

CERN-RRB-2012-040  
20 APRIL 2012

# PRINCIPAL LHCC DELIBERATIONS

---

21<sup>ST</sup> MEETING OF THE COMPUTING RESOURCES REVIEW BOARD

24 APRIL 2012

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 December 2011 and March 2012.

**The W-LCG continues to operate very well. The Committee congratulates the W-LCG team and the experiments on their continued impressive performance.**

---

## 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.	ALICE has taken several measures to reduce its requirements for computing resources. Resources are still tight, but the LHCC is pleased to see that there has been a significant improvement to the long-standing shortfall between the resources available to ALICE and the requirements.

---

## STATUS OF THE W-LCG

---

Since the previous Computing Resources Review Board, much of the proton-proton data from the 2011 run have been reprocessed, the heavy-ion data from the 2011 Pb-Pb run have been processed, large samples of simulated proton-proton collision data at centre-of-mass energies 7 TeV and 8 TeV have been generated and reconstructed, and many physics analyses have been completed and published.

The shortage of disk supplies, which had been anticipated due to flooding in Thailand, has turned out not to be a major problem. However, funding for computing resources is likely to remain under severe pressure, as major equipment purchased some years ago will need to be replaced, in addition to providing any increase in capacity. This could limit the ability of some sites to honour the pledges made to install disk, tape and CPU resources in the future. The LHCC is **concerned** that these pledges, which have been agreed to following careful scrutiny and review of the resources requested, should be fulfilled. This applies as much to the Tier-2 sites as to the larger Tier-1 and Tier-0 centres. The W-LCG was designed to be a distributed computing facility with the Tier-2 sites forming an essential part of the model.

CERN has decided to locate the remote Tier-0 computing centre in Budapest. Prototyping will begin in 2013, ramping up to full production in 2014. Tendering is now being launched for a dedicated high-speed network connection to CERN. Consolidation work to provide additional critical power to the existing CERN Tier-0 building is also in progress.

---

## EXPERIMENTS

---

### GENERAL

The experiments have invested a great deal of effort to improve the speed and efficiency of their reconstruction code and to limit the amount of stored data, by compressing event sizes, reducing the numbers of copies kept on disk and aggressively deleting old versions. This has allowed them to operate within the limits imposed by the resources available, but only by developing flexible strategies that allow data and processing tasks to be moved around to make the most effective use of disks and CPU power.

The LHC run in 2012 is expected to be followed by a period of well over a year without any data-taking. During this long shutdown many of the computing resources that would normally be in use to process new data will be available for other purposes. ATLAS, CMS and LHCb all propose to take advantage of this opportunity by significantly increasing their trigger rates during the 2012 run, opening up their acceptance to record additional physics events that would otherwise fall outside their normal trigger thresholds. This would allow a broader physics programme, increased sensitivity in searches for new physics and more detailed systematic studies in high-precision measurements. The full analysis of these extra data would be delayed until the 2013-2014 long shutdown, when the computing resources would become available to handle the additional processing and storage. Increases of around 10-20% in the requests for resources in 2013-2014 are required in some cases. The details for each experiment are given below.

The LHCC **supports** the strategies for delayed analyses of part of the 2012 data, to maximise the physics potential from the 2012 LHC run, taking advantage of the unique situation provided by the long 2013-2014 shutdown. This is consistent with the Committee's previous recommendations to have a more prioritised approach to data processing, with compromises in the speed at which some channels can be analysed. However, the proposals are for additional data to be recorded on top of the baseline trigger rates. These strategies do require some extra computing resources, but the experiments should understand that funding is restricted. It might prove difficult to satisfy all the requests for storage and CPU resources in 2013-2014, in which case priorities must be set to allow the physics analysis to proceed along with simulation and other work, within the limits set by the resources available.

### ALICE

ALICE has improved its efficiency for production jobs to over 80%, lowered the number of reconstruction passes and the amount of simulated data and had a campaign to delete obsolete files. As a result the requirements for computing resources have been reduced since the previous Computing Resources Review Board. New prototype ALICE Tier-1 centres have been established in South Korea and in Mexico, and with discussions in progress for an Indian Tier-1 site. Resources are still tight, but the LHCC is pleased to see that there has been a significant improvement to the long-standing shortfall between the resources available to ALICE and the requirements.

### ATLAS

ATLAS has compressed the raw event size on disk, reduced the number of reprocessing passes to one per year and steadily improved the software to reconstruct high pile-up events more quickly. In 2012 the baseline rate in physics streams will be at 400 Hz averaged over a fill, with another 15 Hz of zero-bias events for pile-up overlay studies. ATLAS proposes, in addition, to record 75 Hz or more

in a stream for which the processing would be delayed until 2013, in order to allow a broader programme of B physics, for example. This requires 400 TB of tape storage and 200 TB of disk in 2013.

#### **CMS**

The latest version of the CMS reconstruction code shows excellent performance with high pile-up events. It is much faster than the previous version and uses less memory. This allows events to be reconstructed at 300 Hz, with the time between LHC fills being used to complete the processing. In 2012 CMS proposes roughly to double the rate of recording data, for the reasons explained above, with about half the data reconstructed promptly and the remainder “parked” in storage, to be reconstructed later at Tier-1 centres during the long shutdown in 2013-2014. They will need about 20% more Tier-1 CPU than in 2012 and 15% more Tier-2 CPU than estimated in 2011. No further increases in tape or disk requirements over those requested are anticipated, due to reductions in the size of events and the number of raw Monte Carlo events stored and the deletion of old data files.

#### **LHCb**

LHCb suffers from a disk space shortage due to the larger event sizes with high pile-up and the increase of trigger rate from 2 to 3 kHz agreed in 2011 to accommodate a charm physics stream. Installation of the resources pledged for 2012 should improve this situation. LHCb proposes a further increase in trigger rate to 4.5 kHz in 2012. The pledged resources for 2012 are insufficient to exploit the full physics potential at this rate and so they will initially write to disk only the micro-DST for some physics channels. In 2013, if about 20% more storage were available at Tier-1 and Tier-2, then the full DST could be “unlocked” for further selected channels, increasing the output by 50% and allowing enhanced analyses with detailed systematic studies as well as data-mining searches.