

R&D for the expansion of the Tokyo regional analysis center using Google Cloud Platform

Friday, 5 April 2019 10:50 (30 minutes)

Tokyo regional analysis center at the International Center for Elementary Particle Physics (ICEPP), the University of Tokyo, is a computing center for the ATLAS experiment at Large Hadron Collider (LHC) and one of the Worldwide LHC Computing Grid (WLCG) Tier2 sites supporting ATLAS VO. The center provides 10,000 CPU cores and 10 PB disk storage. A part of resources is dedicated to the local usage of the ATLAS Japan member. Hardware devices in the center are supplied by the three years rental. The current system is the 4th system which contract will end in this year. Currently, a migration to the next 5th system is ongoing. In the next system, the number of CPU cores is almost the same as the current system while the performance will be improved about 9% per core based on SPECint. The file storage will be increased to 15TB.

LHC plans the High-Luminosity LHC project starting from 2026. The peak luminosity will be 5 times higher. This will require more than 10 times of computing resources for the experiment. Such a requirement is 5 times higher than expected resources under the assumption of the flat budget scenario. Although many software improvements have been achieved, the gap between the requirement and the expectation is still large. To fill the gap, the availability of computing resources must be improved. One of the possibilities is to use GPGPU, which requires additional software development. There are also some R&D to use external resources of commercial cloud, HPC, or volunteer computing resources.

To expand Tokyo regional analysis center, R&D project using commercial cloud resource was launched. The batch system of the center for WLCG is managed by HTCondor. The first R&D was started to deploy worker nodes of HTCondor on Google Cloud Platform. The system is a hybrid system that worker nodes are on the cloud while header nodes and file storages consist of on-premises resources. To reduce the cost, Google Cloud Platform provides preemptible instances, which has a running time limit of 24 hours. To use preemptible resources, a load balancer called GCP_Condor_Pool_Manager (GCPM) has been developed. GCPM checks HTCondor's waiting queues and create new instances on demand. Instances are deleted after one job is executed. By this procedure, the system can use preemptible instances effectively. The system is set up by Puppet and it is easy to set up a similar system in other places.

In this presentation, the current status of the R&D project and our experiences of usage of Google Cloud Platform will be presented.

Primary author: Dr KANEDA, Michiru (ICEPP, the University of Tokyo)

Co-authors: Prof. TANAKA, Junichi (University of Tokyo); Dr MATSUI, Nagataka (The University of Tokyo); Dr RYU, Sawada (The University of Tokyo); Prof. MASHIMO, Tetsuro (The University of Tokyo); Mr KISHIMOTO, Tomoe (University of Tokyo)

Presenter: Dr KANEDA, Michiru (ICEPP, the University of Tokyo)

Session Classification: Physics & Engineering Application

Track Classification: Physics (including HEP) and Engineering Applications