

# High Performance Computing Cloud - a Platform-as-a-Service Perspective

Pratima Dhuldhule, J. Lakshmi, S. K. Nandy

Supercomputer Education and Research Centre

Indian Institute of Science, Bangalore, India

pratima@cadl.iisc.ernet.in, {j|lakshmi, nandy}@serc.iisc.ernet.in

**Abstract**—HPC applications are widely used in scientific and industrial research, data analytic and visualization, social behavioral studies etc. Most HPC applications require dedicated, available and highly customized resources and environments for computation since they exhibit intense resource utilization. These needs were traditionally provided by clusters and supercomputers which are difficult to setup, manage or operate. While majority of the HPC installations ensure good resource utilization, the reach of these is restricted to few who are members of a specific HPC community. Cloud computing is emerging as a latest computing technology. The on-demand nature of cloud has provoked interest to explore if cloud properties can be useful for HPC setups. This paper is a work in that direction. The prevalent public clouds have accessibility to many and have been explored by the HPC community too. The biggest deterrent identified on these computing platforms for HPC workloads is the virtualization layer used by the cloud systems for resource provisioning.

In this paper we propose a Platform-as-a-Service model to build an HPC cloud setup. The key goals for the architecture design is to include features like on-demand provisioning both for hardware as well as HPC runtime environment for the cloud user and at the same time ensure that the HPC applications do not suffer virtualization overheads. The architecture builds the required HPC platform by providing dedicated node or a group of nodes booted with the desired HPC environment without the virtualization layer. Technologies like Wake-on-LAN and network booting are used to achieve this goal. Once the usage of these resources is relinquished, the same nodes are re-deployed for another HPC platform. Thus this architecture merges cloud properties with HPC platforms for delivering effective performance. We show the results of benchmarks used to evaluate performance difference between a virtualized and non-virtualized environment for this observation.

**Keywords:** Cloud Computing, High Performance Computing, Virtualization, Platform-as-a-Service

## I. INTRODUCTION

Cloud computing is growing as a popular computing technology. Cloud is being extensively used for enterprise or business tasks because there is no need to own an infrastructure, handle the services and maintain it. It provides on-demand resources and services on a metered basis. This is cost efficient as we pay for whatever we are using and once a task is done we give away the resources and stop paying for it. Also the availability, scalability, flexibility are some of the attractive features of cloud computing. It's availability and on-demand nature are some of the major reasons for it's growth. The main idea behind cloud computing is enabling a virtual machine to run on any server. Since there are many

customers and many servers, the management of the infrastructure must be highly automated. A customer can request the creation or removal of a virtual machine and without human intervention a virtual machine is started or stopped on one of the servers. To take advantage of the economic benefits, the cloud providers use multi-tenancy, where virtual machines from multiple customers share a server. This catering to different users at a time is possible due the hypervisors present in virtualization layer. At a physical level, the resources are shared between different users and applications. It caters to different variety of applications such as enterprise workload, social networking, data storage and research computations. The needs of these applications are also different. Resource requirements are mainly in the area of networking, disk IO, CPU etc.

Due to all the benefits of cloud, it seems an attractive option for HPC workloads. Feature like flexibility plays an important role as it is not seen in dedicated machines. If on a dedicated machine, different environment is required, a new cluster has to be set up with the required environment which can take upto several days. On the other hand cloud provides variety of environment on-demand easily. The property of elasticity in the cloud is also useful for HPC applications as an HPC application may require scaling in between the computation. Availability is also one of the property useful for HPC workloads.

In spite of all the above advantages of cloud for HPC applications, there are some serious issues related with it. HPC applications are used in academia and laboratories for research, hence need guarantee of resources and timely results. They tend to consume more than 90% of the available CPU cycles, require low latency and high bandwidth inter-process communication, which is difficult to achieve in a virtualized environment due to I/O overhead and delay in CPU cycles [1]. Previous work shows that Cloud is one or two magnitude worse as compared to Infiniband, which is currently used in supercomputers for interconnect network [2] [3]. The virtualization of interconnect becomes a bottleneck for HPC applications and increases latency of communication intense applications. Thus the use of virtualization layer proves to be a major hindrance in performance of HPC applications [4].

The notion of virtualization was brought in mainly to utilize all the existing resources, as virtual machines share the overall resources instead of having equal share of it. Under-utilization

is not an issue in case of HPC applications as they are capable of utilizing major portion of available resources. Past results [2] [4] [5] on evaluation of HPC applications on Cloud using IaaS platform with virtualization have been pessimistic. [2] tells that not all types of HPC applications are suitable for Cloud. The HPC applications with less intensive communication patterns and having less sensitivity towards interference perform well on Cloud. Applications with performance needs that can be met at small to medium scale execution (in terms of number of cores) are suitable for Cloud under virtualized environment [5] [6].

This paper discusses the idea of converging the benefits of Cloud like elasticity and flexibility with HPC platforms to create an environment for HPC workload. PaaS properties are explored for achieving this goal as PaaS has the ability to provide on-demand platform, abstracts and controls the underlying resources and gives choice of platform and guarantee of performance. Instead of improving the virtualization layer we propose a model where virtualization layer is completely removed. By elimination of virtualization layer and with PaaS model a framework is obtained for provisioning of dedicated platforms. On-demand environment is setup by booting a node with desired operating system and runtime libraries for building user-specific HPC platform. The key contributions of this work are as follows:

- Provision of a non-virtualized platform to obtain contention free and dedicated use of resources.
- Design of a PaaS model to provide on-demand and isolated environment to meet HPC application requirements.

