

15 February 06

Global Grid service for LHC computing succeeds in gigabyte-per-second challenge

Today, at the international Computing for High Energy and Nuclear Physics 2006 conference (CHEP'06) in Mumbai, India, the Worldwide LHC Computing Grid collaboration (WLCG) officially announced the successful completion of a service challenge. This challenge involved sustaining a continuous flow of physics data on a worldwide Grid infrastructure at up to 1 gigabyte per second. The maximum sustained data rates achieved correspond to transferring a DVD worth of scientific data from CERN every five seconds.

The Worldwide LHC Computing Grid collaboration (WLCG) has officially announced the successful completion of a service challenge that involved sustaining a continuous flow of physics data on a worldwide Grid infrastructure at up to 1 gigabyte per second. The maximum sustained data rates achieved in this challenge correspond to transferring a DVD worth of scientific data from CERN every five seconds. The announcement was made today during the international Computing for High Energy and Nuclear Physics 2006 conference (CHEP'06) in Mumbai, India.

The data was transferred from CERN¹ in Geneva, Switzerland to 12 major computer centres² around the globe. Over 20 other computing facilities, were also involved in successful tests of a global Grid service for real-time storage, distribution and analysis of this data. The completion of this service challenge is a key milestone on the way to establishing the necessary computing infrastructure for the Large Hadron Collider (LHC), the world's largest scientific instrument, which is scheduled to startup in 2007 at CERN. The results represent a significant step forward compared to a previous service challenge in early 2005 that had involved just seven centres in Europe and the USA, and achieved sustained rates of 600 megabytes per second.

Commenting from Mumbai on the significance of the results, Jos Engelen, the Chief Scientific Officer of CERN said "Previously, components of a full Grid service have been tested on a limited set of resources, a bit like testing the engines or wings of a plane separately. This latest service challenge was the equivalent of a maiden flight for LHC computing. For the first time, several sites in Asia were also involved in this service challenge, making it truly global in scope. Another first was that real physics data was shipped, stored and processed under conditions similar to those expected when scientists start recording results from the LHC."

The goal of the WLCG is to unite the efforts of established scientific Grid infrastructures to provide sufficient computational, storage and network resources to fully exploit the scientific potential of the four major LHC experiments: ALICE, ATLAS, CMS and LHCb. These experiments will be studying the fundamental properties of subatomic particles and forces, providing insight into the origins of the Universe. They are expected to generate in total some 15 million gigabytes of data each year. WLCG uses a range of national and international Grid infrastructures, including, the Enabling Grids for E-Science (EGEE) project and the Open Science Grid (OSG)³.

LHC scientists designed a series of service challenges to ramp up to the level of computing capacity, reliability and ease of use that will be required by the worldwide community of over 6000 scientists working on the LHC experiments. During LHC operation, the major computing centres involved in the Grid infrastructure, so-called Tier-1 centres, will collectively store the data from all four LHC experiments, in addition to a complete copy being stored at CERN.

Much of the data analysis will be carried out by scientists working at over 100 Tier-2 computing facilities in universities and research laboratories in over 30 countries. These scientists will access the data via the Grid resources that the WLCG is bringing together. Already today, these computing facilities provide a combined computing power of over 20,000 PCs, and this number is expected to reach 50,000 by the time the LHC is operational. During the recent service challenge, the participating computing centres sustained more than 12,000 concurrent computing jobs.

Speaking on behalf of the organizers of CHEP 06, Sabyasachi Bhattacharya, Director of the Tata Institute of Fundamental Research in Mumbai, remarked, “the fact that this announcement is being made in India reflects the truly global significance of these new results. This sort of collaboration, which we in India are delighted to be taking part in, provides an excellent example of what scientists from around the world can achieve together when they have a clear, common goal.”

