

On the Use of Cloud Computing for Scientific Workflows

Christina Hoffa¹, Gaurang Mehta², Timothy Freeman³, Ewa Deelman², Kate Keahey³, Bruce Berriman⁴, John Good⁴

¹Indiana University, ²University of Southern California, ³Argonne National Laboratory, ⁴Caltech

Abstract

This paper explores the use of cloud computing for scientific workflows, focusing on a widely used astronomy application-Montage. The approach is to evaluate from the point of view of a scientific workflow the tradeoffs between running in a local environment, if such is available, and running in a virtual environment via remote, wide-area network resource access. Our results show that for Montage, a workflow with short job runtimes, the virtual environment can provide good compute time performance but it can suffer from resource scheduling delays and wide-area communications.

1. Introduction

Recently, cloud computing has been under a growing spotlight as a possible solution to some of the problems with Grid computing, but the differences between cloud computing and Grid computing are often so diminished and obscured that they become indistinguishable from one another. The term “Grid” computing was coined in the early 1990's to liken a distributed computing infrastructure to the electrical power grid, although the concept has been around a bit longer [19, 23]. Like the electrical power grid, a computational Grid uses resources that are potentially very geographically far apart. These resources can be allotted to combinations of one or more virtual organizations, with the owners of the resources deciding when and to whom they should be allotted. In this manner, collaborations can integrate pools of resources to give supercomputer-class capability to their users.

Just as resources can be spread out geographically, so too can the members of a virtual organization. A virtual organization is a group of individuals or institutions who share direct access to resources such as computers, data, and software according to a set of rules [13]. The problems with dynamic sharing of resources are what one would expect in such an environment. An available resource may not meet the needs of a particular virtual organization or a particular member of a virtual organization. As a result, resources may also become underutilized. Conversely, a specific resource may not be available at the time it is needed, leaving the user to wait or to seek an alternative. In addition, significant overhead is incurred to ensure software compliance on different clusters, and security is always an issue in any kind of networked environment.

In the past decade, many of these issues have begun to be addressed, but several of the problems still exist when working with a Grid. For example, security infrastructures have evolved and have been increasingly implemented as the use of Grid computing has grown, but issues with dynamic scheduling and planning still occur regularly. One solution that has come about in recent years is cloud computing [4, 23]. “The Cloud” refers to the typical depiction of the Internet as a mass of thin, branching lines spreading in all directions. Cloud computing is a set of virtual servers that work together through the Internet and can be dynamically managed, monitored, and maintained. Users are expected to develop their own images or use existing ones as an executable environment on the cloud. Using virtual machines (VMs) that can be configured before deployment has the potential to reduce inefficient resource allocation and excess overhead. A VM can create an environment on a resource that is configured independently from that resource, allowing multiple such environments to be deployed on the same resource at the same time. In this manner of separation, each environment is kept secure from any others. Because sharing can be much more flexible, this also can also increase resource utilization [14].

While reducing or eliminating some sources of overhead, cloud computing introduces a few others. A network must be set up, and an image must be generated or discovered before any computations can take place. In the case of a virtual cluster, several virtual nodes need to be launched within the same time frame and made aware of each other. Once all these requirements are complete, execution overheads must still be taken into account. In addition, monetary costs must be considered. A group can purchase a Grid cluster for a relatively high cost, but it has complete control and sole access to it. The cluster will be sufficient for a certain amount of time until the cluster machines either break or become obsolete and a new cluster must be purchased. In contrast, cloud allotments may be purchased for a relatively small cost and for a relatively small amount of time, but the cost is incurred every time resources are used.

One of the primary obstacles Grid users face today is that while a Grid offers access to many heterogeneous

resources, a user typically needs a very specific environment that is customized to support a legacy application, remains consistent across multiple executions, or supports the idiosyncrasies of many service configurations. Resource providers clearly cannot support the diversity of all the required environments while users are often unable to use what is available. Further, guaranteeing a desired resource allocation to a Grid user is difficult today as most of the available tools support only unstructured sharing which further limits the impact of Grid computing on science. Using virtualization in conjunction with Grid computing, known as “cloud computing” or Infrastructure-as-a-Service (IaaS), has been recognized as a potential solution to these problems [25, 26].

