



e-ISSN: 2319-8753 | p-ISSN: 2347-6710

# IJRSET

International Journal of Innovative Research in  
**SCIENCE | ENGINEERING | TECHNOLOGY**



# INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

Volume 13, Issue 4, April 2024

**ISSN** INTERNATIONAL  
STANDARD  
SERIAL  
NUMBER  
INDIA

Impact Factor: 8.423

9940 572 462

6381 907 438

[ijirset@gmail.com](mailto:ijirset@gmail.com)

[www.ijirset.com](http://www.ijirset.com)

# EC-Aware Adaptive Virtual Machine Migration Techniques for Cloud Data Centre to Optimize Energy Consumption

J.Noorul Ameen M.E., M.Mohamed Apsar, S.Mohamed Fuhadh, M.Mohamed Tharik

Assistant Professor, E.G.S. Pillay Engineering College, Nagapattinam, India

UG Student, Department of Computer Science and Engineering, E.G.S.Pillay Engineering College,  
Nagapattinam, India

UG Student, Department of Computer Science and Engineering, E.G.S.Pillay Engineering College,  
Nagapattinam, India

UG Student, Department of Computer Science and Engineering, E.G.S.Pillay Engineering College,  
Nagapattinam, India

**ABSTRACT:** The rising energy requirements of cloud data centers (CDCs) tasked with extensive data storage and processing pose both environmental and financial hurdles for cloud service providers. To address this, current CDCs utilize virtualization migration and consolidation techniques to decrease energy consumption (EC) and optimize resource utilization. This research examines virtual machine (VM) migration by breaking it down into three key elements: timing migration, selecting VMs for migration, and identifying suitable migration destination hosts.

It proposes an adaptive dynamic threshold approach, leveraging EC levels, to ascertain optimal migration timing. Additionally, it introduces a strategy based on correlation and utilization for selecting VMs, along with an enhanced EC-aware best-fit algorithm for identifying suitable migration destination hosts. Evaluation conducted using the CloudSim toolbox and actual VM workload traces from PlanetLab showcases promising results for these algorithms.

In comparison to established techniques and benchmark algorithms, the proposed methods demonstrate an average decrease in EC, service level agreement violation (SLAV), and the number of VM migrations by 15.49%, 7.85%, and 83.32%, respectively. This indicates superior performance in both optimizing energy consumption and enhancing service quality, even when faced with high VM workload or resource limitations. However, it's important to recognize that the experiments were carried out in a simulated environment, which could produce slightly varied outcomes compared to real-world scenarios.

**KEYWORDS:** Energy consumption optimization, virtual machine migration techniques, dynamic threshold, virtual machine selection, host selection, and cloud data center.

## I. INTRODUCTION

In recent years, cloud computing technology has surged in popularity and widespread adoption. This trend has triggered a rapid expansion of cloud data centers (CDCs) to cater to the burgeoning demands for extensive data storage and computational power. However, this expansion has brought about a corresponding increase in the energy consumption (EC) of servers, posing a significant challenge for large-scale infrastructures such as clusters, grids, and CDCs comprising thousands of diverse servers. The escalation in EC has not only resulted in substantial operational costs for CDCs but also diminished profits for cloud providers and escalated CO<sub>2</sub> emissions. Consequently, both researchers and industries are now actively exploring avenues to utilize electricity more efficiently and sustainably to tackle these pressing challenges.

Servers are typically configured and deployed within cloud data centers (CDCs) to manage peak workloads and achieve optimal performance. However, this practice can lead to underutilization of servers during non-peak periods, resulting in resource wastage and increased energy consumption (EC). To tackle this challenge, virtualization consolidation has

emerged as a common approach in modern data centers (DCs). It enables service consolidation and reduces energy usage through resource multiplexing. By leveraging virtualization technology to create multiple virtual machine (VM) instances on a single physical server, cloud providers can improve resource utilization and enhance their return on investment.

