JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

# Cost-effective Cloud Services for HPC in the Cloud of SaaS,And IaaS, : Exploring the Cloud Computing Paradigm

**Jaya Shukla[1],TusharGiri[2]**

[1]M. Tech Scholar, Department of CSE, Institute of Technology&Management Lucknow UP India
[2]Assistant Professor, Department of CSE, GCRG. Group of Institutions, Lucknow, UP, India

**Abstract-**With the benefits of cloud computing, cloud services like Infrastructure as a Service (IaaS) and Hardware as a Service (HaaS) enable scientists and researchers to run their HPC applications in the cloud without the upfront investment associated with the traditional HPC infrastructure. Traditionally, in the scientific research domain, HPC refers to the use of supercomputers, grid environments, and/or clusters of computers to solve computation-intensive and/or data-intensive problems. The traditional HPC systems are expensive and sometimes require huge start-up investment, technical and administrative support, and job queuing. The computational performance and financial costs of operating HPC applications in the cloud when IaaS or HaaS is leased are examined in this research. We discover that, in comparison to IaaS, HaaS dramatically lowers the cost of running HPC applications in the cloud by 20% without appreciably affecting the application's performance. Additionally, we discovered that HaaS offers significantly better computational performance than IaaS.and HPC-An overview of several cloud computing models, including several kinds of cloud environments like Public Cloud, Private Cloud, is given in this study article. Community, Hybrid, and Cloud clouds. The models of Platform as a Service (PaaS), Infrastructure as a Service (IaaS), and Software as a Service (SaaS) are also covered in detail. There is a thorough discussion of the security risks related to cloud computing, including compromised passwords, data breaches, compromised interfaces and APIs, exploited system vulnerabilities, and account hijacking

**Keywords**: computational performance, cloud computing, computation-intensive applications, HaaS, and HPC

## 1. Introduction

Traditional High Performance Computing (HPC) in the scientific research sector refers to the employment of supercomputers, grid environments, and/or computer clusters to resolve issues requiring a lot of calculation. Weather forecasting, airplane accident simulations, computational fluid dynamics for aerodynamics research, and numerous other computation-intensive applications are some typical uses for HPC systems [14], [20]. These days, HPC systems also present new commercial prospects. For instance, real-time modeling using HPC computers helps financial firms make well-informed investment decisions. Top 500 lists the most potent HPC systems [1]. Research communities find it challenging because acquiring HPC systems requires a significant amount of capital. Most research communities would not have had access to HPC systems until recently. Due to recent developments in computing technology, computation-intensive applications are now run on HPC systems as well as traditional HPC systems. within the cloud. A novel computing paradigm for storing data and executing programs, particularly computation-intensive ones, is cloud computing [2], [3], [8]. One of the many advantages it offers is no up-front costs. Additionally, cloud computing lowers hardware, staff (such as administrators), and development time, which improves service and saves a lot of money. It is anticipated that more cloud-based HPC systems will host and operate computation-intensive applications [5], [3]. Additionally, the recent inclusion of the Amazon Elastic Compute Cloud (Amazon EC2) cluster in the TOP500 list [1] indicates that cloud-based HPC systems have a bright future. Researchers and scientists can lease cloud services like Hardware as a Service (HaaS) and Infrastructure as a Service (IaaS) for computation-intensive applications; these services are relinquished when not in use, avoiding the job queuing that is a common occurrence in traditional HPC systems. The pricing model is also attractive when compared to traditional HPC systems, which call for large capital investments, administrative concerns, and allocation policies. However, the cost of executing HPC applications on the cloud may be high if the cloud services arerecognized and the affordable cloud services selected. If executing HPC applications on the cloud has a higher financial cost than using a traditional HPC machine, then the advantages cloud-based computing-intensive applications might have had a flaw. The cost and computing performance of various cloud services are issues that the HPC research community are interested in. In this

