



Contribution ID: 212

Type: oral presentation

BESIII physical offline data analysis on virtualization platform

Tuesday, 14 April 2015 17:45 (15 minutes)

Mass data processing and analysis contribute much to the development and discoveries of a new generation of High Energy Physics. The BESIII experiment of IHEP(Institute of High Energy Physics, Beijing, China) studies particles in the tau-charm energy region ranges from 2 GeV to 4.6 GeV, and requires massive storage and computing resources, which is a typical kind of data intensive application. With the rapid growth of experimental data, the data processing system encounters many problems, such as low resource utilization, complex migration and so on, which makes it urgent to transplant the data analysis system to a virtualization platform. However, offline software design, resource allocation and job scheduling of BESIII experiment are all based on physical machine. To solve those problems, we bring the virtualization technology of Openstack and KVM to BESIII computing system.

In this contribution we present an ongoing work which aims to make BESIII physical analysis work on virtualized resources to achieve higher resource utilization, dynamic resource management and higher job operating efficiency. Particularly, we discuss the architecture of BESIII offline software and the way to optimize the offline software to reduce the performance loss in virtualized environment by creating event index(event metadata) and do event pre-selection based on index, which significantly reduces the IO throughput and event numbers that need to do analysis, and then greatly improves the job processing efficiency. We also report the optimization of KVM from various factors in hardware and kernel including EPT (Extended Page Tables) and CPU affinity. Experimental results show the CPU performance penalty of KVM can be decreased to about 3%. This work is validated through real use cases of production BESIII jobs by working on physical slots and virtualized slots. In addition, the performance comparison between KVM and physical machines in aspect of CPU, disk IO and network IO is also presented. Finally, we describe our development work of adaptive cloud scheduler, which allocates and reclaims VMs dynamically according to the status of TORQUE queue and the size of resource pool to improve resource utilization and job processing efficiency.

Primary author: Dr HUANG, Qiulan (Chinese Academy of Sciences (CN))

Co-authors: Ms KAN, Bowen (Institute of High Physics Chinese Academy of Sciences); Dr LI, haibo (Chinese Academy of Sciences); Ms LEI, xiaofeng (Chinese Academy of Sciences)

Presenter: Ms KAN, Bowen (Institute of High Physics Chinese Academy of Sciences)

Session Classification: Track 7 Session

Track Classification: Track7: Clouds and virtualization