Recent statistics reveal that idle servers can consume a significant amount of energy, ranging from 50% to 70% of that used by servers operating at full capacity. This underscores the inefficiency of idle servers, which contribute minimally to computational tasks while consuming substantial energy resources. To address this inefficiency, many data centers (DCs) have adopted energy management strategies. These strategies monitor server utilization and initiate VM migration processes to consolidate underutilized servers. During VM migration, VMs are reallocated based on their resource requirements, aiming to minimize the number of active hosts. Additionally, to further reduce energy consumption (EC), idle hosts are transitioned into low-power modes such as sleep or hibernation, effectively eliminating their power consumption until workload demands necessitate their reactivation to prevent performance degradation of applications due to resource scarcity.

In the selection of VMs for migration, various migration metrics such as resource utilization, load balancing, migration time, migrated data volume, application performance, and energy consumption (EC) are evaluated. However, there is a need for a comprehensive method that effectively addresses these metrics while integrating EC considerations efficiently.

## II. RELATED WORKS

Optimizing Within the realm of information technology, optimizing energy consumption (EC) in cloud data centers (CDCs) has emerged as a central research focus. This goal can be attained by formulating a systematic migration strategy that prioritizes both energy efficiency and quality of service (QoS). Current endeavors in EC optimization for CDCs commonly center on three fundamental components: identifying the optimal timing for migration, selecting appropriate VMs for migration, and determining the most suitable destination hosts for migration.

### A. DETERMINING MIGRATION TIMING

At present, migration trigger thresholds are established to govern the timing of migration events. Typically, migration is initiated when the host load exceeds the overload threshold or falls below the low load threshold. While transferring VMs from underutilized servers and transitioning them into sleep mode can mitigate energy consumption, it may also contribute to minimizing service level agreement violations (SLAV). However, setting unreasonable migration thresholds could lead to excessive consolidation and frequent, unnecessary migrations, potentially impacting application performance due to factors such as migration time and downtime. This degradation in service quality could result in SLA breaches and subsequent penalties. Therefore, it is crucial to identify overloaded and underloaded hosts and ascertain the optimal timing for VM migration.

Migration trigger thresholds can be categorized as either static or dynamic. Static thresholds are manually configured and remain unchanged throughout the migration process. Conversely, dynamic thresholds adapt based on real-time resource conditions, cluster load, and time, rendering them more flexible. Various studies have investigated both approaches. For instance, Beloglazov proposed adaptive threshold methods such as interquartile range (IQR) and median absolute deviation (MAD), employing statistical analysis of historical host data. Yadav introduced the GradCent algorithm to dynamically calculate upper limit thresholds for CPU utilization based on historical workload data. Singh proposed a dynamic threshold with enhanced search (DT-ESAR) for VM consolidation systems, while Kulshrestha optimized host overload detection through an exponentially weighted moving average-based threshold formulation method. Real-time load detection and prediction techniques are also utilized to monitor and anticipate host load states, enabling proactive resource allocation and scheduling in CDCs.



### A. SELECTION OF MIGRATED VMs

The selection of VMs for migration involves a comprehensive evaluation of various migration metrics, including resource utilization, load balancing, migration time, migrated data volume, and application performance. Reddy et al. concentrated on resource utilization and migration time metrics and introduced a VM selection algorithm based on memory utilization, bandwidth utilization, and optimization of VM size. Ahmadi et al. took into account migration time, migration risk, VM connectivity, freeable resources, and SLAV rate, presenting a multi-criteria decision-making method grounded in hierarchical analysis. Baskaran utilized the fuzzy soft set method, considering CPU usage, memory usage, RAM usage, and correlation values to identify suitable VMs for migration. Mekala and Viswanathan introduced the energy-efficient resource ranking and utilization factor-based VM selection (ERVS) method, prioritizing EC and resource utilization metrics. Li et al. leveraged content similarity between VM memories to select VMs for migration, with the aim of reducing migration time, data transmission volume, and network traffic pressure. Haghshenas and Mohammadi proposed a regression-based approach to predict VM resource utilization using historical data, selecting VMs with higher utilization prediction results for migration.