The organization of the rest of the paper is as follows: Section 2 discusses about the related work done in the field of HPC cloud which is followed by Section 3 where we discuss the methodology of the framework proposed. Section 4 talks about experimental setup and shows the results of the experiments carried out for performance evaluation between virtualized and non-virtualized environment. We discuss conclusion and future work in section 5.

## II. RELATED WORK

Significant work has been done for achieving high performance in cloud. The merging of the two worlds is not new and has been studied from long time. Scientists have marched towards deploying their HPC workloads on cloud. So far, provision of dedicated resources for HPC applications is largely unexplored on the cloud setup.

There has been work done towards using latest technology for improving performance in various aspects in the cloud. A design of heterogeneous high performance IaaS using OpenStack to address issues regarding virtualization performance, heterogeneous hardware and GPUs, high speed interconnects, advanced scheduling and high performance storage is discussed in [7]. Several science clouds have been established for scientific computations [8]. But no considerable work is done towards completely removing the virtualization layer to avoid the overhead. In [9], strategies have been discussed to remove the hypervisor while still keeping the notion of

VMs. But there are architectural as well as implementation issues which are not properly addressed. [10] discusses about environment aware scheduling of HPC applications on cloud setup to improve performance.

It has been observed that the performance obtained on cloud is much below than the dedicated machines. The major performance hindrance is virtualization. Also security issues are also attached for the scientists while deploying their HPC workload on cloud [11]. Amazon has come up with a HPC cloud with special instances provided for compute intense applications. Different instance types and sizes are provided for deploying HPC workload at low cost. NASA has conducted a performance evaluation test of Amazon EC2 vs Pleiades, a supercomputer in NASA. Their performance is tested using various benchmarks, and considerable drop in performance is shown in [12]. This work done by NASA concludes that Amazon EC2 are not yet ready for HPC workloads but the technology sure is catching up fast. Similar tests have been carried out to evaluate the performance of HPC applications on current Cloud setups in [5] [2] and [4].

[3] discusses about the virtualization impact on the network performance. The unstable network characteristics are caused by virtualization and processor sharing on server hosts. The unstable network performance degrades the performance of and bring new challenges to many applications.

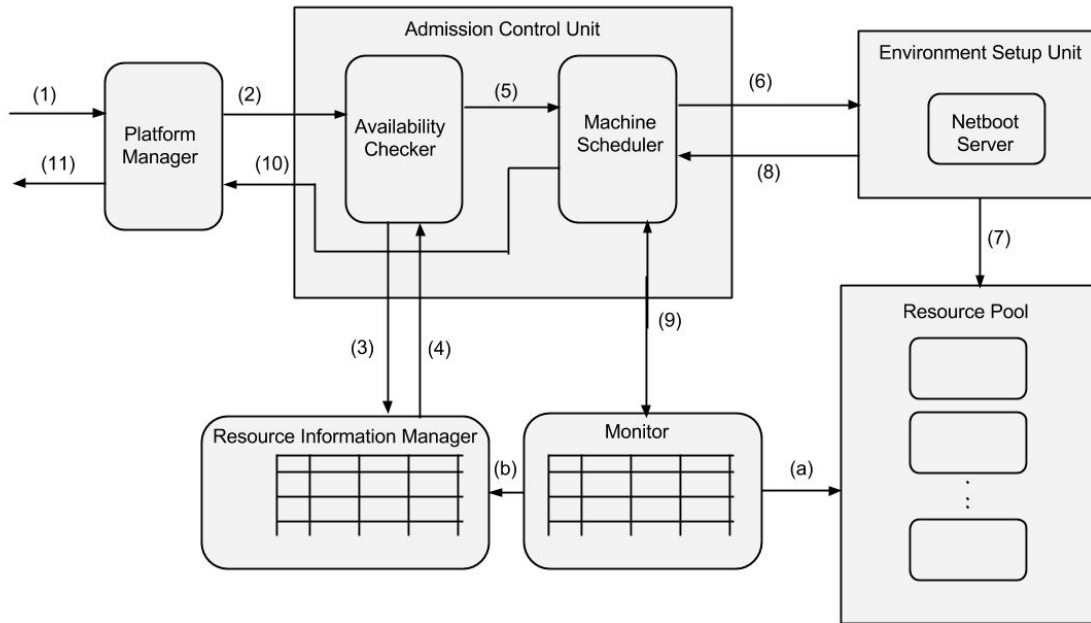
[13] discusses about virtualization impact on HPC applications for IO and makes an effort in improving IO performance with the help of SR-IOV. It concludes that virtualized performance for HPC workloads is competitive with that of native if network-interrupt parameters are tuned to increase the responsiveness of the SR-IOV system.

In [14], linux containers and VMs are discussed. This study compares the performance of linux containers and virtual machines. Performance over different physical resources is measured and it concludes that the performance of linux containers is equal or better than VMs in most cases.

The studies carried out till now talk about improvement of HPC application on the existing architecture of Cloud. But the methodology presented in this paper gives an architecture which addresses all the challenges of deploying HPC on Cloud. The architecture completely removes the virtualization layer and provides dedicated resources in terms of CPU cores, memory bandwidth, network bandwidth and storage.

