workload knowledge transfer in order to speed up the process of model selection based on new workload pattern identification of the structural similarities between the workload pattern witnessed previously and the specific workload type under consideration. The selection process must balance multiple competing objectives, including prediction accuracy, inference latency, memory footprint, and operational costs, often producing Pareto-optimal solution sets rather than single optimal configurations. Studies examining self-learning applications in cloud resource management and scheduling demonstrate that machine learning techniques can effectively address the complexity of multi-objective optimization in cloud environments, with various approaches including supervised learning for workload classification, unsupervised learning for pattern discovery, and reinforcement learning for dynamic decision-making showing promise for different aspects of the resource management problem [5].

The automated deployment orchestration subsystem converts model suggestions into executable infrastructure changes using Infrastructure as Code templates and configuration management tools, and executes gradual rollout plans that incrementally increase model control and perpetually evaluate performance metrics and rollback only. The deployment pipelines should be able to respond to the risks of automated changes to the production infrastructure, such as configuration errors, unforeseen interplay between components, and negative performance effects due to poorly optimized decisions. Canary deployment schemes introduce model configurations to small groups of traffic, allowing model forecasts to be compared to observed responses before the launch of fully deployed systems, and circuit breaker policies allow the detection of aberrant model behaviors and roll back to a stable system state. A combination of continuous integration and continuous deployment methods with container orchestration platforms has made it possible to use more advanced deployment strategies to reduce risk and increase the speed of infrastructure development.

Research examining Kubernetes-based container deployment frameworks demonstrates how containerization technologies combined with CI/CD pipelines enable rapid, reliable deployment of application workloads with built-in health checking, automatic rollback capabilities, and fine-grained resource allocation controls that support efficient multi-tenant infrastructure operation [6].

4. Event-Driven Automation and Continuous Optimization Pipelines

Event automation is the backplane of adaptive cloud infrastructure systems, which allows real-time responsiveness to workload variations and, at the same time, provides stability and predictability of the systems. The event processing architecture is used to realize a directed acyclic graph of processing steps that refine and enrich raw infrastructure events into useful optimization decisions by means of progressive refinement and enrichment operations. Mechanisms of event ingestion receive infrastructure signals at the heterogeneous sources, such as hypervisor metrics, container orchestration systems, application performance monitor systems, and cloud provider APIs, and require standardization transformations to normalize disparate event schemas into unified forms that are cross-platform processed and analyzed. The volume and velocity of telemetry data in a contemporary cloud context pose great challenges to event processing systems, which are required to effectively filter, aggregate, and route events to optimal processing pipelines without compromising the end-to-end latency attribute that is capable of ensuring prompt response to emergent conditions.

Stream processing frameworks enable continuous evaluation of optimization opportunities through windowed computations over telemetry streams, with sliding window operators maintaining rolling statistics that characterize recent workload behaviors and tumbling windows delineating discrete time periods for batch-oriented analyses. Complex event processing patterns identify multi-event sequences that signal emerging performance degradation, capacity constraints, or optimization opportunities requiring intervention, while stateful stream processing maintains contextual information across event sequences to enable detection of long-duration trends and behavioral pattern changes. Research examining tail latency management in datacenter-scale file systems demonstrates the critical importance of understanding performance characteristics at high percentiles, showing that techniques for managing worst-case latencies require careful attention to queueing dynamics, request scheduling policies, and resource allocation strategies that prevent individual slow requests from cascading into broader system-level performance degradation [7]. These insights apply directly to cloud infrastructure optimization, where maintaining consistent performance across varying load conditions requires predictive scaling mechanisms that anticipate demand spikes and proactively provision additional capacity before

quality of service degradation becomes visible to end users.