Kors Bos, WLCG Grid Deployment Board chairman, expressed satisfaction with the recent results: “Not only did we achieve our gigabyte per second goal for this service challenge, but all sites achieved their target data rates and many went well beyond this. The challenge involved interoperation between four different mass storage system technologies and required a big technical push. The staff at all the sites involved deserves credit for putting in the extra effort required.”

EGEE Project Director Bob Jones, noted: “The significance of these results goes well beyond the immediate needs of the high energy physics community. What has been achieved here is nothing less than a breakthrough for scientific Grid computing. The lessons learned from this experience will surely benefit other scientific domains such as biomedicine, nanotechnology and environmental sciences in their future use of Grids.”

OSG Executive Director, Ruth Pordes, was enthusiastic about the progress achieved: “Just as important as the data transfer rate is the fact that the scientists are beginning to test their computing models under realistic conditions, and are interacting closely with the service providers in the computing centres to optimize this. The centres involved have built a strong collaborative spirit, and I am particularly pleased at the progress on interoperability between different Grids, illustrated by the ability we have recently demonstrated to send computing jobs between OSG and EGEE.”

The current service challenge is the third in a series of four leading up to LHC operations in 2007. The next service challenge, due to start in the summer, will extend to many other computing centres and aim at continuous, stable operations. That challenge will allow many of

the scientists involved to refine their computing models for handling and analyzing the data from the LHC experiments, in anticipation of the start of real data taking in 2007.

---Ends---




Figure 1: Histogram monitor by the Bhabha Atomic Research Centre, Mumbai, India, showing data transfer from CERN to 12 major computing centres [see footnote 2 for full names of centres] during the service challenge, demonstrating data rates up to 1 gigabyte per second.

For more information contact:

For CERN

François Grey

CERN

Phone: +41 22 767 1483

Email: Francois.Grey@cern.ch

For CHEP'06

Atul Gurtu

Tata Institute of Fundamental Research, Mumbai

Phone: +91-22-2278 2357

Email: gurtu@tifr.res.in

For EGEE

Joanne Barnett

EGEE External Relations Officer, TERENA Secretariat

Phone: +31 (0)20 530 4488

Email: barnett@terena.nl

For OSG

Katie Yurkewicz

U.S. Grid Communications, Fermilab

Phone: +1 630 840 2877

Email: Katie@fnal.gov

¹ CERN, the European Organization for Nuclear Research, has its headquarters in Geneva. At present, its Member States are Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom. India, Israel, Japan, the Russian Federation, the United States of America, Turkey, the European Commission and UNESCO have Observer status.

² The computing facilities involved in this service challenge were: Academia Sinica Grid Center (ASGC) in Taipei; Brookhaven National Laboratory (BNL) in Brookhaven, NY, USA; CCIN2P3, the Computing Center of the National Institute of Nuclear Physics and Particle Physics (CCIN2P3) in Lyon, France; the German Electron Synchrotron Laboratory (DESY) in Hamburg, Germany; Fermi National Accelerator Laboratory (FNAL) in Batavia, Illinois, USA; Forschungszentrum Karlsruhe (FZK) in Karlsruhe, Germany; the National Center for Research and Development in Technology, Computer Science and Data Transmission (INFN-CNAF) in Bologna, Italy; the Nordic DataGrid Facility (NDGF) a distributed facility in Denmark, Finland, Norway and Sweden; Port d'Informació Científica (PIC) in Barcelona, Spain; the National Center for Computing and Networking Services and the National Institute for Nuclear Physics and High Energy Physics (SARA-NIKHEF) both based in the Netherlands; the Rutherford Appleton Laboratory (RAL) in Oxfordshire, UK; the National Laboratory for Particle and Nuclear Physics (TRIUMF) in Vancouver, Canada.

³ Further information about these Grid infrastructures can be found at:
Worldwide LHC Computing Grid (WLCG) : <http://www.cern.ch/lcg/>
Enabling Grids for E-Science (EGEE): <http://www.eu-egee.org/>
Open Science Grid (OSG): <http://www.opensciencegrid.org/>