

Cost Analysis Comparing HPC Public Versus Private Cloud Computing

Patrick Dreher^(✉), Deepak Nair, Eric Sills, and Mladen Vouk

Department of Computer Science, P.O. Box 890, Raleigh, NC 27606, USA
{padreher,dnair,edsills,vouk}@ncsu.edu

Abstract. The past several years have seen a rapid increase in the number and type of public cloud computing hardware configurations and pricing options offered to customers. In addition public cloud providers have also expanded the number and type of storage options and established incremental price points for storage and network transmission of outbound data from the cloud facility. This has greatly complicated the analysis to determine the most economical option for moving general purpose applications to the cloud. This paper investigates whether this economic analysis for moving general purpose applications to the public cloud can be extended to more computationally intensive HPC type computations. Using an HPC baseline hardware configuration for comparison, the total cost of operations for several HPC private and public cloud providers are analyzed. The analysis shows under what operational conditions the public cloud option may be a more cost effective alternative for HPC type applications.

Keywords: High performance cloud computing · Economic analysis · Public cloud · Private cloud

1 Introduction

Over the past ten years cloud computing technology has capitalized on advances in computing hardware, storage and network technologies. These advances have helped produce new cloud architectures and software environments capable of supporting a variety of cloud computing “pay-as-you-go” service model options. Early economic analysis of the public cloud price points for general purpose applications showed potential savings to move certain types of applications from the private IT to public cloud providers. These early cost comparisons generated interest from industry, government and academia in applying these “pay-as-you-go” service model within their organizations. If such a business paradigm could be implemented, it would allow an organization to convert costly capital line items to less expensive operating costs on their balance sheet while simultaneously continuing to meet some of their users’ computational needs and requirements. Chen et al. [1] investigated the overall performance of private clouds for regular or small scale commercial applications but which are not necessarily

computationally intensive and Walker [2] has done detailed calculations of the total cost of a CPU hour.

In recent years public cloud providers greatly expanded the number and type of public cloud computing hardware configurations along with the base and incremental pricing options for computation, storage and network transmission of outbound data from their cloud computing facilities. This expanded set of choices and options has prompted several new studies and analyses of the public versus private cloud question, taking into account these expanded public cloud hardware platforms and pricing options for enterprise/business based cloud implementations. Many of these new studies have expanded on the original basis for cloud implementation of applications characterized by extremely low system utilization, highly dynamic user demand requirements and critical response times. With these new advances in cloud hardware architectures and software environments the application of cloud systems to areas of scientific and other high performance computing (HPC) type applications needs to be re-analyzed.

Several high performance computing studies compared the difference in performance for an HPC job run on a private HPC cluster and different public cloud provider platforms such as Amazon EC2 [3–5] or Microsoft Azure [6]. Other studies have suggested ways to improve the usage of cloud computing features for the benefit of HPC applications [7] and schedulers [8]. There has also been related work [9,10] specifically focused on how the Infrastructure as a Service (IaaS) cloud option performed when handling scientific workflow simulation environments. While the IaaS study was not specifically confined to analyzing computational performance for HPC clouds, the analysis of scientific workflow simulation environments for Infrastructure as a Services has a similar type of tradeoff criteria between the observed performance and incurred monetary cost for analyzing HPC cloud computing options.

This paper focuses on the cost analysis for HPC cloud computing implementations. A test case cost comparison with a given HPC hardware baseline configuration is established to determine under what conditions will a public cloud rather than a private cloud be more cost effective for computations, storage and network data transfers for HPC type applications. The paper is organized as follows. Section 2 describes the design of a common HPC baseline configuration for a cloud computing system that will serve as a common platform when comparing HPC options offered by various public and private cloud providers. Section 3 summarizes the various pricing models offered by this selected set of cloud providers. Section 4 provides a detailed analysis for the various computation pricing options among the selected private and public cloud providers. Section 5 reviews the detail analysis of the storage options and Sect. 6 summarizes the network transfer pricing options among these private and public cloud providers. Finally, Sect. 7 offers some comments, observations and a summary as to how these costs impact decisions as to when and where it may be advantageous to utilize public and private cloud computing platforms for HPC type applications.

2 The HPC Baseline Configuration

Cloud computing systems have successfully been adopted to service some types of applications characterized by extremely low system utilization, highly dynamic user demand requirements and critical response times. However, these characteristics do not always describe the properties of scientific and other high performance computing (HPC) applications. This paper is focused on examining cost comparisons among several public and private cloud computing providers using a typical HPC type baseline hardware configuration and under what conditions do the cloud computing costs favor a public or private cloud option.

As a starting point, it is important to define this baseline hardware configuration so that the analysis can be standardized when comparing the performance levels across different cloud providers. The baseline configuration must be sufficiently robust to handle the more computationally intense high performance computing demands and requirements and must be capable of effectively delivering similar expected service levels and productivity when compared against these same applications run on a supercomputer.

The detailed technical aspects of utilizing cloud computing for HPC applications has been covered elsewhere [11, 12] and will not be the primary focus here. It is assumed that both the public and private cloud vendors have fully operational cloud systems that can deliver access to HPC level resources and services throughout the entire time interval selected for comparison.

Table 1 lists the types of performance characteristics and capabilities that one would expect from a cloud computing provider offering hardware platforms suitable for HPC and related computationally intensive applications. It was decided to use the parameters listed in this table as the baseline standard when comparing hardware characteristics among the systems supported by the various cloud vendors. These characteristics reflect the specifications of the Intel Xeon E5 2650 v3 processor. This baseline hardware system consists of two CPU sockets on the motherboard each connected to a chip with 10 cores. The baseline total physical

Table 1. Processor baseline specifications.

Number of cores	10
Number of threads	20
Operating frequency	2.3 GHz
Max turbo frequency	3.0 GHz
L2 cache	10 × 256 KB
L3 cache	25 MB
Manufacturing tech	22 nm
64-bit support	Yes
Hyper-threading support	Yes
Virtualization technology support	Yes

core count for these systems will be 20 with an option to enable hyperthreading so that each physical processor can virtualize two cores. Finally, these needs to be some assumption for local storage and network connectivity. For this baseline comparison it will be assumed that there is a total RAM of 160 GB available and a 10 GB network connection to provide support for I/O.

