# Evaluating the Impact of Full Virtualized High-Performance Computing Platform on Large Scale Scientific Data using Quantum Espresso

Michael Ametepe Kattah Kwame Nkrumah University of Science and Technology Department of Computer Science Kumasi – Ghana University of Ghana Computing System University of Ghana Legon - Ghana

Dominic Asamoah Kwame Nkrumah University of Science and Technology Department of Computer Science Kumasi - Ghana Frimpong Twum Kwame Nkrumah University of Science and Technology Department of Computer Science Kumasi - Ghana

# ABSTRACT

High Performance Computing (HPC) applications are becoming vital in scientific research for analyzing large scale scientific data, but there is inadequate knowledge on the impact of a fully virtualized HPC cluster on these applications when they are used to analyzed large scientific data. The main purpose of this research is to carry out a comparative experiment on a virtual HPC cluster and the traditional HPC cluster by executing a benchmarking tool called para-speedup with an input file on both clusters using Quantum Espresso (QE) as an HPC application to determine the impact on the clusters. The research focuses on Central Processing Unit (CPU) utilization, turnaround time of jobs run on the cluster, memory and input/output (I/O) operations. The virtual cluster was setup using VMWare ESXi 5.5.0 as hypervisor of choice and ROCKS was installed on the cluster as an HPC platform of choice. During the experiment, it was observed that the job was not memory and I/O intensive on both clusters, so there was little to discuss on these metrics but generally it was observed that running job using HPC applications like QE on a fully virtualized HPC cluster to analyze large scale scientific data has a negative performance impact on the completion of the job as compared to the traditional cluster.

#### **General Terms**

High Performance Computing, Cluster, Virtualization, VMWare, Quantum Espresso, ROCKS, Hypervisor

# **Keywords**

High Performance Computing, Cluster, Virtualization, VMWare, Quantum Espresso, ROCKS, Hypervisor

# **1. INTRODUCTION**

Scientific computing happens every time where research scientists seek to evaluate their research data in the shortest possible time. This desire gave rise to High Performance Computing (HPC) cluster. The data is analyzed on a traditional HPC cluster thereby flooding datacenters with physical devices leading to high energy consumption and high operating cost.

The HPC community has not taken full advantage of the enormous benefits virtualization brings to institutions, in order to fully virtualized HPC cluster to analyze scientific data. There is this argument that scientists seek to attain performance gains by using all available resources within the traditional HPC cluster, though there is limited research on impact on virtualized HPC be it positive or negative when scientific data is analyzed on the virtual cluster (Ranadive *et al.*, 2008). Also, there is limited information on the comparative cost of setting up the traditional HPC cluster as against the virtual HPC cluster. There is therefore the need for a comparative study to be carried out on both clusters to ascertain their impacts when large scale scientific data are analyzed to enable researchers to make an inform decision on the appropriate platform to analyze their scientific data

# 1.1 Aim and Objectives

This paper implements comparative experimental analysis on a virtual HPC cluster and a traditional HPC cluster by executing a benchmarking tool to evaluate the impact of full virtualized HPC cluster on a large-scale scientific data using Quantum Espresso (QE) as an HPC application by determining the run-time or execution time of a job, the compilation time of QE, and the memory utilization when jobs are been executed of both clusters.

# **1.2 Background**

Scientific HPC applications are increasingly becoming common in scientific research. These applications can analyze large and complex scientific data. However, the applications require huge amount of resources such as CPUs, memory, high speed network etc. to successfully run. This large amount of computing power requirement has been addressed by traditional High-Performance Computing (HPC) clusters which is a cluster of computers made up of multi-core nodes and using high performance devices such as InfiniBand to interconnect them and managed by "CPU scavenger" software such as ROCKS for scientific computing (Netto *et al.*, 2018).

Even though there have been attempts by (Gavrilovska *et al.*, 2007), (Ranadive *et al.*, 2008), (Walters *et al.*, 2008) and (Regola and Ducom, 2010) to fully virtualized HPC cluster, research scientists are unwilling to accept this platform to host scientific applications in order to analyze their large and complex scientific data because of their perceived desire to explore all the available resources in the cluster for performance gains. However, a research by the world's most leading virtualization company revealed that, deploying application on a physical server uses only 15% of its efficiency in terms of resources.

Virtualization is the process of creating a software-based (or virtual) representation of something rather than a physical one. Virtualization can apply to applications, servers, storage, and networks and is the single most effective way to reduce IT expenses while boosting efficiency and agility for all size businesses; bringing simplicity to datacenters.

One of the major scientific software that scientists use in large scale experiment analysis is Quantum ESPRESSO (The name "ESPRESSO" stands for opEnSource Package for Research in Electronic Structure, Simulation, and Optimization). This is an integrated suite of OpenSource application program for electronic structure calculations and materials modeling based on Density-Functional Theory (DFT), Plane waves, and Pseudopotentials which is free and released under the terms of the General Public License (GPL) (Giannozzi and Cavazzoni, 2009).