research, we examine the computational efficiency and financial expenses associated with executing computationally demanding applications on cloud-based HPC systems when IaaS and HaaS are leased. We discover that when HaaS is leased, the cost of running a computation-intensive application is much less than it is with the IaaS model. We demonstrate that when a computation-intensive application is not a network, there is a notable improvement in the computational performance of the program on HaaS. The Message Passing Interface (MPI) implementation is used in our experimental setup [4]. We offer the findings of our tests, but we do not disclose the cloud providers' identities to prevent any confusionhead-to-head contrasts. We do, however, provide the pertinent cloud instance technical specifications.We provide an overview of cloud services in Section 2. for cloud-based HPC systems, and Section 3 details the experimental configuration. In Section 4, MPI applications and benchmarks are displayed. Sections 5 and 6 contain the experimental results and the cost analysis, respectively, while Section 7 covers related work. Finally, Section 8 presents a few findings.

## 2. Overview of Cloud Services for HPC Systems in the cloud

Cloud service providers including Salesforce.com, Amazon [5], Rackspace, Baremetalcloud [6], Microsoft Azure, and SoftLayer have emerged with the development of cloud computing infrastructures. [7] Providers of cloud services to users include Google and IBM. Software as a Service (SaaS), Platform as a Service (PaaS), Infrastructure as a Service (IaaS), Hardware as a Service (HaaS), Network as a Service (NaaS), and Storage as a Service (STaaS) are a few of the services that are provided. Cloud computing services can be divided into four main rival categories according to the capabilities offered by the cloud service provider [8], [9]. Hardware as a service, Infrastructure as a service, Platform as a service, and Application as a service. The architecture of the cloud computing services is depicted in Figure 1

| |
|---|
| Software-as-a-Service (SaaS) |
| Platform-as-a-Service (PaaS) (Developers implementing cloud applications) Infrastructure- |
| Infrastructure-as-a-Service (IaaS) [(Virtualization, Storage Network) as-a-Service] |
| Hardware as a Service (HaaS) |

Fig. 1: Cloud layered architecture [8]

The term "Software as a Service" (SaaS) refers to the highest level of abstraction in the cloud, offering users ready-to-use online applications that have already been deployed in the cloud. This layer is hidden from users and managed by SaaS providers; users simply use these applications without knowing where or how they are deployed. SaaS cloud applications can be accessed via the internet with any internet-ready device, such as laptops, smart phones, or iPads. This allows relatively dumb clients to perform complex tasks by shifting the real work, transparently to the user, into the cloud. A good example of SaaS is Google's popular email services Gmail. Platform as a Service (PaaS) gives cloud users access to a managed and fully configured computing platform that is ready to run custom applications that they have created. Every PaaS platform refers to software that has been written using a certain programming language or software framework (like Java EE) and is prepared to run builds that correspond to it. Without configuring servers and software stacks, considering scalability or clustering, and frequently even without knowing how many machines or CPUs their application will operate on, PaaS cloud users deploy and run their program. Similar to HaaS, Infrastructure as a Service (IaaS) rents out virtual machines rather than physical hardwareUsers using IaaS clouds must install, setup, and manage the operating system and software stack they put in their virtual machines (VMs), which they rent. IaaS customers frequently use a base installation that is a pre-installed and preset virtual machine image from their supplier. The hardware is not accessible to users via root access. Amazon is a great illustration of a cloud provider that provides IaaS for HPC applications [5]. For HPC applications, cluster computing instances are available through the Amazon Elastic computing Cloud (Amazon EC2). Hardware as a Service (HaaS): In this scenario, the cloud provider essentially rents out "bare-metal" hardware (such as servers, hosts, and storage); notable examples of cloud providers offering HaaS are Baremetalcloud [6], Soft-Layer [7]. Cloud users connect to HaaS over the Internet, install, and configure (e.g., virtual machines) the server they leased. Cloud users choose HaaS because it gives them complete control over the server, operating system, and software/hardware stack, as well as the number of virtual machines (VMs) they run on it. Research communities do lease HaaS for computation-intensive and/or data-intensive applications and configure HPC systems according to their needs [8], [9]. As a result, computation-intensive applications that were previously run on HPC systems can now be executed in the cloud.

