3rd annual report
June 2003 – June 2004

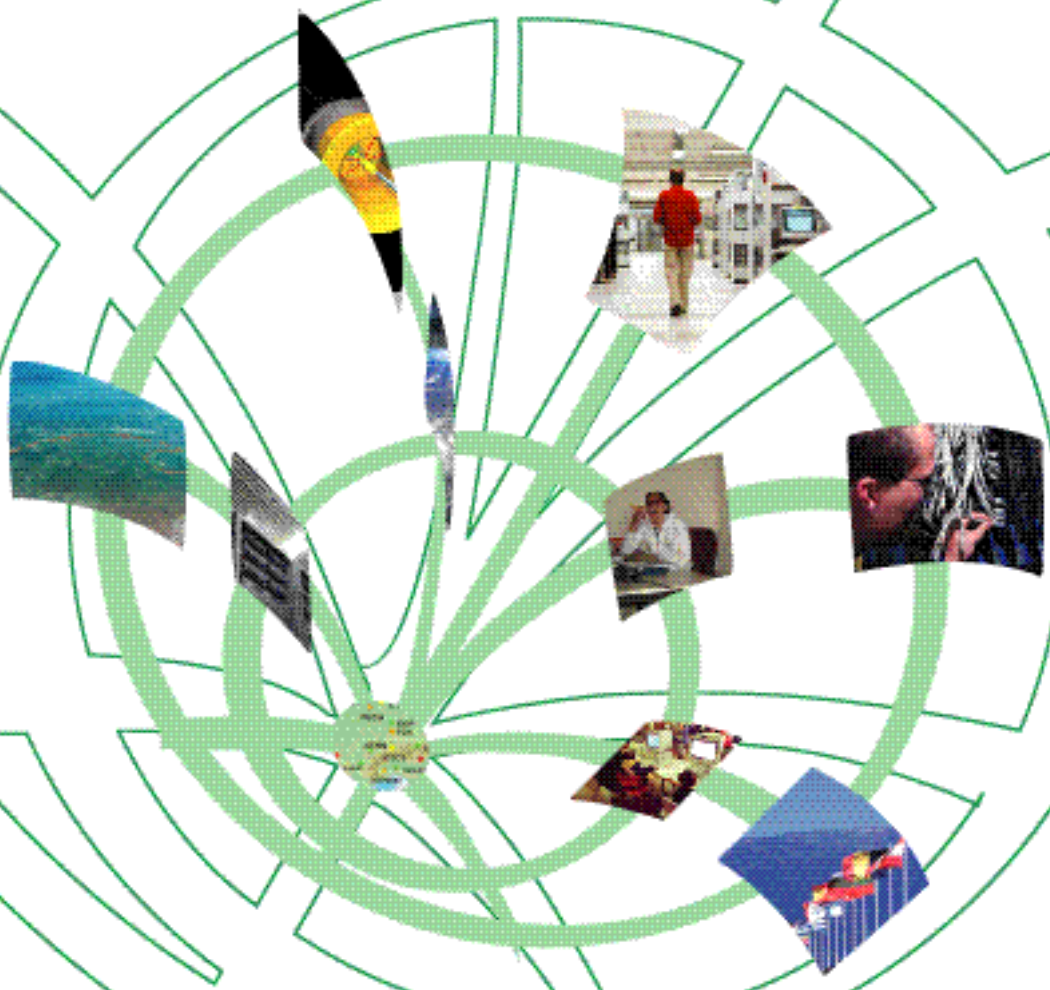

CERN openlab
for DataGrid applications

Table of Contents

Third Annual Report of the CERN openlab for DataGrid Applications
Editor: François Grey
Design and layout: Fabienne Marcastel
Photography: CERN photo service
Production: IT Department CERN
ISBN 92-9083-212-6

CERN openlab: a role model for the future

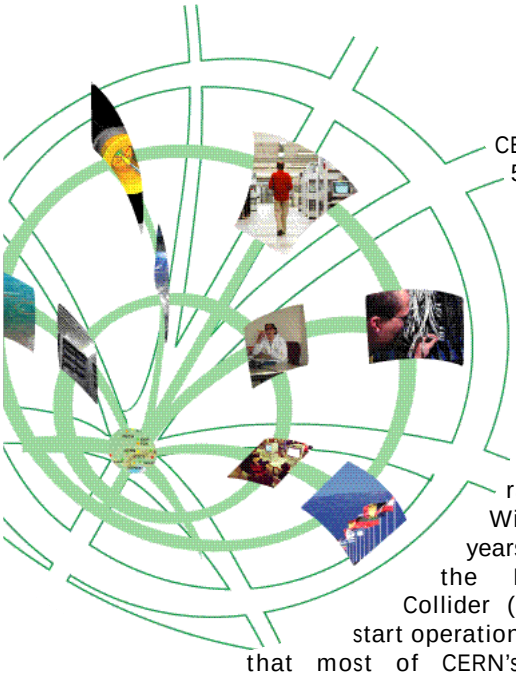

CERN marks its 50th anniversary in 2004, and while this is an occasion to remember CERN's many past achievements, it is also an opportunity to consider the laboratory's future. With barely three years to go before the Large Hadron Collider (LHC) is due to start operation, it is no surprise that most of CERN's energies are focussed on the immediate future, and meeting this deadline. In this context, the CERN openlab for DataGrid applications is making a crucial contribution to the LHC project, by testing and verifying some of the hardware and software solutions that LHC computing will rely on in the years to come.

For the longer term, CERN must continue its efforts to ensure both the financial and political feasibility of future major accelerator-based programmes, which require broad international support. Here, the international collaborations behind the four LHC experiments provide a role model for the future, bringing together hundreds of institutes and thousands of scientists to cooperate on a common goal.

