

CERN openlab adds a new dimension to Grid computing

July 6 2004

Geneva, Switzerland 5 July 2004. The CERN openlab for DataGrid applications, a partnership between CERN , the European Organization for Nuclear Research, and five leading IT companies – Enterasys Networks, HP, IBM, Intel and Oracle – has announced a series of server and storage technical results regarding the first global science Grid – the Large Hadron Collider Computing Grid project, LCG. The announcement was made at the recent annual sponsors meeting of the CERN openlab.

The openlab partners have demonstrated that a cluster of 40 HP servers running 64-bit Intel® Itanium® 2 processors can be successfully integrated with the LCG, which involves over 60 major scientific computing centres in Europe, North America and Asia. The openlab partners have also completed intensive testing of IBM's SAN File System to demonstrate scale-out capabilities of the new storage software.

With this landmark addition of servers, the openlab partners have proven that the LCG, otherwise based on 32-bit processors, can be extended to a truly heterogeneous computing environment. This is crucial for the future evolution of this Grid, as it must grow rapidly in capacity and power to prepare for the tremendous data storage and analysis requirements of CERN's Large Hadron Collider (LHC) project. The LHC is expected to produce some 15 petabytes of data per year after it is switched on in 2007. Thousands of physicists will sift through this data for years to come, analysing it for tell-tale signs of new fundamental particles that will provide insights into the early origins of our Universe.

The CERN openlab, a three-year industrial cooperation formally launched in January 2003, marked its halfway mark at the annual sponsors meeting on June 22nd, and has already tallied a number of impressive technical results. Together, the partners have built the CERN opencluster, a state-of-the-art system for testing prototype Grid applications of increasing power and functionality. The open, collaborative environment of the partnership places an emphasis on a common development programme for data-intensive Grid computing based on open standards. This includes a 28 terabyte high-end storage system and advanced storage management software, supplied by IBM, state-of-the-art switching and routing equipment from Enterasys Networks, and the advanced, grid-enabled Oracle® Database 10g.

CERN's Director General, Dr. Robert Aymar, described the contribution of the CERN openlab as being of crucial importance to the LHC project. Looking ahead, he noted that "The CERN openlab provides a role model for how CERN and its academic partners may in future wish to organise collaboration between the private and public sector, in order to develop the many new technologies that will surely be needed for endeavours beyond the LHC. As the results so far show, CERN openlab has effectively established a framework for collaboration between multiple industrial partners, in a pre-competitive spirit and based on open standards."

Other key results obtained with the CERN opencluster include a data challenge where storage-to-tape rates of over 1GB/s were maintained for hours, corresponding to the maximum rates at which data from the LHC will need to be stored to a primary tape backup. The high-speed switching environment sponsored by Enterasys Networks played a crucial role in that result. Also, some of the HP server nodes with Intel Itanium® 2 dual processors contributed to the Internet-2 landspeed record that was set last October by CERN and partner Caltech, during Telecom 2, which demonstrated the rate – more than 1 terabyte in 30

minutes – at which data from the LHC will need to be distributed to the LCG’s so-called Tier-1 centres around the globe, for local storage and analysis.

Additionally, the researchers completed breakthrough testing on IBM’s Storage Tank™ storage management technology, which is used in IBM’s latest storage software product the IBM TotalStorage® SAN File System. The software is designed without inherent limitations on the amount of storage that can be supported. The project recently managed more than 100 simultaneous SAN File System clients and over 28 terabytes of storage distributed among 10 storage servers. These tests measured data rates exceeding the initial expectations from CERN.

Recently, Oracle-sponsored researchers in the CERN openlab managed to increase the availability of CERN’s grid computing environment by significantly reducing downtime of its catalogue, whose job is to ensure correct mapping of filenames and file IDs and identified a way to upgrade an Oracle database from one release to another whilst keeping the database operational. Both results represent new ways to use existing technology to address the challenges of grid-based computing.

The original press release can be found [here](#).

Citation: CERN openlab adds a new dimension to Grid computing (2004, July 6) retrieved 6 May 2024 from <https://phys.org/news/2004-07-cern-openlab-dimension-grid.html>

<p>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.</p>
--