This paper looks at the differences between running scientific workflows on the cloud and running them on the Grid, using the Montage [7] application and Pegasus WMS, a scientific workflow software [10, 11]. Our perspective is not that of a computer scientist trying to analyze detailed performance metrics, but rather that of a domain scientist who can use a number of different computing environments to perform computations. In our scenario, four different sizes of workflows were run on individual virtual machines and on a virtual machine cluster on the Nimbus science cloud [1] as well as on a local cluster without an overlying virtual environment to explore the differences in overhead. Using these strategies, the usability of cloud computing for large-scale scientific experiments is explored.

2. Motivation

With some of the limitations and inefficiencies of Grid computing exposed, such as heterogeneous execution environment, difficulties of acquiring resources when needed, etc., it is surprising at first glance that large scale computing using a virtual cloud is not more widespread. However, cloud technology is gaining in prominence as the forerunners of mainstream computing are beginning to use their marketing power. Companies such as Google, IBM, Amazon, and Yahoo! are aiming to make supercomputing power available to the masses, not just a relatively tiny number of skilled users [19]. Services like online storage, social networking sites, e-commerce, and web based email are all examples of how cloud computing has already begun to become mainstream. In addition, users in scientific and research communities will have the opportunity to access resources as they need them and only pay for what they use instead of paying for a certain amount of time – scheduled in advance – on a certain number of resources. Large corporations as well as individuals will be able to use cloud resources as they please.

The potential applications of cloud computing are many: financial institutions, health care services, business enterprises, the list is endless. People are beginning to recognize the benefits of cloud computing in the form of services provided over the Internet by the earlier mentioned computer giants, but because of its newness, the extent of the benefits of cloud computing to the scientific community are largely unknown. Very little research has been conducted in comparison with what has been done with Grid technology.

The main goal of this work was to determine the applicability of cloud computing to large scale scientific workflows that are usually run on a Grid. Using the Pegasus WMS software, workflows from the Montage application were run both on a Grid and on a science cloud to test for overall differences in time and overhead costs.

3. Approach

In order to provide a comparison between running on the cloud and on the Grid, Montage workflows of varying sizes were run in three different environments. Combinations of individual virtual machines as well as a virtual cluster were run on a science cloud set up at the University of Chicago [1], and the results from these two environments were compared with the results of running workflows on a local machine and a local Grid cluster. Different methods were implemented to change each environment in order to improve the comparison between the three, such as limiting the number of jobs to be released onto the local Grid cluster since it used more nodes than the science cloud as well as running workflows using several virtual machines that were not set up as a virtual cluster. Much work has been done to date on detailed comparisons of various aspects of virtual machines, including differences between various virtualization approaches, performance of the I/O subsystem, etc [6, 9, 14, 17, 22, 24]. In this paper we aim at providing a comparison from the perspective of a scientist who needs to run the computations and has access to limited local resources and to a potentially much larger pool of virtual resources.

3.1 Experimental Setup

Before any workflows could be run on the Nimbus science cloud, the appropriate software needed to be installed on a virtual machine image. A single virtual machine was launched, and necessities including Globus [2] (GridFTP, and GRAM), xinetd, Condor, and Pegasus WMS were installed. The VM was terminated and saved to an image so that the software would be available for each VM that was deployed using that image.

