INTERVENTIONS FOR SCIENTIFIC AND ENTERPRISE APPLICATIONS BASED ON HIGH PERFORMANCE AND CLOUD COMPUTING

Kitty Ahuja¹, Neetu Sharma²

Abstract

High performance computing refers to the practice of aggregating computing power in a way that delivers much higher performance than one could get out of a typical desktop computer in order to solve large problems in science, engineering or business. While cloud computing is 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.

The scope of HPC is scientific research and engineering as well as the design of enterprise applications. As enterprise applications are data centric, user friendly, complex, scalable and often require software packages, decision support systems, warehouse while scientific applications have need of the availability of a huge number of computers for executing large scale experiments. These needs can be addressed by using high-performance and cloud computing. The goal of HPC is to reduce execution time and accommodate larger and more complicated problems. While cloud computing provides scientists with a completely new model of utilizing the computing infrastructure, computing resources, storage resources, as well as applications can be dynamically provisioned on a pay per use basis. These resources can be released when they are no more needed. This paper focuses on enabling and scaling computing systems to support the execution of scientific and business applications.

Keywords: Scientific computing, enterprise applications, cloud computing, high performance computing

1. INTRODUCTION

High-performance computing (HPC) aims to control supercomputers and computer clusters to solve scientific computation and enterprise problems. The original objective of cloud computing and HPC can be manifestly different which yields different computing paradigms as well as applications. While HPC has been widely used for scientific tasks and cloud computing was set out for business applications. But now HPC and cloud computing can be used collectively for scientific computing and enterprise applications. The users of scientific computing are scientists and engineers which include applications such as climate prediction, oil exploration and gas discovery, problems of biochemistry, drug discovery, automotive design, protein folding simulations, defense and aerospace work, financial forecasting, etc. The enterprise customers encompass examples like the corporate data center that stores customer records, inventory management, and employee details.

Scientific computing revolves around mathematical models, quantitative analysis and numerical solution techniques to solve scientific and engineering problems. These models often require a huge number of computing resources to perform large scale experiments or to cut down the computational complexity into a reasonable time frame. Computational and Data-Driven Sciences are the major pillars of scientific discovery in addition to experimental and theoretical sciences. The support for data intensive scientific computing is critical to advance modern science as storage systems have exposed a widening gap between their capacity and their bandwidth by more than 10-fold over the last decade.

In today's corporate environment, enterprise applications are complex, scalable, distributed, component-based, and missioncritical. They may be deployed on a variety of platforms across corporate networks, intranets, or the internet. They are data-centric, user-friendly, and must meet stringent requirements for security, administration, and maintenance. Designing and developing such enterprise applications means satisfying hundreds or thousands of separate requirements.

Scientific Computing and enterprise applications are the key to solving "grand challenges" in many domains and providing breakthroughs in new knowledge, and it comes in many shapes and forms like **High Performance Computing** and **Cloud Computing** which are heavily focused on computeintensive applications and data-parallel execution.

The remainder of the paper is organized as follows. Section II discusses related work of high performance computing, their applications and also discusses issues with HPC. Section III explains the arising of the cloud and its characteristics. Section IV presents High Performance Computing and Cloud Computing Interventions for Enterprise Applications. Section

V analyzes the High Performance Computing and Cloud Computing Interventions for scientific applications, section VI focuses on challenges of using HPC in cloud. Finally, section VII explains conclusions and future work directions.

2. HIGH PERFORMANCE COMPUTING

HPC is a type of distributed computing, where processors and cores in multiple computers coordinate their actions. High performance computers of interest to small and medium-sized businesses today are clusters of computers which work together as if they are a single system. Each individual computer in a commonly configured small cluster has between one and four processors, and today's processors typically have from two to four cores. Now, we address the bigger and more important question – why we should choose HPC and for what applications?

HPC reduces computation time, which adds quality and value of the product. It also reduces product development time and cost. With HPC, we can test more ideas to help set priorities for work that gets done. HPC helps us to eliminate some testing and evaluate more alternatives. With the use of HPC, design and project run-around is reduced which improves design process to produce better products, innovation ability and reduce cost of development.

Typical HPC applications are in the fields of high-end computational science and engineering like weather forecasting and prediction, computational chemistry and drug discovery, Computational Fluid Dynamics (CFD), Quantum mechanics, computer-aided engineering, rapid prototyping, modeling and complex simulations. Cray Inc is the market leader in HPC. IBM, Intel and SGI are the other major players in this area.

