ORIGINAL RESEARCH PAPER

Volume - 10 | Issue - 05 | May - 2021 | PRINT ISSN No. 2277 - 8179 | DOI : 10.36106/ijsr

INTERNATIONAL JOURNAL OF SCIENTIFIC RESEARCH

ENABLE HIGH PERFORMANCE COMPUTING IN CLOUD: A REVIEW

Information Technology

Reena Saini Department of Computer Science, B. K. Birla Institute of Engineering & Technology, Pilani 333031, India.

Nachiket Sainis* Department of Computer Science, B. K. Birla Institute of Engineering & Technology, Pilani 333031, India. *Corresponding Author

ABSTRACT

High Performance Computing (HPC) resulting whole computing power in a way that delivers much higher performance than it could get in typical desktop computer or workstation. High Performance Computing (HPC) allows scientists and engineers to solve complex science, engineering, and business problems using applications that require high bandwidth, enhanced networking, and very high compute capabilities. HPC's democratization has been driven particularly by cloud computing, which has given scientists access to supercomputing-like features as the pay as you go. This paper will provide an overview on benefits, challenges and future of HPC in cloud.

KEYWORDS

IaaS, SaaS, PaaS, HaaS, HPC

INTRODUCTION

In the coming years, the cloud computing market will grow more efficient. Many organizations have found using the cloud has a clear advantage over in-house hardware. Indeed, shifting what was a capital expense to an operational expense has many advantages, including instant availability and the ability to scale up (and down) rapidly. In HPC, the Globus project has shown the power of grid-based computing and has spawned many successful production computation grids. Cloud computing takes grid computing to a whole new level by using virtualization to encapsulate an operating system (OS) instance. In addition to cloud computing, cloud storage also can be used independently or combined with OS instances. Cloud computing would seem to be an HPC user's dream offering almost unlimited storage and instantly available and scalable computing resources, all at a reasonable metered cost [4].High Performance computing used in parallel processing to run the advanced application program quickly and reliable. The term use especially to systems that function above a teraflop or 10^{12} floating point operations per second. The high performance computing is used as a synonym for supercomputing. The most common users of high performance computing systems are scientific researchers, engineers and academic institutions. Many government agencies, particularly the military, also rely on high performance computing for complex applications. High-performance systems often use custom-made components in addition to so-called commodity components. As demand for processing power and speed grows, high performance computing will likely interest businesses of all sizes, particularly for transaction processing and data warehouses. An occasional techno-fiends might use an HPC system to satisfy an exceptional desire for advanced technology. HPC's democratization has been driven particularly by cloud computing, which has given scientists access to supercomputing-like features at the cost of a few dollars per hour.

Performing HPC jobs on the cloud is a cost effective solution that overcomes and replaces real dedicated cluster. Cloud computing providers use service- driven models. All hardware, platform and software resources are leased to customer on-demand [2]. There are four cloud categories which are: Infrastructure as a service (IaaS), Platform as a service (PaaS), Software as a service (SaaS) [1]. Hardware as a service (HaaS)-on demand access to a specific computation, storage and networking products and/or equipment configuration possibly at a particular site [3].

ANATOMY OF HPC

Performance of a HPC is measured in floating point operations per second (FLOPS). As of, there are HPC which can perform up-to quadrillions of FLOPS. HPC used only a few processors in the beginning but as years passes, HPC with thousands of processors began to appear and, by the end of the 20th century, massively parallel high performance computer with tens of thousands of "off-the-shelf" processors were the norm. Tianhe-2 HPC is currently the fastest in the world at 33.86 petaFLOPS (PFLOPS), or 33.86 quadrillion floating point operations per second.

Systems with massive numbers of processors generally take one of two paths: In one approach (e.g., in distributed computing), a large number of discrete computers (e.g., laptops) distributed across a network (e.g., the Internet) devote some or all of their time to solving a common problem; each individual computer (client) receives and completes many small tasks, reporting the results to a central server which integrates the task results from all the clients into the overall solution. In another approach, a large number of dedicated processors are placed in close proximity to each other (e.g. in a computer cluster); this saves considerable time moving data around and makes it possible for the processors to work together.