Figure 1 below shows the process of creating and running a workflow from start to finish using a Grid cluster, a virtual cluster run on the science cloud, and a single virtual machine run on the science cloud. First, a directed acyclic graph XML file (DAX) was generated for each of the Montage workflows to be run. A DAX is a high-level resource-independent description of the applicable tasks and their dependencies. Using Pegasus WMS software to plan out the workflow and map each job to an appropriate resource or node, a directed acyclic graph (DAG) file was

In order to obtain the workflows that were run, the Pegasus mapper was used to map high-level workflow descriptions onto distributed resources. Pegasus stands for Planning for Execution in Grids and is used for numerous applications in many areas of science [8, 15, 16], including Montage, an astronomy application which creates science-grade image mosaics of the sky. The workflow mapper depends on Condor DAGMan, which launches workflow tasks and maintains the dependencies between them. Pegasus and Condor were both installed on a running virtual machine,




Figure 1: Experimental Setup.

created. This type of file contains specific information needed by Condor to submit each job to the correct resource. After planning, a command given by Pegasus allows the workflow to be submitted to Condor. Condor's directed acyclic graph manager (DAGMan) [3] sends jobs to Condor, which uses Condor-G, to be run on the sites that have been selected during planning.

3.2 Tools

Several tools were used to set up this experiment, the first of which are the Virtual Workspace Service and the workspace cloud client [4]. The Virtual Workspace Service is installed on a Grid, allowing users to create virtual environments using the workspace cloud client. For part of this experiment, Xen virtual machines were deployed using the cloud client on the University of Chicago's sixteen-node TeraPort cluster that had the Virtual Workspace Service installed to create the Nimbus science cloud [1]. The cloud client allows a user to upload virtual machine images, download, modify, delete, or save copies of preexisting images, deploy images as virtual machines, and deploy virtual clusters from virtual images.

and then the VM was saved as an image that could be used to deploy future VMs. A convenient benefit of using the cloud is that once an image is developed, it is not necessary to reinstall software each time a new VM is launched. The configurations that are set while a VM is running can be saved, and the image can be used to launch multiple copies of the original VM.

Parts of the Globus Toolkit [2] were used in combination with Pegasus and Condor to run the workflows, most notably GridFTP [5] and GRAM. GridFTP is a high-speed transfer service that expands the File Transfer Protocol (FTP) to include features needed for Grid computing applications [12]. Globus Resource Allocation Manager, or GRAM, provides a single protocol for communicating with different job schedulers. It is used to locate, submit, monitor, and cancel jobs on remote resources while addressing security, reliability, and performance issues. This experiment used a pre-Web Services GRAM service.

3.3 Environments

We ran our experiments in four different environments: 1) on a local machine, 2) on multiple VMs, where each VM is

treated as an independent computational site, 3) on a local cluster, and 4) on a Virtual cluster. We note that the virtual machines each have 2 CPUs thus we compared the execution not based on machines but based on CPU counts. Table 1 shows the characteristics of the machines we used in our experiments.

The following graphs show the execution times for the workflows of various sizes on resources with 1, 2, and 4 CPUs. The runtimes do not include the set and teardown of the virtual environment. The local execution represents execution of the workflow on the local host. Figure 2 shows the execution of the workflow on 2 CPUs (for

Table 1: The type of machines used in the study.

	Processor Type	Memory	Architecture	Disk space	Schedule r
Local Machine	Pentium, 2.2Ghz	1GB	32 bit	> 100GB	fork
VM	AMD 248, 2.2 Ghz	3.5GB	64 bit	1.0GB free local	fork
Local Cluster	Xeon	2.4GB	32 bit	> 100 GB	pbs
VM Cluster	AMD 248, 2.2 Ghz	3.5GB	64 bit	1.5GB free shared	pbs

The local Grid cluster used was comprised of eight compute nodes, and it had GRAM, GridFTP, and PBS software installed on it. Individual Xen virtual machines were deployed on nodes of the University of Chicago TeraPort cluster. The TeraPort cluster is composed of sixteen nodes, but each VM was only run on one node. GridFTP, GRAM, and the other required software was installed on a template image before deployment. The third environment on which the workflows were run was a virtual cluster. The virtual cluster was composed of two to five VMs: a head node and one to four worker nodes. The VMs each occupied one node of the TeraPort cluster, just as with a single VM.