Dana Holding Corporation is an HPC service provider for US department of defense and DARPA and has identified technologies and partners to support the HPC infrastructure throughout the supply chain of the department of defense. Technical customers usually require the Message Passing Interface (MPI) whereas enterprise customers require an SQL database. Because a number of MPI and database implementations have been created with sockets or DAPL, it is possible for a company to target both classes of customer by delivering a quality software stack. Unfortunately, many smaller HPC companies often target technical customers in the hopes of becoming famous, and in the process neglect enterprise customers who would make them rich. Enterprise is a much bigger market than technical simply because storage is required more than processing.

However progress towards HPC is threatened by two related factors:

(i) Based on current technology, systems with Exascale performance will use excessive amounts of energy.

(ii)The performance of processing elements can no longer be assumed to follow Moore's Law and double every 18 months to 2 years due to physical constraints.

HPC architectures have remained relatively unchanged. Massive investment will be desired to migrate or re-implement these applications to utilize future HPC architectures.

Many HPC applications are so-called legacy codes that have been developed over many years following a model of HPC which has been stable during that period, but which may not persist for much longer. Most of these legacy applications cannot be changed quickly, easily or cheaply.

Robust, real-time or even faster than real-time and reliable HPC is required for applications which deal with constant operation and monitoring of infrastructure such as electricity generation and distribution, transport control, industrial processes and decision support in crisis situations.

The effective use of HPC is often complex because of the diverse skills and expertise that are required to solve specific problems. This complexity has led to slow adoption of the technology by many companies.

The use of HPC is complex and, particularly when a user is new to the domain, may require significant help and support in order to realize its benefits.

HPC users can better quantify the benefits of moving to a cloud and identify which applications are better candidates for the transition from in-house to cloud. Cloud providers can optimize the allocation of applications to their infrastructure to maximize utilization, while offering best-in-class cost and quality of service. There can be three usage modes as shown in Figure 1: 'HPC in the cloud', where researchers outsource entire applications to current public and - or - private cloud platforms; 'HPC plus cloud', which focuses on exploring scenarios where clouds can complement grids and high-performance computers to support science and engineering application workflows; and 'HPC as a cloud', that's focused on parallel resources using elastic on-demand cloud abstractions [10].



Fig 1: Usage Modes of HPC and Cloud

Cloud computing is changing the landscape of many applications. The HPC should move towards exploiting the opportunities associated with the cloud computing.

3. THE RISE OF THE CLOUDS

HPCs can also contribute towards building components for the 'Cloud' as also be part of large data centers and server farms which host cloud-based services. HPC in the cloud is fast catching up as an effective method to manage big data and analytics.

Cloud computing is a form of utility computing, where hardware, software, storage and platform is made available as per need and subscription basis. In this model, clients can access the application through an internet browser. The data can be resident at a remote place also. Complementing the cloud is the usage of server farms and data centers where all applications and data can be stored, shared and accessed on demand using virtualization. Cloud is an eco-friendly green technology as it eliminates the need for enterprises to maintain multiple servers, which takes lots of space and infrastructure.

Cloud computing promises reliable services delivered through next-generation data centers that are built on compute and storage virtualization technologies. Users will be able to access applications and data from a cloud anywhere in the world on demand. In other words, the cloud appears to be a single point of access for all the computing needs of users. The users are assured that the cloud infrastructure is robust and will always be available at any time.

There are two deployment models for cloud, public and private cloud. Public cloud services are available to general public. In public cloud, one can access the services using the internet. Amazon and Google have built their large public clouds. Private cloud is built by enterprises for their internal company operations including communication with their divisions and branches. These could be maintained by a third party though. Of late, there is some movement to architect hybrid clouds which combine features of both private and public clouds. In a hybrid cloud, a company can maintain its own private cloud and on saturation of these resources, use the public cloud [12].

Cloud computing is generally broken down into three primary service levels

Software-as-a-Service''(SaaS): The simplest example to understand is e-mail.

Infrastructure-as-a-Service(IaaS): where we outsource the hardware. In such cases, it's not just the computing power that we rent; it also includes power, cooling, and networking. Furthermore, it's more than likely that we'll need storage as well. Generally IaaS is this combination of compute and cloud storage

Platform-as-a-Service(PaaS) : the vendor takes care of the underlying infrastructure for us, giving us only a platform with which to (build and) host our application(s).