### B. SELECTION OF MIGRATION DESTINATION HOSTS

Selecting the destination host involves choosing hosts from a pool to migrate the selected VMs, presenting a nondeterministic polynomial (NP)-a hard problem that becomes increasingly complex with larger-scale scenarios. Researchers predominantly rely on heuristic, meta-heuristic, and machine learning algorithms to tackle this challenge. Heuristic algorithms like next fit (NF), first fit (FF), best fit (BF), first fit decreasing (FFD), and best fit decreasing (BFD) are commonly employed. Liang et al. utilized the BF algorithm to select a target host with the smallest remaining space to accommodate VMs, while Chhikara et al. implemented the BF algorithm using a heap structure for efficient target host determination. Chen et al. explored FF, BF, and random algorithms for container destination host selection, comparing their performance. Fan et al. utilized the FF algorithm to ensure destination servers have sufficient available resources, while Beloglazov and Buyya used the BFD method to assign each VM to the host with the smallest increase in power consumption, sorting VMs by present CPU usage.

Machine learning algorithms used for determining target hosts mainly include clustering, classification, reinforcement learning, and neural network methods. Liang et al. introduced a dynamic hybrid machine learning-based algorithm for energy-aware resource deployment in CDCs, employing an extended K-means clustering algorithm and extended k-nearest neighbors classification algorithm for VM deployment. Ma et al. developed an online VM scheduling scheme (OSEC) for joint EC and cost optimization based on reinforcement learning theory. Rezakhani et al. applied reinforcement learning and artificial neural networks to propose an integrated algorithm focused on energy-aware QoS for dynamic VM management in CDCs. Liu et al. proposed a deep reinforcement learning model that utilizes QoS feature learning to optimize DC resource scheduling.

Meta-heuristic algorithms encompass various techniques such as simulated annealing, particle swarm optimization, differential evolution, genetic algorithm (GA), ant colony optimization algorithm, and cuckoo search algorithms. Additionally, in practical applications, there's a growing interest in multi-method-based approaches that combine the strengths of multiple metaheuristics. For example, Liu et al. proposed a VM consolidation algorithm for CDCs integrating ant colony systems and extreme learning machines (ELM). Luo et al. utilized improved shuffled frog leaping and extreme value optimization algorithms to address dynamic VM allocation, aiming to reduce power consumption while meeting QoS requirements. He et al. employed GA for VM consolidation tasks. Bibiks et al. introduced an enhanced discrete cuckoo search algorithm to tackle resource-constrained scheduling problems. Meanwhile, Moazeni et al. suggested a dynamic resource allocation strategy using adaptive multi-objective teaching-learning-based optimization (AMO-TLBO), which incorporates features like the number of teachers, adaptive teaching factor, tutorial training, and self-motivated learning. AMO-TLBO aims to minimize makespan and cost while maximizing utilization by ensuring a well-balanced load across virtual machines.

## III. PROPOSED METHODOLOGY

### A. ENERGY CONSUMPTION MODEL AND RESOURCE MODEL OF CDC

In this section, we introduce the CDC-EC model utilized in this study. The overall energy consumption (EC) of a CDC comprises various components, including the power consumed by server hosts, cooling systems, network equipment,



and other auxiliary systems like fire, electrical, lighting, and lightning protection. Since the primary focus is on optimizing the EC of the server hosts, this study primarily addresses host EC optimization. The overall EC of the CDC is modeled as the sum of the EC of all hosts within the center, calculated over a specific time interval from  $t_1$  to  $t_2$  as depicted in Equation (1), where  $E$  represents the EC of the DC,  $H$  denotes the collection of all hosts in the CDC, and  $Ph(t)$  signifies the power consumption of host  $h$  at time  $t$  within the interval  $[t_1, t_2]$ .

