The Scientific Temper (2024) Vol. 15 (spl): 314-320
E-ISSN: 2231-6396, ISSN: 0976-8653
Doi: 10.58414/SCIENTIFICTEMPER.2024.15.spl.37
https://scientifictemper.com/

REVIEW ARTICLE

Expanding the quantity of virtual machines utilized within an open-source cloud infrastructure

D. Jayadurga*, A. Chandrabose

Abstract

As cloud computing continues to evolve, the efficient management and scalability of virtual machines (VMs) have become pivotal for maximizing performance and resource utilization, particularly within open-source cloud infrastructures. This literature review investigates existing approaches and methodologies focused on expanding the number of VMs in open-source cloud environments. Key topics include the impact of VM scaling on resource allocation, load balancing, and energy efficiency, as well as the role of orchestration tools and hypervisor optimization in handling large-scale VM deployments. Furthermore, the review assesses the challenges related to VM density, network latency, and system reliability alongside emerging strategies for enhancing VM elasticity through containerization, microservices, and distributed computing models. This study aims to provide a comprehensive understanding of current trends, innovations, and limitations in VM expansion, offering insights into the future of scalable virtual infrastructures in open-source cloud systems.

Keywords: Cloud computing, Virtual machine allocation, VM migration, VM deployments, Cloud infrastructure.

Introduction

The rapid growth of cloud computing has positioned virtual machine (VM) deployment as a core component in supporting large-scale computing environments, particularly within open-source cloud infrastructures.[1] VMs offer scalable, isolated, and efficient resource utilization, enabling diverse applications to coexist on shared physical infrastructure while ensuring flexibility and cost efficiency. As organizations increasingly adopt cloud services to meet demand spikes, enhance system resilience, and reduce operational costs, expanding the number of VMs in open-source cloud platforms has become an essential challenge. This need for scalability drives the exploration of advanced approaches to accommodate a larger quantity of VMs without compromising performance or system stability, Jalali Moghaddam, M., Esmaeilzadeh, A., Ghavipour, M., & Zadeh, A. K. (2020), Yang, L., Yang, D., Cao, J., Sahni, Y., & Xu, X. (2020) [2].

Open-source cloud infrastructures such as OpenStack and CloudStack provide customizable platforms for deploying, managing, and scaling VMs. However, as VM density increases, several technical challenges emerge, including resource contention, network latency, and energy inefficiency. Addressing these challenges requires an effective balance between system resources, such as CPU, memory, and storage, as well as robust VM orchestration to manage workload distribution. Efficient VM expansion strategies also play a critical role in load balancing and fault tolerance, ensuring cloud systems can scale while maintaining high availability.

Recent advancements in virtualization, containerization, and hypervisor technologies have opened new avenues for optimizing VM deployment. Tools such as Kubernetes for container orchestration and advancements in hypervisor optimizations allow cloud providers to address scalability more dynamically. Additionally, emerging strategies, including microservices architecture and distributed computing models, offer promising solutions to enhance

Edayathangudy G.S Pillay Arts and Science College (Autonomous) (Affiliated to Bharathidasan University, Tiruchirappalli), Nagapattinam, Tamil Nadu, India.

***Corresponding Author:** D. Jayadurga, Edayathangudy G.S Pillay Arts and Science College (Autonomous) (Affiliated to Bharathidasan University, Tiruchirappalli), Nagapattinam, Tamil Nadu, India, E-Mail: jayaddd23@gmail.com

How to cite this article: Jayadurga, D., Chandrabose, A. (2024). Expanding the quantity of virtual machines utilized within an open-sourcecloudinfrastructure.TheScientificTemper,**15**(spl):314-320.
Doi: 10.58414/SCIENTIFICTEMPER.2024.15.spl.37

Source of support: Nil

Conflict of interest: None.

Received: 08/09/2024 **Accepted:** 20/09/2024 **Published:** 10/10/2024

VM scalability in open-source platforms, Singh, S., & Singh, D. (2021).[3]

This literature review examines the current landscape of VM expansion strategies in open-source cloud infrastructures, identifying the strengths and limitations of existing approaches. By synthesizing these findings, the review aims to provide insights into scalable and sustainable VM management practices, ultimately contributing to the development of more resilient and adaptable cloud infrastructures.

