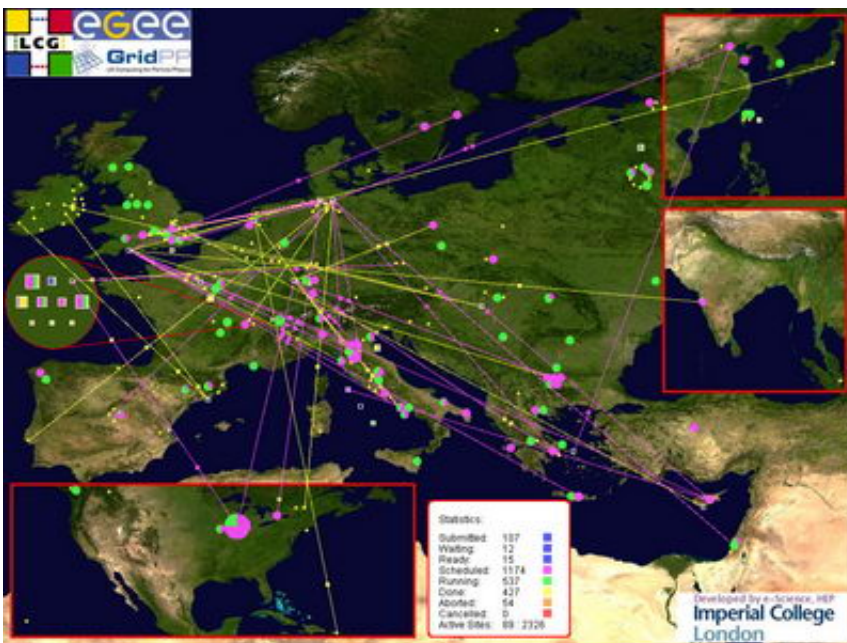


'Maiden Flight' for LHC Computing Grid Breaks Gigabyte-per-Second Barrier

February 15 2006



Map with data transfer rates superimposed.

UK physicists have successfully taken part in a challenge to test an international scientific computing Grid under working conditions. During the week-long challenge, the LHC Computing Grid sustained transfer rates of a gigabyte per second. - a world first for a permanent, international Grid using scientific data.

Completion of the tests was announced on February 15th by the Worldwide LHC Computing Grid collaboration (WLCG) at the

international Computing for High Energy and Nuclear Physics 2006 conference (CHEP'06) in Mumbai, India.

The maximum sustained data rates achieved correspond to transferring a DVD worth of scientific data every five seconds. Professor Tony Doyle, leader of the UK particle physics Grid, commented, "At these rates, it would take 25 days to transfer the nearly 400,000 films listed at IMDB.com and only an hour and a half to transfer the 1000 films produced each year by the Mumbai-based Bollywood. It might take a bit longer to watch them all!"

The data was transferred from CERN in Geneva, Switzerland to 12 major computer centres around the globe. CCLRC Rutherford Appleton Laboratory (RAL) in Oxfordshire represented the UK, receiving data from CERN at nearly 200 megabytes per second via the UKLIGHT network managed by UKERNA. Over 20 other computing facilities worldwide, including those at the UK's Universities of Edinburgh, Lancaster and Imperial College London, were also involved in successful tests of a global Grid service for real-time storage, distribution and analysis of this data.

The LHC Computing Grid is needed to manage the data deluge from CERN's Large Hadron Collider (LHC). When it starts in 2007 the LHC will probe the physics of the Universe at the earliest moments after the Big Bang - and in the process produce 15 million Gigabytes of data a year that need to be shared, stored and analysed around the world.

The results of this challenge represent a significant step forward compared to a previous service challenge in early 2005 that involved just seven centres in Europe and the USA, and achieved sustained rates from CERN 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."

RAL is the UK's national computer centre for the WLCG project. It works with sites at UK universities to form GridPP, the UK particle physics Grid, which currently consists of more than 4,000 CPUs and 250 terabytes of disk storage capacity. Dr Andrew Sansum, the manager of the particle physics computing centre at RAL, was pleased with the results of the tests, "By receiving nearly 200 megabytes per second from CERN, we went well beyond our target data rate. With the installation of the new academic network, SuperJanet5, later this year, we plan to double these data rates - and then we'll really be approaching the speeds we need when the LHC comes on line in 2007".

The next step for GridPP in the UK is to schedule similar tests between RAL and the UK universities, who are grouped into four regional Tier-2s. When the LHC is running a fraction of the data sent to RAL will be duplicated to the Tier-2s. During the next few months GridPP will be testing that these transfers work at the required rates, and then that these rates can be maintained simultaneously with data flowing into RAL from CERN.

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."

The goal of the WLCG is to provide sufficient computational, storage and network resources to fully exploit the scientific potential of the four major LHC experiments. These experiments will be studying the fundamental properties of subatomic particles and forces, providing an 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 GridPP, the Enabling Grids for E-Science (EGEE) project and the Open Science Grid (OSG).

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.

Source: PPARC

Citation: 'Maiden Flight' for LHC Computing Grid Breaks Gigabyte-per-Second Barrier (2006, February 15) retrieved 23 April 2024 from <https://phys.org/news/2006-02-maiden-flight-lhc-grid-gigabyte-per-second.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.