Approaches for assessing energy efficiency in CDCs can be categorized into three main groups: measurement-based methods, simulation-based methods, and analytical modeling-based methods. Measurement-based evaluation is deemed the most accurate, providing real-world data on energy consumption, whereas simulation-based methods, due to inherent limitations and simplifications in the models used, offer lower accuracy.

The power consumption of a host in a CDC is influenced by various factors such as its hardware configuration and processing components like the CPU, memory, hard disk, I/O, and network. Previous studies have primarily utilized linear or square regression models with CPU utilization as the main parameter to estimate power consumption. However, modern servers equipped with larger memory capacities highlight the significance of memory consumption, which is often overlooked in such models, leading to inaccurate estimations. Consequently, accurately modeling power consumption for multi-core CPUs pose a complex and challenging task.

To address these challenges, instead of solely relying on specific quantitative models for server power consumption, this study leverages actual power consumption data generated by the SPECpower benchmark results. The SPECpower committee introduced the SPECpower\_ssj2008 benchmark suite, measuring server power consumption at maximum workload, categorized into 11 discrete load levels ranging from 0% to 100%. Linear interpolation is employed to determine dynamic power consumption between two load levels.

For the purposes of this study, HP ProLiant ML110 G4 (Intel Xeon 3040, 2 core, 1860 MHz, 4 GB) and HP ProLiant ML110 G5 (Intel Xeon 3075, 2 core, 2660 MHz, 4 GB) servers were selected as hosts. Table 1 presents the power consumption properties of the chosen servers, indicating that power consumption increases with higher processor utilization and varies among servers with different CPU frequencies at the same CPU utilization efficiency.

Machine Type	Idle	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
HP G4(Watt)	86.0	89.4	92.6	96.0	99.5	102.0	106.0	108.0	112.0	114.0	117.0
HP G5(Watt)	93.7	97.0	101.0	105.0	110.0	116.0	121.0	125.0	129.0	133.0	135.0

Equation (2) delineates the  $d$ -dimensional resources available on the host  $I$ , where  $P_i$ , CPU,  $P_i$ , mem, and  $P_i,bw$  denote the CPU, memory, and bandwidth resources possessed by host  $I$ , respectively. Equation (3) encapsulates the  $d$ -dimensional resource requirements of VM  $j$ , where  $V_j$ , CPU,  $V_j$ , mem, and  $V_j,bw$  signify the CPU, memory, and bandwidth resource demands of VM  $j$ , respectively.

## B. DETERMINING MIGRATION TIMING: ADAPTIVE DYNAMIC THRESHOLD BASED ON EC LEVELS

The strategic migration of VMs from overloaded hosts plays a crucial role in reducing SLAV, while transitioning VMs from less efficient servers to energy-efficient modes helps mitigate EC. In this study, the approach to migration timing involved establishing a migration trigger threshold. It was recognized that setting unreasonable thresholds could lead to excessive consolidation and frequent, unproductive migrations, resulting in diminished QoS and other adverse outcomes.

To address this challenge, an adaptive dynamic threshold method based on EC levels (ADT-EC) was introduced. This method dynamically adjusts the migration trigger threshold according to varying EC levels, aiming to prevent unnecessary migrations while streamlining EC optimization processes.



Initially, the EC level within the CDC was defined, as outlined in Equation (4). Here, the EC level is categorized as BLUE when the center's EC falls below a certain threshold ( $aE_{max}$ ), indicating an optimal state. If the EC falls within the range of  $aE_{max}$  and  $bE_{max}$ , the level is designated as GREEN, representing a normal state. Conversely, when the EC exceeds  $bE_{max}$ , the level shifts to RED, indicating an undesirable state.