PaaS user is a SaaS developer. an IaaS user could very well be a PaaS or SaaS developer.

Cloud computing exhibits the following key characteristics:

- Agility improves with users' ability to re-provision technological infrastructure resources.
- Application programming interface (API) accessibility to software that enables machines to interact with cloud software in the same way the user interface facilitates interaction between humans and computers. Cloud computing systems typically use REST-based APIs.
- Cost is claimed to be reduced and in a public cloud delivery model capital expenditure is converted to operational expenditure.
- Virtualization technology allows servers and storage devices to be shared and utilization be increased. Applications can be easily migrated from one physical server to another.
- Reliability is improved if multiple redundant sites are used, which makes well-designed cloud computing suitable for business continuity and disaster recovery.
- Scalability and elasticity via dynamic ("on-demand") provisioning of resources on a fine-grained, self-service basis near real-time, without users having to engineer for peak loads.
- Performance is monitored and consistent and loosely coupled architectures are constructed using web services as the system interface.

Clouds are basically high-end high performance data centers, which use virtualization technology creating virtual machines through hypervisors. The cloud computing introduces several benefits for end users, enterprises, service providers, and scientific institutions. Within the public cloud of Amazon, i.e Amazon EC2, in 2011, it launched Cluster Compute EC2 instance, which is a typical application of HPC in cloud. Meanwhile, SGI unveiled on demand cloud HPC service called Cyclone in 2010 [13].

4. HIGH PERFORMANCE COMPUTING AND CLOUD COMPUTING INTERVENTIONS FOR ENTERPRISE APPLICATIONS

Efficient enterprise management requires agility, collaboration, competitiveness, cost benefit, flexibility, resilience, responsiveness, scalability, simplicity, transparency and visibility, among various requirements. With the internet and ICT revolution affecting all aspects of an enterprise, enterprises are now exposed to greater competition as a result

of liberalization, privatization, outsourcing and globalization. Information and Communication Technology (ICT) tools are a great enabler of operations of an enterprise. Its adoption and deployment has become a determinant of competitive advantage for many enterprises [16]. Extensive use of technologies like SCM software packages, decision support systems, warehouse, RFID [15], software agents [17] barcodes, transportation and inventory management systems etc have resulted in better visibility, agility, collaboration and communication both within and among enterprises [15].

As in enterprises, supply chain visibility coupled with quick responsiveness and decision-making will be vital to gain a competitive advantage. This will entail actionable information about various stakeholders and processes to be available realtime. The ICT infrastructure to enable this will need to be seamlessly interconnected, networked, scalable and easily accessible. A major challenge will be lack of homogeneity in the legacy MIS and enterprise applications of various partners in the supply chain. Most of the data that companies need to run their supply chains reside with supplier and partners, who use their own MIS, SCM or ERP solutions. Another issue that prevent companies struggle to get a cohesive picture of their supply chains is that most of these information systems were designed to operate within a single company, not across a network of companies.

The intrinsic features in High Performance Computing systems are of scalability, data integration and interoperability makes it an ideal fit for enterprise applications. Large companies have slowly started adopting HPC for enterprise applications as they were struggling with the twin challenges of data management and data analysis. Pratt and Whitney use HPC for modeling and simulation of inventory positions and situations. They also expect their multitude suppliers to use the same system. However the fact of the matter is that HPC interventions in enterprise applications are limited to some big players. The reasons are primarily the need for all suppliers to migrate from legacy systems to HPC systems.

Cloud computing and associated technologies like Software as a Service (SaaS), Infrastructure as a Service (Iaas) and Platform as a Service (PaaS) have been generating a lot of interest as a game changing technology. Cloud computing can integrate shippers, service providers, distributors, logistic providers, customers, sellers etc in a global enterprise applications. A social network like community could be created with typical data items, placed online on the cloud like information on prices, inventory, schedules, service options, contracts, announcements etc. Timely information updates from all stakeholders renders companies to be demand-driven rather than forecast-driven. Major ERP vendors are also working offering cloud-based solutions.

