

CERN-RRB-2012-092  
23 OCTOBER 2012

# PRINCIPAL LHCC DELIBERATIONS

---

22<sup>ND</sup> MEETING OF THE COMPUTING RESOURCES REVIEW BOARD

30 OCTOBER 2012

EMMANUEL TSESMELIS  
SCIENTIFIC SECRETARY, LHCC

---

**GENERAL**

---

This document summarizes the principal LHCC deliberations concerning the World-wide LHC Computing Grid (WLCG) Project at the Committee's sessions in June and September 2012.

**The WLCG continues to operate very well. The Committee congratulates the W-LCG team and the experiments on their continued impressive performance. This has allowed an extensive harvest of physics results to be presented in a timely manner, highlighted by the Higgs-like particle discovery announcement based on data taken up until only a few weeks earlier.**

---

**CONCERNS FROM THE PREVIOUS COMPUTING RESOURCES REVIEW BOARD**

---

SUB-SYSTEM	CONCERN	STATUS
Computing Resources	Funding for the pledged WLCG computing resources remains under pressure.	There remains a potential danger that LHC physics output could be limited by having insufficient computing resources.

---

**STATUS OF THE W-LCG**

---

In total 19 PB of data have been written in 2012, with a monthly rate close to 4 PB. The scale of the CASTOR storage system is approaching 100 PB, with data transfer rates increased to 3-4 GB/s input and ~15 GB/s output. When LHC is running, CERN exports around 2 GB/s data, and global transfers can exceed 10 GB/s.

WLCG operations over the summer have been smooth, with no particular problems. The heavy user load has led to a slowdown of response from the LSF batch system at CERN. Short-term corrective actions have been taken and the long-term strategy is being reviewed.

The CERN Computer Centre extension is on track for completion in November 2012. Also, the Wigner Centre in Budapest is on schedule, with first installations expected in 2013. The two 100 Gb/s network lines to CERN are being procured. A fraction of *lsbatch* has been run for two months with an artificial 35 ms delay to test the expected CERN-to-Budapest latency, with no observed adverse effects so far.

---

## EXPERIMENTS

---

### GENERAL

For the LHC experimental programme as a whole, the extension of the 2012 run implies a need for  $\sim 20\%$  additional resources, which cannot be satisfied in most cases. Requests for 2013 have been revised to take into account the effects of the extended run, and 2014 requests are close to those of the revised 2013 ones. During Long Shutdown 1 (LS1), computing activities will include a full reprocessing of the complete 2010-2012 data sets, simulations for 13 TeV center-of-mass running, and physics analysis. The experiments have made first rough estimates of the potential requirements for 2015 when data-taking resumes in full. Both ATLAS and CMS aim for an increased trigger rate of 1 kHz, and LHCb has the potential to double its current rate. Only ALICE would expect a modest increase of data volumes, by  $\sim 20\%$ . In total, this translates into a significant step up in CPU, disk and tape requirements, at all Tiers. An annual increase of  $\sim 25\%$  from 2012 levels would be needed to reach this. (Beyond current projections, 2016 and 2017 may require an even steeper rise, as data continue to accumulate.) Taking Moore's Law into account, the required rise up to 2015 might be possible with a roughly flat budget.

Adequate funding cannot be taken for granted, however, and the new computing equipment will require significant annual investments. In addition to the ongoing financial problems in many countries, the EU Framework 7 programmes are coming to an end by 2014 or earlier, while the opportunities for support under the new Horizon 2020 programme are only currently being evaluated. This implies a significant reduction in the number of Grid-related personnel. The remaining resources will have to be prioritised for operations, leaving little or no room for further development effort or for support of the experimental frameworks.

There is thus a potential danger that LHC physics output could be limited by having insufficient computing resources. On the one hand, the LHCC recognises that the struggle to maintain healthy computing budgets must continue. The detrimental impact to physics of inadequate computing provision must be stressed at national and international levels. On the other hand, the collaborations should aggressively pursue ways to make more efficient use of new computing technology, such as multi-core CPUs and GPUs (Graphics Processing Units). Other efforts to speed up existing code should also continue. Otherwise, it might be necessary to restrain trigger rates or to increase data reconstruction and analysis lead times, in order to work within the resources available.

The LHCC requests the collaborations to formulate and present their updated computer models for approval well in advance of the 2015 restart. Demands for significant increases to computing resources, beyond those anticipated by Moore's Law, should be matched by realistic compensatory measures. The LHCC and the Computing Resources Scrutiny Group should cooperate in the review, with the former taking main responsibility for the physics case and the latter for the technical and financial aspects.

### ALICE

The year 2012 has been a standard data-taking year for ALICE. The proton-proton run has an emphasis on rare triggers. The luminosity loss from noise early in the runs is roughly compensated by the prolonged running period. The pilot pPb run gave 2.5 million events, and the Collaboration is looking forward to the long pPb run in February 2013. This will not require any additional computing resources. All Grid sites have provided stable performance. Improved compression of raw data has reduced the data volumes. Until September  $\sim 1$  PB has been recorded. Two copies are

stored, one at Tier-0 and one shared across the Tier-1s. Organised activities now account for ~80% of the CPU usage, a fraction still increasing with the introduction of analysis trains, at the expense of chaotic user analysis.

#### **ATLAS**

The ATLAS computer usage predictions turned out to be accurate, modulo the LHC run extension. ATLAS is CPU-limited for Monte Carlo production but, due to the efforts of the computing sites, more CPU is available than was formally pledged. Disk space will be tight in 2012, but early deployment of 2013 pledges should resolve the issue. To date 7 PB of data have been stored, of which 1.7 PB are raw data and 4.3 PB event summary data. The latter format is optional and could be reduced. ATLAS Tier-1 (and Tier-0) sites are performing very well. There are two upcoming data reprocessings, for 8 TeV centre-of-mass energy proton-proton data and for 2011 heavy-ion data. Also, Tier-2 plays a successful role in digitisation and reconstruction, and in group production.

#### **CMS**

CMS computing has performed well, with live-time and trigger rates close to expectations. Even with the extra parked data to repack, Tier-0 is working well, with a CPU efficiency of 80-90%. Occasionally, Tier-0 can spill over into the public queues, mostly to catch up after long or particularly high-luminosity runs. Utilisation of Tier-1 CPU averaged to 102% of pledged capacity over the last 12 months, and Tier-1 tape volumes agree with plans. To better manage Tier-2 disk space, popularity information is used to make reasonable choices of what should be cleaned up. The speed of the reconstruction code has been improved by a factor of eight since 2010, and memory reduced by ~40%. The next generation of analysis tools has moved into integration testing.

#### **LHCB**

LHCB computing production activities in 2012 have included prompt processing of new data, “swimming” of 2011 stripped data, and Monte Carlo production for both 2011 and 2012 conditions. The reprocessing of the complete 2012 dataset started on 17 September 2012. Up until now resources have been sufficient. The CPU usage has been ~50% of that pledged, as expected for the first half of the year, but with reprocessing ongoing it is now at 100%. Only ~60% of 2012 data will be available for the winter 2013 conferences. The 2012 disk pledge is sufficient, after some aggressive reduction of copies. There is a major shortage of tape, however, projected to reach ~5 PB by March 2013. The key reason is a new reconstructed data format (FULL.DST), which contains a copy of the raw data. Previously, re-stripping required the SDST and raw data types to be staged synchronously, so the new format eases the demand on the tape systems, at the expense of size. The extension of the proton-proton run and the LHCB operation at higher instantaneous luminosity and increased High Level Trigger (HLT) output rate further aggravates the situation. The LHCC hopes that measures can be taken to overcome this shortage.