4 Experimental Results and Analysis

We conducted the experiments using 4 different size mosaics: 0.1 deg², 0.2 deg², 0.5 deg², and 1deg², corresponding to workflows with 50, 74, 150, and 426 application jobs respectively. There are also additional jobs in the workflow, inserted by Pegasus, to perform the data movement if any is necessary, to cleanup the data as the computation progresses [18, 20], and to register the newly created data products. Another important metric in the study is the total footprint of the workflow. This footprint is 58, 122, 338, and 1,090MB for increasing mosaic sizes. The application executables add an additional 11MB to total footprint of each workflow. In all cases input data was available on a machine on the same local area network as the local machines. The virtual environments were connected to the local submit host via a wide area network. As a result, data had to be automatically stage-in by Pegasus-WMS.

comparison we also include the local execution on 1 CPU). The local cluster was configured to use no more than 2 CPUs at any one time. We used 1 VM with 2 CPUs, and a virtual cluster with one headnode that runs PBS and one worker node (with 2 CPUs). We note that the 1 degree square runs conducted on individual VMs did not complete due to lack of disk space in the image. We also notice that the performance of a single processor machine outperforms that of other configurations, even though the local host is less powerful than others. Although we don't see that in these experiments, the size of the mosaics one can run on a single machine is limited by the amount of memory on that machine. Additionally, the number of simultaneous mosaics that can be executed locally is limited. As a result relying on a single machine to do many large-scale mosaics is not feasible in general. Figure 3 shows the execution of the workflow on 4 CPU (and includes the local 1 processor execution for comparison). Again we see that local execution is the shortest. In both the 2CPU and 4 CPU case, the virtual cluster performs better than the local cluster. In the 4CPU case individual VMs perform similarly to the local cluster. In order to understand the reason for good performance of the local host versus the performance of the other platforms, we examined the constituents of the end-to-end runtime. These are composed of the Computation Time (here it also includes data staging in and out), the time it takes DAGMan to release jobs ready to run into the Condor queue, the time jobs wait in the Condor queue, and the time the jobs wait in the queue in the scheduler at the remote resource. Figure 4 shows the breakdown of the quantities for the 0.5 degree square Montage workflow running on 4 CPUs. We see that during local execution, the computational time is




Figure 2: Runtime of the Montage Application on 2 CPUs.




Figure 3: Runtime of the Montage Application on 4 CPUs.




Figure 4: Overheads in the Workflow Execution in Various Environments.

long—because the processor is slow. The Figure also shows that the DAGMan overhead is constant for all the runs. What is different is the overhead associated with the Condor queue. Once we are outside of the local machine, jobs need to go from the Condor queue across a network to the resource. Jobs incur the network delay as well as the delay related to the acknowledgment of the remote resource that the job has been accepted. We can see that this delay is particularly long for the virtual cluster. In our runs, the jobs also wait longer in the remote queue in the virtual cluster case. There is virtually no remote queue time seen in the virtual machine because these are managed by fork. We also analyzed the computational time of the individual VMs because it was longer than that of the virtual cluster. The reason for the difference is that the cluster uses a shared file system to access data during the workflow execution, whereas the individual VMs need to share data via GridFTP transfers. This time constitutes 6% of the overall computation time.

5. Discussion and Conclusions

In this paper, we aimed to quantify the performance differences an application scientist would see while running workflows on various types of resources, including both physical and virtual environments. Although we see that in some cases running in a local environment is sufficient, it is not a scalable solution. However, virtual environments can provide the necessary scalability. In multi-resource environments, we noticed large overheads of jobs waiting in the Condor and resource queues. This is particularly exacerbated in applications such as montage where the average execution time of a job is on the order of a few seconds. However, clustering techniques [21] that cluster small jobs together can greatly reduce the scheduling overheads. In the future, we plan apply these techniques in our virtual environment evaluation studies. Finally, we also noticed that disk space management is critical in virtual environments. When a virtual image is created, the size of the disk is fixed. Having a too small initial virtual disk size can adversely affect the execution of the application.

