

HPC & Cloud Computing



Giuseppe Fiameni - g.fiameni@cineca.it

SuperComputing Application and Innovation - CINECA

INFN CCR GARR Workshop

May 15th, 2012 - Napoli

- Cloud Definition
- Scalability vs. Performance
- Barriers of Adoption
- Performance Evaluation
- CINECA's perspective
- Conclusions

- Self Service Dynamically Scalable Computing Facilities - Cloud computing is not about new technology, it is about new uses of technology
- Main innovation to be associated with clouds consists in the sale of resource capabilities to external customers, as opposed to internal purposes only
- Important step in "*utility computing*" as envisaged by John McCarthy back in 1961
- *In particular the main interests in cloud systems, such as elasticity and availability, do not compare to the main interests in HPC systems, which focus on reliability and performance*

- **Broad network access**
 - Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms
- **Resource pooling**
 - The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
- **Rapid elasticity**
 - Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.
- **Measured service**
 - Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.
- **On-demand self-service**
 - A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

- **Broad network access (OK)**
 - Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms
- **Resource pooling (OK)**
 - The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
- **Rapid elasticity (PARTIALY)**
 - Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.
- **Measured service (OK)**
 - Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.
- **On-demand self-service (PARTIALY)**
 - A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

- **Rapid elasticity (PARTIALY)**
 - Capabilities can be elastically provisioned and released, in some cases automatically, to **scale** rapidly outward and inward commensurate with demand.
- Scalability (Cloud) \neq Performance (HPC)
- **Performance** is the capability of particular component to provide a certain amount of capacity, throughput
- **Scalability**, in contrast, is about the ability of a system to expand to meet demand. This is quite frequently measured by looking at the aggregate performance of the individual components of a particular system and how they function over time
- **Performance** measures the capability of a single part of a large system while **scalability** measures the ability of a large system to grow to meet growing demand. Thus scalable systems may have individual parts that are relatively low performing

- **Rapid elasticity (PARTIALY)**
 - Capabilities can be elastically provisioned and released, in some cases automatically, to **scale** rapidly outward and inward commensurate with demand.
- Most HPC clusters today are built out for very high performance *and* scalability, but with a particular focus on performance of individual components (servers) and the interconnect. The price/performance of the overall system is not as important as aggregate throughput of the entire system
- Cloud & High Scalability Computing Cloud, or HSC, by contrast, focuses on hitting the price/performance sweet spot, using truly commodity components and buying *lots* more of them. This means building very large and scalable systems

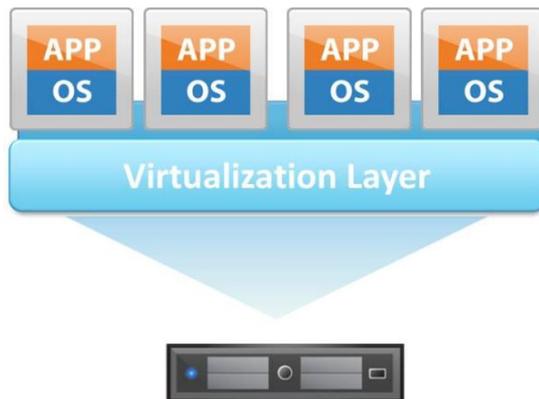
- **Rapid elasticity (PARTIALY)**
 - Capabilities can be elastically provisioned and released, in some cases automatically, to **scale** rapidly outward and inward commensurate with demand.
- Cloud computing offers horizontal scalability – *you have a single application and you replicate it many times, to ensure availability* - it's the same data and application that is replicated over the infrastructure.
- In HPC, we don't want multiple instances of the application, as that doesn't improve their performance. Instead, the application itself distributes over multiple instances – one single application takes a thousand cores (vertical scale)
- Horizontal Scalability (Cloud) vs. Vertical Scalability (Cloud)

- **On-demand self-service**

- A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

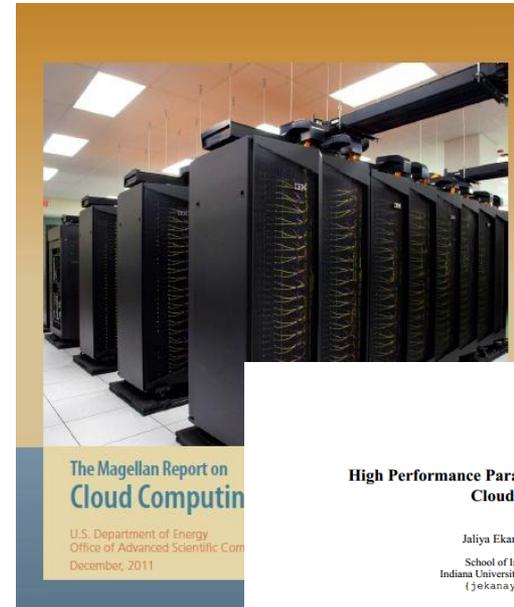
- HPC Systems are not able to provide this capability although it might bring high benefits for improving user's experience
- *"It is too complicated and takes too much effort to port my application on your supercomputer"*
- *"My workflow is not suited for such static environment"*
- *"My application does not run on your Linux ;-("*
- *"I want to choose my own operating system"*
- *"I want to use the Cloud, it is and cool!"*

«The virtualization is the “secret sauce” of cloud and “overbooking” is the “secret sauce” of virtualization»



- **Performance issues**
 - Virtualization tolls CPU and especially I/O
 - Cloud networks are not designed for low-latency communication
- **Data issues**
 - HPC can consume or produce enormous data volumes
 - Need to move them in and out of cloud, which leads to latency & cost
- **Vendor-related issues**
 - Memory caps (currently ~16 GB) limits some shared memory jobs
- **Security issues**
 - Data privacy, availability and integrity
 - Private data moving over WAN

- Hardware contention between VMs causes a compute time increment, this happens when the small amount of individual resources available lead to memory sharing between different virtual machines
- Communication overhead between the hypervisor and the VMs. Communication overhead between the hypervisor and the VMs increases more that proportional with respect to the running VMs. This happens especially in para-virtualized architectures (i.e. Xen), where VMs are not directly able to perform I/O operations but an interaction with the hypervisor is needed
- Different tests have been performed using MPI applications with different characteristics (Matrix Multiplication and Current Wave Equation Solver) and it was discovered that loosely coupled applications perform better than those with a high level of inter-process communication, this effect is even more evident in public clouds



High Performance Parallel Computing with Clouds and Cloud Technologies

Jaliya Ekanayake and Geoffrey Fox
School of Informatics and Computing,
Indiana University, Bloomington, IN 47405, USA
{jekanaya,gcf}@indiana.edu

Abstract. Infrastructure services (Infrastructure-as-a-service), provided by cloud vendors, allow any user to provision a large number of compute instances fairly easily. Whether leased from public clouds or allocated from private clouds, utilizing these virtual resources to perform data/compute intensive analyses requires employing different parallel runtimes to implement such applications. Among many parallelizable problems, most “pleasingly parallel” applications can be performed using MapReduce technologies such as Hadoop, CCL-MapReduce, and Dryad, in a fairly easy manner. However, many scientific applications, which have complex communication patterns, still require low latency communication mechanisms and rich set of communication constructs offered by runtimes such as MPI. In this paper, we first discuss large scale data analysis using different MapReduce implementations and then, we present a performance analysis of high performance parallel applications on virtualized resources.