In a similar way, 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 in this annual report show, CERN openlab has now effectively established a framework for collaboration between multiple industrial partners, in a pre-competitive spirit and based on open standards. This framework goes well beyond the scope of bilateral collaborations with industrial suppliers which have existed in the past, and has the potential to inspire technological development activities outside of the IT sector, where CERN openlab happens to have been conceived.

I therefore thank all CERN openlab partners for their contributions to the success of this partnership, and warmly encourage them to continue to work with CERN in the future – both short term and long term. There is much we can learn from each other, and from the extreme technological challenges that the high-energy physics community sets itself in the process of pushing back the boundaries of mankind's understanding of the Universe.

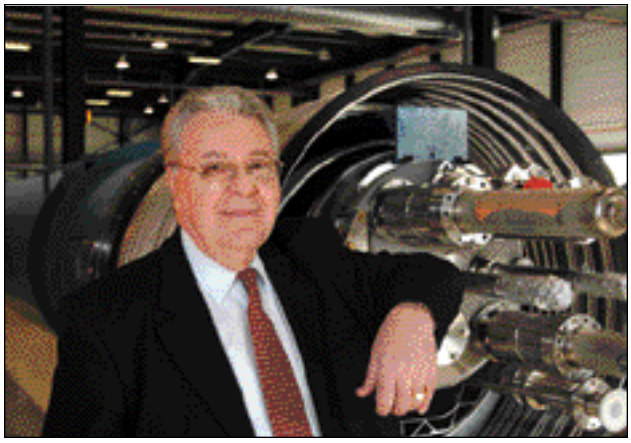

Dr. Robert Aymar
Director General of CERN

THE CONTEXT

CERN, the LHC and the Grid

CERN, the European Organization for Nuclear Research, is the world's largest particle physics centre. Founded 50 years ago, the laboratory was one of Europe's first joint ventures and includes today 20 Member States. Here, some 6500 physicists from universities and laboratories around the world come to explore what matter is made of and what forces hold it together.

CERN, with a staff of 2500 scientists and engineers, exists primarily to provide the international community of high-energy physicists with the necessary tools to do their research. These are accelerators, which accelerate particles to almost the speed of light, detectors to track the results of particle collisions occurring in the accelerators, and computing facilities to support the analysis of the data.

The world's largest and most powerful particle accelerator, the Large Hadron Collider (LHC), is now being constructed at CERN, in a 27km circular tunnel previously used by CERN's LEP experiments. As its name indicates, the LHC will collide beams of hadrons, primarily protons at an energy of 14 TeV. Beams of lead nuclei will be also accelerated, smashing together with a collision energy of 1150 TeV. There will be four detectors on the LHC ring, each a massive instrument several stories high, and each optimised for probing different aspects of the high-energy collisions.

The LHC will start operation in 2007 and run for up to two decades. It will be used to answer some of the most fundamental questions of science. A prime target of the LHC is to pin down the Higgs boson, a particle hypothesised to explain the origin of the different masses of fundamental particles. However, the LHC will address a host of other fundamental questions, too, such as the origin of "dark matter" in the universe, the unification of electroweak and strong forces and the origin of the asymmetry between matter and antimatter.

The Grid is a term coined by Ian Foster of Argonne National Laboratory and his colleagues, to describe the concept of globally distributed computing and data storage. While the notion of distributed computing has been around since the beginning of electronic computing, the idea of extending it to a global scale is recent, and only conceivable thanks to the amazing progress that telecommunications has made over the last few years.

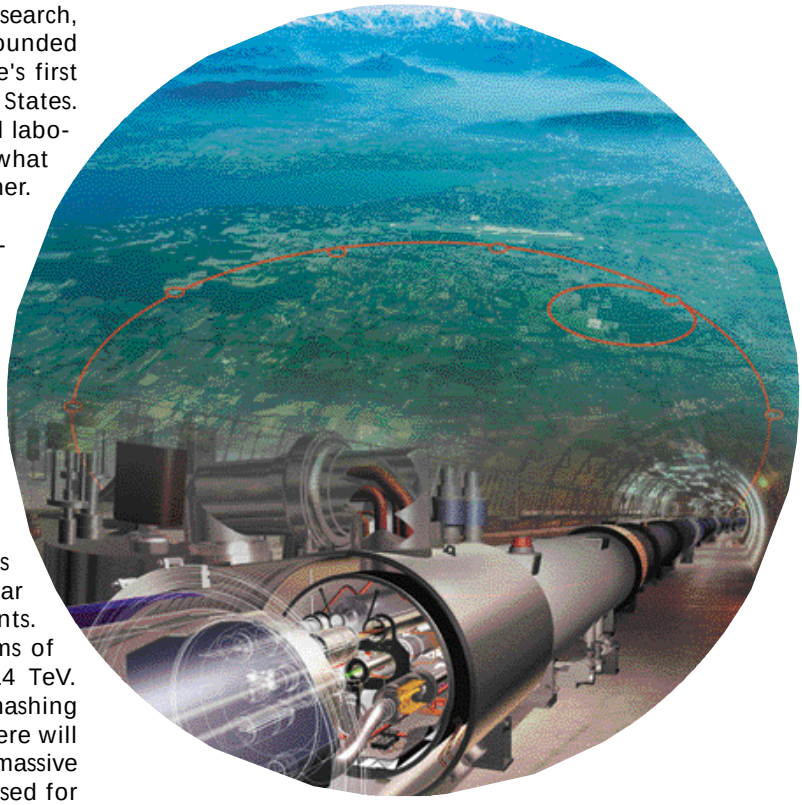

View of the LHC from the air and artist's impression of the finished accelerator inside the tunnel.