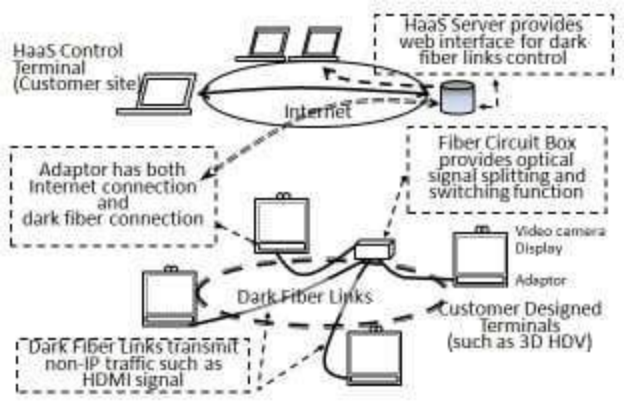


Fig. 2: An example of HaaS architecture with level ofinvolvement of key players.

## Experimental setup

To assess the financial cost and computational performance of executing computationally demanding applications on IaaS and HaaS, we have set up test environments. Our Two services that we have rented from two cloud service providers are part of the experimental setup; we will refer to them as Cloud-A and Cloud-B to avoid comparing the two cloud providers side by side. For HPC applications, Cloud-A provides IaaS in several types of cluster instances, such as cluster compute instances. HaaS, which may be set up to run HPC workloads, is provided by Cloud-B.

## 3.1 Cluster Compute Instances from Cloud-A(IaaS)

One of the main suppliers of cloud services is Cloud-A.They provide IaaS for HPC applications in various instances.A example of cluster compute instances with is displayed in Table 1.
pricing information from cloud providers for on-demand instances.

The cluster compute instances come in 32- and 64-bit platforms, supporting the widely used Operating System (OS) (Windows and Linux). Because the Ref-C virtual instance in Table 1 is frequently used for HPC applications, we selected it for our trials. The instances are fully virtualized with Xen. Ten gigabit Ethernet is used for I/O network connection within the cluster instance. in order to contrast the cost and computational efficiency of executing HPC applications when leasing IaaS and HaaS services. A cluster computing instance was rented by us. having sixteen processors in all. Table 2 displays the specifics of the rented cluster compute instance. On the node, we set up OpenMPI 1.6 [23]. An open-source implementation of the Message Passage Interface (MPI) is called OpenMPI.

Table 1: Virtual and HaaS Instances from Cloud-A andCloud-B

| Instance type | Memory | CPU | Disk | Cost of instance Cost of instancefor Linux | Cost of instance Window |
|---|---|---|---|---|---|
| Ref-A Virtual instance | 30 GB | 2x2.0 GHz (sixteen-core) | 500 GB | $1.600 per hour | $1.800 per hour |
| Ref-B Virtual instance | 244GB | 2 x Intel Xeon E5-2670 (eightcore) | 240 GB | $3.600 per hour | $3.8130 per hour |
| Ref-C Virtual Instance | 22GB | 2 x Intel Xeon X5570 (quadcore) | 1690 GB | $2.100 per hour | $2.600 per hour |
| Ref-D Virtual Instance | 23 GB | 2 x Intel Xeon X5570 (quadcore) | 1690GB | $1.300 per hour | $1.6100 per hour |
| Ref-E Virtual Instance | 96 GB DDR3 - 1333 | 2x2.13 GHz E5606 (eightcore) | 1000.0 GB, 7200RPM | $.990 per hour | $1.900 per hour |
| Ref-F Virtual Instance | 48 GB DDR3 - 1066 | 2x2.66 GHz X5650 (twelvecore) | 300GB, 10000RPM | $.7600 per hour | $0.930 per hour |
| Ref-G Virtual instance | 64 GB | 2x2.0 GHz E5-2650-OctoCore (sixteen-core) | 500 GB | $1.540 per hour | $1.590 per hour |
| Ref-H Virtual instance | 32 GB | 2x2.0 GHz (eightcore) | 250 GB | $1.250 per hour | $1.300 per hour |