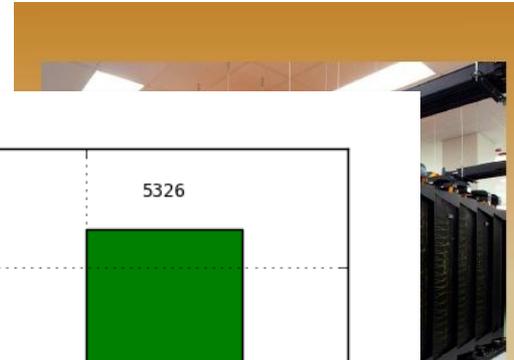
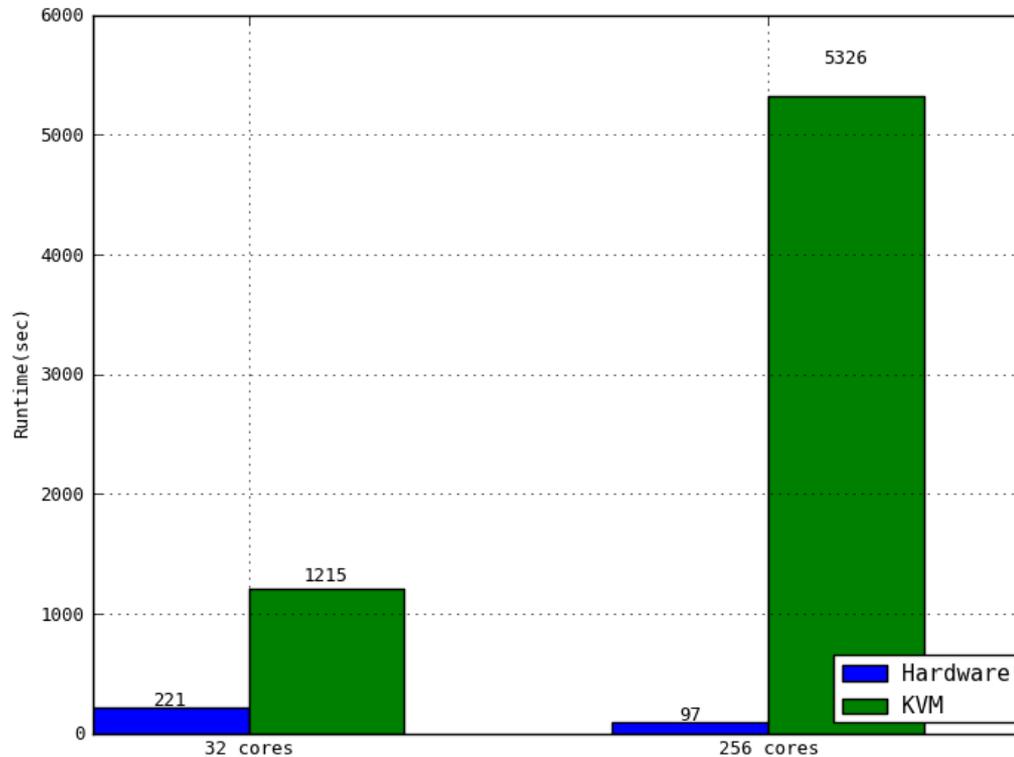
Keywords: Cloud, Virtualization, MapReduce, Dryad, Parallel Computing.

1 Introduction

The introduction of commercial cloud infrastructure services such as Amazon EC2/S3 [1-2] and GoGrid[3] allow users to provision compute clusters fairly easily and quickly by paying a monetary value only for the duration of the usage of resources. The provisioning of resources happens in minutes as opposed to the hours and days required in the case of traditional queue-based scheduling systems. In addition, the use of such virtualized resources allows the user to completely customize the Virtual Machine (VM) images and use them with root/administrative privileges, which is another feature that is hard to achieve with traditional infrastructures. The availability of open source cloud infrastructure software such as Nimbus [4] and Eucalyptus [5], and the open source virtualization software stacks such as Xen Hypervisor[6], allows organizations to build private clouds to improve the resource utilization of the available computation facilities. The possibility of dynamically provisioning additional resources by leasing from commercial cloud infrastructures makes the use of private clouds more promising.

With all the above promising features of cloud, we can assume that the accessibility to computation power is no longer a barrier for the users who need to

- Hardware contention between VMs causes a compute time increment, this happens when individual memory shares in a virtual machine are exceeded.
- Communication overhead in a hypervisor. Communication overhead in a hypervisor that propagates to running VMs (e.g. Xen), when they need to perform interaction with the hardware.
- Different test configurations using MPI characteristics and Current and it was coupled applications than those with a high level of inter-process communication, this effect is even more evident in public clouds.



Performance Parallel Computing with Clouds and Cloud Technologies

Jaliya Ekanayake and Geoffrey Fox
School of Informatics and Computing,
Indiana University, Bloomington, IN 47405, USA
{jekanaya,gcf}@indiana.edu

Infrastructure services (Infrastructure-as-a-service), provided by clouds, allow any user to provision a large number of compute instances. Whether leased from public clouds or allocated from private clouds, provisioning these virtual resources to perform data/compute intensive applications requires employing different parallel runtimes to implement such applications. Among many parallelizable problems, most "pleasingly parallel" can be performed using MapReduce technologies such as Hadoop, Dryad, and Dryad, in a fairly easy manner. However, many applications, which have complex communication patterns, still require communication mechanisms and rich set of communication libraries. In this paper, we first discuss large scale analysis using different MapReduce implementations and then, we present a performance analysis of high performance parallel applications on clouds.

Cloud, Virtualization, MapReduce, Dryad, Parallel Computing.

