CERN openlab Results Summary
January 2003 – December 2005

CERN openlab
for DataGrid applications

enterasys
Networks that Know

hp
invent

IBM

intel

ORACLE

The CERN openlab concept

In January 2003, CERN openlab was launched as a new framework for multilateral projects between CERN and the IT industry. At that time HP, Intel and Enterasys were contributing key equipment to the openlab's first major three-year project, the CERN opencluster. IBM joined the partnership a few months later, and Oracle by the end of that year. These partners have brought state-of-the-art technologies – both hardware and software – to the collaboration, each primarily in a specific technological area: processors for Intel, servers for HP, storage management for IBM, switches and routers for Enterasys and databases for Oracle.

CERN openlab partners committed to a three-year partnership, based on sponsoring of equipment, young researchers at CERN, summer students and various training and outreach events, as well as a considerable investment of time and effort by the partners own engineers. In addition to the CERN openlab partner status, a contributor status was created for shorter-term involvement in specific technological areas, the first contributor being Voltaire in 2004. At the end of 2005, the first phase of CERN openlab came officially to a close.

Wolfgang von Rüden
Head of the CERN openlab
and Head of CERN's IT Department

The focal point of the first three year phase of CERN openlab has been to build a state-of-the-art computing cluster – the CERN opencluster – and test and validate its performance in CERN's uniquely demanding computing environment, as well as on the global computing Grid that CERN and its international partners are managing. Over 20 persons at CERN, including technical experts in a variety of IT-related fields as well as young scientists and engineers, were closely involved in the openlab activities. Annual Board of Sponsors meetings with top managers of the partner companies provided opportunities to review the technical agenda. The list of successes achieved in this first phase of CERN openlab amply justifies the efforts that CERN and its partners have put into this project, and some of these are detailed on the next pages.

Besides technical achievements, a number of joint workshops and press events were organized during the course of the three years. A particularly significant event for CERN openlab was the Computing in High Energy Physics conference, held in Interlaken in September 2004. The level of participation in plenary talks and in terms of sponsorship from the partners was unprecedented for such an event, and contributed directly to the high quality of the event, which hosted a record number of over 500 delegates.

In conclusion, I would like to thank all the openlab members for their generous and unwavering support of the CERN opencluster project, and I look forward to developing new projects of similar scope and ambition in a second phase of CERN openlab.

The CERN opencluster

The CERN opencluster comprises over 100 **Intel** Itanium® dual processor **HP** Integrity servers as well as several Intel Xeon 64-bit HP Proliant servers. Both processors types were used to accelerate the conversion to 64-bit addressability with special emphasis on the validation of High Energy Physics software, middleware and Scientific Linux. Many servers were equipped with Intel's 10 GbE Network Interface Card (NIC) for use in LHC Computing Grid Data Challenges.

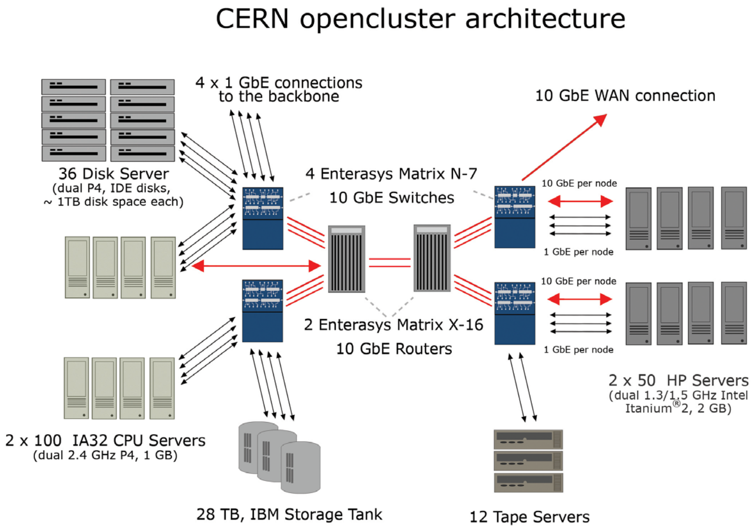

Enterasys Networks contributed three generations of network routers. The initial set of Enterprise Routers, ER16, were followed by four 10 Gb/s Matrix N-7 switches with an additional capacity for 1 Gb/s attachments (in total 20 10-GbE and 300 1-GbE ports). Towards the end of the project, Enterasys also supplied two of their latest generation Matrix X-16 10Gb routers. **Voltaire** supplied a 96-way Infiniband switch and a correspond ing number of host adapters.

IBM sponsored 28TB of high-end SCSI storage together with the IBM TotalStorage SAN File System storage subsystem. In its largest configuration the subsystem was equipped with 15 storage controllers.

Oracle streams enabled consistent replication of physics data between CERN and Tier-1 centres in Europe, North America and Asia.