The continuous optimization pipeline implements closed-loop control mechanisms that automatically adjust infrastructure configurations in response to model recommendations, with feedback controllers comparing predicted outcomes against observed performance metrics to compute error signals that drive incremental configuration adjustments. Proportional-integral-derivative controllers achieve rapid convergence to target operating points while minimizing oscillations and overshoot, with adaptive control strategies modifying controller parameters based on system identification results to enable robust performance across varying workload conditions and infrastructure states. Model predictive control formulations incorporate constraint satisfaction across multiple operational constraints and multi-step lookahead planning to anticipate future optimization opportunities and preemptively adjust configurations, solving optimization problems formulated as mixed-integer programs to identify near-optimal resource allocation policies within computational time budgets compatible with real-time deployment requirements. The application of control theory to cloud resource management draws on decades of research in automated system regulation, adapting classical techniques to address the unique challenges of distributed computing environments, including network delays, partial observability, and the discrete nature of resource allocation decisions. Automated model retraining workflows ensure that deployed models remain aligned with evolving workload characteristics and infrastructure capabilities, executing retraining cycles triggered by drift detection events that identify when statistical properties of incoming telemetry data diverge significantly from training data distributions. When significant drift is detected, automated pipelines trigger model retraining using recent historical data, hyperparameter optimization, and validation against held-out test sets to verify that performance improvements justify deployment of updated models. A/B testing frameworks compare newly trained models against incumbent versions through randomized traffic splitting, collecting performance metrics, and applying sequential statistical tests to verify performance improvements before deployment, with rollout strategies implementing gradual traffic shifting that increases candidate model exposure contingent on continued performance verification. Research examining container-based workload distribution in Kubernetes environments demonstrates practical approaches to implementing canary deployments and gradual rollouts through native platform

capabilities, leveraging service mesh technologies and traffic management policies to control the flow of requests to different model versions while collecting detailed observability data that enables rapid detection of performance regressions or unexpected behaviors [6]. The deployment practices allow optimization models to keep evolving continuously without compromising the high availability and reliability requirements of production infrastructure systems.

5. Multi-Cloud and Hybrid Infrastructure Considerations

The extension of adaptive AI model selection frameworks to multi-cloud and hybrid infrastructure environments introduces additional complexity dimensions related to heterogeneous resource abstractions, cross-platform data movement constraints, and vendor-specific operational characteristics. Contemporary enterprise cloud strategies increasingly embrace multi-cloud approaches that distribute workloads across multiple public cloud providers to achieve risk diversification, avoid vendor lock-in, and exploit geographic distribution for latency optimization and regulatory compliance. Research on the adoption of clouds conducted within the industry indicates that companies run workloads on a combination of multiple cloud environments, with large shares of computing resources being deployed in either private cloud or on-premises data centers as well as using the public clouds, with systems operation under planned strategies to balance control needs, security, performance demands and cost optimization objective across diverse infrastructural portfolios [1]. Nevertheless, every cloud provider has its own resource abstraction, API interfaces, and operational semantics that make it difficult to manage infrastructure holistically, necessitating the existence of abstraction layers that translate provider-specific resources into canonical infrastructure representations while maintaining the flexibility to access provider-specific capabilities where they are useful.

The suggested architecture deals with the issue of multi-cloud heterogeneity by providing cloud-agnostic Infrastructure as Code tools that can specify the infrastructure requirements without considering the implementation details of the provider, and provider-specific adapters can transform them to native provisioning APIs with high fidelity and low manual effort for edge cases.

Hybrid infrastructure environments integrate on-premises data centers with public cloud resources to balance control requirements, compliance constraints mandating on-premises hosting for

certain data types due to regulatory requirements, and economic considerations where total cost of ownership analyses show on-premises infrastructure achieving cost parity with cloud resources at higher utilization rates for predictable workloads. Research examining multi-cloud and hybrid cloud strategies for enterprise architectures emphasizes the importance of careful workload placement decisions that consider data residency requirements, network latency characteristics, security policy constraints, and integration requirements with existing on-premises systems when determining optimal deployment topologies across hybrid environments [8]. Data gravity effects, wherein computational workloads exhibit a preference for colocation with large datasets due to transfer cost and latency constraints, influence optimal workload placement decisions across hybrid topologies, particularly for data-intensive analytics workloads that must process substantial volumes of information. Network latency and bandwidth constraints between on-premises and cloud components introduce additional optimization variables that must be incorporated into model selection criteria, with measured latencies varying substantially based on geographic distance and network path characteristics. Edge computing scenarios extend hybrid architectures to include resource-constrained devices deployed in close physical proximity to data sources, requiring lightweight model variants optimized for limited computational budgets and intermittent connectivity, necessitating local inference capabilities with periodic synchronization to centralized model repositories when network connectivity permits. The architectural implications of edge computing extend beyond simple resource constraints to encompass fundamentally different operational models where autonomous operation during network partitions becomes a critical requirement rather than an exceptional failure mode. Studies examining IoT architectural frameworks identify the importance of hierarchical processing structures that perform initial data