The overload threshold is computed using Equation (5), which adjusts the threshold based on the EC level. Specifically,

$$Level = \begin{cases} BLUE, & E < aE_{max} \\ GREEN, & aE_{max} \leq E \leq bE_{max} \\ RED, & E > bE_{max} \end{cases}$$

when the EC level is BLUE, the threshold is set at 90% of the maximum threshold value. In the GREEN state, it is determined by subtracting a step value from the maximum threshold. In the RED state, it is derived by multiplying the maximum threshold by a coefficient ( $\alpha$ ).

---

**Algorithm 1** ADT-EC

---

**Input:** HostList and EC of CDC

**Output:** OverutilizedHostList and UnderutilizedHostList

```

1: for each host in HostList do
2:   count ← 0
3:   utilizationHistory ← getHostUtilizationHistory(host, T)
//Obtain the m load data of the host in the time interval T
4:   level ← getECLLevel(EC of CDC)
//Obtain the CDC's EC level by (4)
5:   thrhigh ← updateUtilizationThreshold(level)
//Update the threshold by (5)
6:   for each i=1 to m do //m is same as line 3
7:     if utilizationHistory[i] > thrhigh, then
8:       count++
9:     if utilizationHistory[i] < thrlow, then
10:      add host to UnderutilizedHostList
11:     break
12:   if count > n, then //n < m
13:     add host to OverutilizedHostList
14: return OverutilizedHostList and
UnderutilizedHostList
    
```

At times, server utilization may briefly surpass the maximum threshold before quickly declining. In such instances, there's no necessity to migrate VMs from this server to release resources, as the host isn't overloaded at that particular moment. Hence, it's crucial to avoid unnecessary migrations. Within the time interval T, if at least n out of m load data points exceed the high threshold (where n is less than m), it indicates the host may be in an overloaded state and could trigger overload migration. Conversely, for underload migration, the underload threshold is set sufficiently low so that when load data falls below  $thr_{low}$ , it may immediately prompt underload migration.

**C. SELECTING THE VM FORMIGRATION: UTILIZATION AND CORRELATION-BASED SELECTIONSTRATEGY**

The study proposes a selection strategy named SS-CAU (Selection Strategy based on Correlation and Utilization) to effectively choose which VM to migrate when a server host becomes overloaded, aiming to minimize negative impacts on the migration process and swiftly restore the host to a normal load state.



The SS-CAU method operates as follows:

---

**Algorithm 2 SS-CAU**

---

**Input:** OverutilizedHostList

**Output:** SelectedVMList //Selected VM which will be migrated

```

1:  for each host in OverutilizedHostList do
2:    vmList ← getMigratableVms(host) //Get the list of VMs
    on the overloaded host
3:    candidateVmList ← null //The VM that meets the strong
    correlation criterion will add to the candidate list.
4:    for each vm in vmList do
5:      r ← getCorrelationCoefficient(vm, host) //The corre-
    lation coefficient of vm is calculated according to (6)
6:      if r > 0.6, then
7:        add vm to candidateVmList
8:        sortedVMList ← sortByUtilization(candidateVmList)
    //Sort candidateVmList in ascending order by utilization
9:      for each vm in sortedVMList do
10:       add vm to SelectedVMList
11:       if the host returns to the normal load level, then
12:         break
13: return SelectedVMList
    
```

It selects the VM with the lowest CPU utilization and high load correlation from the set of VMs to migrate. Opting for a VM with a high correlation facilitates quicker restoration of the host to a normal load state, as a higher correlation between VM load and host load increases the likelihood of host overload. Therefore, under conditions of high correlation, the VM with the least CPU utilization is chosen. This approach minimizes application performance losses during migration and reduces migration costs.