HPC play an important role in the field of computational science, and are used for a wide range of computationally intensive tasks in various fields, including quantum mechanics, weather forecasting, climate research, oil and gas exploration, molecular modeling.

High performance computers of interest to small and medium-sized business today are really clusters of computers. HPC people often refer to the individual computers in a cluster as nodes. A cluster for small business could have as few as four nodes, or 16 cores. A common cluster size in many businesses is between 16 and 64 nodes, or from 64 to 256 cores.

The point of having a high performance computer is so that the individual nodes can work together to solve a problem larger than any one computer can easily solve. The nodes need to coordinate to one another in order to work meaningfully together [5].

BENEFITS OF HPC IN CLOIUD

The cloud offers applications a range of benefits, including elasticity, small startup and maintenance costs, and economics of scale. Today's rapidly advancing cloud-technology presents significant opportunities for high-performance computing (HPC). HPC system managers can leverage the benefits of the cloud model for their traditional HPC environments for the best of both worlds. Incorporating or transitioning HPC Cloud into your HPC system enables you to better meet workload demand, increase system utilization, and make your high performance computing system accessible to a wider community [6].

- Scale to better support application and job needs with automated workload-optimized node OS provisioning.
- Provide simplified self-service access for broader set of users that also reduces management and training costs.
- Accelerate collaboration or funding by extending HPC resources to community partners without their own HPC systems.
- Enable pay-for-use with show back and chargeback reporting for actual resource usage by user, group, project, or account.
- Enable higher cloud-based efficiency without the cost and disruption of ripping and replacing existing [7].

CHALLENGES

Broader use of HPC in the cloud also presents some key challenges.

Primary among them is the lack of high-speed interconnects and noise-free operating systems to enable tightly coupled High Performance computing applications to scale. New optimized virtualization models (thin VMs, containers, and so on) are reducing virtualization costs, but they're still not noise free. 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.
- Security, regulatory compliance, and various "ilities" (performance, availability, business continuity, service-level agreements, and so on).
- Increasing complexity in product design with extensive use of Monte Carlo simulations with added issues of regulatory requirements, and hedging and risk management needs.
- Asset and Liability Modeling with a Dynamic Hedging Program. Large scale analysis with millions of policies, scenarios and trillions of calculations. All based on sensitive market movements, long time horizons and complex behaviors. 2.5 million core hour runs to target solution over a weekend.
- Disparate internal clusters, lack of resource for large problem execution, and home grown systems pain [9].