3 Cloud Provider Pricing Models

The baseline configuration sets a standard for comparison across all vendors and cloud providers studied in this report. It should be noted that some of the public cloud vendors offer other choices and selections for accessing HPC cloud instances that do not align with their baseline configuration proposed for these comparisons. However, these other choices proved to be more expensive than the option of scaling the per unit hardware and then analyzing and comparing the cost effectiveness for each vendor's cloud instance.

3.1 Public Cloud Options

There are numerous cloud vendors today offering various service options and hardware configurations. The public cloud vendors selected for comparison in this analysis include Amazon [13], Google [14], and Microsoft Azure [15].

Amazon AWS. The Amazon Elastic Compute Cloud (Amazon EC2) is a commercial cloud computing vendor offering a wide selection of different types of options and configurations. These instances are spread across various categories including General Purpose, Memory Optimized, Storage Optimized, and Clusters with options for GPUs with previous or current generation processors.

For the purpose of analyzing the different vendor choices, the M4.10xLarge configuration was selected because it best matches the common baseline configuration. The instance has 40 Virtual Cores with 160 GB of RAM with 10 GB interconnect¹. It has a support for Elastic Block Storage (EBS) and so the I/O for the EBS will be dedicated and will not interfere with the connection to the Internet.

Among the various public cloud vendors Amazon offers some of the most complex pricing structures for cloud computing within the commercial cloud computing marketplace. Pricing structure offerings include:

- On Demand - Reservations initiated by the hour whenever required and can be terminated when the work is done. These reservations will be charged based on the total number of hours used
- Reserved - Reservations for up to a year with monthly payments required for the entire duration of the reservation.

¹ The 'M' in M4.10xLarge denotes the General Purpose intention of the instance and '4' indicates that the processors are current generation processors. Each virtual CPU is a hyperthread of an Intel Xeon core.

- Fully Prepaid Reservation - Payment at the outset for the entire duration independent of the level of monthly usage for that prepaid cloud instance. If the use of the reservation is terminated within the prepaid duration, the unused resources can be sold by the user on Amazon marketplace to other bidders.
 - One Year - One year reservation with full amount paid at the outset
 - Three Years - Three year reservation with full amount paid at the outset
- Partially Pre-paid Reservation - Half of the full amount paid at the outset with hourly rate based on a discounted amount and remaining amount due paid in monthly installments.
 - One Year - One year reservation with partial amount paid at the outset
 - Three Year - Three year reservation with partial amount paid at the outset
- Spot Pricing - Bid price submitted by users on an open market basis for unallocated resources with no guarantee of the computation time user will be able to secure through a spot reservation. The user with the highest bid will be awarded the reservation on the unused cloud resource.²

The On-Demand option provides the most flexibility for scheduling access to cloud resources but that convenience and flexibility also makes it the most expensive option. At the present time the M4.10xLarge has a default quota of 5 instances when booked On-Demand and 20 instances when Reserved. The quota can be increased by contacting the customer service department. Additional servers in the region may be obtained depending on availability at the time of the request.

Users who know at the outset that they will need full utilization of cloud computing resources over the course of a year can save money by electing to purchase a fully pre-paid one year reservation versus the on-demand option. If after purchasing the fully pre-paid one year reservation option, the user determines the entire one year block of reserved computation time will not be needed, the excess time can be sold in the Amazon public cloud marketplace. The seller can set a one-time fee that for the instance and can also pro rate the price depending on the remaining time available if the instance goes unsold for some time. Conversely, if the user determines that more computing capacity will be needed in that one year time window, more computing capacity can be added dynamically to a reservation by either buying from Amazon or from other users who are selling their extra computation time in the marketplace.

The 1 year and 3 year reserved prepaid options show cumulative cost jumps at the appropriate yearly anniversary dates while the 1 year and 3 year partial prepaid options show smaller incremental jumps at the appropriate yearly anniversary dates and upward sloping cumulative costs over time. These various pricing plans provide the users with options to manage their operational costs

² The spot pricing is the cheapest way one can secure access to cloud resources. However if another user initiates a reservation instance in the midst of the spot reservation, the spot reservation gets preempted for the reserved instance and the usage will not be charged to the spot reservation user.

and cash flows but also illustrate the growing complexity in properly choosing the most cost effective pricing option for their cloud computing services requirements.

Amazon also supports an additional option for High Performance Computing jobs by providing compute Intensive Virtual Machines guaranteeing low latency in intra-network communication among virtual machines, but are more expensive and have less RAM per core than the baseline configuration³.

Google Cloud. Google Cloud Engine provides users with services for creating and running virtual machines to utilize the compute power of Google Infrastructure. The types of instances they offer are standard, high memory and high CPU instances. These options have various sizes ranging from 1 virtual core with 1 GB RAM to 32 virtual core machines with 208 GB of RAM.

Google also allows users to create custom machines. A custom machine is a configuration allows both the number of virtual CPUs and the total RAM for the instance to be specified by the user. The cost for the instance will be the sum of the costs for the individual CPUs and RAM based on the rate per CPU and rate per GB of RAM specified by Google. The user can specify 1 vCPU or any even number of vCPUs from 2 to 32 with a maximum of 6.5 GB of RAM per vCPU.

For the purpose of the comparison in this paper an instance which conformed with the baseline configuration was required. There is no single instance allowed by Google which provided the required configuration. Hence two custom instances were used with 20 vCPUs and 80 GB of RAM for each instance. The custom instance constructed consisted of 40 vCPUs⁴ and 160 GB of RAM. The custom instance option can help in reducing the total cost of the instances considering that if a smaller instance is required which is not provided by Google pre configured, the required instance can be created using the custom instance.