## III. METHODOLOGY

There has been work done towards improving the performance of HPC applications on cloud platforms. The work done till date leans towards minimization of the virtualization overhead and building strategies for isolation on existing cloud setups. But the problem of latency caused for HPC workload remains. There is a need of such a strategy which gives performance comparable to the supercomputers and dedicated machines used for HPC workloads. In this work, we present a strategy which brings together the benefits of cloud and the necessity of HPC applications. In this section an architecture is presented which delivers the required platform



- |   |   |
|---|---|
| (1) user sends request to the system                    | (8) ESU sends machine ID to the machine scheduler                                 |
| (2) request in the form of XML file                     | (9) machine scheduler informs monitor about new cluster and gets status in return |
| (3) check if the resources are available and compatible | (10) response to the request is the status  |
| (4) response to the request for resources               | (11) status is informed to the user   |
| (5) resource list made available with machine scheduler | (a) monitor gets status of all the machines                                       |
| (6) list passed to ESU for creating a compute cluster   | (b) monitor updates free list of machines   |
| (7) ESU boots machines with required setups             |   |

Fig. 1. HPC Cloud Architecture

to the user. The main components of the architecture, as shown in Figure 1, are Platform Manager (PM), Admission Control Unit (ACU), Resource Information Manager (RIM), Environment Setup Unit (ESU) and Monitor. User interacts with the system through PM which acts as front end, with the help of which user sends a request for a platform. ACU accepts the request from PM and performs the task of provisioning of requested platform to the user if resources are available. RIM is responsible for maintaining the record of available resources, frequently updating it and providing it to the ACU. The setting up of environment and platform is done by the Environment Setup Unit. After a machine is given to a user Monitor maintains the state of the machines and responds back to the ACU. We call the machine or group of machines created for the user as "compute cluster". Each component of the architecture is explained in detail below.

#### A. Platform Manager

The PM in the architecture acts as an interface between the user and the system. It provides user the ability to create, manage and remove the machines allocated to it. Before specifying resource requirement, user needs to specify the type of request, that is, start a new compute cluster (fresh request), scale an existing compute cluster (scalability request) and stop existing compute cluster (termination request). Then the user specifies

resource requirement which is divided into two parts. One is the Host requirement and other is Environment requirement. The host requirement part consist of type of processor, number of CPU cores, RAM size etc. The environment requirement part consist of type of OS, runtime environment etc. Once this information is obtained, PM forwards it to the ACU in the form of a XML file. PM is the only way to communicate with the user hence the system responds back to the user with the status of request. Also a machine ID is generated to represent the compute cluster and to ensure isolation. Any further communication related to the platform, like scaling or stopping the machine, is done through this interface.

#### B. Admission Control Unit

The Admission Control Unit, shown in Figure 2, is responsible for managing the user requests and sending instruction to other units for provisioning of a platform. As the request comes in the form of a XML file from the PM, the Availability Checker passes this information to RIM and asks to check for the required resources. If the resources are available and compatible with the required environment, RIM sends back the marked resources to ACU. Further the Availability Checker Unit forwards these marked resources and runtime environment requirement to the Machine Scheduler. Machine Scheduler is responsible for giving away the control

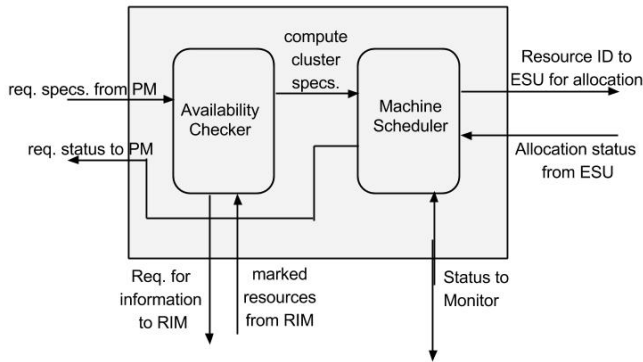


Fig. 2. Admission Control Unit

of compute cluster to the user. It does this by communicating environment requirements to the ESU to obtain a platform of user specifications. Once ESU boots the required machines it sends back a unique machine ID for the user which is used for future communication related to the platform. Once a desired machine with required platform is obtained, Machine Scheduler sends status and machine ID to the PM informing about the successful allocation of the platform. It also informs the Monitor so that it makes an entry in its table to keep track of running systems.

### C. Resource Information Manager

The Resource Information Unit is responsible for maintaining information about the available resources. RIM maintains a list of all free machines in a table populated with information like RAM size, number of CPU cores, size of disk, NICs and runtime environment. These entries are maintained by the Monitor in the table. RIM gets requests from the ACU to check for required resources. RIM matches the request with the available resource. The reason to check the availability is that there is no notion of virtualization present, the resources become finite in number. Hence availability of resources is checked. Apart from availability RIM also checks for compatibility of required runtime environment and hardware as some hardware might not support all the libraries. If they match then those resources are marked and removed from the free list. The information about these marked resources like MAC ID's is passed back to the Availability Checker in the ACU.