#### 2. LITERATURE REVIEW

Virtualization has become popular in improving system utilization and performance by consolidating IT resources, leading to the reduction of power and cooling costs. In addition to the improved utilization, other benefits of virtualization, such as flexible resource management, fault isolation, increased availability, ease of operating system deployment and simplified disaster recovery, have led to the increased interest in the virtualization of computing clusters for high performance computing (Vmware.com, 2015).

(Trangoni and Cabral, 2012), vividly described a modern High-Performance Computing to have diverse stack of hardware and software in which the hardware comprises of servers acting as compute nodes, one or more communication networks, and a storage system, whereas the software includes:

- An operating system, which is usually a Linux distribution,

- A provisioning system that allows a user to install the software stack and configure the cluster, - Tools for monitoring and managing the status of the cluster, configuration and resources,

- A Parallel File System (PFS) which is an optional component, and

- The applications and the programs that the users use to run their jobs.

According to Trangoni & Cabral, 2012, there are a lot of HPC Provisioning platforms available in the field of High Performance Computing. Below are few of them:

- ROCKS is an open-source Linux cluster distribution that enables end users to easily build computational clusters, grid endpoints and visualization tiled-display walls. Hundreds of researchers from around the world have used Rocks to deploy their own cluster

- xCAT is an open source provisioning system from IBM, oriented to large systems (up to 100,000 nodes in a hierarchical routing infrastructure).

- IBM Platform HPC is an easy to use, yet comprehensive technical computing cluster management software. Its robust cluster and workload management capabilities are accessible using the latest design in web-based interfaces, making it powerful and simple to use.



Fig 1: Architecture of an HPC Cluster

According to (Hassani and Luksch, 2014), improving as well as evaluating the performance or the impact of full virtualized HPC platform for large scale scientific computing is considered a critical issue in the field of High-Performance Computing (HPC). HPC employs parallel processing to perform huge computations in shortest possible time. HPC platforms require a cluster setup with massive number of high- end computers, mostly used for scientific research in academia and industries. Virtualizing such 'capacity computing' platforms implies the shared use of not only the nodes and node cores, but also of the cluster interconnect.

Traditional HPC is becoming increasingly popular in the HPC community recently and before recently, researchers in the HPC community have scarcely discussed the comparison between the cost of doing scientific computation in the traditional HPC environments and that of the virtual HPC environment. This comparative cost information is very important to the research scientists to make an informed decision on which environment to use (Carlyle et al. 2010).

In considering the cost of doing scientific computation, several factors must be considered. These include the cost of hardware, software, data center space, energy consumption, Staff to manage the Cluster, consultancy and performance (Wu et al. 2015).

Marathe et al., 2013, also vividly discussed the turnaround time and the total cost of execution in HPC cluster during scientific computation but was unable to decisively conclude on which environment is beneficial in terms of cost. However, in the virtual environment, the benefits of virtualization come to play as resources can be quickly provisioned and released as and when the need arises, hence reducing unnecessary energy consumption cost and the cost of data center space usage as compared to the traditional environment. Hence Wu et al. 2015, concluded that, doing scientific computation in the virtual environment helps to drastically cut down the cost associated with engineering design and manufacturing.

#### **3. METHODOLOGY**

This study seeks to determine the impact of fully virtualized HPC platform on large scale scientific data using Quantum (QE) as an HPC application. Two physical server systems with the same specification were used. One of the physical servers has been fully virtualized by installing VMWare vSphere Hypervisor on it while the other one was used as a compute node in a physical HPC cluster setup. Two virtual machines were created and named Frontend node and Compute node respectively. The frontend node has two network interfaces; internal interface for communicating between the compute nodes within the cluster and external or public interface for communicating with the frontend node and the user community to submit jobs to the compute nodes. On the Frontend node, ROCKS cluster distribution software was installed to help build the HPC compute cluster. After the installation of ROCKS, compute node was added to the cluster from the Frontend node using the "insert-ethers" command and a minimal ROCKS cluster distribution software was automatically deployed to the compute node.

Quantum Espresso was installed and compiled on the Frontend node to be used as an HPC application for large scale scientific data analysis. The configuration on the Frontend note is synchronized with the compute cluster for uniformity.

After all these, an input file is used to benchmark the system and results collected for analysis.

