An Economic and Energy-Aware Analysis of the Viability of Outsourcing Cluster Computing to the Cloud

Carlos de Alfonso[a,*], Miguel Caballer[a], Fernando Alvarruiz[a], Germán Moltó[a]

[a]*Instituto de Instrumentación para Imagen Molecular (I3M). Centro mixto CSIC – Universitat Politècnica de València – CIEMAT, camino de Vera s/n, 46022 Valencia, España*

Abstract

This paper compares the total cost of ownership of a physical cluster with the cost of a virtual Cloud-based cluster. For that purpose, cost models for both the physical cluster and the cluster on the Cloud have been developed. The model for the physical cluster takes into account previous works and incorporates a more detailed study of the costs related to energy consumption and the usage of energy saving strategies. The model for the cluster on the Cloud considers pricing options offered by Amazon EC2, such as reserving instances on a long-term basis, and also considers using tools for powering nodes on and off on demand, in order to avoid the costs associated to keeping idle nodes running. Using these cost models, a comparison is made of physical clusters with Cloud clusters of a similar size and performance. The results show that Cloud clusters are an interesting option for start-ups and other organizations with a high degree of uncertainty with respect to the computational requirements, while physical clusters are still more economically viable for organizations with a high usage rate.

Keywords: Cloud computing, Cluster computing, Green computing, Cost analysis

---

1. Introduction

One of the main problems faced when deploying a cluster of PCs relates to the high Total Cost of Ownership (TCO). This cost involves not only the purchase and installation of the equipment (computational nodes, network components, cables, hard disks, etc.), but also the operating costs. The latter includes the salaries of the personnel in charge of the installation and maintenance, the electricity consumed, and the costs related to rent appropriate housing and its associated cooling systems. The problem is that the usage patterns of these machines are highly dynamic, where peak loads are often restricted to the context of specific experiments or deadlines. In addition to this, the prices of clusters of PCs rapidly decrease (due to the technology obsolescence), thus reducing the value of the initial investment in hardware.

---

*Corresponding author
Email address:* `caralla@.upv.es` (Carlos de Alfonso)

As an alternative, researchers might access the resources at the Computing Centers of national or international institutions. This is the case of the *Spanish Supercomputing Network*, which aggregates several supercomputing centers in Spain. The access to this equipment is supervised by an access committee that grants limited resource access according to the scientific merit of the proposals. Another example is the not-for-profit organization *Partnership for Advanced Computing in Europe* (PRACE), which provides access to a world class computing and data management infrastructure [1].

Another alternative is to use *Cloud computing*, a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [2]. This is the case of Infrastructure as a Service (IaaS), which performs on demand resource provision of computational resources, storage capacity, network access, etc. This is achieved by means of virtual machines that run on the Cloud provider's computing center. In the case of public Cloud providers, a pay-per-use pricing model is typically employed, where users are only charged for the resources that they have consumed.

Virtualization was not considered a feasible approach for High Performance Computing (HPC), due to the penalties involved mainly in the I/O. However, many applications running on PC clusters are CPU-bound, thus their performance is hardly affected by virtualization [3]. In addition, certain public Cloud providers, such as Amazon, offer low latency links among instances, thus leveraging the idea of using virtual clusters on the Cloud. Having a virtual cluster infrastructure on a public Cloud provider has a large number of advantages for the end user, since no hardware costs are involved. However, the sustainability of this infrastructure in the long term might represent a high cost, since the pay-as-you-go model offered by the Cloud providers implies that a running virtual machine costs money regardless of it being used for computations or not.

Due to the increase of the use of virtualisation and Cloud technologies, some initiatives to create HPC clusters over Cloud infrastructures are emerging. One of the first approaches, described in [4], enabled to launch a fully functional Hadoop cluster over Amazon EC2 using a set of simple scripts[1]. Other tools can create HPC clusters in the Cloud using some kind of Local Resource Management System (LRMS) to manage the jobs. StarCluster[2] uses this approach to create a cluster in the Amazon EC2 infrastructure, with a set of predefined installed applications (Sun Grid Engine, OpenMPI, NFS, etc.) to enable launching parallel jobs to the queue system.

Recently Cycle Computing used Cloud infrastructures to create a 30,000-core HPC cluster using Amazon EC2 standard instances[3]. The cluster ran for about seven hours, with 3,809 compute instances and a total of 26.7TB of RAM and 2PB (petabytes) of disk space, with a 10 Gigabit Ethernet network. Amazon itself actually built a supercomputer on its own Cloud that made it onto the list of the world's Top 500 supercomputers. With 7,000 cores, that specific Amazon cluster ranked number 232 in the world in November 2010 with speeds of 41.82 teraflops, falling to number 451 in June of this year[4]. It is estimated that the whole Amazon EC2 infrastructure can be ranked number 42 among

---

[1] http://wiki.apache.org/hadoop/AmazonEC2

[2] http://web.mit.edu/stardev/cluster/

[3] http://arstechnica.com/business/news/2011/09/30000-core-cluster-built-on-amazon-ec2-cloud

[4] http://www.top500.org/system/details/10661

the world's Top 500 supercomputers[5].

These tools and services turn the Cloud into a technologically feasible option for the deployment of clusters of PCs. However, it is important to assess the economic viability of outsourcing the deployment of a cluster on the Cloud, compared to the purchase of a physical cluster. For that, this article analyses and compares the cost of having a physical HPC cluster with that of a similar infrastructure on a public Cloud provider.

The remainder of the paper is structured as follows. First, section 2 describes the related work in the literature comparing physical clusters with virtual clusters deployed on the Cloud, either economically or in terms of performance. Then, section 3 dissects the TCO of an HPC cluster, introducing an energy-aware cost model for physical clusters, and a cost model for virtual clusters on the Cloud. Later, section 4 introduces some simplifying assumptions and considerations in the cost models developed in the previous section, also presenting data related to energy consumption, prices of hardware components and Cloud instance prices. The resulting models are then used in subsection 4.2 in order to compare physical clusters with Cloud clusters. Finally, a discussion of the results is presented in section 5, followed by concluding remarks in section 6.