Allocation and Migration of Virtual Machines

The allocation and migration of VMs are crucial processes in cloud computing environments, enabling efficient utilization of resources and providing flexibility, scalability, and resilience to the underlying infrastructure. VMs are isolated instances of operating systems and applications that run on a shared physical infrastructure, allowing multiple applications and workloads to coexist. By strategically allocating and migrating VMs, cloud providers can optimize the use of hardware resources, balance workloads, and improve energy efficiency, thereby enhancing overall system performance and reliability. This section explores the fundamental concepts, methodologies, and challenges associated with VM allocation and migration in cloud computing, Pande, S. K., Panda, S. K., Das, S., Sahoo, K. S., Luhach, A. K., Jhanjhi, N. Z., ... & Sivanesan, S. (2021), Zolfaghari, R., Sahafi, A., Rahmani, A. M., & Rezaei, R. (2021).[4,5]

Virtual Machine Allocation

VM allocation refers to the process of assigning virtual machines to specific physical hosts based on workload requirements and resource availability. The goal is to distribute VMs across servers in a way that maximizes resource utilization, minimizes performance bottlenecks, and meets service-level objectives. VM allocation is influenced by several factors, including the demand for resources (e.g., CPU, memory, network bandwidth), the characteristics of the applications hosted on the VMs, and the capacity of the physical infrastructure. Allocating VMs efficiently is crucial, as over-provisioning can lead to resource wastage, while under-provisioning can cause performance degradation and potential downtime, Zhuang, H., & Esmaeilpour Ghouchani, B. (2021), Huang, Y., Xu, H., Gao, H., Ma, X., & Hussain, W. (2021).[6,7]

Traditional VM allocation methods include *round-robin*, which evenly distributes VMs across hosts, and *first-fit* or *best-fit* algorithms, which allocate VMs based on resource availability and demand. However, these methods may not always yield optimal results, especially in large-scale cloud environments with dynamic workloads. Consequently, more advanced algorithms have been developed to address the complexities of VM allocation in modern cloud infrastructures.

Advanced Allocation Techniques

To enhance VM allocation efficiency, researchers have developed sophisticated techniques that leverage machine learning, heuristic optimization, and predictive analytics. Machine learning-based allocation algorithms use historical workload data to predict future resource demands, enabling more adaptive and intelligent allocation decisions. For example, supervised learning models, such as decision trees and neural networks, can analyze patterns in workload behavior and recommend allocation strategies that prevent overloading and underutilization of resources, Balaji Naik, B., Singh, D., & Samaddar, A. B. (2020), Masoudi, J., Barzegar, B., & Motameni, H. (2021).[8,9]

Heuristic optimization methods, such as genetic algorithms, simulated annealing, and particle swarm optimization (PSO), provide alternative approaches to VM allocation by searching for optimal or near-optimal solutions within a large solution space. These methods are particularly effective in solving complex, multi-objective allocation problems, where trade-offs between resource utilization, energy efficiency, and latency must be considered. By iteratively exploring potential allocations, heuristic algorithms can identify solutions that balance multiple criteria, thereby improving the scalability and responsiveness of cloud infrastructures.

Virtual Machine Migration

VM migration involves moving a VM from one physical host to another within the same data center or across distributed cloud environments. Migration is typically employed to address issues related to load balancing, energy efficiency, maintenance, and fault tolerance. By migrating VMs to different hosts as needed, cloud providers can prevent server overloading, reduce power consumption, and ensure high availability of services.

There are two primary types of VM migration: *live migration* and *cold migration*. Live migration allows a VM to be moved while it is running, minimizing downtime by transferring the VM's state (e.g., memory, CPU, and I/O states) incrementally. This enables a seamless transition with minimal impact on user applications. Cold migration, in contrast, involves stopping the VM before moving it, which may result in some downtime. Live migration has become the preferred method for dynamic cloud environments where service continuity is critical, Kaur, G., & Sachdeva, R. (2021).[10]

Background Study On Open-Source Cloud Infrastructure

Open-source cloud infrastructure has become an essential component of modern IT environments, allowing organizations to create scalable, flexible, and cost-effective solutions for managing VMs and optimizing resource utilization. Platforms such as OpenStack, CloudStack, and

Eucalyptus provide robust frameworks for deploying and managing cloud resources while giving users the freedom to customize configurations and workflows. Virtual Machine (VM) allocation and migration, key aspects of cloud resource management, are critical for maximizing performance, maintaining service quality, and enhancing fault tolerance within open-source cloud infrastructures.[11,12] This background study examines the characteristics, advantages, and challenges of open-source cloud infrastructures, as well as the specific considerations they present for VM allocation and migration, Yang, C. T., & Wan, T. Y. (2020), Thanh Bui, K., Dac Ho, H., Pham, T. V., & Tran, H. C. (2020).