## 3.2 HPC system on HaaS in the cloud

HaaS gives users complete control over the system and control environment for assessing system performance and other available experiments, as detailed in Section II. Users can use this to calculate how many virtual machines (VMs) should be installed for HPC applications. We have rented a 64GB RAM HaaS instance (Ref-G) from Cloud-B. Some of the HaaS providers' cloud services that are comparable to Cloud-A's cluster compute instances are displayed in Table 1. A summary of HaaS and the cost of the leased service are also provided in the table. One Gigabit Ethernet serves as the communication network between every HaaS unit. Table 2 summarizes the virtual machine (VM) we provisioned on the HaaS. We installed the open-source, industry-standard Xen hypervisor [11] on the host. Ubuntu 12.4 64-bit Linux OS runs on top of the Xen hypervisor. We imported our pre-configured para-virtualized guest OS (Ubuntu 12.4 64-bit) on the HaaS instance. The pre-configured para-virtualized guest OS reduces the time to setup the HPC system

on the HaaS instance. A para-virtualized OS uses a modified kernel and reduces the size of the image. The VM is configured with 16 processors, 60GB of memory, and 200GB of hard drive capacity. We installed OpenMPI on the node. This setup is virtually comparable to the cluster compute instances we leased from Cloud-A. The arrangement also permits us to have a favourable comparative environment for HaaS and IaaS in terms of cash cost and computational performance. The IaaS and HaaS infrastructures we employed are both displayed in Table 2.

**Table 2: Computational environment for IaaS and HaaS**

| Cloud&A,)VM)of)IaaS | Cloud&B,)VM)of)HaaS |
|---|---|
| RAM:24GB | RAM: 60GB% |
| Architecture: x86_64 | Architecture: x86_64 |
| CPU op-mode(s): 32;bit, 64;bit | CPU op-mode(s): 64;bit |
| CPU(s): 16 | CPU(s): 16 |
| Online CPU(s) list: 0;15 | Online CPU(s) list: 16 |
| Thread(s) percore: 2 | Thread(s) per core: 1 |
| Core(s) per socket: 4 | Core(s) per socket: 1 |
| UMA node(s): 1 | UMA node(s): 1 |
| VendorID: GenuineIntel | VendorID: GenuineIntel |
| CPUfamily: 6 | CPU family: 6 |
| Model: 26 | Model: 26 |
| Stepping: 5 | Stepping: 5 |
| CPU MHz: 2933.440 | CPU MHz: 2933.440 |
| Hypervisor vendor: Xen | Hypervisor vendor: Xen |
| Virtualization Type: full | Virtualization Type: full |

## MPI applications and benchmark

MPI benchmarks and applications The MPI applications operating on IaaS and HaaS services were analyzed and evaluated using a widely used HPC benchmark and an actual HPC application. The standard was the ClustalW-MPI was the application, and the benchmark was the High Performance Linpack (HPL) [12, 14]. Below, we describe them.

A popular benchmark for assessing the computational performance of HPC systems is HPL [12], for instance, top500 [1]. Depending on the magnitude of the problem, it calculates the floating execution rate of linear equations. Five distinct issue sizes—2,000, 4,000, 6,000, 8,000, and 10,000—were used to conduct the HPL benchmark on virtual machines (VMs) from both IaaS and HaaS cloud services. Every problem size was carried out twice, and the average of The five different problem sizes allow us to obtain different wall clock execution times of HPL. We recorded the wall clock execution time for each problem size. We then used the wall clock execution time to analyze the computational performance and dollar cost of the two platforms. Figure 3 displays the results obtained on computational-performance. ClustalW-MPI [14] is a parallel implementation of ClustalW [15], which is based on MPI. ClustalW is a tool that is widely used in bioinformatics for multiple alignments of nucleic acid and protein sequences. It uses three alignment steps: pairwise alignment, guide-tree generation, and progressive alignment. We ran a sample of "A full multiple sequence alignment,"

