

Examination of dynamic partitioning for multi-core jobs in the Tokyo Tier-2 center

Friday, March 10, 2017 9:00 AM (30 minutes)

The Tokyo Tier-2 site, which is located in International Center for Elementary Particle Physics (ICEPP) at the University of Tokyo, is providing computer resources for the ATLAS experiment in the Worldwide LHC Computing Grid (WLCG). The official site operation in the WLCG was started in 2007 after the several years development since 2002, and the site has been achieving a stable operation since then. In the current system, which was upgraded in 2016, 6144 CPU cores have been deployed as worker nodes for the WLCG, where each worker node consists of 24 CPU cores.

The ATLAS experiment developed the multi-core implementation of their software framework for reconstruction and simulation jobs, which provides an efficient memory sharing. In 2014, the experiment started to submit eight-core jobs using the above software framework to the Grid. The Tokyo Tier2 site has been processing this multi-core job and normal single-core job (e.g. user analysis job) separately using dedicated worker nodes and computing elements (static partitioning). However, we have often observed idle CPUs of the worker nodes due to this static partitioning when either multi-core or single-core jobs are not assigned to the site. Therefore, we started to evaluate an implementation of the dynamic partitioning of the worker nodes using HTCondor batch scheduler to reduce this idle CPU time, and have deployed a small cluster (1536 CPU cores) in the production.

For the dynamic partitioning, draining of single-core jobs is necessary in order to dispatch a new multi-core job into a worker node when the worker node is filled by only single-core jobs. This draining should be performed until the number of running multi-core jobs reaches a target share. In order to perform an efficient draining, we need to consider several parameters, such as the number of draining machines at the same time, based on properties of the jobs. In this presentation, improvement of CPU efficiencies by introducing the dynamic partitioning and optimization of the drain parameters in the Tokyo Tier-2 center will be reported.

Primary author: Dr KISHIMOTO, Tomoe (The University of Tokyo)

Co-authors: Prof. SAKAMOTO, Hiroshi (The University of Tokyo); Mr MATSUI, Nagataka (The University of Tokyo); Prof. MASHIMO, Tetsuro (The University of Tokyo); Prof. NAKAMURA, Tomoaki (KEK)

Presenter: Dr KISHIMOTO, Tomoe (The University of Tokyo)

Session Classification: Physics & Engineering I

Track Classification: Physics (including HEP) and Engineering Applications