Not since the World Wide Web was developed at CERN, over ten years ago, has a new networking technology held so much promise for both science and society. Once again, CERN is set to play a leading role in making the technology a reality. Core to the success of this venture is the development of "middleware" to run computers distributed around the globe as though they were a monolithic resource. Whereas the Web gives access to distributed information, the Grid aims to do the same for distributed processing power and storage capacity.

There are many varieties of Grid technology. In the commercial arena, Grids that harness the combined power of many workstations within a single organisation are already common. But CERN's objective is altogether more ambitious: storing petabytes of data from the LHC in a distributed fashion, and making it easily accessible to thousands of scientists around the world. This requires much more than any single organisation can provide – a network of major computer centres around the world must provide their resources in a seamless way, in order to ensure the quality of service that scientists need.

A single particle collision, or "event", can generate many megabyte of data in the detectors, in the form of traces through the detector of the paths of particles created by the collision. In a typical experiment at the LHC, there will be some 40 million collisions per second. After on-line filtering, this huge data flow will be reduced to some 100 events of interest per second. This means a recording rate of up to one gigabyte/sec. Thus, over a year – and taking account the need for backup data – the total data stored from all four detectors on the LHC is expected to reach about 15 Petabytes. Since the lifetime of the experiment is expected to be at least 10 years, and researchers will need to analyse the totality of the data with increasingly sophisticated programs, the data challenge will grow considerably with time.

Simulation of an event involving a Higgs boson being produced in the CMS detector at CERN. This result was part of a data challenge that CMS ran in 2004 using the LHC Computing Grid, LCG.

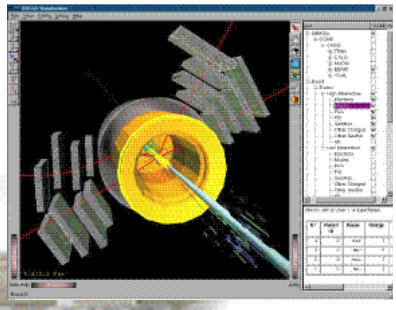

View of the new basement level of the Computing Centre at CERN, which is part of a major upgrade of the facility being carried out in preparation for the LHC.

LCG: A global Grid goes online

Monitoring map of active LCG-2 sites, June 2004.

The year 2003 marked the official launch of the LHC Computing Grid. The first phase of the project, LCG-1, began with a dozen sites operating a series of prototype Grid services, with the ambition that this Grid would gradually increase in scale and complexity as its builders develop an understanding of the functional and operational complexities involved in building a Grid of this scale. LCG-1 uses middleware developed by the European DataGrid project as well as the Virtual Data Toolkit (VDT) in the US, which includes Globus, Condor and related projects.

The LCG-1 system determines what resources and data a computing job requires, arranges for the job to run anywhere in the world that can provide the resources, and moves the data files required and produced by the job. The LCG-2 release was made in February 2004. This represented a significant upgrade of the functionality of the Grid middleware. In 2004, LCG is heavily involved in the so-called data challenges, where the different LHC experiments test their data storage and analysis routines under realistic conditions. Over the next two years, further upgrades are planned in close coordination with the EGEE project, since LCG is expected to form the core of that multi-science Grid.

EGEE: Towards a Grid for all sciences

The European DataGrid (EDG) project, a three year effort to develop and test Grid middleware for a variety of scientific applications which ended in March 2004, was led by CERN, and has provided some of the key software for the LCG project. The success of EDG generated strong support for a follow-up effort to build a permanent European Grid infrastructure that can serve a broad spectrum of applications on a reliable and continuous basis.

Providing such a service requires a major effort, comparable to that which went into establishing the GEANT high-speed network infrastructure for European science. Therefore, the EGEE (Enabling Grids for E-science in Europe) project was launched in April 2004, with the express goal of building a production Grid infrastructure, and providing round the clock Grid service to scientists throughout Europe. A proposal for such a project has been submitted to the EU 6th Framework Programme in 2003.

CERN is the lead partner in this new project, which involves a total of 70 institutional partners from Europe, the US and Russia. The LCG project will be the pioneering application for this Grid, but benefits of such an infrastructure for Europe will extend well beyond the LHC, to nearly all areas of science and commerce where large amounts of data must be managed in a distributed way.

EGEE has established an Industry Forum, to encourage industry both as a potential user of the EGEE infrastructure, and as a provider of Grid-related hardware and software solutions. CERN openlab partners have played an active role in the launch of this industry forum, and are represented on the Forum's steering committee.

Schema of the successive generations of Grid middleware software planned for the EGEE project, starting with current middleware used by the LCG project.

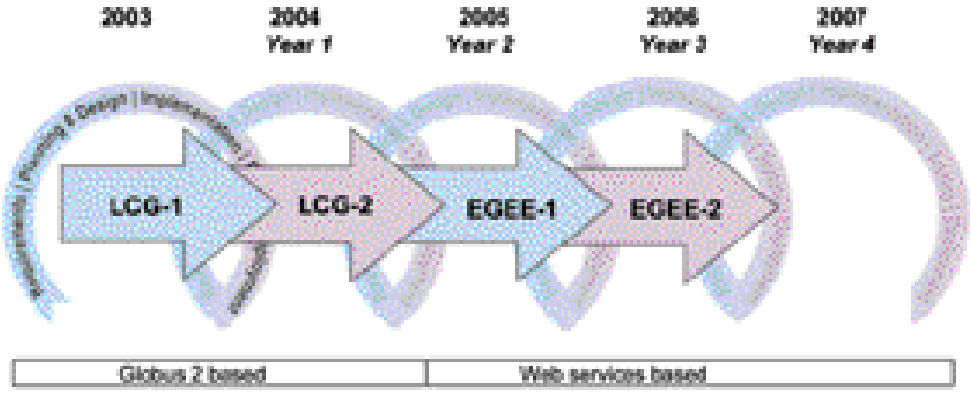

THE STATUS
Reaching cruising altitude

The CERN openlab for DataGrid applications is a new type of partnership between CERN and industry, created in order to tackle some of the Grid computing challenges that CERN faces for the LHC. The LHC experiments necessitate the storage and analysis of unprecedented amounts of data, estimated at about 15 Petabytes per year, or approximately 1% of current world annual production of information. To prepare for this data onslaught, CERN must be able to test industrial technologies that may still be in the pipeline today, in order to anticipate possible commodity solutions that will be available in a few years' time, and adjust CERN's Grid technology roadmap accordingly.

