

CERN-RRB-2013-017

11 APRIL 2013

PRINCIPAL LHCC DELIBERATIONS

23RD MEETING OF THE COMPUTING RESOURCES REVIEW BOARD

16 APRIL 2013

EMMANUEL TSESMELIS

SCIENTIFIC SECRETARY, LHCC

GENERAL

This document summarises the principal LHCC deliberations concerning the World-wide LHC Computing Grid (WLCG) Project at the Committee's sessions in December 2012 and March 2013.

The WLCG continues to operate very well. The LHCC congratulates the WLCG and the experiments for the efficient data processing and for the successful completion of the first LHC running period.

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. A common document by the LHC experiments on the optimisation of their computing models is expected for the September 2013 LHCC session, in time to be presented to the Computing Resources Review Board session in October 2013.

STATUS OF THE WLCG

The completion in April 2013 of the European projects – the European Middleware Initiative EMI (for middleware) and the European Grid Initiative EGI-SA3 (for support to communities of heavy computing users) - will have a significant impact on CERN groups supporting the experiments. In the short term, the measures are oriented towards the optimisation and consolidation of activities at CERN, including re-prioritisation of functions and common solutions on WLCG operations, service and support. These improvements have already started, but there is a gap of about one year between the funding programmes. In the longer term, the overall organisation needs to take into account the new features expected from the continuation of the European Union funding programmes. In particular, the multi-disciplinary aspects are considered as a necessary ingredient of the future European Union funding programmes and they have to be considered during the building of the next round of project proposals.

EXPERIMENTS

GENERAL

The LHC experiments have expressed their satisfaction concerning the availability of WLCG resources in the past year, which made possible the timely processing of the largest LHC data sample collected thus far as well as the completion of a large number of physics publications in 2012.

The LHCC noted some changes in the distribution of resources distribution: The Tier-1 at ASGC (Taiwan) will stop to deliver resources for CMS as a consequence of reduced funding. Two new Tier-1 centres are being implemented, namely KISTI (Korea) as resource provider for ALICE and Russia (JINR and Kurchatov) for all four LHC experiments. New resources are planned for ALICE from Mexico, which together with KISTI will provide 7-8% of the ALICE needs.

In general, the profile of the expected needs for resources in the 2013-2015 period are influenced by the imbalance between the moderate requirements over the shutdown period followed by a significant increase for 2015. The levelling requested by the funding agencies needs to be carefully implemented, such that the longer-term perspectives are not jeopardised and the production of physics results continue with similar efficiency as in the past years. The increase in the request for resources for 2015 are mostly due to the scaling factors induced by the increase in cross-sections, off-time overlays etc. and there are several ideas on how to reduce the usage of resources by optimising the computing models. The LHCC have recommended in its session in December 2012 a coordinated action to update the computing models before the start of LHC running in 2015 and work has started along these lines. The LHCC will monitor closely the technical improvements as well as the physics studies which have to underline the decisions for data taking configurations after Long Shutdown 1 (LS1).

The WLCG resources are complemented during the shutdown by using the available resources from the High-Level Trigger (HLT) farms of the experiments. The LHCb farm is already in production and delivered close to 50% of the CPU for Monte Carlo production in January 2013. Technical developments are actively pursued by ATLAS and CMS using Cloud-like software to aid deployment, integration, and future reconfiguration of farms. These efforts will lead to significant gains, but the farms will not be continually available during LS1 due to various upgrades and interventions.

In addition, ATLAS and CMS have used non-standard resources in an opportunistic mode. CMS profited from a few weeks of running at the San Diego Supercomputer Centre (SDSC), while ATLAS had access to the Amazon Cloud resources. This demonstrates the advantage of being prepared to rapidly make use of such resources (e.g. via Cloud interfaces and smartly packaged and deployable services).

The LHCC **congratulates** the WLCG and the experiments for the efficient data processing and for the successful completion of the first LHC running period. It encourages the experiments to continue the optimisation of their computing models in preparation for the next running period and to pursue the studies on the impact on physics. A common document on computing is expected for the September 2013 LHCC session, in time to be presented to the Computing Resources Review Board session in October 2013. The LHCC also supports the arrangements to continue the middleware and support projects and recommends that a common vision be installed in preparation of the next European Union funding plan. The LHCC acknowledges the progress towards using the available resources from HLT farms during LS1 and recommends the installation of a longer-term solution, also allowing for intensive use of these resources during the short technical stops after LS1.

The usage of opportunistic resources by the experiments (on supercomputing centres or private Clouds), already demonstrated on a few occasions, in conjunction with the need to place the WLCG in a multi-disciplinary context, may open new opportunities for LHC computing.

ALICE

For ALICE, the request for resources is essentially flat for 2014 and 2015, assuming the KISTI and Mexico Tier-1 centres become available. The close collaboration of the experiment with the computing experts at the Tier-1 and Tier-2 centres has led to a dramatic improvement in the efficiency of the CPU usage for analysis (from about 65% in September 2012 to about 85% in February 2013). In addition, the experiment has introduced a new processing scheme, requiring a reduced number of full-processing passes.

ATLAS

The ATLAS needs fit within a flat budget for 2014 and 2015 with the assumption that the event sizes and CPU/event remain at 2012 levels. Since these assumptions are expected to be significantly affected by the new running conditions, new ideas have to be deployed during LS1 to optimise the CPU use, event sizes, memory, etc. These efforts have already started and an ambitious plan for improvements has been presented.

CMS

CMS estimates a large increase in CPU resources in the case that the parameters from 2012 are unchanged. In particular, there is an increase in the reconstruction time due to off-time overlays from the previous bunch occurring in the 25 ns configuration. A few solutions for optimisation are already considered, consisting of software tuning; the use of the Tier-1 centres for prompt reconstruction; only one full reprocessing per year; the commissioning of Tier-2 centres for Monte Carlo reconstruction; and usage of remote data access as a result of data federation work in the last 18 months. The remaining factor of two in CPU seems to be to very hard to gain, but the work on further optimisation has started. The increase in the required storage is more modest and is concentrated on data only.

LHCb

LHCb has sufficient CPU for LS1 needs but is limited by the available disk space. This situation may lead to an under usage of the CPU in 2014 and to disturbances in 2015 since the Monte Carlo production is delayed. A significant effort is needed to reduce disk consumption, so both the size of the Data Summary Tapes (DSTs) and the replication policies are being considered. The software is now under scrutiny, with the ambitious objective to efficiently exploit the new hardware capabilities (parallelisation etc.).