Google prices the level of usage for the instance to the minute for which access was provided, regardless of whether or not the resource was actually used during that reservation time. Google does not provide HPC centric cluster machines, and their compute intensive machines are very low on RAM with only 0.98 GB of RAM per virtual core.

Microsoft Azure. Microsoft, along with a host of cloud services, allows Windows and Linux Virtual Machines to be created and managed on Microsoft Azure Cloud. Azure has many types of instances like General Purpose, Compute Optimized, Memory Optimized and Intranet work Optimized with Infiniband support. For the purpose of comparison to the baseline configuration, the instance

³ For instance the most powerful Compute Intensive instance provided by Amazon is C4.8xLarge which has 32 virtual cores and 60 GB of RAM.

⁴ A virtual CPU is equivalent to a single hyperthread on a 2.6 GHz Intel Xeon E5 (Sandy Bridge), 2.5 GHz Intel Xeon E5 v2(Ivy Bridge), or 2.3 GHz Intel Xeon E5 v3 (Haswell) depending on the processor which makes up the instance.

of Compute and Memory Optimized D15 v2 was chosen because it provided 20 cores and 140 GB of RAM. Thus two instances of the same configuration will have 40 cores and 280 GB of RAM with the latest processor. The instance also provides a 1000 GB of standard storage attached to the node. In Azure each virtual core is equivalent to one hyperthread in the physical core of the machine. Microsoft rounds the amount of time the instance is used to minutes.

Microsoft supports multiple payment options for Azure services (Pay-as-you-go, Prepaid, Microsoft Resellers, and Enterprise Agreements). The pay-as-you-go option allows the user to create on demand reservations with costs calculated on a per minute usage basis. The prepaid reservation is for 12 months and the full amount must be paid regardless of usage pattern during that time period. Microsoft Resellers provide Open License Keys which can be bought and used to access the Azure infrastructure for reserving an instance. Enterprise Agreements involve users making commitments for large amounts of compute minutes in return of a discounted rate.

Microsoft also supports compute intensive virtual machines which have infiniband support and ones without infiniband support which are specific for HPC [16]. They also have closer proximity and hence are more tightly coupled than regular virtual machines which helps HPC applications. However they are also more expensive than the HPC baseline configuration outlined in this paper.

3.2 Private Cloud Option

The hardware and operational pricing for a hypothetical private cloud computing configuration were based on the costs for a proven and tested production level private cloud architecture called the Virtual Computing Laboratory (VCL) [17–20]. For the purpose of building a baseline configuration for fair comparison with the respective public cloud options, an Intel Xeon E5 2650 v3 blade was chosen. This motherboard has two CPU sockets each connected to a chip with 10 cores. The chips have hyperthreading enabled and so the each physical core can virtualize two cores. This provides a total of 40 virtual cores, similar to the public cloud hardware options. The blade comes with a RAM of 128 GB and by adding a 32 GB RAM stick, the total RAM becomes 160 GB and is equivalent to the public cloud hardware configuration. There is also a fast intra-network connection between the cores similar to the public cloud configuration.

4 Cost Analysis of HPC Cloud Computation Options

When reserving an HPC instance there are two major types of reservations from which to choose. All public and private cloud providers offer some form of the guaranteed reservation option. Some cloud providers also offer options to competitively bid for any unused cloud resources in an open market system. Both of these choices offer both advantages and drawbacks in terms of the users' needs, requirements and time constraints for accessing HPC cloud computation resources.

4.1 HPC Cloud Guaranteed Reservation Options

The public cloud providers examined here are Amazon, Google and Azure. The Virtual Computing Laboratory is the private cloud option. For the guaranteed reservation option the user's job will run to completion, although the time required to complete may vary depending on the prioritization level, with higher prioritization options costing more money. With each provider, a cloud option was selected that best matched the HPC baseline cloud configuration discussed in Sect. 2. In order to make a quantitative comparison from among the many choices offered by HPC public cloud providers that can satisfy our example HPC baseline configuration requirements, it is essential to gather and enumerate in the next subsections the detailed pricing structures and options offered by each HPC public cloud provider.

Amazon AWS. For the compute instance which is selected (M4.10xLarge), the configuration matches the Baseline Configuration and so only one of the instances needs to be reserved.

- The hourly rate for the On-Demand reservation of this instance is \$2.394.
- For the one year reservation option, the effective discounted hourly rate is \$1.645.

If the partial amount is prepaid at the outset, the total effective hourly rate becomes \$1.406. The prepaid amount of \$6158.28 is initially costed and the hourly rate of \$0.703 is charged for all reservations made throughout the year.

- For the Fully paid reservation, the full amount of \$12071.28 is costed immediately at the time of the first reservation.
- The 3 year reservations.

For a partial prepayment plan, \$12483 is initially costed and the remaining amount is paid in monthly payments of hourly rate \$0.475. The effective hourly rate becomes \$0.95.

The effective hourly rate for the 3 year full prepaid reservation is \$0.893 with a total amount of \$23468.04 is costed immediately at the time of the first reservation.

Figure 1 shows a graph of cost versus time for these various Amazon pricing plans for cloud computing instances. This figure graphically illustrates the complexities in picking the correct pricing option for the users' time dependent computational requirements.

Google Cloud. For the custom instance that was chosen, multiplying the quantity of virtual cores and RAMs with their respective rates provided by Google and adding them together will give the hourly rate of the custom instance. Google provides built-in incentives in their pricing model. The more the amount of usage in a month, the more the discount is applied to the rate. For our example the full rate for the month with minimum utilization was calculated to be \$2.1456/h,




Fig. 1. Cumulative costs versus time for various Amazon cloud computing pricing plans.

Table 2. Google rates for vCPU and GB memory.

Item	\$/h full usage	Quantity	\$/h full usage
vCPU	0.02444	40	0.9776
Memory (GB)	0.00328	160	0.5248

while for sustained utilization of 100%, the effective monthly rate went down to \$1.5024/h (Table 2).

The itemized list below summarizes the pricing structure.