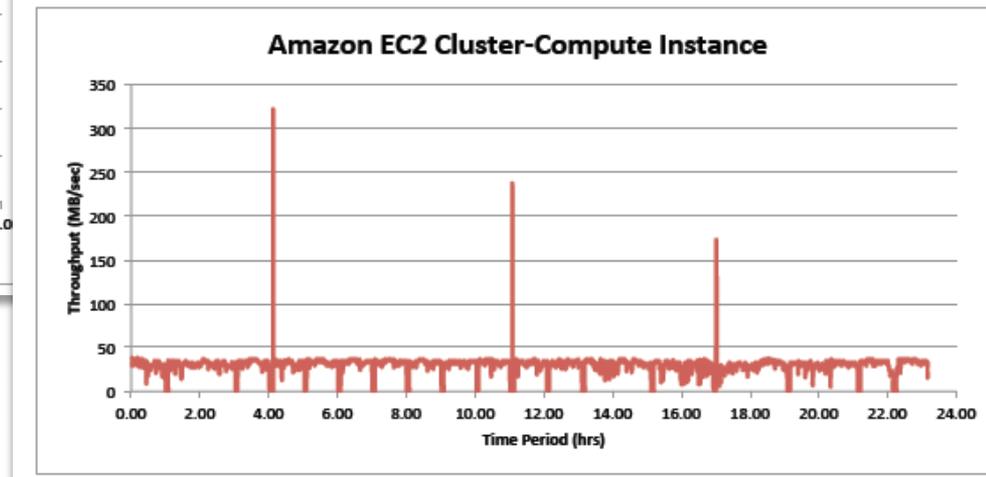
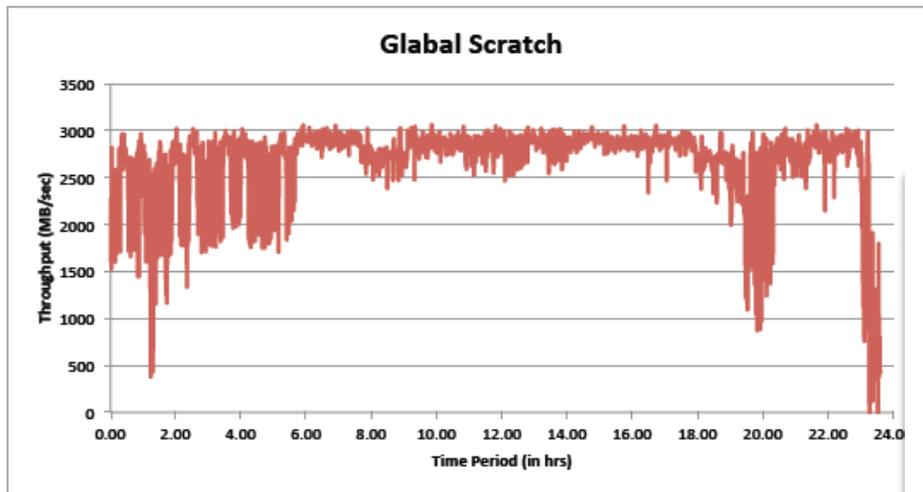
Introduction

Commercial cloud infrastructure services such as Amazon EC2/S3 [3] allow users to provision compute clusters fairly easily and at a monetary value only for the duration of the usage of resources. Provisioning of resources happens in minutes as opposed to the hours and days required in the case of traditional queue-based job scheduling systems. In addition, the use of such virtualized resources allows the user to completely customize the Virtual Machine (VM) images and use them with root/administrative privileges, which is another feature that is hard to achieve with traditional infrastructures.

The availability of open source cloud infrastructure software such as Nimbus [4] and Eucalyptus [5], and the open source virtualization software stacks such as Xen Hypervisor [6], allows organizations to build private clouds to improve the resource utilization of the available computation facilities. The possibility of dynamically provisioning additional resources by leasing from commercial cloud infrastructures makes the use of private clouds more promising.

With all the above promising features of cloud, we can assume that the accessibility to computation power is no longer a barrier for the users who need to

- The I/O performance results clearly highlight that I/O can be one of the causes for performance bottleneck on virtualized cloud environments.
- Performance in VMs is lower than on physical machines, which may be attributed to an additional level of abstraction between the VM and the hardware.



Global Scratch: IBM's GPFS with a peak performance of approximately 15 GB/sec

- **PCI Pass-through** grants a VM direct access to a dedicated HCA (Host Channel Adapter). It requires an I/O Memory Mapping Unit (IOMMU) to ensure memory protection between different VMs. A guest OS uses regular drivers
- **Para-Virtualization** for IB has been proposed by Liu et al. It requires ongoing modifications of drivers in host and guest with respect to changes of the underlying hardware and OS
- **Single Root-I/O Virtualization (SR-IOV)** is a standard for virtualization support in hardware. It allows a PCI Express device to appear as multiple virtual devices which guests can access via PCI Pass-through

I/O Performance of Virtualized Cloud Environments

Devarshi Ghoshal
Indiana University
Bloomington, IN 47405
dghoshal@cs.indiana.edu

R. Shane Canon
Lawrence Berkeley National
Lab
Berkeley, CA 94720
scanon@lbl.gov

Lavanya Ramakrishnan
Lawrence Berkeley National
Lab
Berkeley, CA 94720
lramakrishnan@lbl.gov

ABSTRACT

The scientific community is exploring the suitability of cloud infrastructure to handle High Performance Computing (HPC) applications. The goal of Magellan, a project funded through DOE ASCR, is to investigate the potential role of cloud computing to address the computing needs of the Department of Energy's Office of Science, especially for mid-range computing and data-intensive applications which are not served through existing DOE centers today. Prior work has shown that applications with significant communication or I/O tend to perform poorly in virtualized cloud environments. However, there is a limited understanding of the I/O characteristics of cloud environments. This paper will present our results in benchmarking the I/O performance over different cloud and HPC platforms to identify the major bottlenecks in existing infrastructure. We compare the I/O performance using IOR benchmarks on two cloud platforms - Amazon and the Magellan cloud testbed. We analyze the performance of different storage options available on different instance types in multiple availability zones. Finally, we do some custom benchmarking in order to analyze the variability in the I/O patterns over time and region. Our results highlight the performance of the different storage options enabling applications to make effective storage option choices.

1. INTRODUCTION

Data is a critical component of next-generation scientific processes. Scientific processes are generating and analyzing large data sets to derive scientific insights. Cloud computing technologies have largely evolved to process and store large data volumes of web and log data. In the last few years, there has been an increasing interest in evaluating the use of cloud technologies to meet the needs of scientific applications. Several groups have run both standard benchmark suites such as Linpack and NAS [15, 3, 13, 17, 22], and network performance tests [23]. Previous results have shown that the communication-intensive applications do poorly in these environments. However there is limited understanding of the I/O performance in virtualized cloud environments. Understanding the I/O performance is critical to understand the performance of scientific applications in these environments.

The goal of the Magellan project is to evaluate the ability of cloud computing to meet the needs of DOE workloads. We have previously benchmarked various cloud platforms to quantify their performance for scientific applications [7]. In this paper, we evaluate the various cloud I/O offerings

and understand their performance characteristics relative to current HPC centers where these scientific applications run.