Major application areas where cloud-based solutions are available include demand forecasting, demand planning, e-

procurement, distribution, inventory, warehouse and transportation systems. Every major ERP vendor like SAP, BAAN are coming out with a cloud version of their large software solutions. This reduces the Total Cost of Ownership for enterprises and improves both their top and bottom lines. FedEx has a private cloud deployed in 2011 with CloudX as service-provider with sales processing and CRM as primary activities on the cloud. Cosco Logistics, the largest 3PL company of China and the world's second largest ocean shipping company is using Saas service and integrating all stakeholders like customers, branches, suppliers and distributors in order all of them to use the same logistics management software. Intel was able to replace hundreds of their order clerks using online ordering applications. Several supply chain activities like planning and forecasting, sourcing and procurement, logistics and inventory management were migrated to cloud.

However like HPC, enterprises which go for the cloud solutions will have to face challenges like legacy systems and software applications of various partners. There are also privacy issues in both public and private clouds as the service provider has all the data.

5. HIGH PERFORMANCE COMPUTING AND CLOUD COMPUTING INTERVENTIONS FOR SCIENTIFIC APPLICATIONS

Research and development activities are focused on enabling and scaling scientific applications using distributed computing paradigms, such as cluster, grid, and cloud computing. With the rapid emergence of virtualized environments for accessing software systems and solutions, the volume of users and their data are growing exponentially. According to the IDC, by 2020, when the ICT industry reaches \$5 billion - \$1.7 billion larger than it is today - at least 80% of the industry's growth will driven by third platform technologies, such as cloud services and big data analytics. Existing computing infrastructure, software system designs, and use cases will have to take into account the enormity in volume of requests, size of data, computing load, locality and type of users, and every growing need of all applications.

Scientific institutions can be more concerned with PaaS and IaaS contributions that allow having complete control over the infrastructure used for scientific research and finely customizing their software systems according to the specific needs of the experiments to be performed. Cloud computing also guarantees the desired Quality of Service, which is well-known by means of Service Level Agreements. This aspect comprises of an additional value, which could make scientists prefer computing clouds to traditional grids. For example, different types of analysis can be performed at different costs by optimizing the trade-off between the budget and the expected accuracy of results.

By applying cloud technologies and standards to established scientific computing and HPC infrastructures, costly, inflexible silos can be eliminated and shared computing resources can be fully utilized to:

- Enhance productivity and facilitate collaboration with an easy-to-use, web-based portal
- Improve data access, resource availability and security while boosting performance
- Increase throughput of heterogeneous workloads across diverse resources
- Increase automation and decrease manual effort for improved administrator productivity
- Improve performance and scalability with highperformance clustered systems – factory integrated, tested, single-point of contact
- Reduce operating costs and total cost of ownership with a consolidated high-performance infrastructure [1]

Cloud computing introduces new challenges and new problems, particularly from a legal and a security point of view. In the case of public clouds, systems, applications, and even personal data are hosted into datacenters owned by third parties. These datacenters are often placed into the more expedient geographic location for reducing maintenance and consumption costs. In addition, privacy and confidentiality of data depends on the location of its storage. In order to address this issue some cloud computing vendors have included the geographic location of the hosting as a parameter of the service level agreement made with the customer. For example, Amazon EC2 provides the concept of availability zones that identify the location of the datacenters where applications are hosted. Security is another important issue. This issue is more gripping for enterprises but there could be significant implications even in the case of scientific computing: many scientific projects are often funded by federal bodies or directly by the government that often puts strict restrictions in the use and the management of data.

Even though cloud computing fosters the dynamic discovery of services and a wide variety of runtime environments for applications, in practice a limited set of options are available for scientists, and sometimes they could not be elastic enough to cover their needs.

6. CHALLENGES OF USING HPC WITH THE CLOUD

One of the primary challenges is the lack of high-speed interconnects and noise-free operating systems to enable tightly coupled HPC applications to scale. Other challenges include the costing/pricing model, which is still evolving from the traditional supercomputing approach of grants and quotas toward the pay-as-you-go model typical of cloud-based services; the submission model, which is evolving from job queuing and reservations toward VM deployment; the bringing of data in and out of the cloud, which is costly and results in data lock-in; and security, regulatory compliance, and various other utilities like performance, availability, business continuity, service-level agreements, and so on.