- In a month the usage of 25% of total time is at the full rate of the instance which is \$2.1456/h.
- The second 25% is at a discount of 20% in the full rate of the instance which turns out to be \$1.7164/h.
- The third 25% is at a rate of \$1.2874 which is a discount of 40% of the full rate of the instance.
- And the last 25% at 60% discount of the full rate of the instance which is \$0.8582/h.
- If an instance is used for the entirety of a month, an effective discount hourly rate would be \$1.5012/h which gives an effective discount of 30% on the actual rate by the end of the month.

Microsoft Azure. Microsoft’s Azure provides an instance D15 v2 which has the closest resemblance to the baseline configuration. The D series of virtual machines have an additional feature of Solid State Drives attached to the machine directly which gives high speed Disk IO. The storage of the SSD is limited. The v2 series is the advanced version of the D series with more powerful, latest generation Intel processors coupled with Intel Turbo Boost Technology which can increase the speed of the processors. The hourly rate of one instance

is \$1.853. Since two instances of D15 v2 are considered for matching the baseline configuration, the hourly rate becomes twice the rate of one instance which is \$3.706. This instance comes with 140 GB of RAM and 1000 GB of Solid State Drive support.

The high rates charged by Azure reflect the addition of Solid State Drives and isolation to dedicated hardware for a customer. Both of these additions help increase the speed of the data throughput because of data proximity and high IO and enhance the processing of HPC and compute intensive jobs.

Private Cloud Option. Using historical costs in the private VCL cloud system the following hardware was selected and costs assigned to the baseline configuration defined in this paper. An Intel Xeon E5 2650 v3 blade cost (\$5,500) has two CPU sockets on the motherboard each connected to a chip with 10 cores which brings the total physical core count to 20. The chips are hyperthreading enabled so the each physical core can virtualize two cores. Thus the virtual core count is 40 cores. The blade comes with a RAM of 128 GB. Adding the cost of a 32 GB RAM stick (\$200) makes the total RAM to be 160 GB which confirms to the baseline configuration. The blades will require a rack cost (\$2,000) for its installation. Considering that the blade will be added to an existing and working data center, the entire cost of the rack need not be borne. The total cost of the rack will be divided by the number of blades the rack can house. In this instance the rack held 84 blades, giving a proportionate cost of \$23.81 for each blade.

The blade, RAM and the rack costs will be capital costs which will be paid one time and will be amortized over three years. The rack space rent (\$4,380) will be a recurring fee, a fraction proportionate to the one blade (\$52.14) needs to be borne by the user. The blades will also require power supply which also would be divided by the total number of blades housed by the data center. Each rack will require power circuits to deliver the power. That needs to be bought or paid for in proportionate manner. Each rack could use two 60 A circuit (rent \$3,480 each) or four 30 A circuits (rent \$1,980 each). In case of the example in consideration, four 30 A circuits (total rent \$7,920) are being used which

Table 3. Private cloud baseline costs.

Equipment	Price	Price per blade
Capital costs Intel Xeon E5 2650 v3 blade	\$5,500	-
32 GB RAM	\$200	-
Rack	\$2,000	\$23.81
Switch	\$20,000	\$238.1
SFP	\$200	\$200
Operational costs rack space rent	\$4,380	\$52.14
Four 30 A circuits	\$7,920	\$94.29
Developer	\$100,000	\$100

ultimately cost more than two 60 A circuit (total rent \$6,960). So for one blade's power usage a proportionate fee (rent - \$94.29) needs to be paid by the user. The costs to outfit a private cloud instance are summarized in Table 3.

4.2 Comparison of Cloud Computation Reserved Options

Figure 2 shows a graph of the accumulated cost versus time for on-demand cloud computing access options offered by Amazon, Google, Azure and a private cloud configuration using the HPC baseline hardware configuration referenced in this paper. This sample figure illustrates that for short periods of time the cost for each public cloud option is less than the cost for the private cloud option. However, if there is demand for HPC type cloud computations for extended periods of time, then at some point the cumulative cost for each public cloud intersects and exceeds the cumulative costs for the private cloud. These calculations show that private cloud costs for sustained cloud computing usage are more cost effective over the longer term in a private cloud implementation than in the public cloud sector. However, the cost of operations is only one component in determining the economic effectiveness when comparing public versus private clouds for HPC applications. This will be discussed in more detail in Sect. 7 summarizing HPC cloud computing planning based on cost analysis.




Fig. 2. Azure, Amazon, Google, and private cloud on-demand cumulative costs versus time.

4.3 HPC Cloud Spot Market Options

In 2009 Amazon was the first cloud vendor to introduce the concept of a “spot market” into the mix of options for purchasing cloud services [21]. The basic idea is to allow users to bid on unused CPU cycles under the constraint that their cloud instance would run as long as their bid price exceeded the current spot price. Users can place a bid for the highest amount they wish to pay for usage of their cloud facility instances. Whenever the Amazon current price for

a Spot Instance (SI) is equal to or less than the user's bid, access to those cloud computing facilities is provided to the user. As long as the user's bid price is higher than the instance price, the user gets access to the instance at the market rate of that instance. However, if the spot price increases and/or other users bid prices that are higher than the user's bid price, the instance that was allocated to the user gets preempted and the user does not pay for their partial usage of the instance.

This type of purchasing option has attracted attention within the cloud computing community. Several projects have experimented with this spot market option working from the assumption that the spot market operates on the basis that spot prices are set through a uniform price, sealed-bid, market-driven auction [22,23]. The price for the instance increases or decreases in time depending on the supply and demand for each type of instance.

The argument put forward is that by analyzing the spot market for the historical trends, an informed bid can be placed that has a high probability of giving the user the maximum amount of time for that spot instance without having the instance preempted by higher bids. Those interested in using the spot market for general and commercial computing have studied these historical trends. Mazzucco and Dumas [24] analyzed the spot market and made observations on how to achieve maximum availability on the spot instances and Mattess, Vecchiola and Buyya [25] have published an analysis of the economic benefits of the practice of leasing these spot instances. Although these techniques have improved the success rate for accessing the spot market, there are occasions when the spot price goes even above the reservation price for that instance and preemption of the instance is unavoidable. For these type of situations other techniques such as checkpointing and migration have been suggested as alternate strategies [26].