I/O is commonly used in scientific applications: to store output from simulations for later analysis; to implement algorithms that process more data that can fit in system memory and must page data to and from disk; and for checkpointing to save the state of application in case of system failure. HPC systems are typically equipped with a parallel file system such as Lustre or GPFS that can stripe data across large numbers of spinning disks connected to a number of I/O servers to provide scalable bandwidth and capacity. These file systems also allow multiple clients to concurrently write to a common file while preserving consistency. On systems such as NERSC, often there are two file systems available: local and global. Local file systems accessible on a single platform typically provide the best performance whereas global file systems simplify data sharing between platforms. These filesystems are tuned for achieving high performance that is desired by these scientific applications. Thus it is critical to understand the I/O performance that can be achieved on cloud platforms in order to understand the performance impact on scientific applications that are considering these platforms.

In this paper, we evaluate a public cloud platform and the private cloud platform available on the Magellan testbed. We select Amazon as our public cloud platform since it is currently the most popular Infrastructure-as-a-Service (IaaS) cloud offering. We benchmark three instance types - small, large and Cluster Compute, the specialized HPC offering. The Magellan virtual machine testbed runs the Eucalyptus 2.0 cloud software stack on top of KVM and uses virtio for disk access.

We used IOR [6] benchmarks and a custom timed benchmark for analyzing the I/O performance on clouds. We compare the performance of different instance types, both local and block store and different availability regions on Amazon to understand the spectrum of I/O performance. Specifically, we make the following contributions:

- We compare the I/O performance of various storage options in virtual machines and instance types provided through private cloud software stacks as well as commercial providers such as Amazon,
- We also study the variability of the performance over multiple runs and multiple time periods on Amazon.

The rest of this paper is organized as follows. In Section 2 we discuss related work. We present our methodology for

- Novel architecture for such HPC clouds based on the IB cluster interconnect
- Currently working on completing a prototypic HPC Cloud, which will incorporate SR-IOV IB access. (SR-IOV drivers still not available)
- VMs with access to IB cannot be live-migrated in a transparent way. Anyway it must be considered that providing this capability for virtual machine-based systems could lead to several challenges in terms of security, scalability and stability

Virtual InfiniBand Clusters for HPC Clouds

Marius Hillenbrand
System Architecture Group
KIT, Germany
m.hillen@kit.edu

Viktor Mauch
Steinbuch Centre for
Computing
KIT, Germany
mauch@kit.edu

Jan Stoess
System Architecture Group
KIT, Germany
stoess@kit.edu

Konrad Miller
System Architecture Group
KIT, Germany
miller@kit.edu

Frank Bellosa
System Architecture Group
KIT, Germany
bellosa@kit.edu

ABSTRACT

High Performance Computing (HPC) employs fast interconnect technologies to provide low communication and synchronization latencies for tightly coupled parallel compute jobs. Contemporary HPC clusters have a fixed capacity and static runtime environments; they cannot elastically adapt to dynamic workloads, and provide a limited selection of applications, libraries, and system software. In contrast, a cloud model for HPC clusters promises more flexibility, as it provides elastic virtual clusters to be available on-demand. This is not possible with physically owned clusters.

In this paper, we present an approach that makes it possible to use InfiniBand clusters for HPC cloud computing. We propose a performance-driven design of an HPC IaaS layer for InfiniBand, which provides throughput and latency-aware virtualization of nodes, networks, and network topologies, as well as an approach to an HPC-aware, multi-tenant cloud management system for elastic virtualized HPC compute clusters.

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network Management—*InfiniBand, Isolation*; D.4.4 [Operating Systems]: Communications Management—*InfiniBand Virtualization*

General Terms

Design

Keywords

HPC, InfiniBand, Cluster, Cloud Computing, Virtualization

1. INTRODUCTION

Today's High Performance Computing (HPC) clusters are typically operated and used by single organizations. A cluster

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CloudCP 2012: 2nd International Workshop on Cloud Computing Platforms, Bern, Switzerland
Copyright 2012 ACM 978-1-4503-1161-8...\$10.00.

operator deploys a predefined, fixed runtime and systems environment in a known-good configuration. Such physically owned clusters have two major drawbacks: First, the resource demands of HPC workloads are often fluctuating, leaving the cluster under-utilized or overloaded. Second, application developers are faced with a restricted runtime environment, which allows them at best to change application libraries, but not the underlying operating system or core runtime libraries, such as the job scheduler.