2. Related work

There is recent work in the literature comparing large scale public Cloud infrastructures with PC clusters, especially for HPC. In [5], the authors include a comprehensive evaluation of performance comparing physical HPC clusters to virtual HPC clusters on Amazon EC2, where the larger network latency in the latter introduces a severe performance penalty for parallel applications. In [6], a similar performance comparison is made with workflow applications composed of loosely-coupled parallel applications consisting of computational tasks linked via data and control dependencies. Different EC2 instance types were employed to assess the performance of the applications and a virtualisation overhead below 8% was computed.

Other works have studied the cost or benefit of using Cloud technologies from different points of view: in [7], the authors study the cost of executing the Montage astronomy application in public Cloud environments. In [8], the authors evaluate the cost of expanding a local virtual cluster using a Cloud technology provider, in order to reduce the response time of the user requests. In [9] the authors compare the performance and monetary cost-benefits of Clouds versus desktop Grids (or Volunteer Computing) infrastructures, ranging in size and storage. In [10], the TCO and Utilisation Cost of a Cloud infrastructure are analysed from the point of view of the IaaS service provider. They also developed a web tool where the users can introduce the parameters of their Clouds and obtain the total cost analysis. Finally, other works such as [11], [12] or [13] have tried to compare the cost of owning a datacenter infrastructure versus the pay-per-use costs of Cloud deployments.

Those contributions show that it is crucial to evaluate the economic impact of outsourcing an organisation's HPC computing infrastructure to an external Cloud provider. As opposed to previous works, this paper performs a detailed comparison between physical and virtual HPC clusters from the point of view of the TCO, considering ener-

---

[5]http://www.readwriteweb.com/cloud/2011/11/amazon-ec2-now-42-supercompute.php

getic, management and infrastructural issues, using concepts and estimations from related work, but considering a more detailed analysis. In previous works, the cost of the energy is only estimated, while in this work a detailed model of the energy consumption cost has been defined, where green-aware technologies are a key task to minimize the energetic consumption and the costs in the Cloud deployments. This paper also analyzes how the Amazon EC2 reserved instances can be used to reduce the cost when the users can estimate the average usage of the cluster.

3. The Total Cost of Ownership (TCO) of an HPC Cluster

TCO is generally used as a means of addressing the real costs attributed to owning and managing an IT infrastructure in a business. Therefore, the TCO of owning a HPC cluster not only includes the capital cost, but also the cost of operating the IT infrastructure, and other factors [11].

The cost of owning an HPC cluster can be modelled according to the expression (1), where C_F stands for the fixed costs, which only occur once, as opposed to variable costs (C_V), required during the operation of the equipment. C_F can be detailed as in (2).

$$C = C_F + C_V \tag{1}$$

$$C_F = C_P + C_S + C_{CP} + C_A \tag{2}$$

Concerning the costs related to the purchase and configuration of the equipment (C_P), we have considered the computing nodes, additional components, such as switches, Power Device Units, etc., and auxiliary physical elements (racks, cables, etc.). Besides the costs of purchasing the cluster itself, it is important to consider the costs related to buying or renting the physical space where the cluster will be located, together with appropriate refurbishment (C_S). The expression includes the costs related to the purchase of the cooling system (C_{CP}) and the administrative costs (C_A) involved in the purchase (mortgages, loans, infrastructure documentation, etc.). These costs have already been studied in the literature (see e.g. [14]). Another concept traditionally considered when calculating the TCO is the equipment disposal. This topic has not been included in the equation since due to initiatives such as the European Recycling Platform[6], most computer vendors like Dell[7] or HP[8], etc. offer a free recycling program to their clients. In addition, other tasks such as backing up the data or removing data from hard drive to ensure data privacy, etc, are included as part of the maintenance costs.

In addition to the fixed costs, it is possible to break down the variable costs C_V, which account for periodic costs during the lifetime of the hardware:

$$C_V = C_L + C_M + C_O + C_E \tag{3}$$

Among the variable costs included in (3), one should consider the purchase and update of the software licenses employed (C_L), together with the costs of preventive and

---

[6]http://www.erp-recycling.org/
[7]http://www1.euro.dell.com/content/topics/topic.aspx/emea/topics/services/recycle?c=es&cs=esbsdt1&l=es&s=bsd
[8]http://www8.hp.com/es/es/hp-information/environment/hardware-recycling.html

corrective maintenance (C_M) to repair the machine and to update certain parts. The operation costs of the cluster (C_O) broadly include the costs of the personnel in charge of deploying, updating and securing the cluster.

Finally, we have to consider the energy cost (C_E), which is one of the most complex and highly variable aspects to evaluate. This cost is expressed in (4), where four principal components are included, which correspond to different aspects of the cluster energy consumption. These components are aggregated and multiplied by the cost of the energy unit (C_U), in order to obtain the final cost.

$$C_E = (E_0 + E_I + E_J + E_{CO}) \cdot C_U \tag{4}$$

$$E_0 = t \cdot P_0 \tag{5}$$

$$E_I = P_I \cdot \sum_{i=1}^{n} t_I(i) \tag{6}$$

$$E_J = \sum_{j=1}^{m} (P_P \cdot t_P(j) \cdot n_j + P_U(j) \cdot t_U(j) \cdot n_j) \tag{7}$$

We have considered four energy consumption patterns, related to the different main states in which a cluster can operate:

- **Energy consumption of the essential components** (E_0) for the normal functioning (switches, front-end, network cards, etc.), represented in (5). It is related to the power consumed by those components (P_0) and to the considered time (t).