However, not everyone in the cloud community subscribes to this view of how the cloud spot market operates. There have been studies that suggest that the algorithm for the fluctuation of the spot price of each instance type in this market is entirely controlled by Amazon and that the the spot price is set according to a constantly changing reserve price that discounts actual client bids [27]. In contract Bonacquisto et al. [28] argued that the Amazon spot market was more aligned with the method of resource maximization through allocations and instead proposed a model for a procurement auction market as an alternative to maximize the utilization rate for a provider's data center resources.

Independent of these contrasting ideas, this work has extended the concept of the spot market to HPC instances and has investigated whether spot market strategies and methods can benefit HPC applications. A first step in determining whether there are other strategies that may help maximize the economic benefits for accessing these cloud computing resources is to gain access to the SI bid prices for each region and availability zone. The M4.10xLarge General purpose baseline configuration was used as the basis for this spot market study. Amazon currently provides historical spot price data for this configuration in both the U.S. East and West regions on a three month rolling basis. To obtain this historical data a python code was written using the Boto Python Library [29,30] to connect with

the AWS API [31] and extract the M4.10xLarge General purpose spot pricing data.

The hourly spot prices for the M4.10xLarge General purpose cloud instance from 7th May, 2016 to 18th July, 2016 were downloaded for the US East Region, availability zones 1B and 1C, and for the US West Region, availability zones 1B, 1C, 2A and 2B. The spot price data within each availability zone was verified to be free of misprints or corrupt data. Autocorrelation calculations in each availability zone determined that the daily bid prices were uncorrelated. Using the set of hourly prices spanning a 24 h period, the full correlation matrix for all of the availability zones was calculated for each day between 7th May, 2016 to 18th July, 2016. It was observed that the correlations among availability zones seemed to vary over time.

Among all of the different combinations of regions and availability zones there were certain pairs of availability zones that showed either strongly correlated or strongly anti-correlated correlation matrix values as a function of time during that period. As shown in Fig. 3 it was observed that the spot pricing correlation matrix values for the W1B/W1C and the W2A/W2B availability zones were highly correlated for large portions of time between the May 7th through July 18th time frame. In contrast, Fig. 4 shows periods of strongly anti-correlated spot pricing for the W1B/W2B, W1C/W2A and W1C/W2B availability zones.

These anti-correlated spot pricing patterns between different pairwise combinations of availability zones were unexpected. It should be noted that these anti-correlations do not occur at regular or periodic intervals and so they cannot be anticipated in advance. As a result, in order to potentially take advantage of these favorable anti-correlations in spot pricing positions in the marketplace, the spot price market must be regularly and closely monitored. Regardless of whether or not the spot market operates as a true auction market or if the spot price is controlled by Amazon based on a changing reserve price approach, if these hourly changes and daily spot pricing correlations are carefully monitored, it may be possible to minimize spot pricing costs by migrating applications needing HPC level baseline hardware platforms between different availability zones




Fig. 3. Correlation versus time showing positive correlations for specific pairwise Amazon availability zones.




Fig. 4. Correlation versus time showing intervals of negative correlation for specific pairwise Amazon availability zones.

that are highly anti-correlated during a given time frame. The success of such a program will highly depend on the type of HPC calculations being performed. If the application requires that large quantities of data be pre-staged in a particular region before the application is launched, then it is not practical to attempt to switch availability zones between different regions. However, if the HPC type application is compute bound with little to no need for access to large data sets for the computations, it may be practical to consider launching spot price bids between zones that are highly anti-correlated at a particular time frame in order to get the best probability for uninterrupted running of applications using the spot market option.

5 Cost Analysis of the Cloud Storage Options

The cloud computing storage option can be attached to the compute node itself or located at a remote physical location. The storage options range from slow and cheap archival storage up through options for more expensive solid state drives.

5.1 Amazon AWS

The data associated with EC2 compute nodes must have capabilities to be stored and accessed as needed. The Amazon instances come with attached storage per instance or with Elastic Block Storage (EBS) support. The I/O from the EBS to EC2 has dedicated throughput between 500 to 4000 Mbps. This dedicated throughput minimizes contention between Amazon EBS I/O and other Network I/O from the EC2. The listing below shows the various types of storage including SSDs and HDDs which are provisioned at different rates pro-rated to the hour of its provisioning.

- General purpose SSD costs \$0.10 per GB per month
- Standard HDD costs \$0.045 per GB per month.
- Cold HDD costs \$0.025 per GB per month
- EBS Snapshots storage to S3 costs \$0.05 per GB per month.

5.2 Google Cloud

Storage in Google comes in four main categories- Persistent Disks, Local SSDs, Cloud Storage Buckets and RAM disks.

- The persistent disk storage at Google are of three types Standard Storage (HDD), SSD and Snapshot Storage.
- The Standard Storage of HDD costs \$0.04 per GB per month.
- The SSD costs \$0.17 per GB per month
- The Snapshot Storage costs \$0.026 per GB per month because it is the slowest of them all.
- The maximum storage per instance is 64 TB with data replication and encryption.

Local SSD are physically attached to the machines that are used and hence are the fastest to access. The cost is \$0.218 per GB per month and has a limit of 3TB. The data is not replicated but it is encrypted. The Cloud Storage Buckets are the slowest and cost anywhere from \$0.01 to \$0.026 per GB per month and provide replication and encryption. The RAM Disks are the extensions for the RAM for in memory storage of files. They are the most expensive types of storage costing between \$3.37 to \$3.71 per GB per month with a limit of 208 GB per instance.