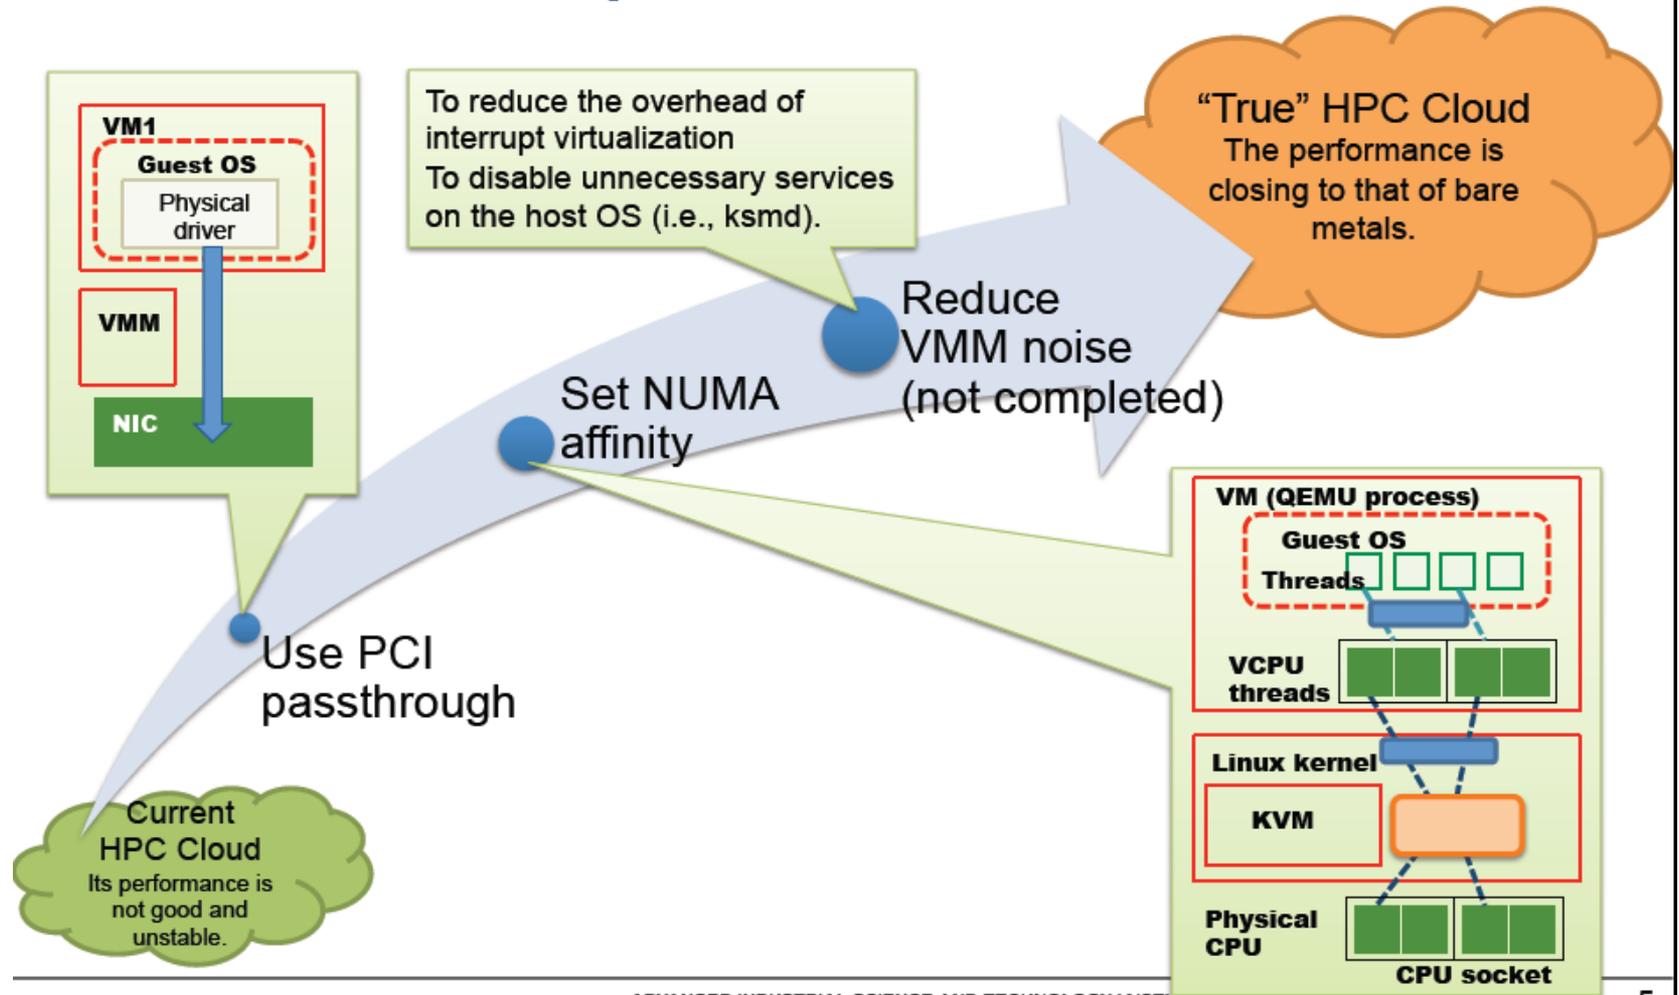
HPC in the Cloud promises increased flexibility and efficiency in terms of cost and energy consumption. Large providers of HPC Infrastructure as a Service (IaaS) can reduce personnel costs through high degrees of automation and a better overall utilization than in privately operated clusters. For end users, elastic virtual clusters provide precisely the capacity that suits demand, without the need to purchase and operate own hard- and software. Since virtual resources can be fully granted to users, they can choose or customize the OS or runtime environment according to their demands.

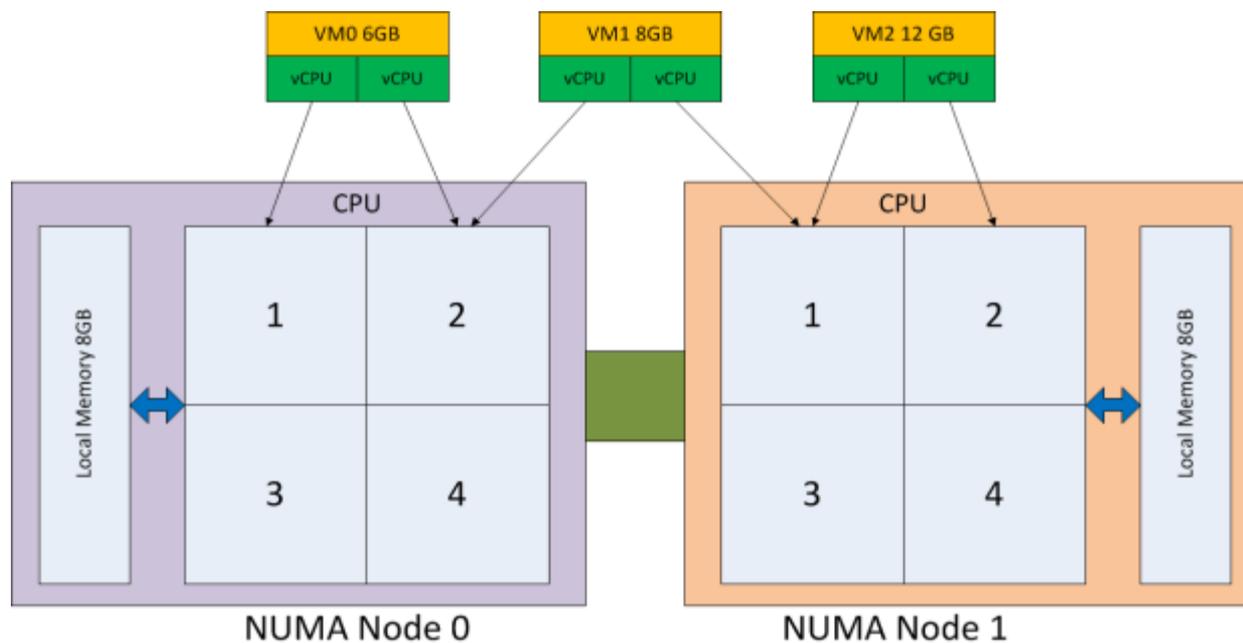
As an example, consider a team of researchers using fluid dynamics simulations with varying demand for compute capacity: They employ coarse test runs at the beginning of their project and highly detailed simulations when completing their publication (continuous high demand). With an HPC cloud, they could scale the capacity of their virtual cluster according to the stage of their project. During periods of little or no use, there would be little or no running expenses. Pay-per-use allows to associate the costs directly with the resource usage (and the project's budget). The researchers use an experimental simulation package with special libraries and therefore customized the software on their virtual nodes.