### D. Environment Setup Unit

The Environment Setup Unit is responsible for setting up the required environment for the user platform. The Machine Scheduler from the ACU sends MAC ids of the marked machines and the runtime environment to setup on those machines. The ESU has a netboot server which acts as NFS, DHCP and FTP server. All the OS images are stored here. All the machines in the resource pool are in low power mode and hence the server in the ESU sends a magic packet to bring up those machines i.e. a notion of Wake-on-LAN is used here. Once the desired machines are up the netboot server configures them for booting over network. Once the machines are booted

with desired OS, required libraries are installed in the selected machines. After environment setup is done, ESU creates an ID associated with the compute cluster for provisioning to the user. This ID is unique and represents a single compute cluster. As the user interacts with the machines using this unique ID, total isolation of resources is achieved. This ID is given to the Machine Scheduler to notify it about the completion of platform creation.

### E. Monitor

Monitor is responsible for keeping track of all the compute clusters and maintaining their state. It maintains the information in a table with fields such as User ID, compute cluster ID, MAC IDs of all machines present in the compute cluster and IP addresses associated with it and finally the status of the compute cluster. The User ID column need not be unique as a single user can own multiple compute clusters. The compute cluster ID field is unique one as each group is given different ID for isolation purpose. Monitor gets information about the new compute cluster from the Machine Scheduler to make an entry in the table. Until the machines are freed, Monitor keeps track of all the activities and maintains the current status of the compute cluster. Monitor is also responsible for updating the free list of machines maintained by the RIM. Once a machine or group of machines is freed by a user, Monitor adds those machines in the free list again.

### F. System Tasks

The architecture presented in this paper provides a platform with specified environment on user demand. The main component, ESU, is responsible for creating a compute cluster. It sets up a platform over the network which gives flexibility to the user to choose various options for the runtime environment. The management of the user machine is done through the user interface that is PM. PM allows user to start, scale or shutdown the machine. The user monitors the system through this interface only. All the status updates from the Monitor are displayed to the user for monitoring of activity and for the statistics of resource usage. During a scale up request, the user sends requirements to ACU specifying that it is a scale up request and provides the ID associated with the machine and required amount of resources to scale. Now again the process goes as before like a fresh request but the only difference now is that rather than associating a new ID to the machine, the old ID sent by the user is given. Hence it becomes a single cluster associated with same ID.

For scaling down the user specifies the number of nodes that are no longer required for the computation task. Hence they are freed and corresponding changes are made in the table managed by the Monitor. When user wants to give away the resources, request to terminate the use is sent and corresponding machines are freed and kept ready for re-deploying. The respective changes are made with the Monitor. In this manner, the architecture provides the functionality of flexibility and elasticity similar to cloud for HPC workload.

### G. Case Study - Machine Failure

A machine in the resource pool can fail or may not respond at times. Failure can occur in two cases. Once when a machine is not responding to the ESU during a netboot operation and second when a machine is given to the user. In the first scenario, failure is less likely to happen because the monitor keeps a free machine list with RIM. During updating the free machine list monitor has to talk to the machine and then it gets added. Hence only those machines are kept which respond to the monitor. But still if the machine fails and does not respond during booting, the ESU send back an error message to the Machine Scheduler conveying failure of machines. The Machine Scheduler collects the information about the required resources from the Monitor and again requests the Availability Checker to check for the availability of those resources. The Availability Checker again gives the required resources if available. In the later scenario, that is when the user owns the machine, the failure is detected by the Monitor. Monitor informs the Machine Scheduler about the failure and gives the information about failed machines. Machine Scheduler asks the Availability Checker for machines with same configurations. Once the machine MAC ids are obtained they are booted with required environment setup. Thus new set of nodes are booted up and are provided to the user.

## IV. EXPERIMENTS AND RESULTS

The proposed architecture in this paper builds a cluster without virtualization of any resources. The resource allocated by the architecture are dedicated to a particular user and no other user intervenes in the working of it. Hence it can be said that the compute cluster provided by this architecture will be same as the non-virtualized dedicated machines used for HPC workloads. Hence the performance measurement on non-virtualized machines will justify the performance of compute cluster and we claim that compute cluster will give better performance than virtualized environment. To support our claim, we have conducted experiments to measure performance difference between virtualized and non-virtualized environment. Experiments were carried out to measure usage of physical resources like CPU, Memory, Disk and Network.

### A. Experiment Setup

The experiments were conducted on a virtualized and a non-virtualized environment. The host machine had quad-core intel core i7 Processor with 8GB RAM and 1TB SATA2 disk. For setting up virtualized environment KVM was used as a hypervisor. All the VMs used in the experiments had single vCPU, 1 GB RAM and 30 GB disk space. All the VMs were pinned to separate CPU cores to obtain isolation and to avoid delay caused by CPU scheduling. The corresponding non-virtualized machine was booted with single CPU core and 1 GB RAM. All the machines used for the experiments were booted with Linux. During the experiment, number of VMs were increased gradually and the results were taken on VMs running the same workload simultaneously. For experiments

running same workload on multiple VMs, only those results were taken into consideration whose completion time differ by negligible amount. Average of the results on all VMs was calculated. The benchmarks considered depict the HPC workload behavior in terms of CPU, memory, disk IO and networking. Stream benchmark was used to measure memory bandwidth and corresponding execution time. NBench benchmark was used to measure CPU speed. IOR (Interleaved or Random) benchmark was used to measure parallel file system IO performance. And Iperf benchmark was used to measure bandwidth, packet loss and jitter during UDP and TCP data streams.