Simulation environments were;

- VMWare ESXi 5.5 and vSphere Client 5.5 ROCKS 6.2
- Quantum Espresso 6.2.1
- HPC Benchmarking tool Para-Speedup

Table 1. Speedup Test Results in Bare Metal Cluster

CPU Cores	CPU-Time(s)	Speedup	I/O-Time(s)
2	17863.14	1.00000	0.00
4	9836.18	1.81606	0.00
6	7966.78	2.24220	0.00
8	7383.11	2.41946	0.00
10	6283.46	2.84288	0.00



Fig 2: Architecture of the proposed virtualized HPC Cluster

The virtual infrastructure of the cluster consists of two (2) virtual machines – Frontendnode with 6GB RAM and 8vCPU core and a Computenode with 10GB RAM and 12vCPU core,

and two virtual network switches (Internal and Public or External) were created using VMware vSphere Hypervisor.

The VMware ESXi hypervisor was selected as a hypervisor of choice because a comparative analysis between virtual platforms carried out by (Pawar and Singh, 2015), indicates that VMware ESXi Hypervisor is far better to meet the demand of an enterprise than any other.

# 3.1 Benchmarking of the systems under study

To achieve the research goals, a benchmarking tool called para-speedup.sh for a small-size cluster was downloaded from the Quantum Espresso (QE) official website and edited to suit the system under study. A python script was written to call computing resources from the computeNode and using QE to run the para-speedup.sh tool to benchmark the system under study. This tool was run on both platforms (virtual and traditional) and the following metrics were recorded; number of CPU cores per execution, run-time or Turnaround time, and speedup values. The Compilation time of QE was also recorded from both environments during the compilation.

The input file used 6 K-point value of  $2 \ 2 \ 2 \ 0 \ 0$  for the benchmark. An automated script was written to run the job with an increment of two (2) CPUs and the results was sent into an output file for discussion.

# 4. EXPERIMENTAL RESULTS AND DISCUSSION

#### 4.1 Speedup Test

Speedup is simply how much faster a parallel version of a program runs as compared to a non-parallel version. The performance of a system is characterized by its speedup.

This benchmarking test used a script to automate the calculation of atomic structure positions using 6 K-points for a number of iterations. The higher the K-Points the longer the execution time of the application hence these K-points were carefully chosen to suit the system under study and also not to exceed the time allocated to the research. At the end of the test, the number of cpu cores used, cpu-time, also known as run-time or turnaround time of each iteration, the I/O time and the Speedup values were recorded. The Results from both experiments are presented in table 1 and table 2 below.

Table 2. Speedup Test Results in Virtual Cluster

CPU Cores	CPU-Time(s)	Speedup	I/O-Time(s)
2	17457.19	1.00000	0.00
4	13193.71	1.32314	0.00
6	13742.54	1.27030	0.00
8	13059.25	1.33677	0.00
10	14047.00	1.24275	0.00

Table 1 and 2 represents the results of both bare-metal and virtual platforms with CPU increment of two (2) for each parallel computation.

Comparing the speedup columns from both tables, the baremetal platform has shown an increase in computational speed as the processing power of the cluster increases. The speedup values indicate the bare-metal cluster is 2.8x faster in handling large scale scientific data in parallel computation than a nonparallel one when the CPU cores are increased to 10 while the speedup values for the virtual platforms were inconsistent with increment in the processing power of the cluster as shown in table 2. Even though, on the average, there has been an insignificant increase in the computational speed of the virtual cluster as compared to the bare-metal. The virtual platform shows an average of 1.25x faster in handling large scale scientific data in a parallel computation compared to a non-parallel. A comparative analysis of the speedup values for both bare-metal and virtual platforms is presented in Fig. 3 in which bare-metal has shown better performance than the virtual cluster.



Fig 3: Comparative graph of Speedup values in Bare metal vs Virtual Cluster



Fig 4: Comparative graph of turnaround time of Bare metal vs Virtual Cluster

Fig. 4 shows a declining curve for the turnaround time of both platforms under study. The graph compares the runtime of a job executed on the bare-metal and the virtual cluster respectively from Table I and II. As the CPU cores are increased by 2 for each job there is a reduction in the runtime. At a glance on the graph, it is clearly observed that there is declining curve. This pattern on the graph shows a reduction in the execution time of job in parallel computation as the computing power of the cluster increases.

Nevertheless, there is a huge gap between bare-metal and virtual cluster. This means there is a better performance in bare-metal cluster with parallel job execution compared to the virtual cluster. Fig. 4 shows the graphical comparison that reveals good turnaround in bare- metal as against virtual throughout the computation.

# 4.2 Compilation Time of Quantum Espresso (QE)

QE version 6.2.1 was downloaded from the original site and install on both traditional and the virtual clusters. The was successfully compiled on the cluster using (tar, ./configure, and make all) commands in oerder to successfully benchmark the clusters. Compilation times for both clusters were recorded. To get a better result, QE was compiled three times and the time for each compilation was recorded. The average of the three was calculated and recorded as the final value. It was noted that the virtual machine recorded approximately 600s against 875s for the bare-metal. This reveals a better performance in the virtual machine compared to the baremetal as shows in Fig. 5 below.