The Challenge of HPC-as-a-Service

HPC applications differ from standard cloud workloads in their much higher demands on the underlying resources and their guaranteed and timely delivery. These requirements, however, are particularly hard to achieve in a virtualized environment with its high I/O-overhead and jitter. It is unclear, whether the elasticity and standardization of cloud environments, which are both achieved by means of virtualization, can be achieved along with the predictability associated with HPC. Further, and in contrast to server workloads, HPC jobs are typically CPU-intensive and task synchronization points (e.g., in the MPI communication library) often require equal-paced CPU resources on all available cores. Therefore,

Toward a practical HPC Cloud





- **VM0** will be fine as each core will have sufficient local memory available.
- **VM1** should never get assigned cores in different NUMA nodes because a NUMA aware hypervisor should only assign a VM to a single NUMA node.
- **VM2** will have NUMA memory fragmentation that could affect performance because there is insufficient local memory to satisfy the 12GB requirement

- Loosely coupled scientific applications with minimal communication and I/O are best suited for clouds
- Providing VMs with direct access to parallel file systems could increase I/O performance in cloud systems, although these kind of solutions require to be carefully evaluated before being adopted in production
- The cloud is set up for non-time-critical processing, it's a way of getting work done quickly without buying a lot of hardware. **It was never designed to replace High Performance Computing**
- Scientific applications have special requirements that require cloud solutions that are tailored to these needs
- Clouds can require significant programming and system administration support
- Public clouds can be more expensive than in-house large systems. Many of the cost benefits from clouds result from the increased consolidation and higher average utilization
- Cloud is a business model and can be applied at any supercomputing centers

Type	Examples	Requirements
Compute Intensive	Monte Carlo simulations, Embarrassing parallel	CPU Cycles
Data Intensive	Signal/Image processing, Pattern matching	Fast I/O to data (SAN File Servers)
Communication Intensive	Particle Physics, Fluid Dynamics, MPI, etc.	Fast interconnect network, low latency
Memory Intensive	DNA assembly, Image reconstruction, etc.	Large Memory
Continuous services	Databases, web servers, web services	Dynamically scalable

Type	Examples	Requirements
Compute Intensive 	Monte Carlo simulations, Embarrassing parallel	CPU Cycles
Data Intensive	Signal/Image processing, Pattern matching	Fast I/O to data (SAN File Servers)
Communication Intensive	Particle Physics, Fluid Dynamics, MPI, etc.	Fast interconnect network, low latency
Memory Intensive 	DNA assembly, Image reconstruction, etc.	Large Memory
Continuous services 	Databases, web servers, web services	Dynamically scalable

“Given the unalloyed good that is HPC, how come everybody isn’t using it?”

Council on Competitiveness:

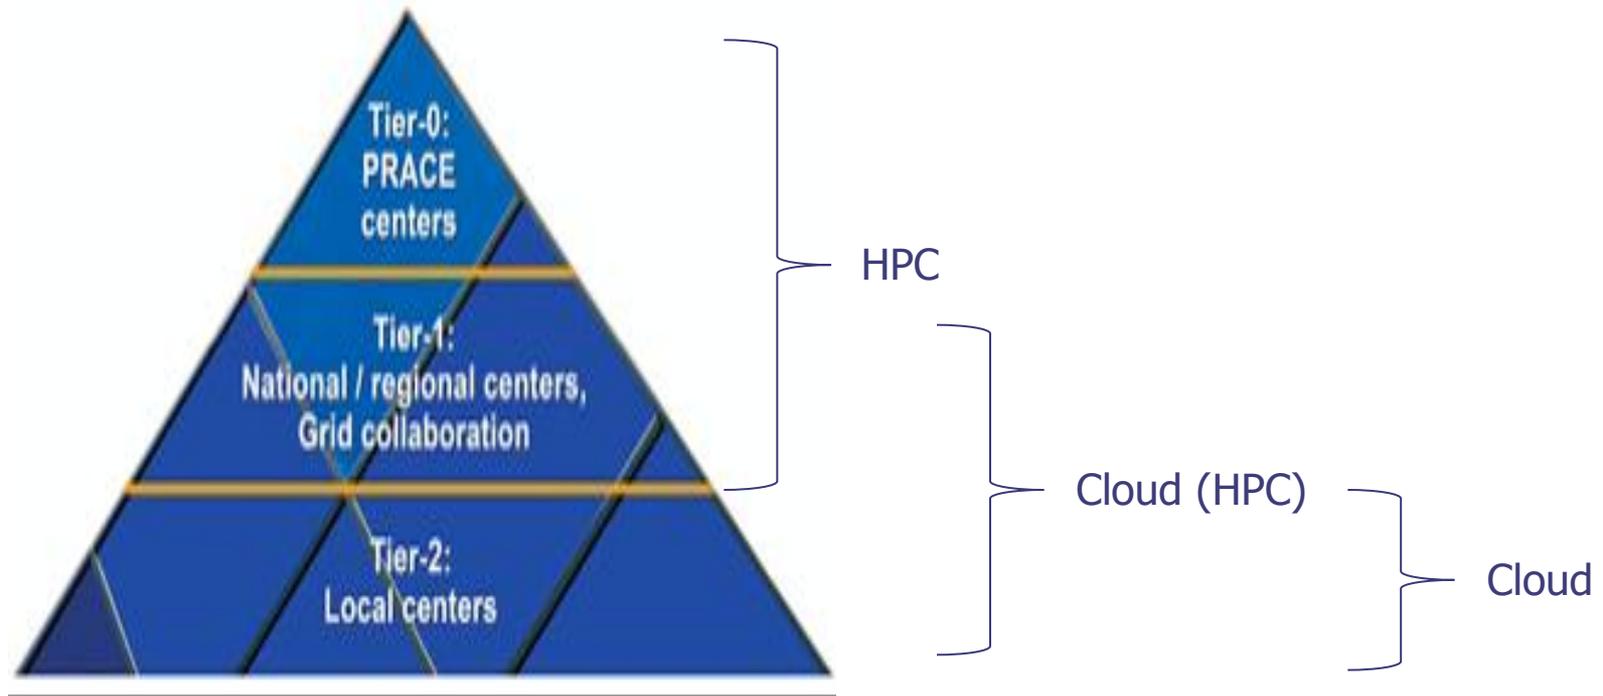
- **“missing middle”** refers to: the group of HPC users between low-end, mostly workstation-bound HPC users, and the kind of high-end HPC typically performed at national labs and some universities. The problem is that transitioning from desktop HPC to server-based HPC is filled with roadblocks, especially for commercial users looking to make the leap to small clusters.