5.3 Microsoft Azure

Azure provides standard disks for storage in multiple forms of data redundancy that include Local Redundancy Storage (LRS), Zone Redundancy Storage (ZRS), Geo Redundancy Storage (GRS) and Read access Geo Redundancy Storage (RA GRS). Azure storage option have at least 3 replications of data. LRS stores all the 3 replicas of the data in the same region as the storage account location and is the least expensive. Zone Redundancy Storage replicates the data within 2 or 3 facilities within the same region or multiple regions and it is available only for block data. Geo Redundancy replicates data 5 times into facilities which are very far away from each other with a total of 6 copies of the data. GRS costs more than LRS because of maximum redundancy. The characteristics for each data redundancy are listed below and summarized in Table 4 below.

5.4 Private Cloud

Private cloud storage options cover wide ranges of both price and performance. A capacity storage system from a top-tier vendor is used as the storage baseline.

Table 4. Microsoft Azure LRS and GRS rates per GB per month.

Storage	LRS	GRS
1 TB	\$0.05	\$0.095
50 TB	\$0.05	\$0.08
500 TB	\$0.05	\$0.07
1000 TB	\$0.05	\$0.065
5000 TB	\$0.045	\$0.06

This type of storage system is typically used for HPC storage where there is requirement to maximize storage capacity. The baseline storage system has 240 6 TB near-line SAS disk drives with redundant controllers supporting various attachment methods including iSCSI, fibre channel, and NFS. Raw capacity is about 1.4 PB and usable capacity is about 1.1 PB. The list price for the storage system is \$1,350,000 including three years of maintenance. Discounts vary by organization and circumstance of the purchase but generally would be expected to be in range of 15–60%. Annualized capital expense over the three year warranty period is between \$0.030 per GB/month and \$0.014 per GB/month. Operating expense for the baseline system is about \$8.60/TB/year or about \$0.001 per GB/month. Overall baseline storage expense is between \$0.015 and \$0.031 per GB/month.

5.5 Comparison of the Cloud Storage Options

The types of storage options vary by the provider and each vendor has different pricing algorithms. The data show that Google offers the most inexpensive regular storage options such as hard drives, solid state drives and snapshot storage drives. Although Microsoft provides multi region replication for data recovery as one of the storage options, it may not be as useful for HPC jobs.

6 Cost Analysis of Cloud Network Transfer

Network transfer of data can be a potentially expensive when using public clouds. Cloud providers do not charge for network ingress of data but data extraction may incur substantial costs from public cloud providers.

6.1 Amazon AWS

Amazon provides free network transfer into its EC2 compute nodes or storage nodes. It also provides free data transfer between EC2 nodes in the same region. However Amazon charges for data transfer into another AWS region and for removal of data from their cloud system into the internet. The incremental costs for data transfer are summarized below.

- The first 1 GB of data transfer from the compute node to the internet is free each month.
- Above that, the first 10 TB will cost \$0.9 per GB.
- The next 40 TB of data transferred out will cost \$0.085 per GB.
- The next 100 TB of data transfer in a month will cost \$0.07 per GB.
- Data transfer between 150 TB and 500 TB will cost \$0.05 per GB per month.
- If there is more than 500 TB of data transferred per month from the EC2 compute node to the Internet, the customer service should be called to get a better deal on the data.

6.2 Google Cloud

The data transfer into the compute engine is free. The data transfer from Google to external Internet locations is charged in tiers listed below.

- The data transfer up to 1 TB from the compute engine to the internet has a cost of \$0.12 per GB.
- The next 9 TB of data transfer has a cost of \$0.11 per GB.
- All the data transfer more 10 TB has a cost of \$0.08 per GB.

6.3 Microsoft Azure

Azure also does not charge for inbound data transfer to the compute nodes. Azure does have an outbound data transfer rate after the initial transfer of 5 GB of data from the compute instance to the Internet with the incremental costs for data transfer summarized below.

- For the data transfer of 5 GB to 10 TB of data, the rate is \$0.087 per GB per month.
- For the transfer of next 40 TB of data the rate is \$0.083 per GB per month.
- For the next 100 TB the rate is \$0.07 per GB per month.
- For the next 350 TB of data from 150 to 500 TB, the rate is \$0.05 per GB per month.

6.4 Private Cloud

For the network transfer of data among the blades and to the Internet, a network switch (\$20,000) must be included. Assuming 84 blades per rack, the apportioned cost per blade is approximately \$238. In addition, each port of the switch will require an SFP (cost - \$200) for fast communicating with the blade over fiber optic cables. The SFP and the switch will be capital costs and incur a one time fee. The private cloud will also require personnel to operate and maintain the cluster and network. Salaries for personnel can vary and are apportioned to users proportionate to the number of blades the developer handles (salary - \$100-\$250 per blade and network connection).

6.5 Comparison of the Cloud Network Transfer Options

Table 5 gives a tabulated summary of network transfer incremental pricing for each of the public cloud providers as a function of the total size of the data to be transferred⁵.

Table 5. Network transfer rate costs.

Total data transfer	Amazon AWS	Google cloud	Microsoft Azure
1 GB	Free	\$0.12	Free
5 GB	\$0.09	\$0.12	Free
1 TB	\$0.09	\$0.12	\$0.087
10 TB	\$0.09	\$0.11	\$0.087
50 TB	\$0.085	\$0.08	\$0.083
150 TB	\$0.07	\$0.08	\$0.07
500 TB	\$0.05	\$0.08	\$0.05




Fig. 5. Costs Charged by Cloud Providers to Transfer Data to an external location.

These data transfer costs are illustrated in Fig. 5 showing a graph of the cumulative network transfer costs for Amazon, Google, and Azure versus the amount of data transferred (in TB). To further illustrate the impact of data transfer on the cost comparisons of public versus private cloud options, a three-dimensional plot was constructed with axes of cost in dollars, time in days and data transferred in terabytes. Figure 6 shows planes in the figure representing the cumulative on-demand computation costs for a private cloud provider and an Amazon on-demand computation instance. The additional plane on the graph shows the Amazon add-on costs for transfer of data up to 50 TB over and above the Amazon computation on demand instance. These calculations show the sensitivity of adding data transfer requirements when performing a costs analysis

⁵ The Private Cloud does not have a fee for the data transfer over the network.




