A Test-bed to Evaluate The Impact of Dynamic Monitoring Interval Adjustment on Power Consumption in Virtualized Data Centers

Mark White, Hugh Melvin, Michael Schukat

Discipline of Information Technology, NUI Galway, Galway, Ireland

m.white1, hugh.melvin, michael.schukat{@nuigalway.ie}

ABSTRACT

Virtualization is one of the principle Data Center (DC) technologies increasingly deployed in recent years to meet the challenges of escalating costs, industry standards and the search for a competitive edge. In this poster we describe a novel approach to management of the virtualized system which dynamically adjusts the monitoring interval with respect to the average CPU utilization for the DC. We propose this will facilitate identification of cost opportunities which may otherwise have remained hidden mid-interval. We explain our adjustment algorithm and outline our CloudSim [1] test-bed with which we aim to evaluate our approach.

INTRODUCTION

In a virtualized data center, multiple Virtual Machines (VMs) are typically co-located on a single physical server, sharing the processing capacity of the server's CPU between them. When, for example, increased demands on the CPU result in reduced performance of one of the VMs to the point where a Service Level Agreement (SLA) may be violated, virtualization technology can facilitate a migration [2]. Migration relocates the services being provided by the VM on this 'over-utilized' host to a similar VM on another physical server, where sufficient capacity (e.g. CPU) is available to maintain SLA performance.

By dynamically adjusting the monitoring interval with respect to the incoming workload (indicated by a weighted CPU utilization average for the data center), we aim to identify performance issues at an earlier stage than systems which use a static interval and thus facilitate improved decision making.

PREDICTION OF UTILIZATION

We calculate the weighted mean of the next predicted CPU utilization value for all operational hosts in the data center using:

$$x = \frac{\sum_{n=1}^{n} w_i x_i}{\sum_{n=1}^{n} w_i}$$

where w is the weight applied to the range within which the predicted utilization value x for each operational host falls and n is the number of operational hosts. This value is then used to set the monitoring interval such that the interval is decreased when overall utilization is higher and increased for periods of low utilization. See Fig. 1 for an illustration.

TEST-BED VALIDATION PHASE

In order to validate our approach, we have performed some preliminary tests based on comparison of the default LR/MMT CloudSim simulation with a simulation which includes our dynamic interval adjustment algorithm. Our test-bed servers and VM models are provided by the CloudSim framework. Workload data was gathered by CloudSim from over 500 DC locations worldwide. We modified the relevant modules in CloudSim to implement our algorithm. Figure 1 depicts the intervals (727) calculated by our dynamic adjustment algorithm during a 24-hour simulation. Initial test results suggest a reduction of 2-3kWh in power consumption. However, a reduced number of migrations is also reported by CloudSim. This is counter-intuitive. We are currently carrying out a full review of the CloudSim architecture before completing detailed testing.

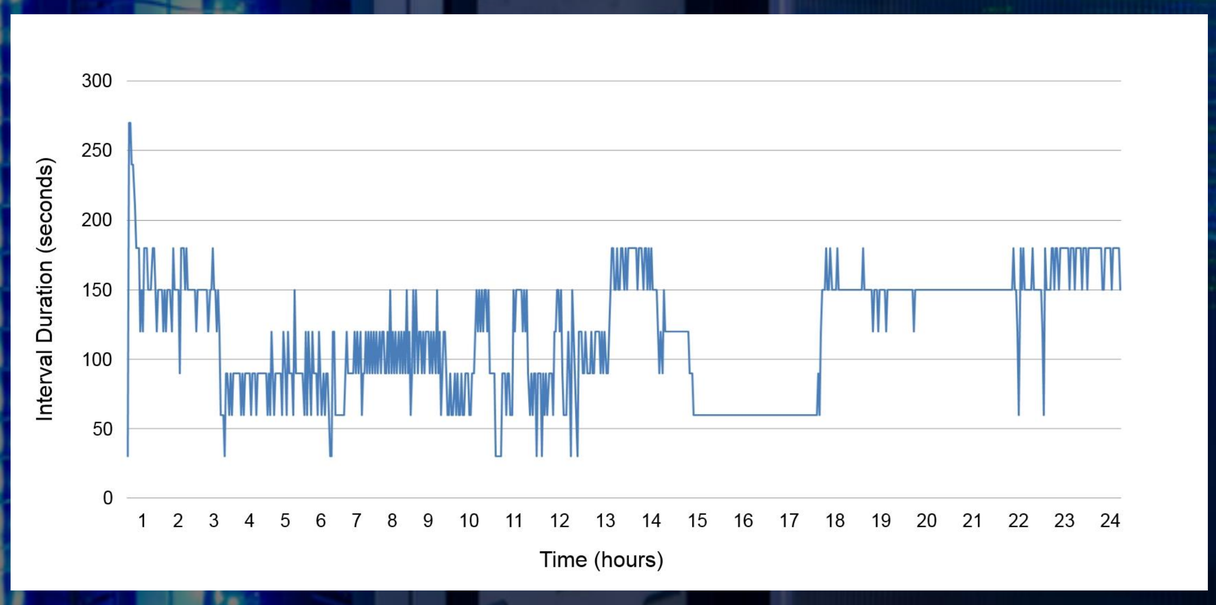

Figure 1: Intervals calculated by the adjustment algorithm

REFERENCES

[1] Buyya, R., Ranjan, R., Calheiros, R. N.: Modeling and Simulation of Scalable Cloud Computing Environments and the CloudSim Toolkit: Challenges and Opportunities. In: High Performance Computing & Simulation, 2009. HPCS'09. International Conference on, pp. 1-11. IEEE (2009)

[2] Clark, C., Fraser, K., Hand, S., Hansen, J. G., Jul, E., Limpach, Pratt, I., Wareld, A.: Live Migration of Virtual Machines. In: Proceedings of the 2nd Conference on Symposium on Networked Systems Design & Implementation, vol .2, pp. 273-286. USENIX Association (2005)