- **Energy consumption of idle nodes** (E_I), represented in (6). It depends on P_I (power consumed by an idle node), $t_I(i)$ (amount of time node i is idle), and n (number of nodes in the cluster).

- **Energy consumption dedicated to workload computation** (E_J), as expressed in (7), which can be split in two. On the one hand, the energy consumed while preparing the nodes for the job (data staging, file transfers, environment setup, etc.), which depends on P_P (power consumed by a node while it is under preparation), $t_P(j)$ (time invested in that preparation for job j) and n_j (number of nodes used by the job). On the other hand, the energy consumed while the job uses the resources, which depends on the function $P_U(j)$ (power consumed), $t_U(j)$ (job duration) and n_j (number of nodes used by the job). The total energy is the sum of these two components for all the m jobs.

- **Energy consumption of the cooling system**, denoted by E_{CO}.

These states have been considered to be different because $P_I < P_P << P_U(j)$. In our case we have represented the power consumed by a node while it is computing as a function $P_U(j)$, since that consumption largely depends on the number of cores and processors being simultaneously employed.

$$C = C_P + C_S + C_{CP} + C_A + C_L + C_M + C_O +$$

$$\left(t \cdot P_0 + P_I \cdot \sum_{i=1}^{n} t_I(i) + \sum_{j=1}^{m} \left(P_P \cdot t_P(j) \cdot n_j + P_U(j) \cdot t_U(j) \cdot n_j \right) + E_{CO} \right) \cdot C_U \quad (8)$$

Expression (8) includes the TCO of an HPC cluster during a certain amount of time. This expression collects the principles and concepts employed in other studies as covered in the related work section. In addition, it introduces the energy consumption breakdown. The previous studies focus on aspects related to the purchase and maintenance of hardware, while our approach puts more emphasis on analysing the costs that arise when operating the cluster.

Some of the aforementioned costs are constant, while others are typically covered by the economic resources of the organisations. In research centers or universities a space reorganisation can be performed in order to reduce C_S to the bare minimum. This is precisely the case with C_A since this cost might be reduced if the existing administrative personnel take responsibility for the administrative management of the cluster. In some cases, costs such as C_L can also be neglected since licences could be included in the purchase contract of the equipment.

Finally, the energy required for the cooling system tends to be estimated as proportional to the energy consumed by the cluster components, using the Power Usage Efficiency (PUE) ratio. This energy consumption increment can involve from 30% to 200% of the energy consumed by the cluster components alone [15].

Taking into account the aforementioned considerations, the simplified total cost of the cluster mainly depends on the purchase of the hardware, the maintenance and operation of the cluster and its energy consumption. We are not considering C_S and C_{CP} costs, which have been studied earlier in some papers detailed in the related work section. For the discussion, we assume that the owner of the cluster hosts the hardware in an available data center and, thus, the analysis focuses on the costs related to the purchase and operation of the cluster. Buying or building the appropriate infrastructure to host the cluster represents a larger cost than the cluster itself. In this last case, other more exhaustive studies should be carried out, as in [14], in order to decide the features of the installation, the amount of hardware to host, the facilities included, together with the related cost.

Expression (8) presents a term related to the energy consumption of idle nodes. In some cases, the amount of time in which the nodes remain idle can represent a high percentage of time of the hardware lifetime. This is due to over-provisioning of resources in order to better cope with peak workloads. Therefore, it is not uncommon to see usage rates of clusters of PCs in the order of 20% [11]. Under these conditions, we could power off the idle nodes in order to save energy and reduce the total cost, as pointed out in [16]. Powering off the idle nodes tries to reduce $\sum_{i=1}^{n} t_I(i)$ from (6) and (8) to zero, thus eliminating the corresponding term. In this case, the time involved in starting up the nodes would be included in the preparation phase, $t_P(j)$. Considering all these assumptions, the expression reads as follows:

$$C = C_P + C_M + C_O + \left(t \cdot P_0 + \sum_{j=1}^{m} (P_P \cdot t_P(j) \cdot n_j + P_U(j) \cdot t_U(j) \cdot n_j) + E_{CO} \right) \cdot PUE \cdot C_U$$

$$(9)$$

3.1. The Cost of an HPC Cluster on the Cloud

For the case of clusters on the Cloud, the economic analysis is based on the pay-per-use model that Cloud platforms introduce. Moving the cluster to the Cloud discards the fixed costs, such as the purchase of the equipment and the supporting infrastructure, the cooling system, etc. In addition, the administrative costs are substantially reduced since there is no longer need to perform such an upfront investment to purchase the hardware. In any case, the energy consumption related to the cluster and the cooling system is entirely covered by the Cloud provider. However, in this case there are usage costs, represented by the pay-as-you-go model of Cloud computing.

For the discussion, we consider the pricing model proposed by Amazon EC2, since it is one of the pioneer Cloud providers and it has the biggest market share in the provision of computational resources in Cloud. In addition, many providers are adopting the pricing model proposed by Amazon. In this case, an "instance", which corresponds to a virtual machine running with specific virtual hardware features (processors, disk, memory, etc.), is charged per hours running, regardless of its utilization. The different instances offered by Amazon EC2, with the corresponding features and cost, are shown in [17].

There are also additional costs that should be considered, such as the data storage cost and the network transfer cost. Both Amazon and other IaaS-supplier companies charge for using storage space apart from the one required for the proper functioning of the virtual machine, i.e., databases, persistent volumes, etc. The network bandwidth usage for inbound and outbound connectivity of virtual machines is also considered in the pricing model.