filtering and aggregation at edge devices, intermediate analytics and aggregation at fog computing layers, and comprehensive analysis and long-term storage at centralized cloud resources, with intelligent data routing policies determining which processing tasks execute at which hierarchical levels based on latency requirements, bandwidth constraints, and computational complexity [3]. Cross-platform telemetry aggregation presents technical challenges related to metric schema alignment, timestamp synchronization, and data federation across administrative boundaries, with federated monitoring architectures maintaining local telemetry repositories at each infrastructure site while supporting cross-site query capabilities for global optimization decisions. Cost optimization across multi-cloud environments requires sophisticated economic modeling that accounts for heterogeneous pricing structures, including on-demand, reserved, and spot instance pricing models, data transfer charges, and storage tiering costs, with the model selection framework incorporating total cost of ownership calculations that evaluate long-term economic implications of infrastructure decisions beyond immediate resource expenses. Research examining intelligent workload factoring for hybrid cloud computing models demonstrates practical approaches to workload partitioning that consider both technical constraints and economic factors, with optimization formulations that jointly minimize execution costs while satisfying performance requirements and data locality constraints across hybrid infrastructure topologies [9]. Multi-objective optimization formulations balance cost minimization against performance guarantees, reliability requirements, and other operational objectives, producing Pareto frontiers that inform human decision-making in scenarios requiring explicit policy trade-offs involving budget-performance-reliability triangulation across complex multi-cloud deployment scenarios.[10].

Table 1: Cloud Infrastructure Management Challenges and Optimization Strategies [1,2]

Management Challenge	Traditional Approach Limitation	Adaptive AI Strategy	Expected Outcome
Cost Optimization	Manual resource allocation	Dynamic workload-aware provisioning	Reduced waste from overprovisioning
Visibility Control	Static monitoring dashboards	Real-time telemetry aggregation	Enhanced multi-cloud visibility
Energy Efficiency	Fixed capacity allocation	Dynamic consolidation with prediction	Lower power consumption during idle periods
Workload Placement	Rule-based assignment	Intelligent placement considering constraints	Balanced performance and cost
Resource Utilization	Reactive scaling policies	Predictive scaling mechanisms	Maintained quality of service

Table 2: IoT and Hybrid Cloud Architectural Framework Components [3,4]

Architectural Layer	Processing Capability	Deployment Consideration	Integration Requirement
Edge Devices	Initial data filtering and aggregation	Low-latency local processing	Autonomous operation capability
Fog Computing	Intermediate analytics and aggregation	Distributed processing nodes	Hierarchical data routing
Centralized Cloud	Comprehensive analysis and storage	Scalable computational resources	Long-term data retention
Hybrid Infrastructure	Balanced on-premises and cloud workloads	Data locality and compliance	Security policy alignment
Multi-Cloud Distribution	Geographic workload distribution	Vendor lock-in avoidance	Cross-platform integration

Table 3: Machine Learning Model Categories for Cloud Optimization [5,6]

Model Category	Optimization Task	Learning Paradigm	Deployment Strategy
Time-Series Forecasting	Resource demand prediction	Supervised learning	Historical data training
Anomaly Detection	Reliability monitoring	Unsupervised learning	Pattern deviation identification
Reinforcement Learning	Dynamic resource allocation	Trial-and-error optimization	Policy-based decision making
Workload Classification	Application type identification	Supervised learning	Feature-based categorization
Container Orchestration	Automated deployment management	Rule-based and adaptive	Progressive rollout with health checks

Table 4: Event-Driven Automation and Performance Management Mechanisms [7,8]

Automation Component	Processing Function	Control Mechanism	Quality Assurance
Stream Processing	Continuous optimization evaluation	Windowed computations	Real-time pattern detection
Complex Event Processing	Multi-event sequence identification	Stateful context tracking	Behavioral change detection
Tail Latency Management	High-percentile performance control	Request scheduling policies	System-level degradation prevention
Feedback Controllers	Configuration adjustment automation	Error signal computation	Target convergence with stability
A/B Testing Framework	Model performance validation	Randomized traffic splitting	Statistical significance