Fig 5: Comparative chart of QE compilation time

# 4.3 Memory Utilization

During the experiment, as jobs are run on both platforms, their memory utilizations were monitored and recorded. It was noted after a successfully completion of the job that, the virtual cluster recorded 6493 MB memory utilization while the bare metal recorded 6495 MB memory utilization. This reveal a higher memory utilization in the bare metal as compared to that of the virtual cluster. However, the difference is insignificant to have any effect on performance. This is because the job being run is not memory intensive.

# 5. CONCLUSION

In this study, a comparative experimental impact analysis was performed in a full virtualized HPC infrastructure and a traditional or physical HPC infrastructure to determine the impact if any of full virtualized HPC cluster on large scale scientific data using QE as an HPC Application. This is intended to determine the suitability of analyzing Large scale scientific data in a fully virtualized HPC platform compared to the traditional platforms.

In conclusion, running jobs using HPC Application like Quantum Espresso in a full virtualized HPC platform to analyze large scale scientific data has a negative performance impact on the completion of the job as compared to that of physical platform. This is due to the abstraction layer added by virtualization which results in CPU overheads coupled with the high network latency at the virtualization layer.

#### 6. REFERENCES

- Ranadive, A. *et al.* (2008) 'Performance implications of virtualizing multicore cluster machines', in *Proceedings* of the 2nd workshop on System-level virtualization for high performance computing - HPCVirt '08. New York, New York, USA: ACM Press, pp. 1–8. doi: 10.1145/1435452.1435453.
- [2] Netto, M. A. S. *et al.* (2018) 'HPC Cloud for Scientific and Business Applications: Taxonomy, Vision, and Research Challenges', 51(1), pp. 1–29. doi: 10.1145/3150224.
- [3] Gavrilovska, A. et al. (2007) 'High-performance hypervisor architectures: Virtualization in hpc systems', virtualization for HPC ( .... Available at: http://www.cc.gatech.edu/~adit262/docs/HPHA-HPCVirt2007.pdf (Accessed: 4 June 2017).
- [4] Walters, J. P. et al. (2008) 'A comparison of virtualization technologies for HPC', in Proceedings International Conference on Advanced Information Networking and Applications, AINA, pp. 861–868. Available at: http://ieeexplore.ieee.org/xpls/abs\_all.jsp?arnumber=448 2796 (Accessed: 25 October 2016).
- [5] Regola, N. and Ducom, J.-C. J. (2010) 'Recommendations for virtualization technologies in high performance computing', *Cloud Computing Technology and*, pp. 409–416. Available at: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arn umber=5708479 (Accessed: 25 October 2016).
- [6] Giannozzi, P. and Cavazzoni, C. (2009) 'Large-scale computing with Quantum-Espresso', *Nuovo Cimento C*, Vol. 32 C,. doi: 10.1393/ncc/i2009-10368-9.

- [7] Vmware.com (2015) Virtualization Technology & Virtual Machine Software / United States, VMware, Inc. Available at: https://www.vmware.com/solutions/virtualization.html (Accessed: 8 September 2017).
- [8] Trangoni, M. and Cabral, M. (2012) 'A Comparison of Provisioning Systems for Beowulf Clusters', (3).
- [9] Hassani, R. and Luksch, P. (2014) 'Performance Analysis of HPC Applications Running in Public Cloud', *Future Generation*, 00(2012). Available at: http://www.sciencedirect.com/science/article/pii/S01677 39X12001458 (Accessed: 25 October 2016).
- [10] Carlyle, A. G., Harrell, S. L. and Smith, P. M. (2010) 'Cost-effective HPC: The community or the cloud?', *Proceedings - 2nd IEEE International Conference on Cloud Computing Technology and Science, CloudCom* 2010, pp. 169–176. doi: 10.1109/CloudCom.2010.115.
- [11] Wu, D., Terpenny, J. and Gentzsch, W. (2015) 'Cloud-Based Design, Engineering Analysis, and Manufacturing: A Cost-Benefit Analysis', *Proceedia Manufacturing*. Elsevier B.V., 1, pp. 64–76. doi: 10.1016/j.promfg.2015.09.061.
- [12] Marathe, A. et al. (2013) 'A comparative study of highperformance computing on the cloud', Proceedings of the 22nd international symposium on High-performance parallel and distributed computing SE - HPDC '13, pp. 239–250. doi: 10.1145/2462902.2462919.
- [13] Pawar, S. and Singh, S. (2015) 'Performance Comparison of VMware and Xen Hypervisor on Guest OS', International Journal of Innovative Computer Science & Engineering Issue, 2(3), pp. 56–60. Available at: http://ijicse.in/wp-content/uploads/2015/07/v2i3-14.pdf.