“John West, Director for DoD HPC Modernization Program, discussing “The Missing Middle” at the 26th Annual HPCC Conference (March 26-28 2012, Newport, RI, USA)

- Increasing hardware complexity (including divergent technologies like CPUs and GPUs) – ***this is both an advantage and a disadvantage*** 😊
- Fixed software (OS) environments
- Fixed user environments
- Software portability and usability
- Software scaling (a non-problem, vertical vs horizontal scaling)
- Software licensing
- Lack of expertise
- Lack of human resource, time, etc.
- Applications with different requirements can co-exist on the same physical host



-  HPC Users
-  Potential HPC Users
(Academic, SMEs)
"missing middle"



- Cloud is not HPC, although now it can certainly support some HPC workloads
- Great opportunity for HPC in science and business ('**missing middle**')
- HPC Cloud is complementary to other services like supercomputers (PRACE Tier-0 and Tier-1)
- Users can configure their resources according to their own needs
- Operators can grant (virtual) administrator privileges to users, while retaining full ownership of the physical nodes and networks
- Virtual machines can be monitored without modification of the observed VM
- Reliability can be improved with checkpoint/restart schemes for VMs ☺
- Users can debug and test run their algorithms on VNodes
- You are not requested to exactly estimate the number of needed resources and this is a big advantage in the cost factor for a start-up company, or for a science application



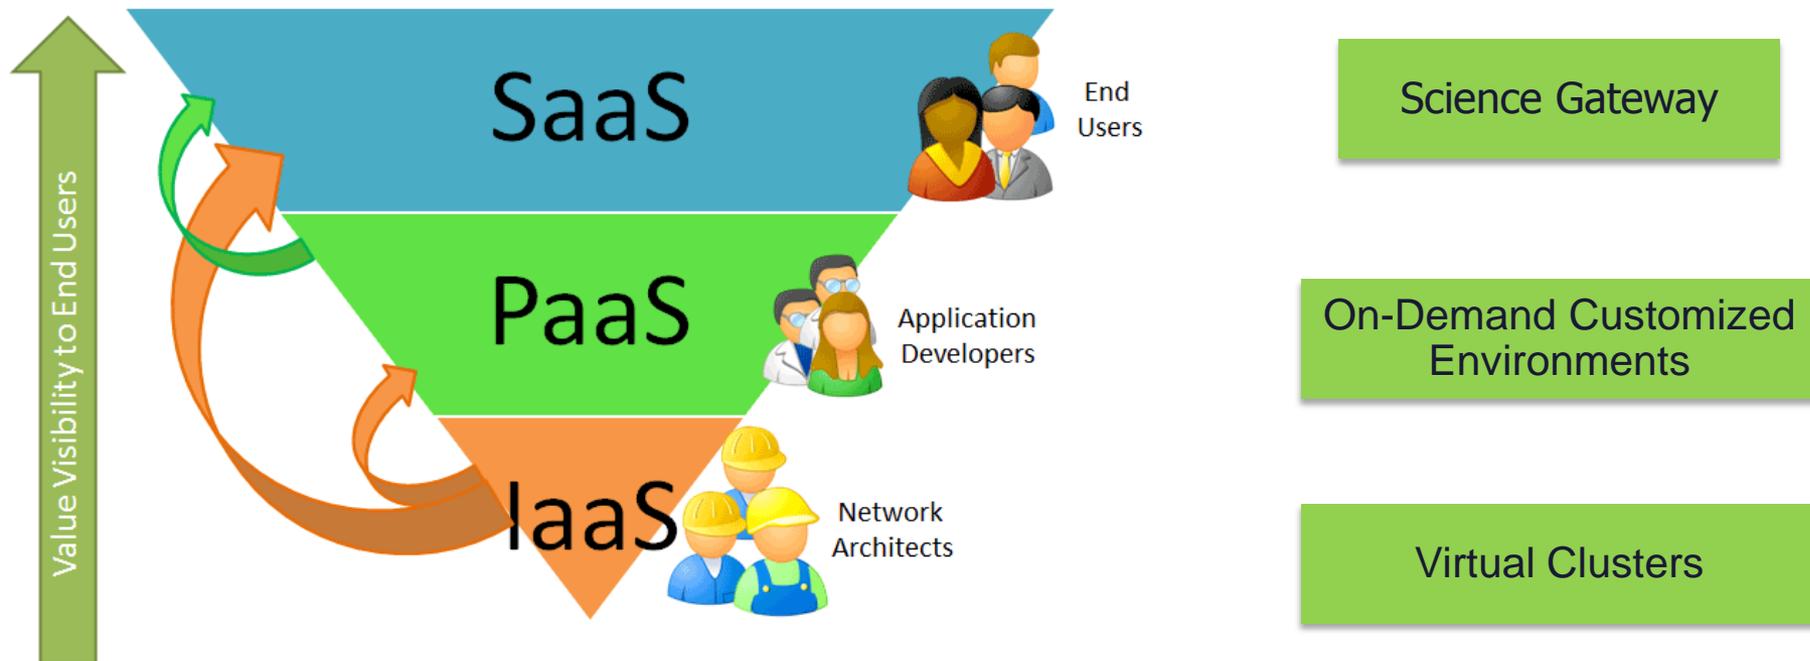


Image courtesy of <http://www.saasblogs.com>



Science Gateway

High interest

- *Does not necessarily imply virtualization of resources (not all HPC resources can be virtualized!)*
- *Well defined applications environment, easy to use and exploit*
- *Dedicated pool of resources to reduce queue waiting time*

On-Demand Customized Environments

Medium interest

- *Implies resources virtualization*
- *Much effort in managing the infrastructure could be required*

Virtual Clusters

Low interest

- *User would rather prefer bare metal*
- *Implies resources virtualization*
- *Does this really bring a big advantage?*

Testbed environment

- *On demand customized environment*
- *Application Prototyping for resource grant assignment*
- *Rapid prototyping and Proof-of-Concept in pilots with concrete use cases*
- *Pre-production with limited number of users*

Engineering applications

- *Flow assurance software like **Olga** or **Leda Flow**, which run serial jobs do not need too much communication overhead and can perform I/O locally on the node*
- ***CAE** (Computing Aiding Engineering) software separate a CAD model in small parts, solving a set of algebraic equations to obtain desired results (**Finite Elements Method**)*
- *As opposite, **FFT (Fast Fourier Transform)**, which could be used in **vibration analysis**, is probably an unsuitable algorithm to run on a cloud system because the processes involved in the computation are highly interdependent and the FFT stresses a lot the global communication bandwidth*

MapReduce data analysis algorithms

- *simulation as well as Big Data analysis*

- Stratuslab** - <http://stratuslab.eu/doku.php/start>
- Venus-C** - <http://www.venus-c.eu/Pages/Home.aspx>
- BigGrid** - <http://www.biggrid.nl/>
- Helix Nebula** - <http://cdsweb.cern.ch/record/1434708>
- White House** - www.apps.gov
- Cloud Book** - <http://www.cloudbook.net/directories/gov-clouds/government-cloud-computing.php>
- Google Exacycle Initiative**



Federal CIO Council

Cloud computing plays a key role in the President's initiative to modernize Information Technology (IT) by identifying enterprise-wide common services and solutions and adopting a new cloud computing business model. The Federal CIO Council under the guidance of the Office of Management and Budget (OMB) and the Federal Chief Information Officer (CIO), Vivek Kundra, established the Cloud Computing Initiative to fulfill the President's objectives for cloud computing.

