

Running WLCG and HPC Workload on Single Openstack Cloud

Our world is changing so fast. We have to adopt this change and act according to the new technologies quickly. Many of the biggest world recognized companies are no longer in the business just because they did not act quickly to adopt the new changing in technologies. Nokia is the example of this. Today, cloud computing has revolutionized Information Technology field. It is the emerging concept and change that every Data Center will have to adopt in near future or even today.

HPC and WLCG applications are widely used in scientific and industrial research, data analytic and visualization, social behavioral studies etc. Most HPC applications and WLCG setup require dedicated, available and highly customized resources and environments for computation since they exhibit intense resource utilization. As stated above, Cloud computing is emerging as a latest computing technology. The on-demand nature of cloud has provoked interest to explore if cloud properties can be useful for HPC and WLCG setups. This paper is a work around in that direction.

Currently, user community of our cluster comprises but not limited to researchers working in the fields of Computational Condensed Matter Physics, Climate Modeling, Density Functional Theory, Computational Plasma Physics, and Semi-conductor Nanostructures. Software packages/compilers support is available for Intel Fortran, gcc, Quantum Espresso, Gaussian09, Mopac2012, Turbo mole, Wien2k, and Geant4 etc. The results produced making use of this facility by submitting user jobs over it looks comparable to the those obtain in the past using similar computational clusters hosted by other international institutes. Whereas, our WLCG setup has 524 CPU Cores and 441 TB of raw Storage. DPM based WLCG Tier-2 site has currently 59 worker nodes dedicated machines. The main innovation towards running HPC and WLCG workload on cloud was to better utilization of IT resources by running WLCG worker nodes on cloud, as well as to test the behavior of jobs being run on cluster nodes and WLCG Worker nodes simultaneously. Private cloud was setup by using Openstack latest release Mitaka. Initially, three Octa Core Sun Servers, each having 16GB of RAM and 160 GB of local storage were installed on CentOS 7. One node was made as controller node while other two nodes were made as compute node. Different subnets were used for cluster and WLCG instances to avoid network latency. Finally, Rocks Cluster and WLCG worker nodes instances were created and added in our production environment and real jobs were submitted, ran and monitored.

We observed that jobs are being run smoothly on cluster instances as well as on our WLCG worker nodes. We also observed that running HPC and WLCG workload on single cloud was more easy, scalable and flexible than running these workload on physical hosts because we could add new nodes in both HPC and WLCG very easily and allocate hardware resources according to the job requirements.

Primary author: Mr AHMAD, Manzoor (National Centre for Physics, Islamabad)

Presenter: Mr AHMAD, Manzoor (National Centre for Physics, Islamabad)

Track Classification: Infrastructure Clouds and Virtualisation