With the aforementioned considerations, it is possible to propose a cost model of a cluster of PCs based on the pricing model of Amazon EC2, as shown in expression (10).

$$C = C_A + C_L + C_O + t \cdot (C_H + n \cdot C_W) + C_{ST} + C_N \qquad (10)$$

This model includes the administrative, license and operation costs, together with the usage costs of the virtual machines, where we differentiate the cost of the front-end instance (a.k.a the head node, C_H) and the cost of the working node instance (C_W). The usage cost also depends on the number n of computational nodes of the cluster and the amount of time that these nodes are running. Following the same considerations as in the case of a physical cluster, we will assume that administrative and license costs are negligible. In this case the prices of virtual machine images can also include the license costs of preinstalled applications.

Finally, the model also considers the costs related to data storage (C_{ST}) and network bandwidth usage (C_N). For the former, a virtual machine that does not require additional storage space, other than the one provided by the virtual machine image, will not incur in further charges. For the latter, if a virtual machine does not perform or receive an outbound connection (external to the Cloud provider's network) it will not produce any

economic charges. In the case of an HPC cluster, it is important to notice that, since the computational nodes are deployed inside the infrastructure of a Cloud provider, the network communications among the nodes will not cause any charge.

Therefore, the network cost would be caused by the data transfers from the user network to the cluster (and to retrieve the results of the computations), as well as the downloads that the working nodes could eventually perform (if they are allowed). In the case of Amazon EC2, uploading data to the Cloud incurs in no additional charges.

Note that working nodes are charged regardless of whether they are busy or idle. Considering the cost model for a physical cluster (9) and the ability to power off nodes in order to reduce energy consumption, it is easy to realize that an approach for energy reduction would also be valid for reducing the economic cost in the Cloud. In order to reduce this cost there is software such as StarCluster, Hadoop, Globus Provision[9] and elasticwulf[10], which enable to deploy a cluster with a given number of nodes on a Cloud infrastructure, such as Amazon EC2. In these cases, the resources are released when the computations finish, in order to save money. Due to the growth of this kind of tools, different computer services providers are offering the HPC cluster creation as a service to their clients. As an example SGI[11], Sabalcore[12] or Penguin Computing[13] provide access to HPC clusters that can scale on demand. Also some research centres such as SARA[14] offer services to create "Virtual Private HPC Clusters".

With these considerations, it is possible to break down the economic cost of (10), applying the criteria from (7) to the pay-per-use model, which results in expression (11). In this case, the startup time of the virtual nodes is included in the preparation phase of the instances $t_P(j)$.

$$C = C_O + t \cdot C_H + C_W \cdot \sum_{j=1}^{m} (t_P(j) + t_U(j)) \cdot n_j + C_{ST} + C_N \qquad (11)$$

4. Cost Analysis of Moving HPC to the Cloud

This section thoroughly analyses the tradeoffs of moving HPC to the Cloud, in the shape of a virtual cluster deployed on Amazon EC2, compared to a traditional in-house physical cluster, from an economic perspective. The expressions in this section no longer include the C_O cost, because we assume this cost is of the same order of magnitude for both the physical and the virtual cluster, and thus can be omitted.

The cost of a cluster throughout a period of time largely depends on its usage rate and the workload distribution among the nodes, determined by its LRMS. In our case, using the expression (9), we compute the lower and upper bounds to the TCO of a physical cluster. For that, according to [12], we estimate the maintenance cost of the machines (C_M), including out of warranty repairs, to be an annual 10% of the initial cost of the hardware. Concerning the PUE, choosing the right value depends on the specific

---

installation but, according to [15], it typically varies from 1.3 to 3.0. For the analysis we choose the average value of that scale, i.e. 2.15, although other cases might involve adjusting that ratio, such as in [13, 18].

The lower bound for the TCO of a physical cluster is shown in expression (12) and it corresponds to the situation where a cluster is purchased but no computations are ever performed. Therefore, all the internal nodes are powered off and only the front-end and the essential components remain powered on. The upper bound expression is shown in (13), and it corresponds to the theoretical state in which the cluster is working at full workload, where all the computing nodes are running and using 100% of all the computing cores (c). In (13), the amount of power consumed by the n nodes is estimated using the power of an idle node (P_I) and the additional power consumed by each of the c cores in the cluster (P_C). Y corresponds to the duration of one year.

$$C = C_P + 0.1 \cdot C_P \cdot \frac{t}{Y} + t \cdot P_0 \cdot 2.15 \cdot C_U \tag{12}$$

$$C = C_P + 0.1 \cdot C_P \cdot \frac{t}{Y} + [t \cdot P_0 + t \cdot (n \cdot P_I + c \cdot P_C)] \cdot 2.15 \cdot C_U \tag{13}$$

Just like the physical cluster, the cost of a cluster in the Amazon EC2 Cloud depends on its usage, since this determines the number of computational nodes to provision. In our case, instead of creating a synthetic benchmark to obtain the average cluster activity, we use a percentage of usage (u) of the equipment through the time of the study:

$$C = C_H \cdot t + C_W \cdot n \cdot u \cdot t + C_{ST} + C_N \tag{14}$$

However, to increase the diversity of cases in the study, we have included the reserved instances model offered by EC2. This allows users to make a one-time payment to reserve an instance for a three-year period, and in turn receive a significant discount on the hourly charge for that instance. Considering our approach, a number of p reserved instances would be kept up and running for the reservation period t_R, and the other instances could be fired up and down on demand up to the maximum number of n nodes. The expression now reads as in (15).

$$C = (C_{HR} + C_{WR} \cdot p) \cdot \lceil \frac{t}{t_R} \rceil + C_W \cdot (n - p) \cdot u' \cdot t + (C_{HU} + C_{WU} \cdot p) \cdot t + C_{ST} + C_N \tag{15}$$

This expression introduces the head node instance reservation price (C_{HR}) and the hourly price for using the reserved instance (C_{HU}). In a similar way, it also considers the cost of the reserved working nodes (C_{WR} and C_{WU}). It also considers u', which corresponds to the usage rate of nonreserved instances, which we will refer to as *nonreserved usage*.