[View the Federal CIO Council Profile >>](#)



Argonne National Laboratory

To test cloud computing for scientific capability, DOE centers at the Argonne Leadership Computing Facility (ALCF) and the National Energy Research Scientific Computing Center (NERSC) will install similar mid-range computing hardware, but will offer different computing environments. The combined set of systems will create a cloud tested that scientists can use for their computations while also testing the effectiveness of cloud computing for their particular research problems.

[View the Argonne National Laboratory Profile >>](#)



National Energy Research Scientific Computing Center (NERSC)

To test cloud computing for scientific capability, DOE centers at the Argonne Leadership Computing Facility (ALCF) and the National Energy Research Scientific Computing Center (NERSC) will install similar mid-range computing hardware, but will offer different computing environments. The combined set of systems will create a cloud tested that scientists can use for their computations while also testing the effectiveness of cloud computing for their particular research problems.

[View the National Energy Research Scientific Computing Center \(NERSC\) Profile >>](#)



General Services Administration (GSA)

The GSA is focusing on implementing projects for planning, acquiring, deploying and utilizing cloud computing solutions for the Federal Government that increase operational efficiencies, optimize common services and solutions across organizational boundaries and enable transparent, collaborative and participatory government. This includes Apps.gov, the official Cloud Computing Storefront for the Federal Government. The site features a complete listing of all GSA approved Cloud services available to federal agencies.

[View the General Services Administration \(GSA\) Profile >>](#)

Defense Information Systems Agency (DISA)

The DISA is developing a number of Cloud Computing solutions available to US military, DoD government civilians and DoD contractors for Government authorized use. They include: Forge.mil, a system that currently enables the collaborative development and use of open source and DoD community source software; GCDS, a commercially owned, globally distributed computing platform that provides a reliable and secure content and application distribution services solution that delivers applications to dispersed user communities; and RACE, a quick-turn computing solution that uses cloud computing to deliver platforms that are quick, inexpensive and secure.

[View the Defense Information Systems Agency \(DISA\) Profile >>](#)



National Institute of Standards & Technology (NIST)

NIST's role in cloud computing is to promote the effective and secure use of the technology within government and industry by providing technical guidance and promoting standards.

[View the National Institute of Standards & Technology \(NIST\) Profile >>](#)



National Science Foundation (NSF)

In 2008 the ACCI partnered with the National Science Foundation to provide grant funding to academic researchers interested in exploring large-data applications that could take advantage of this infrastructure. This resulted in the creation of the Cluster Exploratory (CluE) program led by Dr Jim French, which currently funds 14 University projects.

[View the National Science Foundation \(NSF\) Profile >>](#)



NASA

NASA has several cloud programs including Nebula, a Cloud Computing pilot under development at NASA Ames Research Center. It integrates a set of open-source components into a seamless, self-service platform, providing high-capacity computing, storage and network connectivity using a virtualized, scalable approach to achieve cost and energy efficiencies. The fully-integrated nature of the Nebula components provide for extremely rapid development of policy-compliant and secure web applications, fosters and encourages code reuse, and improves the coherence and cohesiveness of NASA's collaborative web applications.

[View the NASA Profile >>](#)



Department of Energy (DOE)

The DOE is exploring the cloud concept with its federal partners to identify opportunities to provide better service at lower cost through cloud services. The DOE National Laboratories are exploring the use of cloud services for scientific computing. They are also developing high bandwidth networking to transport the high volumes of data between DOE and cloud service providers required by scientific computing.

[View the Department of Energy \(DOE\) Profile >>](#)



Department of the Interior National Business Center (NBC)

The Department of the Interior's National Business Center (NBC) plans to bring the benefits of cloud computing to both NBC's business services clients and data center hosting clients alike through advancements to the highly efficient NBC shared infrastructure. The NBC is offering six cloud computing products: NBCGrid (IaaS), NBCFiles (Cloud Storage), NBCStage (PaaS), NBC Hybrid Cloud, NBCCApps (SaaS Marketplace), & NBCCAuth.



Public Works and Government Services Canada (PWGSC)

The Canadian Government's CTO of Public Works Government Services presented a paper on Cloud Computing and the Canadian Environment. This paper essentially outlines the Canadian Government's considerations of cloud computing by outlining the advantage of their cold landscape (among other things) as a prime location for the construction of large energy efficient Cloud Computing data centers.

[View the Public Works and Government Services Canada \(PWGSC\) Profile >>](#)

Conclusions



- Big Data requirements will start to transform the HPC market
 - Server designs, interconnects, storage, software, and services
 - Minimizing data movement will become increasingly important
 - The HPC storage market will continue to outpace server market growth
- There will be more processor choices
 - X86 will remain dominant, but GPUs are proliferating rapidly
 - ARM and other power-stingy embedded processors will begin showing up
 - Programmability and software availability will drive processor sales
- Energy efficiency will move from buzzword to critical design goal
- HPC private cloud computing will continue ramping up nicely, while public clouds grow more slowly
 - Cloud computing will mainly add to, not replace, existing HPC work



IDC Top 10 HPC Predictions for 2011

1. The Worldwide Economic Recovery Will Restore Healthy HPC Growth
2. The Battle for HPC Leadership Will Become Increasingly Global and Strategic
3. More Real-World Applications Will Be Run at Trans-Petaflop Speeds
4. There May Be More Emphasis on Software...Finally
5. The Alternative Processor Wars Will Keep Heating Up
6. The HPC Staffing Shortage Will Grow More Acute
7. Cloud Computing Will Ramp Up Slowly in HPC, But Will Find a Number of Strong Niches
8. The HPC Storage Market Will Continue to Grow Faster than the HPC Server Market
9. InfiniBand Will Continue to Take Market Share from Proprietary Interconnects While Ethernet Remains the Leader
10. Power and Cooling Will Become an Even Greater Concern

Technology: IBM Bluegene
Model: BlueGene/Q
Processor: IBM PowerA2, 1.6 GHz
Computing nodes: 10240
Computing cores: 163840
RAM: 1GB/core
Peak performance: 2PFlop/s
Users: academic

Moving to production
Aug. 13th, Sept. 3th