### B. Results

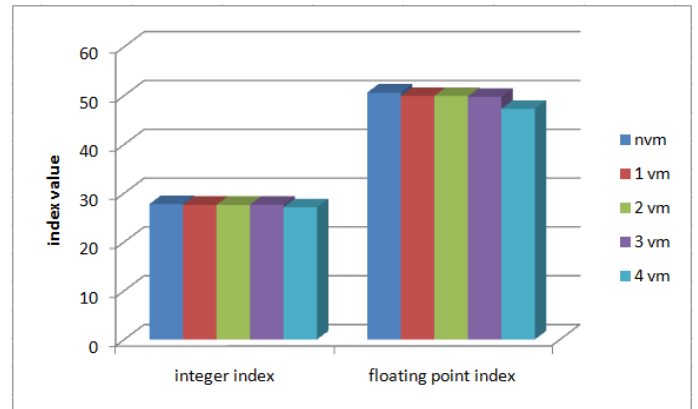


Fig. 3. NBench : CPU performance

1) *Effect of virtualization on CPU performance:* For measuring CPU performance NBench benchmark was used. NBench algorithm consist ten tasks which are a mixture of integer and floating point operations. Due to this mix, this benchmark resembles the HPC compute intense workload. The iterations per cycle are calculated for each algorithm. The integer index and the floating point index plotted in the graph are nothing but geometric mean of tests that require integer and floating point processing respectively. The benchmark was run on non-virtualized and virtualized machines and the number of VMs was increased gradually. It is clear from the graph shown in figure 3, that CPU performance is not affected by the virtualization layer and the performance is comparable with the non-virtualized machine.

2) *Effect of virtualization on Memory performance:* Stream benchmark is a popular benchmark in HPC community. It measures the memory bandwidth and the corresponding rate for four simple vector kernels. During the experiment, the data handled by Stream was kept more than the size of cache to eliminate any caching effect. The four vector kernels run multiple times and average bandwidth is calculated for those runs. The graph in figure 4 clearly shows that the bandwidth of non-virtualized machine is greater than the single virtual machine case. As the memory instruction goes through an additional address translation layer of hypervisor, loss in bandwidth is observed. Further as we increase the number

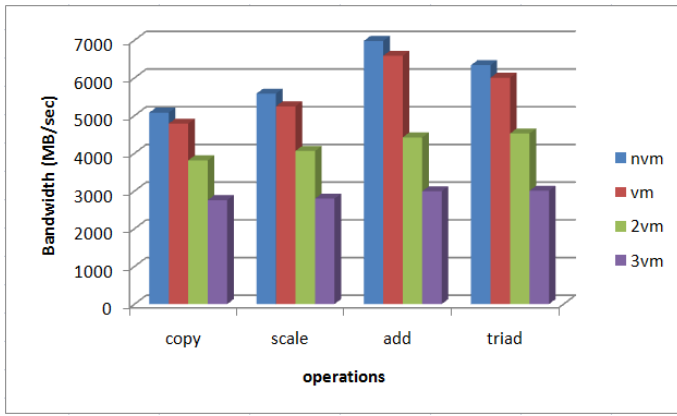


Fig. 4. Stream : Memory Bandwidth

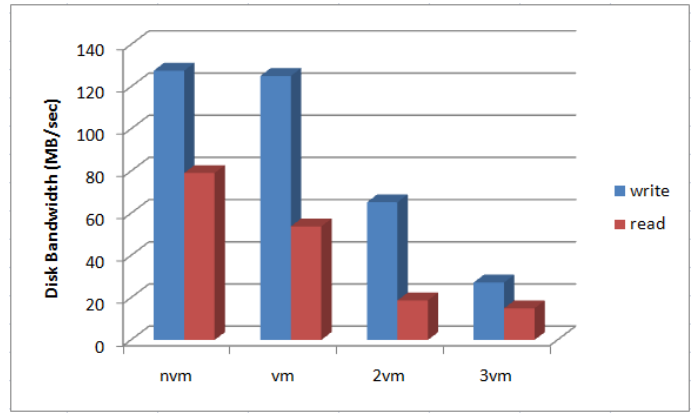


Fig. 6. IOR : Disk Bandwidth

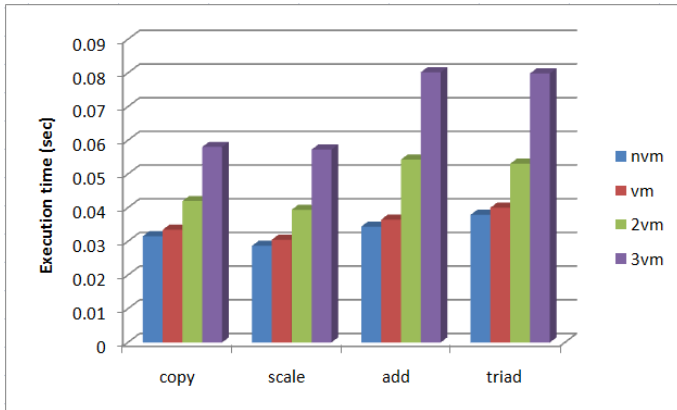


Fig. 5. Stream : Execution Time for Memory Operations

of VMs, the individual bandwidth for each VM decreases. The VMs share the memory channels and thus the placement of memory intense application affect the performance. The virtualization layer becomes a bottleneck for multiple VM case. The other graph in figure 5 shows execution rate for completing a task for each machine. As the bandwidth decreases, the corresponding latency raises. This effect is due to overhead caused by virtualization layer as VMs need extra time for going through extra abstraction layer.