6. Conclusions

The integration of adaptive AI model selection mechanisms with event-driven cloud infrastructure automation represents a transformative advancement toward self-optimizing, resilient computing environments capable of managing diverse workloads while satisfying stringent service-level agreements and cost efficiency objectives. The comprehensive framework addresses fundamental challenges of workload heterogeneity, real-time responsiveness, and continuous optimization through the systematic application of machine learning operations principles integrated with cloud engineering best practices. Real-time telemetry analysis and patterns of past workload help select intelligent models, which can be used to proactively manage the infrastructure to predict the opportunities for optimization, and react in response to the changing

operational conditions, shifting the process of managing the infrastructure between the reactive problem resolution paradigm and predictive and prescriptive management model. The full model lifecycle of training, evaluation, deployment, and rollback across hybrid and multi-cloud environments is automated to ensure that humans are minimally involved in the process without sacrificing the strict compliance with service-level agreements and cost limits to overcome the visibility and control challenges posed by the complex multi-cloud environments. The architectural implementation of event-driven makes it temporal asynchronous processing pipelines decouple monitoring, analysis, and action stages, and the Infrastructure as Code provides reproducibility and auditability of changes in infrastructure. Important contributions are systematic workload-conscious model selection protocols, continuous optimisation pipelines that

enforce closed-loop control models, and practical extensions of the multi-cloud and mixed infrastructure environment that maintain end-to-end optimisation targets whilst supporting platform heterogeneity. Future opportunities include federated learning to support collaborative model training across organizational boundaries, causal inference techniques to facilitate a better understanding of infrastructure performance relationships, and explainable AI techniques that can give insight into model selection choices. The emergence of quantum computing resources and specialized AI accelerators necessitates framework extensions accommodating novel computational paradigms with distinct optimization characteristics, while sustainability metrics integration represents crucial avenues for reducing the environmental impacts of cloud computing operations. Practical implications extend beyond technical implementation to organizational transformation toward intelligent, autonomous infrastructure operations, with demonstrated reductions in operational overhead, improvements in resource utilization efficiency, and enhanced service reliability positioning organizations to extract maximum value from cloud investments while maintaining agility necessary for responding to evolving business requirements in increasingly dynamic competitive landscapes where digital infrastructure serves as foundational enabler of innovation and competitive differentiation.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

- [1] Tanner Luxner, "Cloud Computing Trends: Flexera 2023 State of the Cloud Report," Flexera 2023. [Online]. Available: <https://www.flexera.com/blog/finops/cloud-computing-trends-flexera-2023-state-of-the-cloud-report/>
- [2] Anton Beloglazo, et al., "Energy-aware resource allocation heuristics for efficient management of data centers for Cloud computing," ScienceDirect, 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0167739X11000689>
- [3] Jayavardhana Gubbi, et al., "Internet of Things (IoT): A vision, architectural elements, and future directions," arxiv, 2012. [Online]. Available: <https://arxiv.org/abs/1207.0203>
- [4] Hui Zhang, et al., "Intelligent Workload Factoring for a Hybrid Cloud Computing Model," IEEE, 2009. [Online]. Available: <https://ieeexplore.ieee.org/document/5190708>
- [5] Babak Ravandi; Ioannis Papapanagiotou, "A Self-Learning Scheduling in Cloud Software Defined Block Storage," IEEE, 2017. [Online]. Available: <https://ieeexplore.ieee.org/document/8030616>
- [6] Manish Kumar Abhishek, et al., "Framework to Deploy Containers using Kubernetes and CI/CD Pipeline," International Journal of Advanced Computer Science and Applications, 2022. [Online]. Available: https://thesai.org/Downloads/Volume13No4/Paper_60-Framework_to_Deploy_Containers_using_Kuberne tes_and_CICD_Pipeline.pdf
- [7] Pulkit A. Misra, et al., "Managing Tail Latency in Datacenter-Scale File Systems Under Production Constraints," ACM Digital Library, 2019. [Online]. Available: <https://dl.acm.org/doi/10.1145/3302424.3303973>
- [8] Vamsi Krishna Reddy Munnangi, "Multi-Cloud and Hybrid Cloud Strategies for Enterprise API Architectures," ResearchGate, 2025. [Online]. Available: https://www.researchgate.net/publication/391630699_Multi-Cloud_and_Hybrid_Cloud_Strategies_for_Enterprise_API_Architectures
- [9] Nicola Capodieci, et al., "Deadline-Based Scheduling for GPU with Preemption Support," IEEE, 2009 [Online]. Available: <https://ieeexplore.ieee.org/document/8603197>
- [10] Omdia, "Global cloud infrastructure spending rose 21% in Q1 2025", 2025. <https://omdia.tech.informa.com/pr/2025/jun/global-cloud-infrastructure-spending-rose-21percent-in-q1-2025>