Contribution ID: **490** Type: **presentation**

Elastic resource allocation for multi-queue HEP workloads in cloud based environments

Tuesday, 10 July 2018 12:00 (15 minutes)

As the development of cloud computing, more and more clouds are widely applied in the high-energy physics fields. OpenStack is generally considered as the future of cloud computing. However in OpenStack, the resource allocation model assigns a fixed number of resources to each group. It is not very suitable for scientific computing such as high energy physics applications whose demands of resource various, especially with strong peak demands. In a traditional static cluster, a fixed number of virtual machines are pre-allocated to the job queue of different experiments. What happens often is some queues are queued while some queues are idle. As a result, the overall efficacy of virtual cluster is rather low. To solve this problem, we developed a resource allocation service to provide OpenStack elastic scheduling. In this implement, each experiment queue in OpenStack will have a pair of {Min, Max} quota, which Min represents the minimum number of resources for this experiment and Max represents the highest available resources. The final resource allocation for each experiment is determined by the fair resource scheduling algorithm.

In the paper, we will discuss development work of elastic resource allocation. An use case with VCondor (Virtual cluster with HTCondor implemented by IHEP) is given. The result shows it can greatly improve cloud resource efficiency.

Primary author: LI, Haibo (Institute of High Energy Physics Chinese Academy of Science)

Presenter: LI, Haibo (Institute of High Energy Physics Chinese Academy of Science)

Session Classification: T7 - Clouds, virtualization and containers

Track Classification: Track 7 –Clouds, virtualization and containers