From now on, we will assume that the costs related to storage (C_{ST}) are zero, since our study focuses on the operating costs of the cluster. This assumption is reasonable considering that the front-end instance has enough free storage space to operate the cluster. For example, an *m1.small* instance in Amazon EC2 has a 160 GBytes disk size. The front-end will probably be shared with the internal computing nodes via the network (using NFS, for example).
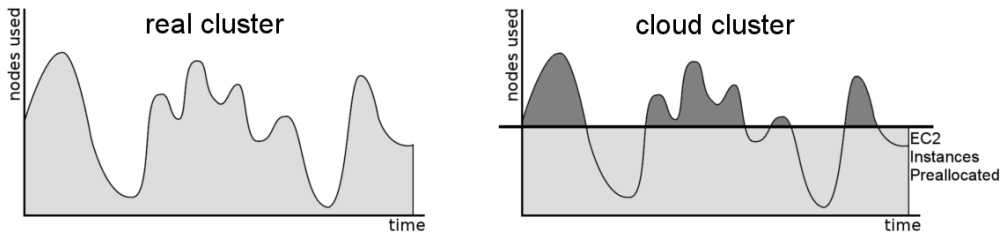
Figure 1: Usage rate beyond the amount of preallocated nodes (reserved instances)

Concerning the costs of the network bandwidth C_N, this will depend on the applications to be executed on the cluster, the amount of data required to start the computations, and the generated output data of the executions. In the case of Amazon EC2, internal transfers do not incur in additional charges. However, outer transfers, i.e., the data movement between the Cloud provider's network and the client, are billed. Considering a reference value of 100 Gb of data transfer per month (more than 3 Gb of daily results by the users), this would represent a monthly cost of 12\$. Since this study handles values in the order of thousands of euros, the cost of the data transfers is negligible.

It is important to point out that expression (15) is an estimation that includes simplifications and, therefore, the actual results will depend on the usage patterns of the cluster and the LRMS. In addition, the expression represents an upper bound to the cost of the cluster in the Cloud, since we have assumed that the reserved instances will remain active all the time. This is because we have assumed that the usage rate of the reserved nodes is very high, in order to match the typical workload of the system. Notice that additional savings could be made if those reserved nodes were powered off if they remain idle for a certain amount of time.

An important issue to be considered is that the variable u' included in expression (15) is not the usage rate of the cluster, but the usage rate of the extra $(n-p)$ unreserved nodes. If we assume that the reserved nodes are always used, the global usage rate of the cluster would be $p/n + ((n-p)/n) \cdot u'$.

Figure 1 shows a conceptual view of the costs of the cluster related to its workload, according to the proposed model. In the left hand side of the figure, an example of the workload of a physical cluster is shown, which is characterized by the number of nodes being used during the time study. Considering the proposed model, this workload results in an economic cost due to the energy consumption. In the right hand side of the figure, the same workload is depicted for the case of a cluster in the Cloud. Assuming the initial purchase of a set of reserved instances in EC2, the final cost would consist of the cost of these reserved instances (lower part of the figure, in light gray), together with the cost of the dynamically provisioned instances, which cope with the excess of workload that cannot be executed by the reserved instances (upper part of the figure, in dark gray). In the expression (15), the dark gray part corresponds to $(n-p) \cdot u' \cdot t$.

As an example, in the case of a 64-node cluster where 40 nodes are reserved instances, there are 24 nodes that should be powered on and off depending on the workload of the cluster. Therefore, the nonreserved usage u' refers to those 24 nodes. Thus, a nonreserved usage of 0.5 would represent that an average of 12 nodes (out of the 24) are powered on during all the time of study, or equivalently, the 24 nodes are powered on during half of

Parameter	Value
C_U ($ per kWh)	0.0988
P_0 (kW)	0.7659
P_I (kW)	0.0966
P_C (kW)	0.01075

Table 1: Energy consumption a cluster node (Intel Xeon E5520 2.3 GHz)

Instances	1	2	3	4	5	6	7	8
Increment (W)	34	44	57	68	71	73	78	86
Increment / core (W)	34.0	22.0	19.0	17.0	14.20	12.17	10.71	10.75

Table 2: Power consumption of a modern cluster node

the time of study.

Introducing the nonreserved usage and the previous assumptions, we will be able to identify a broad spectrum of cases when using a cluster infrastructure, thus performing a thorough study. The following section analyses the cost for two clusters, a small-sized one with 64 cores and a larger one with 1024 cores.

4.1. Supporting Data for the Case Study

The cost of a physical cluster related to energy consumption has been computed considering the values in Table 1. The energy price C_U comes from the average value of the energy in the USA [19]. In some European countries (like Spain), the energy price can reach 0.20$ per kWh [20].

In order to obtain the actual power consumed by the essential components of the cluster (P_0) and by an idle node (P_I), we have relied on the monitoring systems of a physical cluster of PCs. To estimate the power consumption of an individual core we have used the well-known Linpack [21] benchmark, which has been executed with a different number of instances ranging from 1 to the maximum number of cores of a node whose features are described in Table 3. The average power consumption has been analysed and the increment of power with respect to the idle state (P_I) has been computed, taking into account the number of instances employed. Since the benchmark only stresses the CPU, this power consumption increment can only be attributed to the CPU. The resulting data is shown on Table 2.

As shown in the table, the power consumption per core is not linearly related to the number of cores. If we want to compute an upper bound of the power consumption, we can use the increment of power per core corresponding to the case when all the cores are being employed.

Concerning the price of purchasing a cluster, we have assumed a standard rack-based configuration, with the components detailed on Table 3, for which prices have been obtained through different vendors in early 2012.