Overview of Open-Source Cloud Infrastructure

Open-source cloud platforms allow users to deploy cloud environments in a customizable, scalable manner without relying on proprietary software. Such platforms enable organizations to implement private, public, or hybrid cloud solutions tailored to their specific operational needs. Among the most widely adopted open-source cloud infrastructures are:

Openstack

A widely used platform that supports infrastructure as a Service (IaaS), OpenStack provides tools for managing computing, storage, and networking resources in a scalable environment. Its modular architecture, with components like Nova for computing and Neutron for networking, enables administrators to customize deployments for optimal VM management and migration, Prameela, P. K., Gadagi, P., Gudi, R., Patil, S., & Narayan, D. G. (2021).[13]

Apache cloudstack

This platform is designed for simplicity and ease of use, supporting IaaS deployments across diverse hypervisors, including Xen, KVM, and VMware. CloudStack offers built-in support for VM management and migration, allowing for seamless allocation and balancing across resources, Fard, M. V., Sahafi, A., Rahmani, A. M., & Mashhadi, P. S. (2020) [14].

Eucalyptus

Focusing on hybrid cloud environments, Eucalyptus is designed to integrate with AWS, providing users with an open-source solution for creating AWS-compatible private clouds. VM allocation and migration in Eucalyptus leverage the platform's AWS compatibility, simplifying hybrid cloud integration, Madakatte, B. K., & Nagesh, H. R. (2021).

The modularity and extensibility of these open-source platforms make them ideal for organizations that require high flexibility and customization. By using open-source cloud infrastructure, organizations avoid vendor lock-in, enabling the adoption of evolving technologies and community-driven innovations in VM management.

Importance of VM Allocation and Migration in Open-Source Cloud Infrastructure

In cloud computing, VM allocation and migration are critical processes that ensure optimal resource utilization, performance, and reliability. Allocation refers to the initial placement of VMs on specific physical hosts based on workload requirements and available resources, while migration involves moving VMs between hosts to balance workloads, prevent server overload, and maintain system stability. Open-source cloud infrastructure presents unique opportunities and challenges for VM allocation and migration, as users have control over the system's configuration and can implement customized resource management policies.

VM allocation and migration strategies are particularly significant in open-source environments for several reasons, Li, N., Liu, X., Wang, Y., & Mojarad, M. (2023), Talwani, S., Alhazmi, K., Singla, J., Alyamani, H. J., & Bashir, A. K. (2022):

Resource optimization

Open-source cloud platforms allow administrators to develop customized allocation algorithms that align with specific resource optimization goals, such as minimizing latency, balancing CPU and memory loads, and maximizing energy efficiency.

Load balancing and scalability

As open-source platforms support large-scale VM deployments, effective allocation and migration are essential for balancing loads across servers, avoiding resource hotspots, and ensuring scalability.

Fault tolerance and maintenance

Open-source platforms support proactive migration strategies that enhance fault tolerance, allowing VMs to be moved away from potentially failing servers or during maintenance windows without service interruption.

By enabling dynamic allocation and migration, open-source cloud infrastructure can adapt to fluctuating workloads and changing operational requirements, enhancing overall system resilience and flexibility.

Key VM Allocation and Migration Strategies in Open-Source Platforms

Several VM allocation and migration strategies have been developed to address the challenges of managing resources in open-source cloud environments, Seyyedsalehi, S. M., & Khansari, M. (2022):

Static vs. dynamic allocation

In static allocation, VMs are assigned to hosts based on predefined criteria, with minimal adjustments during runtime. Conversely, dynamic allocation uses real-time data to reassign VMs based on current workload demands,

enhancing flexibility. Open-source platforms enable administrators to implement dynamic allocation policies using custom scripts or external orchestration tools, such as Kubernetes, which are often integrated with platforms like OpenStack.