The best way CERN can gain this foresight is by working in close partnership with leading IT companies which will ultimately commercialise these technologies. And with the LHC project already under considerable budgetary pressure, the substantial sponsorship provided by the industrial partners really makes a difference to what can be achieved. Each partner in the CERN openlab is generously sponsoring 2.5MCHF over three years. This sponsorship is primarily in the form of in-kind donations of equipment and software, but also includes significant technical support from leading in-house experts, funding for postdoctoral fellows, and other training and dissemination activities.

The CERN openlab concept was launched in late 2001, and it achieved a critical mass of partners and technical activity by about January 2003, which represents the official start date of the partnership. By this time HP, Intel and Enterasys Networks were contributing key equipment to CERN openlab's first major project, the CERN opencluster. CERN had in turn established a dedicated team of in-house experts to build and test this advanced cluster. Six months later, by the time of the second annual meeting in June 2003, IBM had also joined the CERN openlab, bringing advanced storage technology to the partnership.

This annual report summarises progress from the last annual meeting until June 2004, which represents the halfway mark of the three-year collaboration. Significant milestones during the last year include the integration of the CERN opencluster into the LHC Computing Grid (LCG), as well as contributions to new landspeed records for data transfer across the Atlantic. With Oracle joining the CERN openlab in December 2003, there are now complementary partners in most of the key areas relevant to Grid computing, from processors to databases.

While there is a limit to the number of major IT companies that can usefully collaborate on a single project in the CERN openlab framework, we feel that smaller, more specialised firms can play a significant contributing role to the project. Thus, the announcement in May 2004 of Voltaire as the first company to become an official contributor to the CERN opencluster project opens the way for other possible contributions. In addition, the increasing emphasis on security, mobility and "user-centric" solutions – both in the Grid context and in the commercial sector – suggests opportunities for creating a new project within the openlab framework. We look forward to pursuing these opportunities in the coming year.

Finally, I would like to reiterate my thanks to all the industrial partners for their generous sponsorship and unwavering support of CERN openlab. Your commitment remains the key to our common success.

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

The openlab team: a growing pool of talent

CERN openlab now represents a significant investment in manpower for the IT Division, with a team of eight senior technical staff at CERN, as well as four sponsored CERN fellows (postdoctoral researchers) and two masters-level technical students contributing to the opencluster project. There has also been significant technical contributions from technical liaisons in the partner companies. In addition 11 undergraduate summer students were active for periods of 2-3 months in the CERN openlab student programme during June-September 2003.

CERN openlab Board of Sponsors

Robert Aymar	CERN [Chairman of Board]
Wolfgang von Rüden	Head of the CERN openlab
John Roese	Enterasys Networks
Jim Duley	HP
Tom Hawk	IBM
Tom Gibbs	Intel
Sergio Giacoletto	Oracle

CERN openlab Management Unit

Wolfgang von Rüden	Head of the CERN openlab
François Fluckiger	CERN openlab Manager
Sverre Jarp	Chief Technology Officer

CERN openlab Technical & Administrative Team

Sverre Jarp	Chief Technology Officer
Jamie Shiers	Databases
Maria Girone	Databases
Eric Grancher	Databases
Jean-Michel Jouanigot	Networking
Marc Collignon	Networking
Rainer Többicke	Storage
Jean-Damien Durand	Storage
Andras Horvath	Infiniband
Jarek Polok	Linux support
Jackie Franco-Turner	Administrative support
François Grey	Press communications
Matti Heikkurinen	Student programme

CERN openlab Fellows and openlab Students (opencluster project)

Andreas Hirstius	CERN openlab fellow (HP co-funded)
Andreas Unterkircher	CERN openlab fellow (HP co-funded)
Eva Dafonte Perez	CERN openlab fellow (Oracle funded)
Katarzyna Pokorska	CERN openlab fellow (Oracle funded)
Bárdur Árantsson	openlab Technical student (Jul.03 – Aug.03)
Michal Kapalka	openlab Technical student (Jul.03 – Aug.03)
Glenn Hisdal	openlab Technical student (Jul.03 – Jul.04)
Stephen Eccles	openlab Technical student (Jul.03 – Feb.04)

Industry Partner Technical Liaison with CERN

Markus Nispel	Enterasys Networks
John Manley	HP
Brian Carpenter	IBM
Jai Menon	IBM
Herbert Cornelius	Intel
Bjørn Engsig	Oracle

Industry Partner Organizational Liaison CERN

Christoph Kälin	Enterasys Networks
Michel Benard	HP
Pasquale Di Cesare	IBM
Pierre Mirjolet	Intel
Russel R. Beutler	Intel
Christiane Schlüter	Oracle
Monica Marinucci Lopez	Oracle

Industry Contributor Liaison with CERN

Patrick Chevaux	Voltaire

CERN Technical Liaison with Industry Partners

Sverre Jarp	HP and Intel
Jean-Michel Jouanigot	Enterasys Networks
Rainer Többicke	IBM
Jamie Shiers	Oracle