To compute the cost of the cluster in the Cloud, the values in Table 4 have been employed, which reflect the pricing policies of Amazon EC2 as of early 2012 [17]. To make the comparison, the "Cluster Compute Eight Extra Large" instance (named cc2.8xlarge in Amazon EC2) has been used. It has been selected because it gets the best price/

Component	Unit Price
Computational nodes with two Quad-core Intel Xeon E5620 2.4Ghz processors, 16 GB of RAM and 146GB of hard disk	3600 $
Front-end node with similar features than the computational node	5000 $
Network by means of Gigabit Ethernet switches	2500 $
Rack	4000 $

Table 3: Configuration of a physical cluster

Item	Value
Front-end instance price ($/3 years)	300 $
Front-end instance price ($/hour)	0.013 $
Internal reserved instance price ($/3 years)	10490 $
Internal instance price ($/hour)	2.4 $
Internal reserved instance price ($/hour)	0.494 $

Table 4: Considered values to compute the cost of a cluster on Amazon EC2

performance ratio. This instance type, with 16 cores, provides 88 EC2 Compute Units (ECU). ECU has been used to compare the performance of the different instances, as it has become the "de facto" standard unit to measure the performance of Cloud systems. Other type of instance could be used depending on the concrete cluster needs and the specific required resources for each virtual node. Concerning the reservation of instances, Amazon considers different prices for low, medium and high use instances. With respect to low usage instances, high usage instances have a higher one-time reservation fee and lower hourly price. For simplicity, in this study we consider only high use instances, although results for low or medium use instances would not be substantially different.

Two different cases have been considered. On the one hand, a small-size cluster composed of 1 rack, 1 Gigabit Ethernet switch, 1 front-end, 8 nodes and 64 cores, with an approximate cost of 40,300 $. On the other hand, a large cluster composed of 4 racks, 6 Gigabit Ethernet switches, 1 front-end, 128 nodes and a total 1024 cores, with a total approximate cost of 496,800 $.

4.2. Comparing clusters

Using the previous data, a cost comparison analysis has been performed between a physical cluster (considering the lower and upper bounds in (12) and (13), respectively) and an EC2 virtual cluster of the same size, using the model in (15), from which lower and upper bounds can be derived using the minimum and maximum usage of the cluster. We consider not only the cluster size for the comparison, but also the performance. For that purpose we have run the Linpack benchmark both in an EC2 instance and in a physical cluster node equivalent to the computation node described in the previous section. The results and their normalized values are shown in Table 5, where we see that the performance of the physical cluster node and the cc2.8xlarge EC2 instance are quite similar, both obtaining a result about three times greater than the m1.small instance.

For the comparison, a period of 4 years has been considered, following the recommendations of [22]. The evolution of the cost with respect to the time is shown on Figure 2.

	MFlops (Linpack)	Normalized
Amazon EC2 (m1.small)	522.26	1
Real Node Xeon E5620 2.4Ghz	1516.52	2.903758
Amazon EC2 (cc2.8xlarge)	1614.90	3.092139

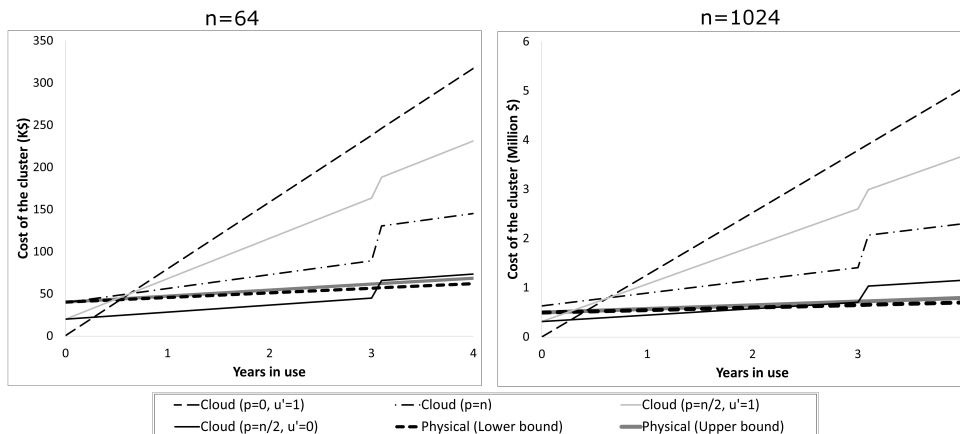Table 5: Performance value of the studied nodes (per core)

Figure 2: Cost comparison between the physical cluster and the cluster in the Cloud for a 64-core cluster (on the left), and for a 1024-core cluster (on the right) for 4 years

The left hand graph corresponds to the 64-core cluster, while the righ hand graph shows the data for the 1024-core cluster.

For each case, four different cases are shown, which correspond to different combinations of reserved nodes (p) and nonreserved usage (u') in the virtual cluster. The combinations illustrate the upper limit and lower limit of the virtual cluster cost for different cases of reserved instances ($p = 0$, $p = n/2$ and $p = n$). There are two special cases: a cluster that has not reserved instances and it is not used at all ($p = 0$ and $u' = 0$), and a cluster that has all its nodes reserved ($p = n$). In the first case, the cost is exactly 0 and it is not shown in the graph because it coincides with the horizontal axis. In the second case, both the upper and lower bounds of the costs are the same since there is no variable part in the cost.

In both graphs, the thick lines correspond to the lower bound (dashed black line) and the upper bound (solid gray line) of the physical cluster cost. The increment in cost with respect to time corresponds to the energy cost and the hardware maintenance. This cost can represent more than 70% of the initial cluster price after 4 years.

We can see that the virtual cluster cost on the left graph is approximately proportional to the same cost on the right graph, as is to be expected according to the expression (15). The differences between the graphs correspond to the scale of the prices and to the price of the physical cluster.