## 1. II. COMPARATIVE EXAMINATION OF PAAS, IAAS, AND SAAS DESIGNS

In the field of information technology, cloud computing is a revolutionary paradigm that is changing how people and businesses access, store, process, and control apps and data. Fundamentally, cloud computing is the provision of computer services via the Internet, which enables customers to take advantage of a shared pool of resources, including servers, storage, and apps, without requiring a significant amount of local equipment.Scalability is a fundamental characteristic of cloud computing, providing users with the ability to adjust the amount of resources they utilize. This demand-driven dynamic. This demand-driven dynamic. Allocating resources effectively and economically guarantees peak performance, enabling businesses to adjust to varyingworkloads.without interruption.The three primary forms of cloud computing services are usually Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). Whereas SaaS delivers software programs on a subscription basis, doing away with the requirement for local installations, PaaS gives a platform for developers to create, deploy, and manage applications, and IaaS offers virtualized computing resources via the internet. Businesses of all sizes are adopting cloud computing due to its cost-effectiveness, scalability, and agility in the ever changing digital landscape. Navigating the multitude of cloud service models, however, can be challenging. Three well-known Three unique value propositions are offered by the models: Platform as a Service (PaaS), Infrastructure as a Service (IaaS), and Software as a Service (SaaS). This leaves enterprises withthe crucial query: Which model best meets their particular requirements? This study compares and contrasts SaaS, IaaS, and PaaS, analyzing each service's key features, advantages, disadvantages, and best uses.We must be able to clearly demonstrate the basis of these three categories of cloud-based services before we can begin the comparison. Thus, this is merely the outline for these:

### 2. Software as a Service:

**A .SaaS** offers web-based apps that are ready to use and accessible on all devices. SaaS removes the requirement for on-premises software, including email platforms and CRM systems.
Software installs easily and is expandable, with automatic updates available.Under this arrangement, the user can access the software through the Internet, which is housed on the cloud. The user is relieved of the burden of

maintaining the program and the infrastructurethat support it. SaaS examples include Microsoft Office 365, Salesforce, and Google Workspace.

### B. IaaS:

**IaaS** provides virtualized resources, such as networking, storage, and servers, on demand. Consider it as leasing a virtual data center, where companies have total authority over their
infrastructure, providing the greatest degree of flexibility and bespoke configurations.Under this arrangement, the user leases servers, storage, and networking from a cloud provider. The infrastructure is entirely under the user's control, and they are free to put whatever software on it. AWS (Amazon Web Services), Azure from Microsoft, and Google Cloud Platform are a few examples ofIaaS.

### C. PaaS:

**PaaS** offers a ready-made cloud environment for development and deployment. Databases, development tools, and runtime environments fall under this category. In essence, PaaS fills the void left by IaaS and SaaS, freeing developers to concentrate on creating apps rather than managing infrastructure.Under this concept, an application developer and user hires a platform. The operating system, middleware, and infrastructure needed to operate the application are all included in the platform. It is only the application code that needs to concern the user. Microsoft Azure, Google App Engine, and Heroku are a few PaaS examples

## Results and Discussion

The fact that setting up the clusters is comparatively simpler with the Cloud-A cluster compute instance than it is with a HaaS cluster is one of its main draws. But someBecause HPC applications have different requirements, a certain amount of technical knowledge is needed to build up a cluster on Cloud-A that will execute HPC applications. We uploaded our pre-configured para-virtualized image to the cloud to shorten the setup time for an HPC system using HaaS instances. Similar virtual machine images are also available for download from various websites. This method, according to our estimates, can cut the setup time by up to 80%. We did not contrast how long it took to build up an HPC system in Cloud-A (IaaS) withFigure 3's computational performance result for the HPL test indicates that the benchmark's wall clock execution time on a supplied Compared to Cloud-A's IaaS, an instance on HaaS is shorter. We were able to accomplish this because the virtual instances running on HaaS have 60GB of memory. Since we can anticipate the memory requirements, we decided to give our virtual instance this much memory. This option isnot available for the IaaS instance (users cannot change the memory of the virtual instance chosen). We also have full control of the Hardware instance and virtual instances.