The openlab partners: a constellation of capabilities

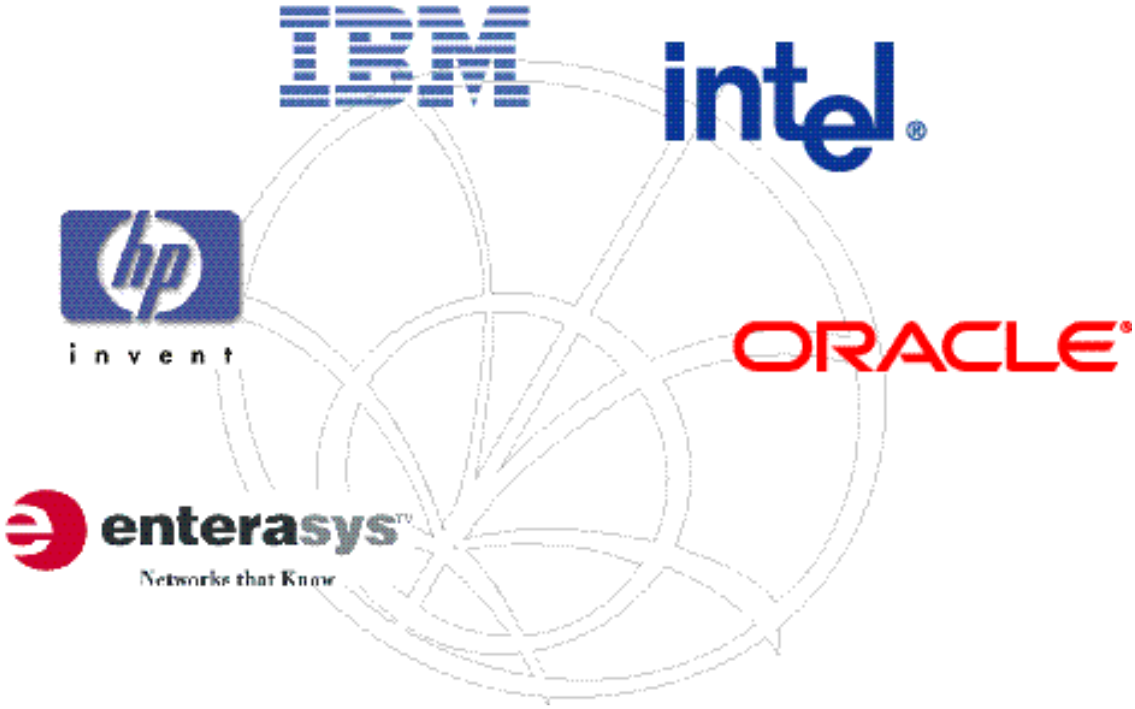

The CERN Rationale

CERN openlab has selected leading IT partners for five key technologies that support advanced clusters working on a Grid: processors, servers, high-speed switching, storage and database software. The CERN opencluster combines 64-bit processor technology and 10Gb/s network cards from Intel, compute servers in a cluster from HP, and a 10 gigabit switching environment from Enterasys Networks. IBM provides 28TB of high-end storage to the project, as well as its Storage Tank data management technology, and Oracle contributes a wide variety of Grid-enabling functionalities.

Enterasys Networks

Enterasys Networks is a leading worldwide provider of broadline intelligent data networking infrastructures for enterprise-class customers. Enterasys' networking hardware and software offerings deliver the innovative security, availability and mobility solutions required by Global 2000 organizations coupled with the industry's strongest service and support. For more information on Enterasys and its products, visit www.enterasys.com.

"The aggregate data throughput for LHC will exceed one terabit per second. Enterasys is confident that its 10-Gigabit Ethernet Technology will enable CERN to unlock the full potential of its DataGrid." John Roese, Chief Technology Officer for Enterasys Networks.

HP

HP is a leading global provider of products, technologies, solutions and services to consumers and businesses. The company's offerings span IT infrastructure, personal computing and access devices, global services and imaging and printing. HP completed its merger transaction involving Compaq Computer Corp. on May 3, 2002. More information about HP is available at www.hp.com.

"Through this collaboration with the CERN DataGrid, HP's researchers and engineers will be put to the test to truly push the envelope in developing advanced Grid computing technologies." Jim Duley, Director for Technology Programs, HP University Relations.

IBM

IBM is the world's largest information technology company, with 80 years of leadership in helping businesses innovate. Drawing on resources from across IBM and key Business Partners, IBM offers a wide range of services, solutions and technologies that enable customers, large and small, to take full advantage of the new era of e-business. Additional information about IBM is available at www.ibm.com/us/

"This is the perfect environment for us to enhance our Storage Tank technology to meet the demanding requirements of large scale Grid computing systems." Jai Menon, IBM Fellow and co-director of IBM's Storage Systems Institute

Intel

Intel, the world's largest chip maker, is also a leading manufacturer of computer, networking and communications products. Additional information about Intel is available at www.intel.com/pressroom.

"CERN's DataGrid project is an ideal application for Intel's most powerful processor yet, the Itanium. The awesome computer power required will find a formidable engine in the Itanium." Steve Chase, Director, Business and Communication Solution Group of Intel.

Oracle

For 27 years, Oracle has been helping customers manage critical information. The company's goal is to make sure that customers spend less money on their systems while getting the most up-to-date and accurate information from them. Oracle does this by simplifying or outsourcing IT infrastructure to reduce costs, and by integrating disparate systems to create a single, global view of the customer's business. www.oracle.com.

"Leading-edge grid technologies developed at CERN will be road-tested as part of its Large Hadron Collider project. As these technologies then come into the commercial mainstream, both we and our customers will benefit even further." Sergio Giacoletto, Executive Vice President, Oracle Europe, Middle East and Africa.