Live vs. cold migration

Live migration, in which VMs are transferred without downtime, is preferred for dynamic cloud environments, as it minimizes service disruption. Open-source platforms support live migration, though it requires significant network and computational resources. Cold migration, involving VM downtime, is used when live migration is not feasible, such as during system upgrades or reboots.

Energy-aware allocation and migration

Many open-source platforms support energy-efficient allocation and migration policies, often through integration with power management tools. Techniques such as VM consolidation and load shifting help reduce energy consumption by migrating VMs to fewer servers during off-peak hours, allowing underutilized servers to enter low-power states.

Literature Review

The authors proposed a new multi-objective optimization method of dynamic resource allocation for multi-virtual machine distribution stability. Combining the current state and future predicted data of each application load, the cost of virtual machine relocation and the stability of new virtual machine placement state are considered comprehensively. A multi-objective optimization genetic algorithm (MOGANS) was designed to solve the problem. The simulation results show that compared with the genetic algorithm (GA-NN) for energy saving and multi-virtual machine redistribution overhead, the virtual machine distribution method obtained by MOGANS has a longer stability time. Aiming at this shortage, this paper proposes a multi-objective optimization dynamic resource allocation method (MOGA-C) based on MOEA/D for virtual machine distribution. It is illustrated by experimental simulation that moGA-D can converge faster and obtain similar multi-objective optimization results at the same calculation scale, Shi, F., & Lin, J. (2022).

The authors proposed a bespoke virtual machine orchestrator (BVMO) as a method for constructing a virtual machine. The BVMO builds resource volumes as core assets to meet user requirements and builds virtual machines by reusing and combining these resource volumes. This can increase the reusability and flexibility of virtual machine construction. A case study was conducted to build a virtual machine by applying the proposed BVMO to an actual OpenStack cloud platform, and it was confirmed that the construction time of the virtual machine was reduced compared with that of the existing method, Park, J., Jeong, S., & Yeom, K. (2023).

The authors a methodology to design VMP algorithms using arbitrary PEs. Moreover, a novel algorithm to address the VMP problem using RF elements in cloud infrastructure is proposed. The methodology includes discovering evaluation the environment, models, parameters extraction, limitations, adaptation, problem formulation and heuristic. Among those, parameter extraction has a critical role in the overall performance. The extracted parameters are employed to make a decision about which PM is more appropriate for hosting the desired VM. According to simulation results on synthetic workloads, our proposed VMP algorithm outperforms others in operation with our proposed cloud architecture model, Farzaneh, S. M., & Fatemi, O. (2022).

The authors proposed an Open-Source Development Model Algorithm (ODMA) as a meta-heuristic approach to solve the VMP problem, which is named VMP-ODMA. VMP-ODMA seeks to dynamically consolidate VMs into a minimum number of active hosts by migrating VMs over cloud data centers. VMP-ODMA can perform the placement process dynamically and periodically by finding the best sequence of VMs for migration. In addition to minimizing the number of active hosts, a load-balancing strategy is included in VMP-ODMA to improve the quality of service without violating the SLA. We demonstrate the effectiveness of the proposed scheme with extensive simulations. Experimental results show that VMP-ODMA can efficiently improve system performance and outperform the best results of existing methods ranging from 11 to 27%, Li, N., Liu, X., Wang, Y., & Mojarad, M. (2023).

The authors presented an energy aware VM allocation and migration approach to meet the challenges faced by the growing number of cloud data centres. Machine Learning (ML) based Artificial Bee Colony (ABC) is used to rank the VM with respect to the load while considering energy efficiency as a crucial parameter. The most efficient virtual machines are further selected and thus, depending on the dynamics of the load and energy, applications are migrated from one VM to another. The simulation analysis is performed in Matlab and it shows that this research work results in more reduction in energy consumption as compared to existing studies, Talwani, S., Alhazmi, K., Singla, J., Alyamani, H. J., & Bashir, A. K. (2022).

