CERN-RRB-2011-096
14 OCTOBER 2011

PRINCIPAL LHCC DELIBERATIONS

20TH MEETING OF THE COMPUTING RESOURCES REVIEW BOARD

18 OCTOBER 2011

EMMANUEL TSESMELIS

SCIENTIFIC SECRETARY, LHCC

GENERAL

This document summarizes the principal LHCC deliberations concerning the World-wide LHC Computing Grid (W-LCG) Project at the Committee's sessions in June 2011 and September 2011.

The W-LCG continues to operate very well. The LHCC congratulates the W-LCG project team and the experiments in their successes at processing and analyzing the LHC data.

CONCERNS FROM THE PREVIOUS COMPUTING RESOURCES REVIEW BOARD

SUB-SYSTEM	CONCERN	STATUS
Resources for ALICE	The pledged resources do not fully match the ALICE requirements for the long term.	Computing resource availability remains a long-standing problem for ALICE.

STATUS OF THE W-LCG

The W-LCG has been running remarkably smoothly, allowing the experiments to record and process data at high rates and produce large samples of simulated data. At the same time many analyses have been completed to give results for the summer conferences, in several cases including data recorded only a few weeks before the conference.

EXPERIMENTS

GENERAL

The large increases in resources requested by the experiments at Tier-0 mean that in 2012, for the first time, the total Tier-0 budget will not be sufficient to satisfy all the demands. This is not just a financial constraint. The power and space available in Tier-0 are also limited, so a higher rate of growth has implications for the timescale when a new Tier-0 will be required. If the increased resource requests are justified then the CERN IT Department will need guidance on how to share the available resources among the experiments.

The Committee encourages the experiments to reconsider their requests and to reduce the demands on Tier-0 when and where possible. In cases where requests greatly exceed the available funding a more radical change in computing model may be required in order to remain viable. This might require a more prioritized approach to data processing, with compromises in the speed at which some channels can be analyzed.

Evolution of the LHC experiment computing models, and in particular of the strategies for data placement, require changes to the network infrastructure. The new scheme is called LHCONE. It is based on Open Exchange Points that enable any Tier-2 or Tier-3 site to connect easily to any Tier-1 or Tier-2 site, without overloading the general network. This requires each Tier-2 site to have a

minimum level of connectivity, depending on its size, which is sufficient to allow the resources and data at that site to be used effectively. These requirements will be included in the site Memoranda of Understanding and networking needs will be re-included as part of the computing resource requests and accounting.

Working groups have been established to plan the evolution of the W-LCG, taking into account the experience and understanding gained so far, the developments in technology and the needs to reduce operational effort and maintain long term support.

ALICE

ALICE has recorded 1.69 PB of data in 2011. Efficiency of CPU usage continues to be a problem because of spikes in the required memory allocation per core. While organized simulation and reconstruction jobs are under control, user analysis jobs have an average efficiency of only 35%. There are several causes and remedies have been developed to mitigate the problem. Basic data parameters are as expected, except for a substantial reduction in the reconstruction time and the raw data size for the PbPb data, with further reduction possible. The projected data storage requirements for 2011 are close to what were requested for disk and somewhat lower for tape. Assuming a similar LHC running schedule and duty cycle in 2012, the requests for resources in 2012 and 2013 are considerably higher. Compared to the requests made in April 2011, the tape storage at Tier-1 has decreased and the disk storage requests have increased, including at Tier-0. Since ALICE computing is dominated by the heavy-ion data recorded near the end of each year, the Tier-0 resources must be available in January of the following year for the first reconstruction of these data. There is a large shortfall in resources pledged and available compared to the ALICE requests. To live within their means ALICE could rely more on tape storage, with data copied to disk when required, and ALICE continues to use the GEANT3 software for simulation, because the more accurate GEANT4 is four times slower. This has an impact on data analysis and the Monte Carlo systematic uncertainties.

ATLAS

ATLAS has made changes to cope with the demands of high pile-up and limited computing resources, driven by the physics goals. The mean number of interactions per bunch crossing, $<\mu>$, has increased from 6.3 earlier in 2011, with $\beta^*=1.5$m, to $<\mu>=11.6$ with $\beta^*=1.0$m. The reconstruction code has been optimized and the event size reduced to lower the rate of increase with $<\mu>$, and work continues to reduce the CPU time required for simulation. The data distribution policy has been modified to use dynamic data placement, and physics analyses now use only derived and reduced data formats. Compared to what was expected for 2011, event sizes and disk usage have been lower, allowing a mean trigger rate of 275 Hz to be maintained (higher than the baseline 200 Hz), but up to 400-475 Hz at the start of fills. ATLAS is motivated to carry on running at 300-400 Hz at higher luminosities, with continued optimization to meet the physics goals. The Collaboration requests to increase its CPU resources at Tier-0 in 2012, to handle the increased luminosity and pile-up expected and to adapt to fit within the other resource limitations.

CMS

CMS trigger rates and event sizes in 2011 are much as expected, except that the reconstruction time is about 20% higher and increasing with pile-up. The memory requirements of the reconstruction code are higher because of the pile-up, resulting in a low CPU utilization efficiency when some nodes run out of memory. Installation of new CPU resources has mitigated this problem, and a new software release should improve the situation. Simulation production is a little ahead of schedule and data reprocessing will begin soon. The transition to using Analysis Object Data (AOD) samples for

user analysis is going well. The CMS request for CERN disk in 2012 has been increased to allow a higher fraction of the more recent AOD to be kept on disk. CMS believes this would allow more efficient use of its CERN CPU resources for analysis during periods when the LHC is not running. The details of the requests depend on the parameters for the 2012 LHC running, which are not yet known.

LHCB

In 2011 LHCb has been running with luminosity leveling and recording data at a fairly constant trigger rate of 3 kHz as expected. Data volumes and processing times have also been as expected, making full use of the LHCb computing resources. Some Tier-2 sites have been commissioned for reprocessing data, so that this can be done in parallel with the processing of new data at Tier-1 sites. A continuous, aggressive policy of cleaning up old data files from disk is required to make room for new and reprocessed data samples. There is a shortfall in the available disk resources relative to that pledged. Unless this can be found (not necessarily at CERN) then the new simulation production will have to be delayed until there is room to store the data. A problem with high memory demands in stripping jobs has been resolved by dropping the use of the POOL persistency package and using ROOT directly.