The CERN openlab facilities: at the heart of the Grid

CERN opencluster

The opencluster facility has been moved to an even more prominent position in the middle of CERN's Computer Centre, directly below the visitor gallery.

In October 2003 the first phase of upgrading of the cluster involved adding 16 nodes rated at 1.3 GHz. The second phase, undertaken in February 2004, involved speed upgrades of 81 processors to 1.5 GHz. The third phase in May/June 2004 doubled the number of nodes. The cluster now comprises just over 100 Itanium 2 dual processor HP rx2600 nodes, with about 10 10-GbE Network Interface Cards (NICs) delivered by Intel.

On the software side, the new opencluster nodes were initially tested with RH ES 3.0 (with kernel 2.4.21), and successful testing was also carried out of the 2.6 kernel. A decision was taken to switch CERN opencluster to "CERN-recompiled Enterprise Linux" (CEL3).

Enterasys Networks delivered in total four 10 Gbps "N7" switches with additional port capacity for 1 Gbps (in total 20 10-GbE ports and 300 1-GbE ports). IBM is sponsoring 28TB of high-end storage, with hardware delivery being completed in 2003.

Openlab openspace

This VIP meeting room, built with access directly on the Computer Centre, provides industrial partners and their clients with a prime location for technical and promotional meetings, where the partnership is prominently featured. It is also widely used by the IT Department and increasingly by CERN's Press and VIP services, for meetings with visiting dignitaries and journalists.

CERN openlab offices

The CERN openlab team has established offices in building 513, along the corridor from the openlab openspace. Most of the core technical team and the openlab management unit is installed there.

The IBM Storage Tank facility, with 28TB of high-end storage, in the CERN Computer Centre.

CERN OPENCLUSTER RESULTS
Rising to the data challenge

In the run up to the start of the LHC, it is imperative that each of the four experiments, ALICE, ATLAS, CMS and LHCb run data challenges to simulate all aspects of data collection and analysis. In particular, with LCG becoming operational as of September 2003, testing the emerging Grid infrastructure is now a prime motivation of such data challenges.

During the last year, CERN opencluster has contributed to several of these data challenges. For example, it contributed 20 servers from August 2003 to January 2004 to the ALICE Data Challenge V, a contribution which was noted by the experiment in its summary report for "running the whole GDC (Global Data Collector) environment flawlessly during the whole exercise". High speed links with Ottawa were also tested with Itanium-2 servers, in the context of the ATLAS data challenges.

In addition to the experimental data challenges, a major openlab focus is the 10 Gigabit challenge: a common effort by Enterasys Networks, HP, Intel and CERN to achieve data transfer rates within the opencluster of 10Gb/s. This involves multiple technological issues: network switch tuning, Linux Kernel and TCP/IP parameter tuning, and so on. In this connection, a number of key tests were performed on the components of the opencluster. For example, 10Gb Ethernet LAN/WAN tests were carried out successfully.

Back-to-back tests between Itanium systems were carried out by openlab summer students, and detailed comparisons between Itanium and Xeon systems were also made. PCI-X hardware investigations were carried out showing that at speeds much above 6 Gb/s PCI-X starts becoming the major bottleneck. A faster bus based on PCI-X2 or PCI-Express, is therefore required to saturate a 10 Gb/s link.

Oracle learns new tricks

Although Oracle only joined CERN openlab six months ago, it has already obtained tangible benefits from the efforts of the two Oracle-funded research fellows working in CERN openlab, which have important implications for Grid availability.

The first project greatly increases the availability of CERN's grid computing environment by significantly reducing the downtime of its catalogue, which ensures correct mapping of filenames and file identifications. This catalogue runs on Oracle Database and Oracle Application Server. If it is unavailable for any reason, the Grid itself becomes unavailable. Previously, applying security patches necessitated taking the catalogue out of service, hence impacting the whole Grid environment. One of the CERN openlab fellows pioneered the use of Oracle DataGuard to reduce this unavailability from hours to minutes. Oracle DataGuard was originally devised for handling catastrophic failures in data centres, but it turns out that it can effectively maintain a hot backup of the catalogue, keep the Grid environment available while patches are applied.

The second openlab project identified a way to upgrade an Oracle database from one release to another while keeping the database operational, using Oracle Streams, a function originally developed to help create powerful information sharing solutions. Both of these novel results could prove very useful not just for CERN's Grid, but for Oracle Database Administrators (DBAs) everywhere.

CERN openlab fellows Eva Dafonte Perez and Katarzyna Pokorska, who are funded by Oracle.

Into the record books

The CERN opencluster contributed to an event which made waves around the world. Announced during Telecom 2003 in Geneva, CERN and the California Institute of Technology received an award for transferring over one Terabyte of data across 7000 km of network at 5.44 Gbps, using Itanium-2, Linux, and ipv4-based single stream from nodes of the CERN opencluster specially configured for this challenge. This beat the previous land speed record by almost a factor of 2. Records for ipv6 were also established.

Such a data rate is nearly 10,000 times faster than a typical home broadband connection, and is equivalent to transferring a full CD in 1 second or a DVD in approximately 7 seconds. This record was set as part of the EU-funded DataTAG project, the aim of which was to push the limits for high-speed data transfer between the US and Europe. Such high-speed data transfer is an essential requirement for the LCG to be able to cope with the data rates from the LHC experiments. The result also garnered CERN a mention in the Guinness Book of Records. Flexible access to the CERN opencluster nodes, and the support of the CERN openlab technical team, played an important part in these technical successes.

In February 2004, the same collaboration announced a new record, and was able to achieve slightly over 6 Gbps with the ipv4 internet protocol, multiple streams and Microsoft's WindowsXP over the same 10,000km distance between Los Angeles and Geneva.