The authors proposed a new approach for VM placement in a multi data center (DC) cloud environment. The aware genetic algorithm first fit (AGAFF) is a context-aware algorithm that distinguishes big data tasks with an input tag and uses a structure to minimize the traffic between MapReduce nodes. This multi-objective algorithm is based on the genetic algorithm, which is incorporated with the first fit methodology. The algorithm minimizes energy usage by minimizing the number of used servers, intra-DC traffic of big data tasks, and VMs' live migration while maximizing relevant usage of CPU and RAM in every server.

Furthermore, it improves job execution time, especially in big data processing, and reduces service level agreement (SLA) violations. A comparison between the results of AGAFF and four other algorithms shows by about 61% energy consumption reduction on average on different scales and approves a decrease in the number of needed PMs, intra-DC traffic of big data processing, and the number of live migrations, Seyyedsalehi, S. M., & Khansari, M. (2022).

The authors combined resource allocation security with efficient task scheduling in cloud computing using a hybrid machine learning (RATS-HM) technique is proposed to overcome those problems. The proposed RATS-HM techniques are given as follows: First, an improved cat swarm optimization algorithm-based short scheduler for task scheduling (ICS-TS) minimizes the make-span time and maximizes throughput. Second, a group optimization-based deep neural network (GO-DNN) for efficient resource allocation using different design constraints, including bandwidth and resource load. Third, a lightweight authentication scheme, i.e., NSUPREME is proposed for data encryption to provide security to data storage. Finally, the proposed RATS-HM technique is simulated with a different simulation setup, and the results are compared with state-of-art techniques to prove the effectiveness. The results regarding resource utilization, energy consumption, response time, etc., show that the proposed technique is superior to the existing one, Bal, P. K., Mohapatra, S. K., Das, T. K., Srinivasan, K., & Hu, Y. C. (2022).

The objective is to find the best allocation of VMs to PMs. Due to the NP-hardness of the VM allocation problem, nature-inspired meta-heuristic algorithms have become commonly used to solve it. However, there are no comprehensive and in-depth review papers on this specific area. This paper aims to bridge a knowledge gap by providing an understanding of the significance of metaheuristic methods to address the VM allocation issue effectively. It not only highlights the role played by these algorithms but also examines the existing methods, provides comprehensive comparisons of strategies based on key parameters, and concludes with valuable recommendations for future research, Xiaoqing, Y. A. N. G. (2023).

Virtual machines play an important role for the improvement of data center performance therefore, different approaches are used for the improvement of virtual machine efficiency (i-e) load balancing of resources and task. For the improvement of this section, different parameters of VM improve like makespan, quality of service, energy, data accuracy and network utilization. Improvement of different parameters in VM directly improve the performance of cloud computing. Therefore, we conducting this review paper that we can discuss about various improvements that took place in VM from 2015 to 20,201. This review paper also contains information about various parameters of cloud computing and the final section of the paper presents the role of machine learning algorithms in VM as well load balancing approach along with the future direction of VM in cloud data center, Ullah, A., Nawi, N. M., & Ouhame, S. (2022).

With the rapid growth of IAAS (Infrastructure as a Service) cloud solutions, its flexibility, energy savings, cost, and optimal resource utilization, most organizations have turned to the right cloud to deploy the application that users need. There are a variety of open-source cloud platforms such as OpenStack, Eucalyptus, Opennebula, etc. However, it is difficult for organizations to select an appropriate platform to host Web applications or deploy virtual machines for high availability. From this perspective, researchers have tended to compare the performance of cloud platforms in order to obtain the most suitable platform for cloud deployment. In this article, the performance of OpenNebula and Eucalyptus open-source clouds was evaluated through a series of hardware and software experiments ending in a user-friendly selection, Khair, Y., Dennai, A., & Elmir, Y. (2022).

The authors presented a comparative study that evaluated three open-source; CloudStack, OpenStack and Eucalyptus to assess their viability as open-source cloud solutions for business and research purposes, also the limitations and strengths were considered. The performance of the individual platforms was evaluated based on Network performance, Memory, CPU utilization and disk I/O of the individual platforms independently deployed as a virtual machine on different underlying hardware of identical configurations. Four benchmarking tools namely LINPACK, Bonnie++, Stream and Iperf were used for the evaluation of CPU utilization, Disk I/O throughput, memory, and Network utilization respectively. Experimental results revealed the strengths and limitations of OpenStack, CloudStack, and Eucalyptus. CloudStack outperformed OpenStack and Eucalyptus on the CPU, Memory, and Network utilization metrics. OpenStack however outperformed Cloudstack and Eucalyptus on the Disk I/O metric. CloudStack recorded a better network (Iperf) compared with Eucalyptus Cloud and OpenStack. Recommendations for possible optimization configuration and direction for further work have equally been presented, Aregbesola, M. K., Aro, T. O., Aiyeniko, O., Olukiran, O. O., & Akanni, O. O. (2022).