Regression analysis is employed to compute load correlation, where the load sequences of the VM and host are denoted as  $\{x_1, x_2, \dots, x_n\}$  and  $\{y_1, y_2, \dots, y_n\}$ , respectively. The correlation coefficient is expressed as shown in Equation (6).


$$R^2 = \frac{\sum_{i=1}^n (x_i - \bar{x})^2 (y_i - \bar{y})^2}{\sum_{i=1}^n (x_i - \bar{x})^2 \sum_{i=1}^n (y_i - \bar{y})^2}$$

The load correlation between the VM and host is denoted by  $(R^2)$ , where  $(\bar{x})$  and  $(\bar{y})$  represent the average values of the load data for the VM and host, respectively.  $(R^2)$  ranges between 0 and 1, with higher values indicating a stronger correlation between the VM and the host. When  $(0.6 < R^2 \leq 1)$ , it indicates a robust correlation between the VM and the host; VMs meeting this criterion are included in the candidate set. Subsequently, the VM with the lowest CPU utilization is selected from this set and migrated to another host until the overloaded host returns to its normal load level.

Algorithm 2 presents the pseudo-code of the SS-CAU. Given overloaded hosts as inputs, Algorithm 2 produces selected VMs for migration as outputs. Initially, the list of VMs on the overloaded host is obtained (line 2), followed by the calculation of the correlation coefficient for each VM (lines 4–5) using Formula (6). VMs with a correlation



coefficient exceeding 0.6 are added to the candidate set (lines 6–7), which is then sorted in ascending order by utilization (line 8). Finally, VMs are sequentially migrated out until the overloaded host returns to its normal load level (lines 9–12).



#### IV. CONCLUSION

This study is dedicated to optimizing the energy consumption (EC) of Cloud Data Centers (CDCs) by decomposing VM migration tasks into three components: determining migration timing, selecting VMs for migration, and identifying destination hosts for migration. To achieve this, three algorithms—ADT-EC, SS-CAU, and IEABF—were proposed and evaluated through four simulation experiments using the Cloud Sim toolbox.

The experimental findings unequivocally demonstrate the superiority of our proposed algorithms over benchmark algorithms in three key aspects. Firstly, they achieve an average EC reduction of 15.49%, indicating their effectiveness in minimizing energy consumption. Secondly, there is an average reduction in Service Level Agreement Violation (SLAV) of 7.85%, demonstrating improved adherence to desired Quality of Service (QoS) levels. Lastly, the algorithms successfully reduce the number of VM migrations by an average of 83.32%, which is crucial for minimizing system overhead and migration costs. These results affirm the superiority of our approaches over benchmark algorithms, highlighting their potential for enhancing energy efficiency and overall system performance.

Our future endeavors will focus on three main aspects:

Exploring the application of multi-objective optimization techniques to address VM consolidation migration, which involves multiple objectives such as load balancing, performance optimization, and cost reduction.

Further evaluation in large-scale real-world deployments is imperative to authenticate the effectiveness and feasibility of our proposed approaches. Such real-world testing can offer valuable insights into the challenges and advantages associated with implementing these approaches in production environments.

With the increasing popularity of container technology for resource virtualization, we intend to augment the migration algorithm developed in this study to ensure its compatibility with containers.

#### REFERENCES

[1] M. Zakarya's paper, "Energy, performance, and cost-efficient data Centers: A Survey," published in Renewable and Sustainable Energy Reviews in October 2018, offers an extensive overview of the topic, exploring various aspects such as energy efficiency, performance optimization, and cost-effectiveness in data center operations. Zakaryaprovides insights into emerging trends and technologies aimed at enhancing the sustainability and economic viability of data centers, making it a valuable resource for researchers and practitioners in the field.