Wolfgang von Rüden congratulates Olivier Martin (CERN-IT, right) and Harvey Newman (Caltech, left) for their award of the Internet2 land-speed record at Telecom 2003.

Entering the Grid

From the outset of the CERN opencluster project, a key goal has been to participate fully in a major Grid, in order to test the performance of this advanced facility in a Grid environment. The LCG project provides the perfect match for such tests, and, after a significant effort, the openlab team has demonstrated complete port to the Itanium-2 platform of the LCG-1 software (launched in October) and the LCG-2 upgrade launched in February.

Originally, the LCG software had only been targeted the x86 platform, but thanks to the openlab effort a new level of heterogeneity was added to LCG. The LCG software includes elements from key Grid software projects CONDOR, VDT, and EDG, the latter contributing software for managing Worker Nodes and Compute Elements. An openlab Web page has been established to guide people who want to download the Itanium-2 binary distribution of the software.

In addition to this, CERN opencluster now runs the QUATTOR client, which is part of the suite of middleware developed by the European DataGrid project. QUATTOR manages automated installation of Grid middleware on cluster nodes. A related software component, LEMON was also installed for performance monitoring of clusters in a Grid environment.

HP signs up for LCG

In January 2004, HP announced that it would be the first commercial member of LCG, contributing computing resources from its HP Labs locations in Palo Alto and Bristol, as well as from HP computing centres in Brazil and Puerto Rico. In total, over 150 HP servers are expected to contribute to LCG, and the CERN openlab technical team has been active in assisting HP with the initial deployment of the necessary Grid middleware, including a site visit to Puerto Rico.

Commenting on HP's decision, Dick Lampman, senior vice president of research and director, HP Labs said that "[the] opportunity to participate in the LCG will provide us with unique insight into the functionality and complexity of large-scale Grid environments. Ultimately HP's customers around the world will reap the benefit of this collaboration".

View of HP Puerto Rico, which is contributing to the LCG project.

Lessons for Industry

CERN Openlab student Jarno Laitinen, co-funded by Enterasys Networks, with the new N7 routers for the CERN opencluster.

A key benefit of the CERN openlab for the industrial partners is the highly qualified feedback that the CERN openlab technical team can provide by in-depth testing of partner hardware and software in CERN's challengingly heterogeneous computing environment. Examples of such testing and validation during the last year include work on HP's SmartFrog, which is a technology for describing, activating and managing distributed software systems. A technical student in CERN openlab completed a thesis on SLP library implementation for SmartFrog, and a test implementation of SmartFrog was carried out, which involved dynamically adding PBS clients as Grid nodes. SmartFrog has capabilities that will allow deployment in both Grid and cluster management scenarios.

Another example is the work that the CERN openlab team has done on Intel compilers. Bi-weekly coordination meetings between Intel and CERN openlab centre on the performance around four ROOT benchmarks, ROOT being an object-oriented data analysis framework developed at CERN. Some 50 performance-related issues were reviewed this way with the Intel compiler team over the course of the year. This regular feedback was supplemented by intense testing of all relevant compiler options and code variants, carried out by an openlab summer student and continued by the regular openlab team. A benchmarking report (based on CERN software packages GEANT4 and ROOT using their own and the GNU compilers was delivered to Intel.

The feedback was mutual, since a ROOT bug was discovered in compression. A result well exceeding 1000 ROOTmarks was achieved with the benchmark called "stress", and the openlab technical team has recently achieved an average of 1000 ROOTmarks for all four ROOT benchmarks in use (stress, stressLinear, stressgeom, bench). Both ROOT and GEANT4 have been submitted to SPEC2004 in order to encourage compiler developers to focus more on large C++ applications.
Porting of CERN physics applications to the opencluster's 64-bit environment continued during the last year in close collaboration with the Physics Department. This included Castor, CLHEP, GEANT4 and ROOT. Other groups from the individual experiments also ported their applications to this environment including ALIROOT by ALICE, CMSIM by CMS US. Results of scalability tests with PROOF on 32 servers were reported at the CHEP2003 conference and new tests with 64 servers are underway. Software porting results were also described at the Gelato Federation meeting in autumn 2003 and spring 2004.

For the IBM StorageTank, initial hardware tests of the speed of the 200i disk servers and the corresponding ServRaid cards were carried out by an openlab student. Simulated pile-up tests for data from the CMS experiment were carried out in March. A peak speed of 70 MB/s was achieved per server with low number of clients (705 MB/s across 10 servers). 484 MB/s was achieved with 100 clients (each with 96 threads) across 8 servers for 40 hours.

With Enterasys Networks, an upgrade of the routers was completed over the past year. The three ER16 routers that were part of the original contribution were successfully integrated into the backbone of CERN's internal communications network. Four new N7 type routers have been installed in the CERN opencluster, replacing the ER16s.

Inside the opencluster, showing the distinctive heat sinks of the Itanium 2 chips.

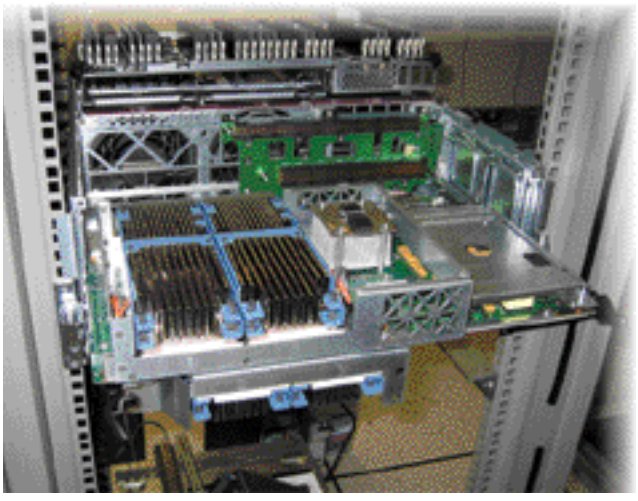

VOLTAIRE

Voltaire's philosophy

A few kilometres from CERN lies the small French town where renowned 18th century Enlightenment philosopher Voltaire lived most of his life. So it came as a surprise to many at CERN when a company of the same name became the first contributor to the CERN openlab in May 2004, a status created explicitly so that small and medium-sized high-tech companies could also contribute to CERN openlab. For the company Voltaire, the specific interest in joining CERN openlab is to test and verify of the InfiniBand solutions that the company provide. These enable high-performance Grid-computing applications to run on commodity servers and storage.