The authors presented an overview of the OpenStack software, with a focus on constructing a scalable hosting architecture tailored to educational settings. We discuss functional and architectural details that are essential to implementing unique cloud computing models for virtualization purposes. We describe an experimental virtualization setup that was implemented within an educational scenario, along with a guideline for configuring OpenStack. Overall, this study offers insights into the potential of OpenStack for virtualizing large-scale educational setups, paving the way for cost-effective and

efficient resource utilization, Abbasi, M., Cardoso, F., Silva, J., & Martins, P. (2023, June).

The authors aimed to improve network management in OpenStack Clouds by taking advantage of the combination of software-defined networking (SDN), Network Function Virtualization (NFV), and machine learning/artificial intelligence (ML/AI) and for making networks more predictable, reliable, and secure. Artificial intelligence is being used to monitor the behavior of the virtual machines and applications running in the OpenStack SDN cloud so that when any issues or degradations are noticed, the decision can be quickly made on how to handle that issue, being able to analyze data in motion, starting at the edge. The *OpenStackDP* framework comprises lightweight monitoring, anomaly-detecting intelligent sensors embedded in the data plane, a threat analytics engine based on ML/AI algorithms running inside switch hardware/network co-processor, and defensive actions deployed as virtual network functions (VNFs). This network data plane-based architecture makes high-speed threat detection and rapid response possible and enables a much higher degree of security. We have built the framework with advanced streaming analytics technologies, algorithms, and machine learning to draw knowledge from this data that is in motion before the malicious traffic goes to the tenant compute nodes or long-term data store. Cloud providers and users will benefit from improved Quality-of-Services (QoS) and faster recovery from cyber-attacks and compromised switches. The multi-phase collaborative anomaly detection scheme demonstrates an accuracy of 99.81%, average latencies of 0.27 ms, and response speed within 9 s. The simulations and analysis show that the OpenStackDP network analytics framework substantially secures and outperforms prior SDN-based OpenStack solutions for Cloud architectures, Krishnan, P., Jain, K., Aldweesh, A., Prabu, P., & Buyya, R. (2023).

The authors presented an accurate energy-conscious method for initial resource allocation by casting the issue of energy-efficient resource allocation as a bin-packing model. The available VMs employ a modified version of the max-min scheduling technique, which saves money and resources. The results of this study give a framework for comparing and contrasting the many different resource distribution approaches that other researchers have proposed. The importance of efficient data centers for the cloud is growing. Power consumption has been a major problem due to its expanding size and widespread usage. The overarching purpose of this effort is to create models and algorithms for resource allocation that are both energy-efficient and take into account a variety of relevant factors, Singh, K., & Kaur, J. (2023).

The authors present the need for improvement in the current security monitoring platform to support more types of monitoring devices and cover the consequences of multi-tenant setups. Future work includes incorporating log collectors and aggregators and addressing the needs of a super-tenant in the security monitoring architecture. The equitable sharing of monitoring resources between tenants and the provider should be established with an adjustable threshold mentioned in the SLA. The results of experiments show that Enhanced Round-Robin uses less energy compared to other methods, and the Fusion Method outperforms other techniques by reducing the number of Physical Machines turned on and increasing power efficiency, Mahmood, S., Yahaya, N. A., Hasan, R., Hussain, S., Malik, M. H., & Sarker, K. U. (2023).

Problem Statements

Despite their advantages, VM allocation and migration processes face several challenges that can impact performance and scalability. One of the main challenges is migration overhead, which refers to the resource consumption and latency associated with migrating VMs. Live migration, in particular, requires substantial bandwidth and computational resources, which can lead to temporary performance degradation. To mitigate migration overhead, techniques such as compression and delta encoding are often employed, which reduce the volume of data transferred during migration.