There can be the increase of HPC's adoption in the cloud with the two prospects. The first prospect is in optical networking, which will improve hardware manageability and interconnect performance while also reducing power consumption, and software-defined networks, which will further improve the manageability of cloud networking. Furthermore, developments in nonvolatile memory will improve check pointing performance and, in the long term, address the data deluge and new programming models leveraging nonvolatility. Both technologies will create disruptive improvements and increase HPC's adoption in the cloud.

As the scientific community moves toward exascale computing, the evolution of new asynchronous programming models and an emphasis on redesigning algorithms can make the cloud even more suitable for HPC [9].

CONCLUSIONS

HPC is a type of distributed computing, where processors and cores in multiple computers coordinate their actions. HPC is an important technology in which further investment is essential if its economic competitiveness is to be maintained. HPC systems have the advantages of scalability, data integration and interoperability. Cloud computing can integrate all partners into an online social network like community with real-time information on all elements. Cloud interventions in the enterprise management space are rapidly increasing with several vendors offering Software as a Service while for scientific computing Platform as a Service and Infrastructure as a Service are being offered. HPC interventions on the other hand are limited to some big companies. A challenge can be faced during deployment of both cloud and HPC systems that can be the legacy applications especially of the suppliers [11].

The technological challenges must be overcome for the full potential of HPC and cloud computing to be realized. There can be numerous opportunities for using HPC in new applications and the possibilities are opening up by using HPC with other computing domains such as embedded systems and mobile computing.

REFERENCES

[1].http://www.03.ibm.com/systems/technicalcomputing/soluti ons/hpccloud.html

[2]. P. Mell and T. Grance, "Draft nist working definition of cloud computing - v15," 21. Aug 2009.

[3]. Tharam Dillon, Chen Wu and Elizabeth ChangCloud Computing: Issues and Challenges, 24th IEEE International Conference on Advanced Information Networking and Applications, 2010.

[4]. Christian Vecchiola, Suraj Pandey, and Rajkumar Buyya,High Peformance computing:cloud computing,10th International Symposium on Pervasive Systems, Algorithms, and Networks,2009

[5]. M. J. Chin, S. Harvey, S. Jha, and P. V. Coveney, "Scientific Grid Computing: The First Generation," Computing in Science and Engineering, vol. 7, 2005, . pp. 24– 32

[6]. D. Thain, T. Tannenbaum, and M. Livny, "Distributed computing in practice: The condor experience," Concurrency and Computation: Practice and Experience, vol. 17, pp. 323–356, February, 2005.

[7]. Victor R. Basili, Jeffrey C. Carver, Daniela Cruzes, Lorin M. Hochstein et.al., Understanding the High- Performance-Computing Community: A Software Engineer's Perspective Victor, IEEE, Aug 2008

[8]. Abhishek Gupta, Filippo Gioachin, Paolo Faraboschi, The Who, What, Why and How of High Performance Computing Applications in the Cloud," Hewlett-Packard Development Company, L.P., 2013

[9].http://www.computer.org/portal/web/computingnow/archive/september2012

[10] http://www.isgtw.org/feature/can-clouds-transformscience

[11].http://cordis.europa.eu/fp7/ict/computing/documents/plan ethpc-strategy.pdf

[12]. Sujay R (2011), "Hybrid Cloud: A new Era",

International Journal of Computer Science and Technology (IJCST), Vol. 2, Issue 2, pp. 323-326

[13]. Betts, Bryan (2012), "HPC Cloud; Supercomputing to go", E&T Magazine, Vol. 7, No. 2.

[14]. Nair, Prashant R (2012), "RFID for Supply Chain Management", CSI Communications, November 2012, Vol. 36, No. 8, pp. 14-18, ISSN # 0970-647X.

[15]. Nair, Prashant R and OA Balasubramaniam (2012), "IT Enabled Supply Chain Management using Decision Support Systems", CSI Communications, May, 2010, Vol. 34, No. 2, pp. 34-40, ISSN # 0970-647X.

[16]. Nair, Prashant R (2010), "Benefits of Information Technology Applications for Supply Chain Management: An Explorative study of progressive Indian companies", Edited Book "Enterprise Information Systems and Implementing IT Infrastructures: Challenges and Issues", IGI Global, March 2010, ISBN13: 9781615206254.

[17]. Nair, Prashant R, "E-Supply Chain Management using Software Agents", CSI Communications, July 2013, Vol. 37, No. 4, pp. 13-16, ISSN # 0970-647X.