3) *Effect of virtualization on Disk IO performance:* IOR is designed to measure parallel file system I/O performance at both the POSIX and MPI-IO level. This parallel program performs writes and reads to/from files under several sets of conditions and reports the resulting throughput rates. The api used here for conducting experiments is MPI-IO with number of clients set to 16. To avoid the possibility of caching, the file size was kept more than the free memory on each client hence the file size used here is 8 GB. While performing the operation, each participating task opens a shared file, transfers the data and closes the file. The time elapsed for transferring the data and total amount of data transferred is calculated from which bandwidth is obtained. As seen from the graph shown in figure 6, there is slight drop in write bandwidth while 30% drop in read bandwidth for a single VM case. As the number

of VMs are increased the write and read bandwidth starts to degrade further. This drop in bandwidth is caused due to the longer path a VM has to take to perform the operation. The disk request has to go from guest OS to the hypervisor and from the hypervisor to the host OS. The host OS then generates a call to the disk drivers to perform the operation. This process is not as speedy as direct physical hardware call by the host system. Due to this delay the bandwidth of the VM suffers. Thus virtualization of disk leads to degradation of IO performance.

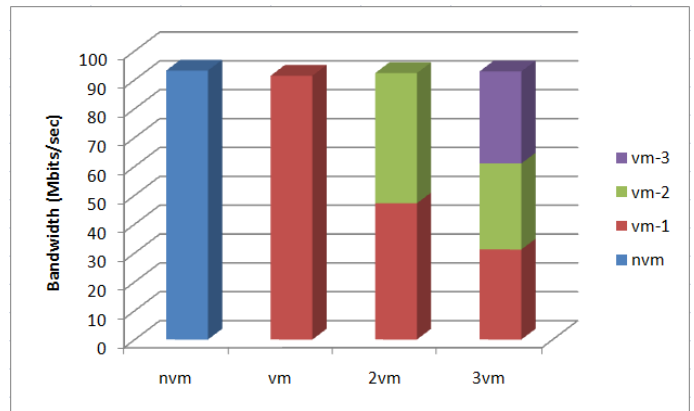


Fig. 7. Iperf : Bandwidth Measurement for UDP

4) *Effect of virtualization on Network performance:* Iperf benchmark creates Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) data streams and measure the throughput of a network that is carrying them. For conducting this experiment, server and client were setup on different machines present in the same network. For virtualization case, server was setup on a VM while client was setup on different machine in a non-virtualized environment. The number of server was increased gradually. Each VM was running as a server while clients were connected to each server separately. To get accurate results, experiment was conducted during a time when network usage was minimum. Network performance was evaluated in terms of bandwidth, jitter and