Network latency is another significant issue, especially in geographically distributed cloud environments. Migrating VMs across long distances may introduce network delays, which can impact application performance. To address this, cloud providers are exploring edge computing models, where computing resources are located closer to end-users, thereby reducing latency for latency-sensitive applications.

Additionally, security and privacy concerns arise during VM migration, as sensitive data may be exposed during transfer. Encryption and secure transfer protocols are essential to protecting data integrity and confidentiality during migration. However, implementing these security measures may increase migration overhead, creating a trade-off between security and performance.

Future Research Directions

To address these challenges, emerging innovations are being integrated into open-source cloud infrastructure:

Edge computing

By extending cloud resources to the edge of the network, platforms such as OpenStack and CloudStack support distributed VM allocation, reducing latency and improving responsiveness for latency-sensitive applications.

Container orchestration

Containers, managed through tools like Kubernetes, offer a lightweight alternative to VMs, enabling faster allocation and migration across nodes. Containers' compatibility with open-source platforms enables organizations to implement microservices and distributed architectures, facilitating resource scaling and efficient migration.

AI and machine learning in allocation

AI-driven resource management is transforming VM allocation and migration by predicting workload patterns and automatically adjusting resources. Integrated within open-source platforms, AI and machine learning algorithms can enhance decision-making for dynamic allocation and proactive migration.

References

Abbasi, M., Cardoso, F., Silva, J., & Martins, P. (2023, June). Exploring OpenStack for Scalable and Cost-Effective Virtualization in Education. *In International Conference on Disruptive Technologies, Tech Ethics and Artificial Intelligence* (pp. 135-146). Cham: Springer Nature Switzerland.

Aregbesola, M. K., Aro, T. O., Aiyeniko, O., Olukiran, O. O., & Akanni, O. O. (2022). Open-source cloud computing platforms: comparative study. *Adeleke University Journal of Engineering and Technology*, 5(2), 125-136.

Bal, P. K., Mohapatra, S. K., Das, T. K., Srinivasan, K., & Hu, Y. C. (2022). A joint resource allocation, security with efficient task scheduling in cloud computing using hybrid machine learning techniques. *Sensors*, 22(3), 1242.

Balaji Naik, B., Singh, D., & Samaddar, A. B. (2020). FHCS: Hybridised optimisation for virtual machine migration and task scheduling in cloud data center. *IET Communications*, 14(12), 1942-1948.

Fard, M. V., Sahafi, A., Rahmani, A. M., & Mashhadi, P. S. (2020). Resource allocation mechanisms in cloud computing: a systematic literature review. *IET Software*, 14(6), 638-653.

Farzaneh, S. M., & Fatemi, O. (2022). A novel virtual machine placement algorithm using RF element in cloud infrastructure. *The Journal of Supercomputing*, 78(1), 1288-1329.

Huang, Y., Xu, H., Gao, H., Ma, X., & Hussain, W. (2021). SSUR: an approach to optimizing virtual machine allocation strategy based on user requirements for cloud data center. *IEEE Transactions on Green Communications and Networking*, 5(2), 670-681.

Jalali Moghaddam, M., Esmaeilzadeh, A., Ghavipour, M., & Zadeh, A. K. (2020). Minimizing virtual machine migration probability in cloud computing environments. *Cluster Computing*, 23, 3029-3038.

Kaur, G., & Sachdeva, R. (2021). Virtual machine migration approach in cloud computing using genetic algorithm. *In Advances in Information Communication Technology and Computing: Proceedings of AICTC 2019* (pp. 195-204). Springer Singapore.

Khair, Y., Dennai, A., & Elmir, Y. (2022). An experimental performance evaluation of OpenNebula and eucalyptus cloud platform solutions. *In Artificial Intelligence and Heuristics for Smart Energy Efficiency in Smart Cities: Case Study: Tipasa*, Algeria (pp. 450-457). Springer International Publishing.

Krishnan, P., Jain, K., Aldweesh, A., Prabu, P., & Buyya, R. (2023). OpenStackDP: a scalable network security framework for SDN-based OpenStack cloud infrastructure. *Journal of Cloud Computing*, 12(1), 26.