The graphs point out the influence of the saving plans of the Cloud provider (i.e. reserved instances) in the total cost of the cluster. In particular, if no reserved instances are pre-purchased, the price of a cluster in the Cloud shoots up to more than 350% (with

respect to the physical cluster) in the case of the small cluster ($n = 64$) and more than 500% in the case of the large cluster ($n = 1024$) for a maximum usage rate. However, using no reserved nodes could be better if the cluster usage rate is sufficiently low. This option is certainly interesting in the case of internet-based start-ups in which estimating the workload of the computing infrastructure is difficult, since this might depend on the success of the product.

Using reserved instances, it is possible to gain an economic advantage in the long term with respect to provisioning all the computing resources on demand. However, sometimes this means no real advantage with respect to the physical cluster. It is true that the hourly rates of reserved resources are lower than those of non-reserved ones. However, the initial investment in the reserved instances makes the total cost of the virtual cluster to easily surpass the cost of the physical cluster. In addition, the slope related to the hourly prices of the instances is steeper than the slope corresponding to the electricity and maintenance costs for the physical cluster. In this case we face the additional problem of having to decide the number of instances to purchase in advance This requires a careful planning and estimation of the computational resources to be needed in the near future. Otherwise, the user would incur in a penalty cost for having an overdimensioned cluster.

It is worth noting that the lines corresponding to the upper bound of the virtual cluster cost intersect in one single point that corresponds to about 7.6 months. In fact, using expression (15) and considering the costs corresponding to any two different values of p, we obtain that both costs are equal for time t_E given in (16), which depends only on u' and on the prices of reserved and nonreserved instances for the cluster working nodes.

$$t_E = \frac{C_{WR}}{C_W \cdot u' - C_{WU}} \tag{16}$$

In particular, the intersection point mentioned above corresponds to the case where $u' = 1$, which substituting in (16) yields $t = 5503.67$ hours, or approximately 7.64 months.

In general, expression (16) could help us decide if the option of reserving nodes is preferable for a particular case, assuming that we have enough information about the cluster usage rate.

5. Discussion

Deciding whether it is still convenient to purchase hardware to enable scientific computations in the shape of HPC (High Performance Computing), HTC (High Throughput Computing) or MTC (Many Task Computing) instead of outsourcing computations to an IaaS Cloud provider depends on several factors. Obviously, a cluster in the Cloud provides the user with the inherent benefits of the Cloud technology, such as avoiding the upfront investment in hardware, maintenance, cooling, etc. In addition, it also frees the user from setting up the space dedicated to host the cluster together with its refurbishing costs. This represents both a considerable investment, studied in publications such as [23], and an entry barrier that might delay the start of the operations of the cluster. In many cases, these delays are not admissible. Therefore, the Cloud could

be employed as the final computing infrastructure or as a transition solution until the physical infrastructure is set up.

We found that, in some cases, a cluster in the Cloud can deliver an amount of computational power comparable to a physical cluster with a similar cost. If the usage of the cluster is going to be very high during all its life-cycle, the physical cluster is the best option. But if the average usage of the cluster is going to be moderate or low, only having some periods with peak workloads, the Cloud can be a very good option. Note that it is important to correctly estimate the usage of the cluster in order to purchase the correct number of reserved instances to take advantage from commercial IaaS.

Another important aspect to be considered is that the hardware depreciates with time, due to the rapid obsolescence of the equipment and its delivered performance. Therefore, a physical cluster cannot be considered a middle or long term investment unless a plan of Return of Investment (ROI) exists. In this sense, one should expect public Cloud providers to periodically upgrade the underlying hardware (or reduce the price) so that users can benefit from the performance improvements. In the case of Amazon EC2 the hardware upgrade produced in the last years enabled an important reduction of the price per ECU (more than 75%) in the last years [24]. So it is logical to think that the price fall could continue in the next years. However, this fact has not been included in the proposed model due to the related uncertainty concerning a future price fall. Nevertheless, this should be considered by a customer when deciding among different Cloud providers.

It is also important to point out that estimating the cluster size is far from being a trivial task. If the cluster is underdimensioned, we face the risk of being unable to fulfill the computational requirements of the users. However, if the cluster is overdimensioned, we face an unnecessary cost for unused resources. In the case of a physical cluster, this decision is critical, since the dimension of the cluster relates to space for housing, energy supply, cooling systems, etc. In the case of a cluster in the Cloud, an inappropriate dimensioning of the cluster is less problematic since computational resources can be provisioned and released on demand, in order to satisfy unexpected peak workloads.

Another aspect to be considered is that one might think that a usage rate of 20% or 40% (73 or 146 days in a year) in a cluster is significantly low. Even though there are scientific applications that require sustained computing power for weeks or months, not many applications used in clusters have usage patterns over that usage ratio. This represents a total of 224,256 and 439,296 CPU hours for a 128-node cluster. For example, a node that our research group dedicates as part of the Spanish NGI (National Grid Initiative) has a usage rate of 33%.

One of the target users of a cluster in the Cloud would be a start-up (an enterprise with a high degree of uncertainty with respect to the computing requirements) or an organisation whose workloads match a high level rate during medium or small time periods (e.g., without exceeding a 20-40% sustained usage rate across the year). In these cases, the organizations avoid the upfront investment, reduce the Time to Market (TTM) and can postpone the decision of investing in a physical infrastructure, possibly depending on the ROI of the business activity. Such organisations could objectively map the Cloud costs to the price charged to the customers, and adjust the use of reserved instances according to the changing needs.

6. Conclusions

Nowadays, there are several IaaS public Cloud providers, which represent an alternative to the traditional purchase of computing infrastructure. However, the users of computer clusters have traditionally relied on physical clusters. This article has focused on the convenience of outsourcing cluster-based computations (HPC, HTC, MTC, etc.) to the Cloud.

For that purpose, a cost model of a physical cluster has been developed, which considers the different aspects studied in the literature, but also includes the cost related to energy consumption and the usage of energy saving strategies. We have used the pricing options of Amazon EC2 to create a cost model of a cluster in the Cloud. This model considers the reserved instances approach offered by the provider and it also considers the usage of tools to power on and off instances on demand, in order to avoid idle nodes on the Cloud and its associated cost.