- [2] N. Hamdi and W. Chainbi's paper, "A survey on energy-aware VM consolidation strategies," published in Sustainable Computing: Informatics and Systems in September 2019, offers an overview of energy-aware Virtual Machine (VM) consolidation strategies in cloud data centers. The authors analyze various techniques aimed at optimizing energy consumption through efficient VM consolidation, providing valuable insights for researchers and practitioners working on energy-efficient cloud computing solutions.
- [3] A. H. T. Dias, L. H. A. Correia, and N. Malheiros conducted a systematic literature review on virtual machine consolidation, published in ACM Computing Surveys in November 2022. Their study offers a comprehensive analysis of existing research on VM consolidation techniques, identifying trends, challenges, and future research directions. This systematic review provides a valuable reference for researchers seeking to understand the current state of the art in VM consolidation.
- [4] In their paper titled "Distributed virtual machine consolidation: A systematic mapping study," published in Computer Science Review in May 2018, A. Ashraf, B. Byholm, and I. Porres present a systematic mapping study of distributed virtual machine consolidation techniques. The study categorizes existing approaches, identifies research gaps, and provides insights into future directions for distributed VM consolidation research for researchers and practitioners in the field.
- [5] R. Yadav et al.'s paper, "Managing overloaded hosts for energy-efficiency in cloud data centers," published in Cluster Computing in September 2021, proposes strategies for managing overloaded hosts to improve energy efficiency in cloud data centers. The authors present an approach to dynamically adjusting resource allocation to alleviate host overload conditions, thereby reducing energy consumption while maintaining performance levels. Their work contributes to the development of energy-efficient resource management techniques in cloud environments.
- [6] A. Beloglazov, J. Abawajy, and R. Buyya's paper, "Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing," published in Future Generation Computer Systems in May 2012, presents energy-aware resource allocation heuristics for efficient management of data centers in cloud computing environments. The authors propose algorithms to optimize resource allocation while considering energy consumption, thereby improving the overall efficiency of cloud data center operations.
- [7] F. D. Rossi et al.'s paper, "E-eco: Performance-aware energy-efficient cloud data center orchestration," published in the Journal of Network and Computer Applications in January 2017, introduces E-eco, a performance-aware energy-efficient cloud data center orchestration framework. The authors propose algorithms for dynamically allocating resources based on workload characteristics to improve energy efficiency while maintaining performance levels in cloud data centers.
- [8] R. Bradford et al.'s paper, "Live wide-area migration of virtual machines including local persistent state," presented at the 3rd International Conference on Virtual Execution Environments in June 2007, discusses techniques for live wide-area migration of virtual machines, including mechanisms for handling local persistent state. The authors propose approaches for the efficient migration of VMs across geographically distributed data centers, enabling dynamic resource management and load balancing in cloud environments.
- [9] X. Ma et al.'s paper, "Real-time virtual machine scheduling in industry IoT network: A reinforcement learning method," published in IEEE Transactions on Industrial Informatics in February 2023, presents a reinforcement learning-based approach for real-time virtual machine scheduling in Industrial Internet of Things (IIoT) networks. The authors propose algorithms to optimize VM scheduling decisions based on dynamic workload conditions, thereby improving resource utilization and overall system performance in IIoT environments.
- [10] U. Deshpande and K. Keahey's paper, "Traffic-sensitive live migration of virtual machines," published in Future Generation Computer Systems in July 2017, discusses techniques for traffic-sensitive live migration of virtual machines in cloud data centers. The authors propose approaches for minimizing disruption and optimizing resource utilization during VM migration, thereby improving the efficiency and reliability of cloud data center operations.
- [11] M. Awad, N. Kara, and A. Leivadreas's paper, "Utilization prediction-based VM consolidation approach," published in the Journal of Parallel and Distributed Computing in December 2022, introduces a utilization prediction-based VM consolidation approach.



INTERNATIONAL  
STANDARD  
SERIAL  
NUMBER  
INDIA



# INTERNATIONAL JOURNAL OF INNOVATIVE RESEARCH

IN SCIENCE | ENGINEERING | TECHNOLOGY

 9940 572 462  6381 907 438  [ijirset@gmail.com](mailto:ijirset@gmail.com)



[www.ijirset.com](http://www.ijirset.com)

Scan to save the contact details