

Building an elastic batch system with private and public clouds

Wednesday, 3 April 2019 16:00 (30 minutes)

The Computing Research Center at High Energy Accelerator Research Organization (KEK) provides a Linux cluster consisting of 350 physical servers with 10000 CPU cores for the scientific data analysis and numerical simulation. All of the computing resources are shared by using a batch system among the various projects supporting at KEK. We have adopted IBM Spectrum LSF as the batch system for many years, and users in KEK are well familiar with the system. However, when the computing cluster is becoming congested by the lack of resource, user jobs make a long stay in a job queue. As a result, the turnaround time often becomes longer for each job to be finished. Furthermore, we have to consider the different requirements for processing environments depending on users/groups.

Providing flexible resources both in terms of computing power and environment, we have investigated the possibility of Cloud computing technology. IBM Spectrum LSF has an optional functionality so-called Resource Connector which enables us to utilize additional computing resources in external cloud providers, such as Amazon Web Service (AWS), Microsoft Azure, OpenStack. By using that functionality, we have succeeded to integrate our batch system with on-premise OpenStack cloud and Amazon Web Service by the collaboration with National Institute of Informatics (NII), Japan.

Users jobs submitted into the dedicated job queue are going to be dispatched to dynamically provisioned instances at on-premise OpenStack, AWS, as well as static local computing node. Any kind of job processing environments can be utilized by choosing a different virtual machine according to the user's request. The OpenStack based private cloud has been integrated with our LDAP service and GPFS (General Parallel File System) for user authentication and data sharing on our batch system, respectively. Amazon Simple Storage Service is utilized for the data exchange between KEK and AWS. Both the physical servers on the batch system at KEK and provisioned instances on AWS mount the S3 bucket via FUSE for transparent data access. We conducted some performance tests on our hybrid batch system for investigation of the scalability and succeeded to execute a Deep Learning job using 3500 cores on AWS.

In this talk, we would like to present the detailed configuration of the hybrid system and some results of the performance tests.

Primary author: TAKASE, Wataru (KEK)

Co-authors: Dr MURAKAMI, Koichi (KEK); Prof. SASAKI, Takashi (KEK); Dr NAKAMURA, Tomoaki (KEK)

Presenter: TAKASE, Wataru (KEK)

Session Classification: Infrastructure Clouds and Virtualisation

Track Classification: Infrastructure Clouds and Virtualisation