



Contribution ID: 243

Type: Poster

Storage Element performance optimization for CMS analysis jobs

Tuesday 22 May 2012 13:30 (4h 45m)

Tier-2 computing sites in the Worldwide Large Hadron Collider Computing Grid (WLCG) host CPU-resources (Compute Element, CE) and storage resources (Storage Element, SE). The vast amount of data that needs to be processed from the Large Hadron Collider (LHC) experiments requires good and efficient use of the available resources. Having a good CPU efficiency for the end users analysis jobs requires that the performance of the storage system is able to scale with I/O requests from hundreds or even thousands of simultaneous jobs.

In this presentation we report on the work on improving the SE performance at the Helsinki Institute of Physics (HIP) Tier-2 used for the Compact Muon Experiment (CMS) at the LHC. Statistics from CMS grid jobs are collected and stored in the CMS Dashboard for further analysis, which allows for easy performance monitoring by the sites and by the CMS collaboration. As part of the monitoring framework CMS uses the JobRobot which sends every four hours 100 analysis jobs to each site. CMS also uses the HammerCloud (HC) tool for site monitoring and stress testing and HC will replace soon the JobRobot. The performance of the analysis workflow submitted with JobRobot or HC can be used to track the performance due to site configuration changes, since the analysis workflow is kept the same for all sites and for months in time. The CPU efficiency of the JobRobot jobs at HIP was increased approximately by 50 % to more than 90 %, by tuning the SE and by improvements in the CMSSW and dCache software. The performance of the CMS analysis jobs improved significantly too. Similar work has been done on other CMS Tier-sites, since on average the CPU efficiency for CMSSW jobs has increased during 2011. Better monitoring of the SE allows faster detection of problems, so that the performance level can be kept high. The next storage upgrade at HIP will consist of SAS disk enclosures which can be stress tested on demand with HC workflows, to make sure that the I/O-performance is good.

Author: Dr LINDEN, Tomas (Helsinki Institute of Physics (FI))

Co-authors: Dr BEHRMANN, Gerd (Nordic Data Grid Facility); Mr JOHAN, Guldmyr (CSC —IT Center for Science Ltd); Mr DAHLBLOM, Jonas (CSC —IT Center for Science Ltd); Mr HAPPONEN, Kalle (Helsinki Institute of Physics HIP)

Presenter: Dr LINDEN, Tomas (Helsinki Institute of Physics (FI))

Session Classification: Poster Session

Track Classification: Computer Facilities, Production Grids and Networking (track 4)