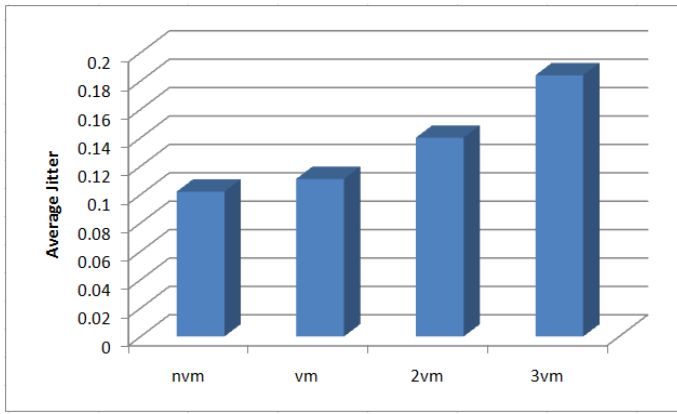


Fig. 8. Iperf : Jitter during UDP Transmission

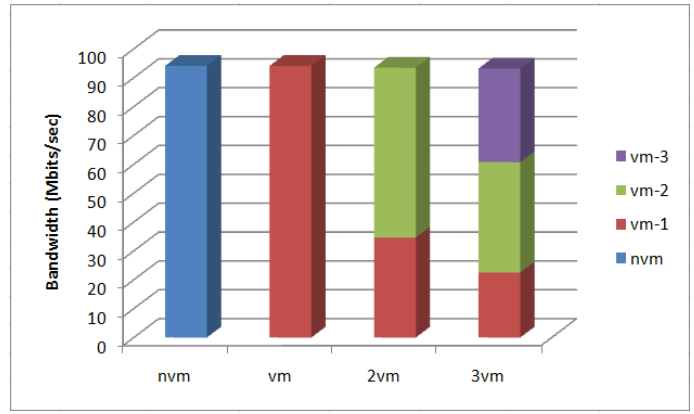


Fig. 10. Iperf : Bandwidth Measurement for TCP

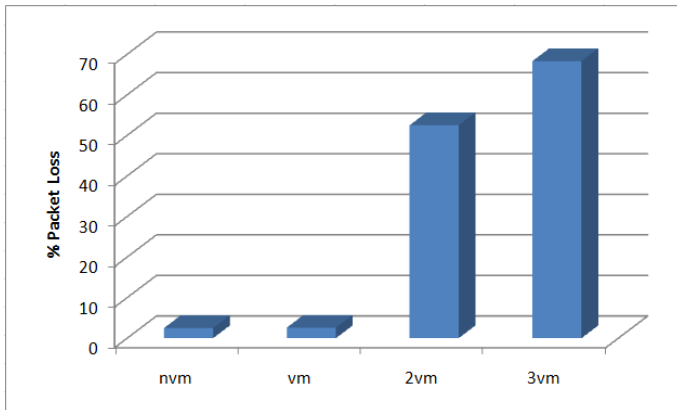


Fig. 9. Iperf : Packet Loss during UDP Transmission

packet loss during communication. Two separate tests for UDP and TCP were conducted. UDP transmission gives information about bandwidth, jitter and packet loss. As seen from the graph in figure 7, the total bandwidth during UDP transmission remains same while individual VM bandwidth is decreased as we increase the number of VMs. It can be seen clearly that the bandwidth is uniformly distributed among all the VMs. The graph in figure 8 shows jitter i.e. delay for response, which increases with number of VMs. As the VMs share the same NIC, the communication overhead increases with increased number of VMs. During the UDP transmission, packet loss is seen in the network. This packet loss also differs from machine to machine. Packet loss can be seen in the third graph from figure 9, where there is nearly 50% packet loss when 2 VMs are running simultaneously and 70% packet loss when 3 VMs run simultaneously.

The TCP transmission gives information about the bandwidth alone. TCP waits for an acknowledgement for certain amount of time and if it doesn't get one it considers that the packet was lost and again transmits the packet. Hence it ensures that all packet reach destination. Hence there is no notion of packet loss in TCP transmission. Figure 10 shows bandwidth of a network during TCP transmission. Same experimental setup as UDP was repeated for conducting this

experiment. It can be seen from the graph that the total bandwidth used remains same while individual bandwidth for each VM differ. Also it is observed from the graph that the individual bandwidth of VMs is not as uniform as in case of UDP. This is due to the re-transmissions caused when a packet or its acknowledgement is lost. When a re-transmission is caused the host resets its TCP window and size of buffer to minimum default size. And thus the throughput suffers due to this. Thus sharing of NIC causes each VM to wait for its share of packets and this delay causes re-transmission of packets. Thus virtualization of NIC is a bottleneck for VMs.

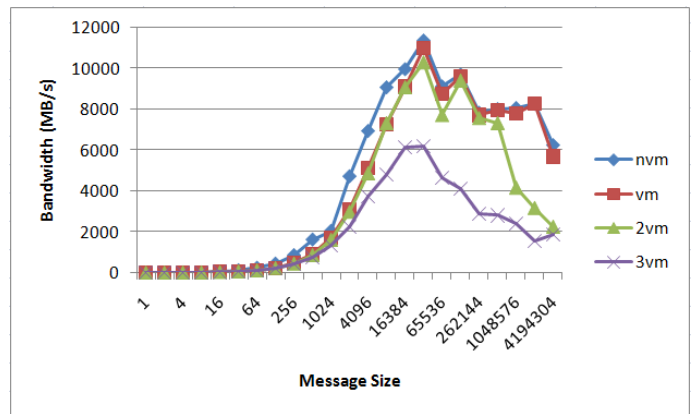


Fig. 11. OSU micro benchmarks : Bandwidth for point-to-point communication

We have also conducted experiments with the MPI version of OSU micro benchmarks to check the communication bandwidth and communication latency among the MPI processes. The algorithms checked the bandwidth, bi-directional bandwidth and latency for point-to-point communication. Measures uni-directional bandwidth from one to another MPI rank. Several Isends are started followed by a Waitall. The receiving side uses matching Ircvs with Waitall. One iteration ends when the sending side gets the receive of all messages acknowledged. The graph in figure 11 shows the bandwidth measured for different message sizes. It is observed that we get a peak performance for message size of approximately 64 KB. Also

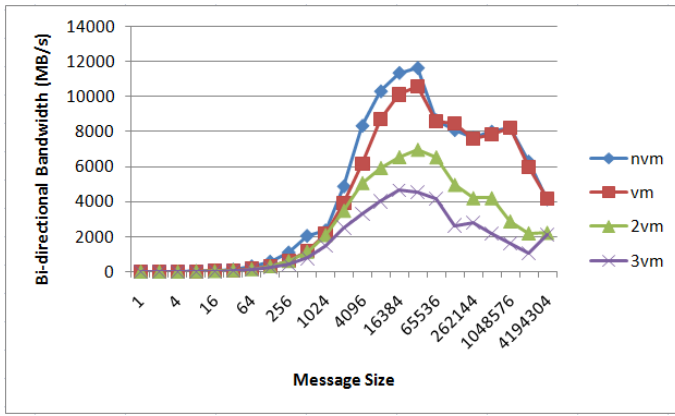


Fig. 12. OSU micro benchmarks : Bi-directional bandwidth for point-to-point communication

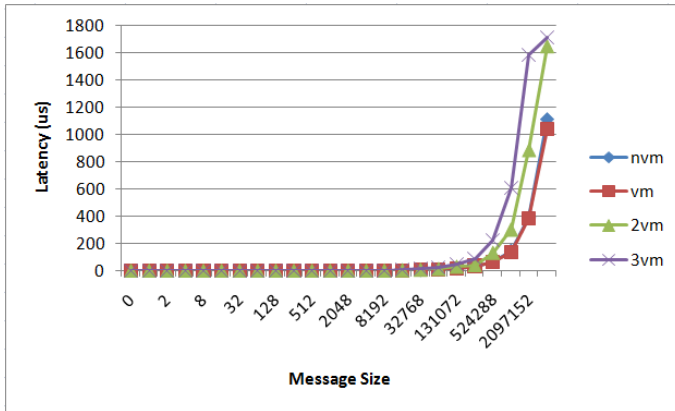


Fig. 13. OSU micro benchmarks : Latency for point-to-point communication

as the message size increases bandwidth for multiple VM case decreases.