The cluster was used extensively for benchmarking and development. After a few months, it became clear that the I/O capacity of the HP servers was excellent, and an Internet2 Land Speed record was achieved in time for Telecom 2003. Subsequently many tests were carried out on a transatlantic scale, together with Datatag, the ATLAS experiment, Caltech (CMS) and others. After multiple tests of this nature (which were done memory to memory) the focus shifted to "next-generation" disk servers that could drive both the 10Gb network and a storage subsystem at a compatible speed. Two 4-way Itanium servers (HP Infinity RX4640), each with three RAID controllers were deployed for this purpose.

The Itanium servers participated in several ALICE Data Challenges and LCG Service Challenges. The excellent floating-point computing capability of the Itanium processors led to the agreement to use between 20 and 40 servers as a service for Computational Fluid Dynamics codes for verification of the cooling setup in the LHC detector caverns. A separate IB switch was purchased from Voltaire for this purpose.

Key CERN openlab Partner Results

Enterasys

A wide range of tests of the Enterasys networking equipment were carried out. For example, the Matrix N-7 switches were interconnected between themselves with 10Gbps Ethernet, to demonstrate the scalability of the technology. Comprehensive tests of the Matrix X-16 demonstrated that the hardware is fully capable of providing the expected bandwidth for such a system the size of the opencluster. In particular, Spirent 10Gbps Ethernet Smartbits tests verified the 10Gbps performance of the X16 at all frame sizes, and demonstrated that performance was unaffected by the number of flows. In addition, the Enterasys equipment played a central role in a series of data and service challenges carried out by the LHC Computing Grid (LCG) in which the CERN opencluster took part.

HP

One of the most significant results from the CERN openlab partnership was the decision of HP to allocate a substantial number of Itanium processors to LHC Computing Grid from its Bristol and Puerto Rico computer centres. HP thus became the first industrial partner on this Grid. This move was based on experience gained with the CERN opencluster, and the openlab technical team ported the LCG middleware stack to Linux/64-bit and played an active part in helping the HP sites connect to LCG. Other significant software work has involved comprehensive testing HP's SmartFrog solution for automatic deployment on clusters, as well as Xen, an open-source virtualization package, on the opencluster. For Xen, work has focused on extending the Itanium port after its initial release by HP Labs, in particular by adding multiple domain support. The excellent I/O capabilities of the HP Integrity servers was key to achieving record results in multiple areas.

IBM

As part of the openlab partnership, IBM achieved breakthrough performance results for storage virtualisation software at CERN. Using Storage Tank, an IBM Research product for storage virtualisation which is marketed under the name IBM TotalStorage SAN File System, the internal tests shattered performance records by reading and writing data to disk at rates in excess of 1GB/second for a total I/O of over 1 Petabyte in a 13-day period. A significant goal of the tests was to monitor the IBM virtualization solution under a range of failure scenarios. The results confirmed the robust performance of this data management solution. Another area of investigation was integration of Storage Tank with CERN's storage management system for the LHC experiments, CASTOR, where promising results were obtained, helping both the IBM software and CASTOR to be further developed.

Intel

It is anticipated that 64-bit computing will be crucial to the future of scientific computing, and thanks in particular to Intel's contribution to the CERN openlab initiative, the HEP community now has the chance to test its applications for 64-bit compatibility on the LHC Computing Grid infrastructure. Of the four LHC experiments, ALICE is the first to be 64-bit compatible, and the CERN opencluster has contributed to several ALICE data challenges over the last two years. The HEP community got an early opportunity to test its applications. In addition to the LCG middleware, several scientific packages, such as ROOT, Geant4, CLHEP and CASTOR have been ported by their authors. Compiler optimization has been another key theme of the Intel collaboration with CERN openlab, a particular focus being on C++ performance. The Intel Itanium servers equipped with an Intel 10-Gb NIC was used for CERN's claim of the Internet-2 Land Speed Record in 2003.

Oracle

One of the themes of database technology evaluation in the CERN openlab context has been to extend tests of specific technologies beyond CERN, to some of the Tier-1 sites of LHC Computing Grid. For example, thanks to CERN openlab, work on Oracle Streams, a technology to connect databases and ensure that data is replicated consistently was expanded to the propagation of physics data from databases at CERN to six Tier-1 sites in Europe, the US and Taiwan. As a result, Oracle Streams is being integrated in the Distributed Database Deployment project (3D) as a baseline replication technique. For the database services at CERN, work has focused on using a mixture of Oracle DataGuard to ensure storage redundancy and another Oracle product, Real Application Clusters (RAC), for computing redundancy. Oracle DataGuard turns out to be effective in maintaining a hot backup of a catalogue, and keeping the Grid environment available while patches are applied.

Impact of CERN openlab

CERN openlab and the Grid

From the early stages of planning the Large Hadron Collider (LHC), CERN's flagship scientific instrument which is due to start operation in 2007, it was clear that the storage and processing requirements of the LHC experiments would dwarf any facility that could reasonably be assembled at CERN. As a result, a distributed solution – a data and computing Grid – was chosen, based on the fact that CERN has over 250 institutional partners in Europe and 200 more in the rest of the world, most of which have significant computing resources to contribute to such a Grid.