A comparison has been made of the cost of physical clusters and their virtual counterparts, with the same number of nodes and similar performance. The result is that from the point of view of a data center, which expects a high usage rate for their clusters, it is still economically preferable to purchase a physical cluster hosted on its facilities. The fact is that these centers expect a ROI related to the usage and renting of their equipment and, therefore, can benefit from the economies of scale to turn themselves into infrastructure providers. But if the goal is obtaining high performance computing, and the sustained usage rate is moderate or low, the Cloud can provide similar equipment at a competitive price. Having a good estimation of the workload is necessary in order to select the correct number of reserved instances. We have also obtained an expression to help us decide whether the option of reserving nodes is economically preferrable, depending on the estimated usage rate of the cluster.

Our cost model does not include additional features such as spot instances, which enable the user to bid for unused Amazon EC2 capacity, since their cost is not deterministic. However, it could be an interesting option when performing HTC computations. A scenario could be envisaged in which the cluster grows and shrinks opportunistically according to the instance prices.

References

[1] PRACE Research Infrastructure.
    URL http://www.prace-ri.eu
[2] P. Mell, T. Grance, The NIST Definition of Cloud Computing. NIST Special Publication 800-145 (Final), Tech. rep. (2011).
    URL http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf
[3] L. Nussbaum, F. Anhalt, O. Mornard, J. Gelas, Linux-based virtualization for HPC clusters, in: Proceedings of the Linux Symposium, 2009.

[4] A. Bialecki, M. Cafarella, D. Cutting, O. O'Malley, Hadoop: a framework for running applications on large clusters built of commodity hardware, Tech. rep. (2005).
URL http://hadoop.apache.org

[5] P. Marshall, K. Keahey, T. Freeman, Elastic site: Using clouds to elastically extend site resources, in: 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing, IEEE Computer Society, 2010, pp. 43–52.

[6] G. Juve, E. Deelman, K. Vahi, Scientific workflow applications on Amazon EC2, Workshop on Cloud-based Services and Applications (5th IEEE International Conference on e-Science (e-Science 2009) (2010) 59–66.
URL http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5408002

[7] E. Deelman, G. Singh, M. Livny, B. Berriman, J. Good, The cost of doing science on the cloud: the montage example, in: 2008 ACM/IEEE conference on Supercomputing, Austin, Texas, 2008.

[8] M. De Assuncao, A. Di Costanzo, R. Buyya, Evaluating the cost-benefit of using cloud computing to extend the capacity of clusters, in: Proceedings of the 18th ACM international symposium on High performance distributed computing, ACM, Garching, Germany, 2009, pp. 141–150.
URL http://portal.acm.org/citation.cfm?id=1551635

[9] D. Kondo, B. Javadi, P. Malecot, F. Cappello, D. P. Anderson, Cost-benefit analysis of cloud computing versus desktop grids, in: 2009 IEEE International Symposium on Parallel & Distributed Processing, 2009.

[10] X. Li, Y. Li, T. Liu, J. Qiu, F. Wang, The method and tool of cost analysis for cloud computing, in: IEEE International Conference on Cloud Computing, 2009, 2009, pp. 93–100.

[11] M. Armbrust, A. Fox, R. Griffith, A. Joseph, Above the clouds: A berkeley view of cloud computing, Tech. rep., UC Berkeley Reliable Adaptive Distributed Systems Laboratory (2009).
URL http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.149.7163\&rep=rep1\&type=pdf

[12] Amazon, User Guide: Amazon EC2 Cost Comparison Calculator.
URL http://awsmedia.s3.amazonaws.com/User_Guide_Amazon_EC2_Cost_Comparison_Calculator.pdf

[13] M. Woitaszek, H. Tufo, Developing a cloud computing charging model for high-performance computing resources, in: IEEE 10th International Conference on Computer and Information Technology (CIT), 2010, pp. 210–217.

[14] J. Koomey, K. Brill, P. Turner, J. Stanley, B. Taylor, A simple model for determining true total cost of ownership for data centers, Tech. rep. (2008).
URL http://www.uptimeinstitute.org/wp_pdf/(TUI3011C)SimpleModelDetermingTrueTCO.pdf

[15] C. Belady, A. Rawson, Green grid data center power efficiency metrics: PUE and DCiE, Tech. rep., The Green Grid (2008).

[16] C. De Alfonso, M. Caballer, V. Hernández, Efficient Power Management in High Performance Computer Clusters, in: Proceedings of the International Conference on Green Computing 2010 (ICGreen 2010), 2010.

[17] Amazon, Amazon EC2 Pricing.
URL http://aws.amazon.com/ec2/pricing/

[18] Amazon, The Economics of the AWS Cloud vs. Owned IT Infrastructure, Tech. rep.
URL http://aws.amazon.com/economics

[19] U. E. I. Administration, Electricity Explained: Factors Affecting Electricity Prices.
URL http://www.eia.gov/energyexplained/index.cfm?page=electricity_factors_affecting_prices

[20] Energy Price in Spain.
URL http://www.boe.es/boe/dias/2011/12/31/pdfs/BOE-A-2011-20650.pdf

[21] J. Dongarra, LINPACK: users' guide, SIAM, 1979.

[22] J. Mahvi, A. Zarfaty, Using TCO to Determine PC Upgrade Cycles, Tech. rep., Intel (2009).
URL http://communities.intel.com/docs/DOC-3172/version/1

[23] W. Pitt, K. Brill, Cost Model: Dollars per kW plus Dollars per Square Foot of Computer Floor, Tech. rep., Uptime Institute (2008).

[24] K. Vermeersch, A Broker for Cost-efficient QoS aware Resource Allocation in EC2, Tech. rep., University of Antwerp (2011).

17