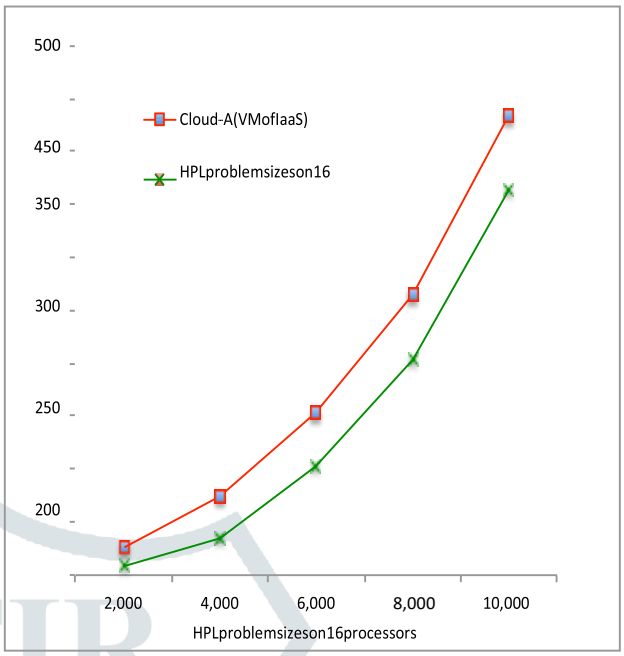


Fig. 3: Computational performance of High Performance

### Linpack on 1 node with 16 processors

Figure 3 illustrates how the bandwidth inequality on both providers is eliminated when the HPL is run on a single node with 16 processors. HaaS surpasses IaaS in virtual instance performance.
This is a result of our allocation of more memory to virtual machines on HaaS and our complete control over the applications operating on our HaaS instance. Other virtual machine instances that were hosted on the hardware during our IaaS lease might have had an impact on the application's performance. Apps operating in virtual machines are impacted by large resource allocations on infrastructure, as demonstrated in [19]. Although we might have benchmarks with identical bandwidths, 10 gigabit Ethernet is not available from the two main HaaS providers. network. Figure 4's results demonstrate that there is no discernible difference in the impact on applications operating on IaaS and virtual instances running on HaaS.
Figure 4 displays the findings of the ClustalW-MPI. Whereas the HaaS we rented uses a 1 gigabit Ethernet network, Cloud-A IaaS employs a 10 gigabit Ethernet network.

## 4. Cost Analysis

As of this writing, Cloud-A provides its cluster compute instance clients with a variety of pricing options; The most popular primary pricing strategy is referred to as "ondemand."

examples. While alternative pricing models could call for upfront payments and/or contracts, the on-demand instances price model permits consumers to pay on an hourly basis

without obligation. Pay-as-you-go pricing, which is comparable to Cloud-A's on-demand instance fees, is what Cloud-B provides its clients. As a result, we employ on-demand pricing instances.

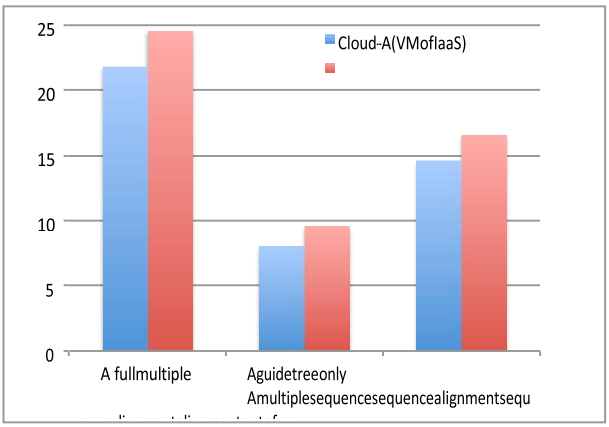


*Fig. 4: Performance of ClustalW-MPI application on 16Processors*

to evaluate how much it costs to operate computationally demanding apps on each cloud provider. There are fees in addition to the on-demand/pay-as-you-go instance rates. which are paid for certain cloud services—like IP addresses and network bandwidth—that we disregard in order to simplify things. To analyze the cost, we used the HPL benchmarking findings. As previously employed in a comparable cost analysis [18], we make the assumption that the hourly rates that both cloud providers supply are equivalent to one second. Additionally, we may conduct the analysis without having to pay for the hours that the experiment would have required. 5. We made use of the leasing service pricing listed in Tables 1 and 2. Figure 5 displays the cost analysis computation for both IaaS and HaaS. It seems more cost-effective to lease HaaS and configure the HPC equipment based on the computational performance and cost evaluations. HaaS cloud service consumers had total authority over the VMs and hardware that they supplied. Measuring application performance and other metrics is simple. Based on the results, it appears that leasing HaaS can lower the cost of running HPC applications by 20%.

