ABSTRACT

Cloud computing is a model in which computer resources are provided as paid services to users. It has enjoyed wide spread acceptance in recent times due to its numerous advantages, particularly in terms of cost savings. Despite its advantages, it still has some challenges, suchas Quality of Service (QoS) adherence, efficient resource utilization, and energy conservation. From the Cloud users' perspective, service provisioning and adherence to QoS are vital; while to the Cloud Service Providers (CSPs), efficient resource utilization and conservation of the energy consumed by Cloud data centres are key. Meeting these requirements collectively is a major Cloud computing challenge. Virtual Machine Consolidation (VMC) has become the *de facto* technique for addressing resource utilization and energy conservation, while efficient scheduling and migration techniques are used to address QoS. Power-Aware Best Fit Descending (PABFD) and Virtual Machine Consolidation with Multiple Usage Prediction (VMCUP-M) are examples of recent research works that usedthe combination of both approaches to address this challenge. However, in these works like many others, users were treated equally without regard to class of requirements. In a bid to address this shortcoming, this study proposed an approach that groupedworkloads into classes, then used a class-based VMC scheme to ensure end-to-end adherence to QoS.It also improved on resource utilization and energy conservation by combining resource utilization prediction with a half-interval workload allocation scheme. The proposed approach is called Multi-Class Load Balancing (MC-BAL). MC-BAL was tested against PABFD and VMCUP-M, using three datasets and across static and dynamic PM threshold schemes. Obtained results show that MC-BAL met the set objectives by performing better than PABFD and VMCUP-M in terms of QoS adherence, by an average of 13 % and 26 % for both static and dynamic thresholds respectively. With respect to resource utilization, MC-BAL used at least 19 % less resources to accomplish the same tasks versus the other two approaches. In terms of energy conservation, MC-BAL consumed an average of 24.7 % less energy than both PABFD and VMCUP-M. MC-BAL is able to address critical requirements that are vital to Cloud stakeholders. For Cloud users, MC-BAL ensures satisfactory (quality)service delivery. Using MC-BAL, CSPs can manage their resources better,creating room to take on new customers and increase profit margins. To the society at large, MC-BALlowers energy consumption of data centres and by extension carbon emission, thereby addressing one of the major concerns in the world today.

Keywords: Cloud Computing, Energy Conservation, Quality of Service, Resource Management, Virtual Machine Consolidation.