Li, N., Liu, X., Wang, Y., & Mojarad, M. (2023). Improving dynamic placement of virtual machines in cloud data centers based on open-source development model algorithm. *Journal of Grid Computing*, 21(1), 13.

Madakatte, B. K., & Nagesh, H. R. (2021). Arr And Fusion Based Virtual Machine Scheduling Techniques For Eucalyptus Cloud. *Webology* (ISSN: 1735-188X), 18(6).

Mahmood, S., Yahaya, N. A., Hasan, R., Hussain, S., Malik, M. H., & Sarker, K. U. (2023). Self-adapting Security Monitoring in Eucalyptus Cloud Environment. *International Journal of Advanced Computer Science and Applications*, 14(3).

Masoudi, J., Barzegar, B., & Motameni, H. (2021). Energy-aware virtual machine allocation in DVFS-enabled cloud data centers. *IEEE Access*, 10, 3617-3630.

Pande, S. K., Panda, S. K., Das, S., Sahoo, K. S., Luhach, A. K., Jhanjhi, N. Z., ... & Sivanesan, S. (2021). A Resource Management Algorithm for Virtual Machine Migration in Vehicular Cloud Computing. *Computers, Materials & Continua*, 67(2).

Park, J., Jeong, S., & Yeom, K. (2023). Bespoke Virtual Machine Orchestrator: An Approach for Constructing and Reconfiguring Bespoke Virtual Machine in Private Cloud Environment. *Applied Sciences*, 13(16), 9161.

Prameela, P. K., Gadagi, P., Gudi, R., Patil, S., & Narayan, D. G. (2021). Energy-efficient VM management in OpenStack-based private cloud. *In Advances in Computing and Network Communications: Proceedings of CoCoNet 2020*, Volume 1 (pp. 541-556). Springer Singapore.

Seyyedsalehi, S. M., & Khansari, M. (2022). Virtual machine placement optimization for big data applications in cloud computing. *IEEE Access*, 10, 96112-96127.

Shi, F., & Lin, J. (2022). Virtual machine resource allocation optimization in cloud computing based on multiobjective genetic algorithm. *Computational Intelligence and Neuroscience*, 2022(1), 7873131.

Singh, K., & Kaur, J. (2023). MECHANISMS AND TOOLS USED FOR RESOURCE ALLOCATION IN THE CLOUD. *International Journal of Advanced Research in Computer Science*, 14(3).

Singh, S., & Singh, D. (2021). A virtual machine migration mechanism based on firefly optimization for cloud computing. *Recent Patents on Engineering*, 15(4), 9-20.

Talwani, S., Alhazmi, K., Singla, J., Alyamani, H. J., & Bashir, A. K. (2022). Allocation and migration of virtual machines using machine learning. *Computers, Materials and Continua*, 70(2), 3349-3364.

Thanh Bui, K., Dac Ho, H., Pham, T. V., & Tran, H. C. (2020). Virtual machines migration game approach for multi-tier application in infrastructure as a service cloud computing. *IET Networks*, 9(6), 326-337.

Ullah, A., Nawi, N. M., & Ouhame, S. (2022). Recent advancement in VM task allocation system for cloud computing: review from 2015 to2021. *Artificial Intelligence Review*, 55(3), 2529-2573.

Xiaoqing, Y. A. N. G. (2023). Nature-Inspired Optimization for Virtual Machine Allocation in Cloud Computing: Current Methods and Future Directions. *International Journal of Advanced Computer Science & Applications*, 14(11).

Yang, C. T., & Wan, T. Y. (2020). Implementation of an energy saving cloud infrastructure with virtual machine power usage monitoring and live migration on OpenStack. *Computing*, 102(6), 1547-1566.

Yang, L., Yang, D., Cao, J., Sahni, Y., & Xu, X. (2020). QoS guaranteed resource allocation for live virtual machine migration in edge clouds. *IEEE access*, 8, 78441-78451.

Zhuang, H., & Esmaeilpour Ghouchani, B. (2021). Virtual machine placement mechanisms in the cloud environments: a systematic review. *Kybernetes*, 50(2), 333-368.

Zolfaghari, R., Sahafi, A., Rahmani, A. M., & Rezaei, R. (2021). Application of virtual machine consolidation in cloud computing systems. *Sustainable Computing: Informatics and Systems*, 30, 100524.