## 6. Related work

A cutting-edge computing paradigm for data storage and application execution, particularly for computationally demanding workloads, is cloud computing. Cloud-based computing offers a host of advantages, one of which is the absence of upfront costs for HPC applications—a compelling contrast to conventional HPC systems. Numerous analyses of the viability of HPC systems in the cloud have demonstrated that it is anticipated that a greater number of computation-intensive HPC applications will be executed in cloud HPC than

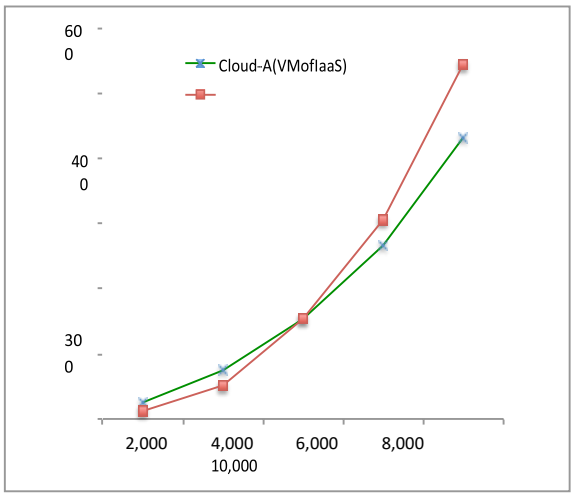


*Fig. 5: The cost analysis VM of IaaS and VM of HaaS*

conventional HPC architectures [16].

Additionally, the Amazon Elastic Compute Cloud (Amazon EC2) cluster was just included in the 2010 TOP500 list [1], indicating that HPC systems have a promising future on the cloud. Numerous previous studies have been conducted assessing HPC applications on cloud-based HPC systems, with a focus on Amazon EC2. The performance of traditional HPC systems and Amazon EC2 is the main emphasis of these studies [16], [17], [18], and [19]. Carlyle et al.'s study [17] focused on the economical HPC system. They demonstrate that running a community or traditional cluster is more affordable for organizations like Purdue University than renting HPC resources from Amazon EC2. This study unequivocally demonstrates that for particular HPC applications, the costs of Amazon on-demand cluster compute instances are not reasonable. organizations. Traditional HPC systems and the Amazon EC2 service (IaaS) are the main topics of their work. In their study titled "The Cost of Doing Science on the Cloud: The Montage examples," Deelman et al. [18] demonstrate that, with appropriate storage and computing resource provisioning, cloud service costs may be drastically decreased without materially affecting application performance. They did not, however, take into account alternative platforms like HaaS. We expanded on their research to show that HaaS can drastically lower the expense of executing computation-intensive applications on cloud-based HPC. Ekanayake and Fox [19] compared the performance of HPC programs with and without latency, based on varying criteria. But they didn't contrast the price of running computationally demanding applications on various services, including IaaS and HaaS. According to Yao et al. [21], the right cloud instance can help achieve the ideal cost-performance ratio. They did not, however, take into account the expense and computational efficiency of leasing IaaS and HaaS. To the best of our knowledge, what sets our work apart from past research is that we examine the financial costs and computational performance of operating computation-intensive applications in cloud-based HPC when IaaS and HaaS are rented. Through experimentation, we demonstrate that using HaaS can lower the cost of running computationally intensive

applications by up to 20% without compromising performance noticeably.