References

- [1] "Nimbus Science Cloud." <http://workspace.globus.org/clouds/nimbus.html>
- [2] "Globus." <http://www.globus.org>
- [3] "DAGMAN (Directed Acyclic Graph Manager)," 2004. <http://www.cs.wisc.edu/condor/dagman/>
- [4] Virtual Workspaces <http://workspace.globus.org/>
- [5] W. Allcock, et al., "Data Management and Transfer in High-Performance Computational Grid Environments," *Parallel Computing*, 2001.
- [6] P. Barham, et al., "Xen and the art of virtualization," *Proceedings of the nineteenth ACM symposium on Operating systems principles*, pp. 164-177, 2003.
- [7] G. B. Berriman, et al., "Montage: A Grid Enabled Engine for Delivering Custom Science-Grade Mosaics On Demand," in *SPIE Conference 5487*., 2004.
- [8] D. A. Brown, et al., "A Case Study on the Use of Workflow Technologies for Scientific Analysis: Gravitational Wave Data Analysis," in *Workflows for e-Science*, Springer, 2006.
- [9] B. Clark, et al., "Xen and the art of repeated research," *USENIX Annual Technical Conference, FREENIX Track*, pp. 135–144, 2004.
- [10] E. Deelman, et al., "Pegasus: a Framework for Mapping Complex Scientific Workflows onto Distributed Systems," *Sci. Prog. J.*, vol. 13, pp. 219-237, 2005.
- [11] E. Deelman, et al., "Pegasus: Mapping Large-Scale Workflows to Distributed Resources," in *Workflows in e-Science*, Springer, 2006.
- [12] M. Feller, et al., "GT4 GRAM: A Functionality and Performance Study," *TeraGrid Conference*, 2007.
- [13] I. Foster, et al., "The Anatomy of the Grid: Enabling Scalable Virtual Organizations," *IJHPCA*, vol. 15, pp. 200-222, 2001.
- [14] I. Foster, et al., "Virtual Clusters for Grid Communities," *Cluster Computing and Grid (CCGRID)*, 2006.
- [15] P. Maechling, et al., "SCEC CyberShake Workflows---Automating Probabilistic Seismic Hazard Analysis Calculations," *Workflows for e-Science*, Springer, 2006.
- [16] V. Nefedova, et al., "Automating Climate Science: Large Ensemble Simulations on the TeraGrid with the GriPhyN Virtual Data System," in *e-Science*, 2006.
- [17] A. Norton, et al., "Evaluation of Virtual Machines for HEP Grids," *Computing in High Energy Physics, Mumbai*, 2006.
- [18] A. Ramakrishnan, et al., "Scheduling Data-Intensive Workflows onto Storage-Constrained Distributed Resources," in *CCGrid 2007* 2007.
- [19] A. Ricalde, "Computing Heads for the Clouds," in *Business Week*, November 16, 2007. http://www.businessweek.com/technology/content/nov2007/tc20071116_379585.htm
- [20] G. Singh, et al., "Optimizing Workflow Data Footprint " *Sci Prog. J.*, vol. 15, 2007
- [21] G. Singh, et al., "Workflow task clustering for best effort systems with Pegasus," *ACM Mardi Gras conference: 2008*.
- [22] VMWare, "A Performance Comparison of Hypervisors."
- [23] A. X3J3/S8.115, "Fortran 90," 1990.
- [24] L. Youseff, et al., "Gorda, and Chandra Krintz. Paravirtualization for HPC systems," *ISPA Workshops*, vol. 4331, pp. 474–486.
- [25] R. Figueiredo, P. Dinda, J. Fortes. "A Case for Grid Computing on Virtual Machines." *23rd International Conference on Distributed Computing Systems*. 2003.
- [26] K. Keahey, K. Doering, I. Foster. "From Sandbox to Playground: Dynamic Virtual Environments in the Grid." *5th International Workshop in Grid Computing*. 2004.