



Contribution ID: 332

Type: Poster

Cosmos : A Unified Accounting System both for HTCondor and Slurm Clusters at IHEP

Thursday, 7 November 2019 16:15 (15 minutes)

HTCondor is adopted to manage the High Throughput Computing (HTC) cluster at IHEP since 2017. Two months later in the same year, a Slurm cluster is set up to run High Performance Computing (HPC) jobs. To provide accounting service both for HTCondor and Slurm clusters, a unified accounting system named Cosmos is necessary to develop.

However, different job workload brings different accounting requirements. Most jobs from the HTCondor cluster are single-core jobs, and more than 30 million jobs are submitted each year. In addition to these jobs, a legacy HTCondor Virtual Machine (VM) cluster is treated as the second HTCondor pool, which means that VM jobs have to be accounted as well. On the other side, most jobs run in the Slurm cluster are parallelism jobs, and some jobs are run on GPU worker nodes to accelerate computing. Besides, some qualified HTC jobs are going to migrate from the HTCondor cluster to the Slurm cluster for research purpose.

To satisfy all the mentioned requirements, Cosmos is designed as a four-layer system, and layers from bottom to top are: data acquisition layer, data integration layer, data statistics layer and data presentation layer. In this proceeding, we present the background and development requirements of Cosmos system, the four-layer system architecture, the issues of each layer, and technical solutions to these issues. Cosmos is running as a production system for more than one year, and running results show that it's a well-functioning system and fulfills requirements from both HTCondor and Slurm clusters.

Consider for promotion

No

Primary authors: DU, Ran; SHI, Jingyan (IHEP); ZOU, Jiaheng (IHEP); JIANG, Xiaowei (IHEP (中国科学院高能物理研究所))

Presenter: JIANG, Xiaowei (IHEP (中国科学院高能物理研究所))

Session Classification: Posters

Track Classification: Track 7 –Facilities, Clouds and Containers