Fig. 6. Three dimensional plot showing cumulative costs for an on-demand private cloud, an Amazon on-demand public cloud and the additional add-on Amazon costs for transferring up to 50 TB of data over and above the Amazon computation on demand cost.

among public and private cloud providers. These results will be discussed in more detail in Sect. 7.

7 Summary

This paper studied the question of determining if and when it may be more cost effective to utilize public or private clouds for HPC type instances. Using a baseline hardware configuration reflective of HPC system requirements as a test case, a systematic sample analysis was done for computation, storage and network transfer options using several typical cloud providers.

The first major observation involved the time dependency and compute requirements of the HPC application in determining whether it would be more cost effective to use a public or private cloud. As discussed in Subsect. 4.2 and illustrated in shown in Fig. 2, the public cloud option is shown to be more cost effective than the private cloud for HPC type applications of short duration or intermittent need for cloud computation resources. However, as the need for extended or prolonged access to computation resources grows, at some point in time the cumulative costs for HPC applications are more cost effective being assigned to a private cloud provider.

The second major observation showed the cost sensitivity when large quantities of data must be transferred from public cloud providers to external Internet locations. As shown earlier in Subsect. 6.5 the three dimensional plot of cumulative cost, time and data transferred (Fig. 6) has strong implications for choosing a cost effective HPC cloud computing option. This plot compared the on-demand cumulative computation costs for an Amazon public cloud option and a typical private cloud HPC option. Previously discussed in Subsect. 4.2, Fig. 2 showed that the intersection point in time for a crossover of cumulative costs between

an on-demand Amazon baseline instance and a private cloud instance is approximately 110 days. The arrows in Fig. 6 pointing to the time axis illustrate how the additional costs of data transfer will shorten the crossover point for determining the cost effective choice between public and private cloud computing options from approximately 110 days to approximately 40 days. The arrows in the same figure pointing to the cost axis show an additional cost of approximately \$5,000 that will be incurred if the Amazon HPC baseline cloud instance user must also pay for data transfer.

The third major observation involved the introduction of the ideas of Spot Market Instances into the mix of cloud computing reservations. Subsection 4.3 studied the hourly spot prices for the M4.10xLarge sample baseline configuration in two Amazon regions and five availability zones. The results of the correlation matrix calculations showed that during the time period when the spot prices were tracked, certain pairwise correlations and anti-correlations were detected among various availability zones. This indicated that any spot price bid submitted to one zone may be highly correlated to the price in the other zone. However, there were periods of time when pairs of availability zones showed strong anti-correlated behavior of their spot instances. In effect, the spot price in one zone would be substantively different than the spot price in the other zone.

This observation has potentially interesting consequences for users dependent on the spot market for accessing the most cost effective HPC level computational resources. The strong anti-correlation signal would indicate when and where to switch availability zones and bid prices to assure the best access to the cloud resources at the most cost effective rate. By closely tracking the spot prices among the different availability zones and calculating the correlation matrix for a given hardware configuration, it might be possible to shift the workload from the availability zone with the higher spot price to the other availability zone with the lower spot price⁶. In effect, the user would be monitoring the spot market for inefficiencies in the pricing structure in order to maximize access to the public cloud resources in the most cost effective manner.

The cost analysis here shows that it is important to determine the full duration needed for a cloud computing system at the outset. If the cost analysis is focused on just the immediate short term or incremental need, then the public cloud is the option of choice. However, if the overall total duration for use of that cloud for the project ultimately extends over the longer term, then repeatedly using the public cloud for each incremental calculation can end up costing the user substantially more money than if the decision was made at the outset to commit the entire project to a private cloud.

In summary, to answer the question when it becomes more cost effective to use a private HPC cloud versus public HPC cloud requires a detailed analysis of computation, storage and network transfer costs associated with the specific HPC type hardware configuration. The evaluation is sensitive as to whether the

⁶ This potential economic strategy may work if the user's job only involves computations and is not dependent on staging large quantities of data in a particular region and zone before submitting a bid price for access.

type of HPC application needs continuous CPU resources for long periods of time, the quantity of storage that application will require in the cloud, and the amount of data that must be ultimately be transferred from the cloud computing vendor to an external location. The sample process described here for the HPC baseline configuration would need to be repeated for the specific HPC hardware configuration platform being considered as well as the public and private cloud computing vendors being considered. It is also recognized that cost is not the only consideration when determining the most economically effective public or private cloud computing configuration for an HPC application. The throughput performance of an HPC application must also be factored into the overall analysis to determine the most economical option for an HPC application in the cloud. The authors are analyzing HPC level benchmark performance of each vendor's baseline configuration and the results will be reported in a future publication.

Acknowledgments. This work is supported in part through NSF grant 0910767, 1318564, 1330553, the U.S. Army Research Office (ARO) grant W911NF-08-1-0105 managed by the NCSU Science of Security Initiative and the Science of Security Labet, by the IBM Share University Research and Fellowships program funding, and the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357. One of us (Patrick Dreher) gratefully acknowledges support with an IBM Faculty award.