Figure 12 shows results for bi-directional bandwidth for point-to-point communication. This benchmark works as the uni-directional bandwidth benchmark only that both sides issue first Irecv followed by Isends and Waitalls. Here also a peak performance is observed for message size of 64 KB. The VM performance keeps decreasing as the message size is increased.

Figure 13 shows the graph for latency during point-to-point communication. Latency denotes the time taken to transfer a message of a certain size from one MPI rank to another. Here the communication is point-to-point and thus two processes communicate with each other. From the graph it can be seen that the difference in latency for non-virtualized machine and a virtualized machine is very less. Also as the number of VMs are increased the latency increases. The amount of increase is not drastic but for the realtime applications this loss can be severe if the application is moved from a dedicated machine to a cloud setup.

## V. CONCLUSIONS AND FUTURE WORK

In this paper we studied different performance tradeoffs of HPC applications. We evaluated performance across different

physical resources and studied the effect of virtualization on each of them and found that HPC applications do not give fair performance on virtualized environment. We presented an architecture to overcome this performance degradation caused by the virtualization layer. The key goal of providing a dedicated and on-demand platform for HPC workload was achieved through the design of the architecture.

The current architecture has the functionality of starting, scaling and terminating the compute cluster. The future work will consist of improving the current architecture for better functionality and adding more features to it such as pausing/resuming and saving the state of compute cluster to make it a full fledged framework which acts as a cloud setup for HPC applications.

## REFERENCES

- [1] V. Mauch, M. Kunze, and M. Hillenbrand, "High performance cloud computing," *Future Generation Computer Systems*, vol. 29, no. 6, pp. 1408–1416, 2013.
- [2] A. Gupta and D. Milojicic, "Evaluation of hpc applications on cloud," in *Open Cirrus Summit (OCS), 2011 Sixth*. IEEE, 2011, pp. 22–26.
- [3] G. Wang and T. E. Ng, "The impact of virtualization on network performance of amazon ec2 data center," in *INFOCOM, 2010 Proceedings IEEE*. IEEE, 2010, pp. 1–9.
- [4] C. Evangelinos and C. Hill, "Cloud computing for parallel scientific hpc applications: Feasibility of running coupled atmosphere-ocean climate models on amazons ec2," *ratio*, vol. 2, no. 2.40, pp. 2–34, 2008.
- [5] A. Gupta, P. Faraboschi, F. Gioachin, L. V. Kale, R. Kaufmann, B.-S. Lee, V. March, D. Milojicic, and C. H. Suen, "Evaluating and improving the performance and scheduling of hpc applications in cloud," 2014.
- [6] K. R. Jackson, L. Ramakrishnan, K. Muriki, S. Canon, S. Cholia, J. Shalf, H. J. Wasserman, and N. J. Wright, "Performance analysis of high performance computing applications on the amazon web services cloud," in *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*. IEEE, 2010, pp. 159–168.
- [7] A. J. Younge, J. P. Walters, J. Suh, D.-I. D. Kang, Y. Park, S. P. Crago, and G. C. Fox, "Towards a high performance virtualized iaas deployment."
- [8] S. Srirama, O. Batrashev, and E. Vainikko, "SciCloud: scientific computing on the cloud," in *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*. IEEE Computer Society, 2010, pp. 579–580.
- [9] E. Keller, J. Szefer, J. Rexford, and R. B. Lee, "Nohype: virtualized cloud infrastructure without the virtualization," in *ACM SIGARCH Computer Architecture News*, vol. 38, no. 3. ACM, 2010, pp. 350–361.
- [10] S. K. Garg, C. S. Yeo, A. Anandasivam, and R. Buyya, "Environment-conscious scheduling of hpc applications on distributed cloud-oriented data centers," *Journal of Parallel and Distributed Computing*, vol. 71, no. 6, pp. 732–749, 2011.
- [11] K. Yelick, S. Coghlan, B. Draney, R. S. Canon *et al.*, "The magellan report on cloud computing for science," *US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR)*, 2011.
- [12] P. Mehrotra, J. Djomehri, S. Heistand, R. Hood, H. Jin, A. Lazanoff, S. Saini, and R. Biswas, "Performance evaluation of amazon ec2 for nasa hpc applications," in *Proceedings of the 3rd workshop on Scientific Computing Date*. ACM, 2012, pp. 41–50.
- [13] M. Musleh, V. Pai, J. P. Walters, A. Younge, and S. Crago, "Bridging the virtualization performance gap for hpc using sr-iov for infiniband," in *Cloud Computing (CLOUD), 2014 IEEE 7th International Conference on*. IEEE, 2014, pp. 627–635.
- [14] W. Felter, A. Ferreira, R. Rajamony, and J. Rubio, "An updated performance comparison of virtual machines and linux containers," *technology*, vol. 28, p. 32, 2014.