## 5. Conclusions and FutureWork

Using an HPC system on the cloud is a viable option because traditional HPC systems require a large capital investment and usually include job queuing. Substitute. Applications requiring a lot of compute can be performed on cluster instances by utilizing IaaS and HaaS provided by cloud computing. IaaS offers clusters that are nearly ready for use and require little installation or deployment work. Virtual machines can be supplied to run computationally demanding applications using HaaS. To find out how well and how much it costs to run computation-intensive applications on cloud services like IaaS and HaaS, we have carried out experimental studies. We demonstrated that leasing HaaS can lower the total cost of operating computation-intensive applications in the cloud by up to 20%.We demonstrated that running the applications on a rented HaaS has no appreciable effect on their performance.

## References

[1]Visit www.top500.org.

[2]The NIST definition of cloud computing draft, Mell, P., and Grance, T. NIST special publication, 800, (2011), pp. 145.

[3]"A view of cloud computing," Armbrust, M., Fox, A., et al., Communications of the ACM, 53(4), pp. 50-58, 2010.

[4]Message Passing Interface Forum, International Journal of Supercomputer Applications, 8(3/4):165–414, 1994. "MPI: A message-passing interface standard."

[5] http://aws.amazon.com/ec2/ Amazon (2013), [Online],

[6] "A taxonomy and survey of cloud computing systems," Rimal, B. P., Choi, E., and Lumb, I. Fifth International Joint Conference on NCM'09, INC, IMS, and IDC, 2009, pp. 44-51). 2009 IEEE.

[7] A proactive fault-tolerance approach to high performance computing (HPC) on the cloud, in The 2nd International Conference on Cloud Computing, Egwutuoha, I. P., Chen, S., Levy, D., Selic, B., and Calvo, R. & Green Computing, pp. 268-273, Xiangtan, Hunan, China, 2012.

[8] Hypervisor Xen, [Internet], available at http://www.xen.org/products/xenhyp.html

[9] "HPL Benchmark," [Online], http://www.netlib.org/benchmark/hpl/, Petitet, A., Whaley, C., Dongarra, J., and Cleary, A. (2008, Sept.).

[10] A survey of fault tolerance techniques and checkpoint/restart implementations for high-performance computing systems was published in The Journal of Supercomputing by I. P. Egwutuoha, D. Levy, B. Selic, and S. Chen. 10.1007/s11227-013-0884-0, February 2013.

[11] Li, Kuo-Bin. "ClustalW-MPI: Distributed and Parallel Computing for ClustalW analysis." Bioinformatics 19, no. 12 (2003), pp. 1585–1586.

[12] Julie D. Thompson, Toby J. Gibson, and Desmond G. Higgins.
"CLUSTAL W: enhancing progressive multiple sequence alignment sensitivity via weight matrix selection, position-specific gap penalties, and sequence weighting" Research on Nucleic Acids 22, no. 22 (1994): 4673–4680.

[13]The paper "Cloud Computing for parallel Scientific HPC Applications: Feasibility of running Coupled Atmosphere-Ocean Climate Models on Amazon's EC2" by C. Evangelinos and C. N. Hill was published in ratio 2, no. 2.40 (2008). Pages 234–236.

[14] Carlyle, A. G., Stephen, L. H., and Preston, M. S., "HPC at a reasonable cost: The cloud or the community?" 2010 IEEE Second International Conference on Cloud Computing Technology and Science (CloudCom), pp. 169–176. 2010 IEEE

[15]"The cost of doing science on the cloud: the montage example," by Deelman, E., Singh, G., Livny, M., Berriman, B., and Good, J. In November 2008, IEEE Press, Proceedings of the 2008 ACM/IEEE Conference on Supercomputing, p. 50.

[16]High performance parallel computing with clouds and cloud technologies," by J. Ekanayake and G. Fox, Cloud Computing (2010): 20–38.

[17] "An overview of high performance computing for the physical sciences," by G. C. Fox and P. D. Coddington, in Proceedings of the Mardi Gras Conference: High Performance Computing and Its Applications in the 1993; Physical Sciences.

[18] "A Performance Evaluation of Public Cloud Using TPC-C Benchmark," Yao, J., Ng, A., Chen, S., et al., The 1st International Workshop on Analytics Services on the Cloud (ASC 2012), in conjunction with ICSOC 2012: 1–7.