References

1. Chen, Y., Sion, R.: To cloud or not to cloud? musings on costs and viability. In: Proceedings of the 2nd ACM Symposium on Cloud Computing, pp. 29:1–29:7 (2011)
2. Walker, E.: The real cost of a CPU hour. *IEEE Comput.* **42**, 3541 (2009)
3. Zhai, Y., Liu, M., Zhai, J.: Cloud versus in-house cluster: evaluating amazon cluster compute instances for running MPI applications. In: State of the Practice Reports, pp. 11:1–11:10 (2011)
4. Gupta, A., Milojevic, D.: Evaluation of HPC applications on cloud. In: Fifth Open Cirrus Summit, pp. 22–26 (2011)
5. Ostermann, S., Iosup, A., Yigitbasi, N., Prodan, R., Fahringer, T., Epema, D.: A performance analysis of EC2 cloud computing services for scientific computing. In: Avresky, D.R., Diaz, M., Bode, A., Ciciani, B., Dekel, E. (eds.) *CloudComp 2009*. LNCSSTE, vol. 34, pp. 115–131. Springer, Heidelberg (2010). doi:[10.1007/978-3-642-12636-9_9](https://doi.org/10.1007/978-3-642-12636-9_9)
6. Ding, F., Mey, D., Wienke, S., Zhang, R., Li, L.: A study on today's cloud environments for HPC applications. In: Helfert, M., Desprez, F., Ferguson, D., Leymann, F. (eds.) *CLOSER 2013*. CCIS, vol. 453, pp. 114–127. Springer, Cham (2014). doi:[10.1007/978-3-319-11561-0_8](https://doi.org/10.1007/978-3-319-11561-0_8)
7. Brandt, J., Gentile, A., Mayo, J., Pebay, P., Roe, D., Thompson, D., Wong, M.: Resource monitoring and management with OVIS to enable HPC in cloud computing environments. In: *IEEE International Symposium Parallel Distributed Processing, IPDPS 2009*, pp. 1–8 (2009)

8. Gupta, A., Kal, L.V.: Towards efficient mapping, scheduling, and execution of HPC applications on platforms in cloud. In: Parallel and Distributed Processing Symposium Workshops Ph.D. Forum, pp. 2294–2297 (2013)
9. Gómez Sáez, S., Andrikopoulos, V., Hahn, M., Karastoyanova, D., Leymann, F., Skouradaki, M., Vukojevic-Haupt, K.: Performance and cost trade-off in IaaS environments: a scientific workflow simulation environment case study. In: Helfert, M., Méndez Muñoz, V., Ferguson, D. (eds.) CLOSER 2015. CCIS, vol. 581, pp. 153–170. Springer, Cham (2016). doi:[10.1007/978-3-319-29582-4_9](https://doi.org/10.1007/978-3-319-29582-4_9)
10. Saez, S., Andrikopoulos, V., Hahn, M., Karastoyanova, D., Leymann, F., Skouradaki, M., Vukojevic-Haupt, K.: Performance and cost evaluation for the migration of a scientific workflow infrastructure to the cloud. In: Proceedings of the 5th International Conference on Cloud Computing and Service Science, CLOSER 2015, p. 110. SciTePress (2015)
11. Coghlan, S., Yelick, K., Draney, B., Canon, R.S.: The Magellan report on cloud computing. In: Office of Advanced Scientific Computing Research (ASCR), US Department of Energy (2011). http://science.energy.gov/~media/ascr/pdf/programdocuments/docsMagellan_Final_Report.pdf
12. Vouk, M., Sills, E., Dreher, P.: Integration of high-performance computing into cloud computing services. In: Handbook of Cloud Computing, pp. 255–276 (2010). Chap. 11
13. Amazon High Performance Computing (2016). <https://aws.amazon.com/hpc/>
14. Google Compute Engine (2016). <https://cloud.google.com/compute/>
15. Microsoft Azure (2016). <https://azure.microsoft.com/en-us/>
16. Microsoft Big Compute: HPC & Batch (2016). <https://azure.microsoft.com/en-us/solutions/big-compute/>
17. Vouk, M.: Cloud computing issues, research and implementations. *J. Comput. Inf. Technol.* **16**(4), 235–246 (2008)
18. Dreher, P., Vouk, M., Sills, E., Averitt, S.: Evidence for a cost effective cloud computing implementation based upon the NC state virtual computing laboratory model. In: Advances in Parallel Computing, High Speed and Large Scale Scientific Computing, vol. 18, pp. 236–250 (2009)
19. Schaffer, H.E., Averitt, S.F., Hoit, M.I., Peeler, A., Sills, E.D., Vouk, M.A.: NCSUs virtual computing laboratory: a cloud computing solution. In: *IEEE Computer*, pp. 94–97 (2009)
20. Apache VCL (2016). <https://vcl.apache.org/>
21. Amazon EC2 Spot Instances. <http://aws.amazon.com/ec2/spot-instances/>
22. Zhang, Q., Gurses, E., Boutaba, R., and Xiao, J., Dynamic resource allocation for spot markets in clouds. In: Workshop on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services Hot-ICE (2011)
23. Chen, J., Wang, C., Zhou, B.B., Sun, L., Lee, Y.C., Zomaya, A.Y.: Tradeoffs between profit and customer satisfaction for service provisioning in the cloud. In: HPDC (2011)
24. Mazzucco, M., Dumas, M.: Achieving performance and availability guarantees with spot instances. In: IEEE International Conference on High Performance Computing and Communications, pp. 296–303 (2011)
25. Mattess, M., Vecchiola, C., Buyya, R.: Managing peak loads by leasing cloud infrastructure services from a spot market. In: IEEE International Conference on High Performance Computing and Communications, pp. 180–188 (2010)
26. Yi, S., Andrzejak, A., Kondo, D.: Monetary cost-aware checkpointing and migration on amazon cloud spot instances. *IEEE Trans. Serv. Comput.* **5**(4), 512–524 (2012)

27. Agmon Ben-Yehuda, O., Ben-Yehuda, M., Schuster, A., Tsafir, D.: Deconstructing amazon ec2 spot instance pricing. In: 2011 IEEE Third International Conference on Cloud Computing Technology and Science (CloudCom), pp. 304–311 (2011)
28. Bonacquisto, P., Di Modica, G., Petralia, G., Tomarchio, O.: Dynamic pricing in cloud markets: evaluation of procurement auctions. In: CLOSER 2014. CCIS, vol. 512, pp. 31–46 (2015)
29. GitHub Repository for Boto Python Library (2016). <https://github.com/boto/boto>
30. Amazon SDK for Python to access Amazon public data (2016). <https://aws.amazon.com/sdk-for-python/>
31. Bhatia, K.: The data science of AWS Spot Pricing (2015). <https://medium.com/cloud-uprising/the-data-science-of-aws-spot-pricing-8bed655caed2#.f9w14i4iq>