The Grid architecture for the LHC also builds on a definite hierarchy of sites. CERN, the Tier-0 centre, is a central hub, storing a copy of all data on tape, and also distributing them to eleven Tier-1 centres. The Tier-1 centres will store data from one or more experiments, for access by regional Tier-2 centres, of which there are already over 100.

CERN has a three-fold Grid strategy. The LHC Computing Grid (LCG) project provides the core Grid services for the LHC experiments. The Enabling Grids for E-sciencE (EGEE) project, co-funded by the EU, runs a Grid infrastructure for a wide range of scientific and industrial applications. CERN is the lead partner of EGEE, and LCG is the flagship scientific application for this Grid. CERN openlab complements these two projects with longer-term test and validation of new technologies from industrial partners, to assess their merits for the future of the Grid.

The CERN opencluster contributed to many data challenges and service challenges run by LCG over the past three years. These include: contributing to a storage-to-tape challenge that achieved 1.1GB/s sustained in 2003; playing a key role in a data transfer challenge which achieved 5.44GB/s in 2003, a feat which won CERN and the California Institute of Technology the Internet-2 landspeed record and a place in the Guinness Book of Records; facilitating service challenges carried out in 2005, where CERN sustained continuous flows of data at 600MB/s to seven Tier-1 centres. In addition, joint workshops between CERN openlab partners and LCG and EGEE experts, on a wide range of technical themes, provided opportunities for all to share information about the rapidly evolving Grid technology that CERN is helping to pioneer.

CERN openlab and the world

CERN openlab results have been disseminated in wide range of international conferences and publications, as well as at partner-specific events. These results and publications are documented in four annual reports of CERN openlab, which are available on the CERN openlab website, www.cern.ch/openlab. In addition, key results of CERN openlab have been the subject of over a dozen press releases from CERN, as well as a similar number of press releases from industry, generating a large number of articles in the printed press and on the Web. CERN is regularly visited by top delegations from governments and industry, as well as customer and press visits organised by openlab partners. These groups are briefed about CERN openlab in a dedicated VIP meeting room known as the CERN openlab openspace.

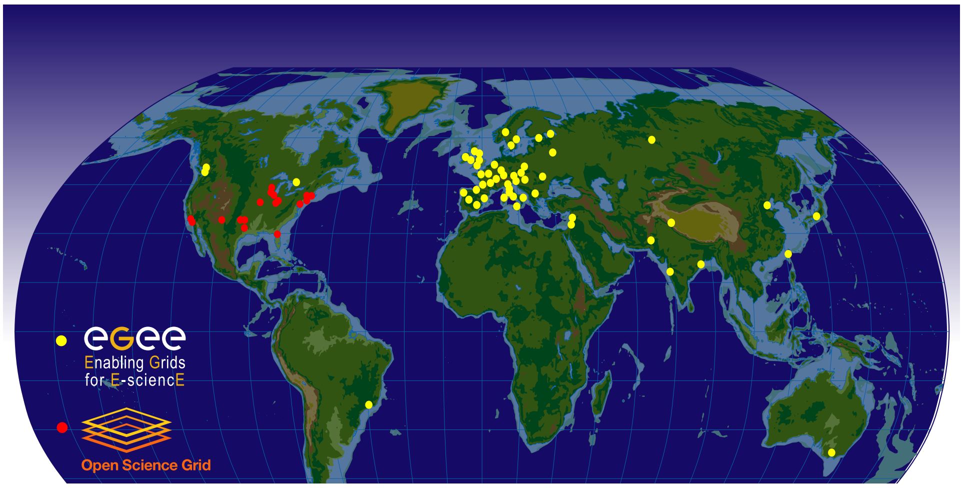

A map of the worldwide LCG infrastructure operated by EGEE and OSG.

The CERN openlab student programme was launched in 2002, and brings some 15 IT and physics students from around the world to CERN each year, to work in teams on Grid-related projects. Some of these students have been co-funded by openlab partners, and part of the programme involved study tours to partner facilities in the region. In addition, two masters students and two Marie-Curie fellows have completed their studies in CERN openlab. The creation of a public educational website about Grid technology, the GridCafé, was co-sponsored through the student programme, as were several international events such as the Role of Science in the Information Society conference, which was part of the World Summit on the Information Society in Geneva in 2003, and the Computing for High Energy and Nuclear Physics (CHEP) conference in 2004.

As well as the many excellent technical results that CERN openlab has provided, the partnership has given CERN a means to share a vision of the future of scientific computing with leading IT companies and gain deep insights into how industry sees computing hardware and IT services evolving in the future. This exchange of information is valuable to all involved. For CERN, it helps chart a path for LHC computing in years to come. For the industry partners, the global Grid being pioneered for the LHC holds significant promise for industrial applications – everyone is aware that this may be a new World Wide Web in the making. As the CERN opencluster project has demonstrated, the CERN openlab is a novel and effective framework for collaboration between multiple industrial partners, in a pre-competitive spirit and based on open standards.