DIFFERENT HPC CLOUD PROVIDER

- **Penguin Computing** One of the first vendors to introduce a true HPC cloud was Penguin Computing. The Penguin On Demand, or POD cloud, was one of the first remote HPC services. From the beginning, POD has been a bare-metal compute model similar to an in-house cluster. Each user is given a virtualized login node that does not play a role in code execution. The standard compute node has a range of options, including dual four-core Xeon, dual six-core Xeon, or quad 12-core AMD processors ranging in speed from 2.2 to 2.9GHz with 24 to 128GB of RAM per server and up to 1TB of scratch local storage per node.
- **R-HPC** R-HPC offers R-Cloud, wherein clients can "rent" HPC resources. R-Cloud offers two distinct computing environments. The first is a Shared Cluster, which offers a login to shared nodes and a work queue. This environment offers a classic cluster environment and is essentially a "shared cluster in the sky." Users are billed by the job, creating a pay-as-you go HPC service. The second environment comprises virtual private clusters that are carved out of a shared configuration. Use can be on-demand with VLAN access. These systems are billed on a 24/7 basis. R-HPC can provide new 3.4GHz quad core Sandy Bridge-based systems with 16GB of RAM/node (4GB/core), DDR 2:1 blocking InfiniBand, and 1TB of local disk.
- **Amazon EC2 HPC** the most well-known cloud provider. Originally, the EC2 service was found not suitable for many HPC applications. Amazon has since created dedicated "cluster instances" that offer better performance to HPC users. Several virtualized HPC instances are available on the basis of users' needs. Their first offering is two Cluster Compute instances that provide a very large amount of CPU coupled with increased network performance (10GigE). Instances come in two sizes, a Nehalem-based "Quadruple Extra Large Instance" (eight cores/node, 23GB of RAM, 1.7TB of local storage) and a Sandy Bridge-based "Eight Extra Large Instance" (16 cores/node, 60.5GB of RAM, 3.4TB of local storage).Pricing can vary depending upon on-demand, scheduled, or spot purchase of resources [8].
- **Univa** Univa is known as the place Sun Grid Engine landed after the Oracle purchase. The companies offer two distinct ways to build clouds. The first is an internal cloud infrastructure system that provides "raw" virtual machines using one of several methods, including VMware, KVM, Xen, or something more sophisticated like OpenStack or VMware vSphere. Their UniCloud product uses an Infrastructure-as-a-Service (IaaS) layer to create the machines and then directs the IaaS to provision a raw machine on the fly. Once the machine is provisioned, UniCloud takes over as a Platform-as-a-Service (PaaS) layer and does additional software provisioning and configuration to turn the raw instances into nodes that can then be part of an HPC cluster.
- **Sabalcore Computing** Sabalcore offers HPC systems over the Internet for both on-demand and dedicated solutions. On-demand access is accomplished using an SSH client or secure remote desktop interface from any Windows or Linux laptop, desktop, or workstation. Users are provided with an NFS-exported home directory with persistent storage for results, code, applications, and data. As part of their service, Sabalcore offers accounting tools that allow administrators to assign and track usage by users, projects, or resources. One interesting feature offered by Sabalcore is called Dynamic Cluster Isolation (DCI), which builds an optional firewall around the compute nodes of a particular job.
- **SGI** SGI's cloud computing service is called Cyclone and is specifically dedicated for technical applications. Through Cyclone, SGI offers both a software stack and HPC hardware in a non-virtualized environment together with HPC applications. Pulling from both open source and commercial vendors, they have applications in Computational Biology, Computational Chemistry and Materials, Computational Fluid Dynamics, Finite Element Analysis, and Computational Electromagnetics.
- **Gompute** Gompute provides CPU hours, storage, system administration, and support for HPC applications, including both open source and commercial. For commercial applications, users may either start a private license server at Gompute using the license key they get from an authorized reseller or allow the connection to their local license server. Gompute uses a pay-per-use model, whereby users get real-time information about their consumption for each job. [4].

CONCLUSION

In this paper we are discussing high performance computing on cloud. It is the best way to compute complex program easily on cloud. High performance computing on cloud increase the efficiency, reliability of computation of the cluster or we can say speed of the system on cloud in the range of teraflop or petaflop. There are some complexity in high performance computing on cloud i.e. disparate internal cluster, lack of resource for large problem execution and dynamic hedging problem. Some providers that provide high performance computing on cloud Penguin Computing Amazon EC2 HPC, SGI etc.

REFERENCES:

[1] A. Marathe, D. Lowenthal, B. Rountree, X. Yuan, M. Schulz and B. de Supinski, "A User Perspective of High-Performance Computing on the Cloud", The University of Arizona, Lawrance Livermore National laboratory, March, 2012.

[2] Portal.azure.com.(2017).MicrosoftAzure[online]Availableat:https://portal.azure.com/. [Accessed 1 Dec. 2017].

[3] Amit Batra, Arvind Kumar. "High Performance Computing into Cloud Services". International Conference on Computing Sciences, 2012.

[4] Admin magazine: http://www.admin- magazine.com/HPC/Articles/Moving-HPC- to-the-Cloud

[5] Insidehpc: http://insidehpc.com/hpc- basic- training/what-is-hpc/

[6] T.Groves, "Characterizing and Improving Power And Performance of HPC Networks", University of New Mexico and Sandia National Laboratories, 2016.

[7] Adaptivecomputing:http://www.adaptivecom puting.com/products/hpc-products/moab- hpc-suite-enterprise-edition/hpc-cloud- overview/

[8] Amazon Web Services, Inc. (2017). Amazon Web Services (AWS), Cloud Computing Services. [online] Available at: https://aws.amazon.com/ [Accessed 1 Dec. 2017].

[9] Morgan Eldred, Dr. Carl Adams, Dr. Alice Good. "Trust Challenge in a High Performance Cloud Computing Project". IEEE 6th International Conference on Cloud Computing Technology & Science, 2014.