Initially a 12-way switch will be evalutated between Xeon systems and between Itanium systems. A 96-way Infiniband switch is planned for the summer of 2004. With InfiniBand, data can be streamed into the CERN opencluster very quickly with minimal loss of CPU cycles, so that these cycles can be retained as far as possible for the data analysis itself. Ronnie Kenneth, Chairman and CEO of Voltaire, emphasizes that "It is a privilege to work in the CERN openlab with our partners and industry leaders such as IBM, HP and Oracle. Together with CERN and these companies, we aim to build, test and verify cutting-edge solutions for the future of cluster-based Grid computing."

OTHER ACTIVITIES
CERN openlab dissemination

The mission of CERN openlab includes a strong emphasis on the public relations opportunities that this novel partnership provides. Therefore, the CERN openlab staff support a spectrum of activities that promote CERN openlab in different contexts, and provide added benefits for the partners.

The openlab Technical Team ensures dissemination of CERN openlab results through traditional technical channels such as conference presentations, technical publications and thematic workshops. In addition, the cutting-edge nature of the CERN openlab data challenges lends itself well to Press releases and other non-technical publications, as well as to presentations to visiting VIPs at CERN and to the general public.

A description of CERN openlab for the general public can be found on the external website for the general public, at http://www.cern.ch/openlab. A detailed description of CERN openlab events and publications can be found on the internal website at http://openlab-mu-internal.web.cern.ch . This site includes an archive of all publicly available documents and presentations made by CERN openlab, as well as a dedicated sponsors area.

Publications and conference presentations

"The CERN openlab: a novel testbed for the Grid", CERN Courier, p31-33, Oct 2003
"Rising to the data challenge", CERN annual report '03, p27, CERN 2004
External presentations, covering LHC computing as well as openlab activities, were made at:
ISC 2003 in Heidelberg
LCSC2003 in Linköping.
Gelato Spring 2004 in Champaign, Il.

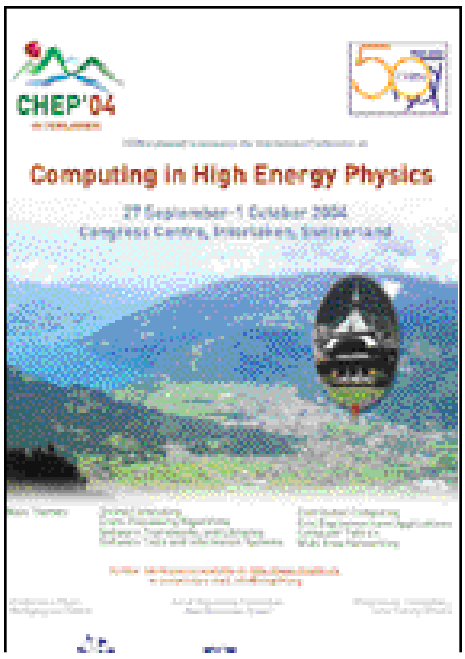

CHEP Conference sponsorship

CERN openlab partners have agreed to be prime sponsors of CHEP '04 (Computing for High Energy and Nuclear Physics), which CERN's IT Department is organising in Interlaken from September 27 to October 1st 2004, as part of the CERN 50th anniversary festivities. CERN openlab partners will contribute keynote speeches and each partner will benefit from the opportunity to organise a dedicated session with clients on the last afternoon of the conference

Press releases

"CERN openlab partners celebrates halfway mark" 06/04
"CERN recognizes UK's outstanding contribution to Grid computing" (openlab involved in UK industry event) 06/04
"CERN collaborates with Voltaire on Grid Technology Project" 05/04
"HP becomes first commercial member of CERN Large Hadron Collider Computing Grid" (HP release) 27/01
"Oracle Joins CERN openlab to advance Grid Computing" (Oracle co-release) 12/03
"CERN and Caltech join forces to smash Internet speed record" (opencluster involved)10/03
"CERN at Telecom World 2003" (openlab prominent on stand) 10/03
"LHC Computing Grid goes online" (Enterasys co-release) 09/03

Thematic workshops

As part of the CERN openlab activities, Thematic Technical Workshops are organized. These are open exclusively to representatives of openlab industrial partners as well as CERN staff. The standard format of Thematic Workshops is as follows: on day one, CERN experts present specific CERN requirements and the status of their technical development; on day two, CERN meets sponsors on a one-to-one basis. The themes of the Workshops are usually proposed by the industrial partners.

In the period June 2003 – June 2004, three Thematic Technical Workshops were organized:

- 2nd thematic workshop on Fabric Management (July 03)
- 3rd Thematic Workshop on Total Cost of Ownership (Nov 03)
- 4th Thematic workshop on Security (April 04)

The workshop on TCO involved extensive follow-up discussions with several partners, who provided further analysis of CERN's TCO situation.

Brian Carpenter (IBM), Chris Kaelin (Enterasys Networks), Lionel Cons (CERN IT Communication Systems Group) and Sverre Jarp (CERN openlab CTO) during a break in the CERN openlab Security Workshop, 27-28 April 2004.

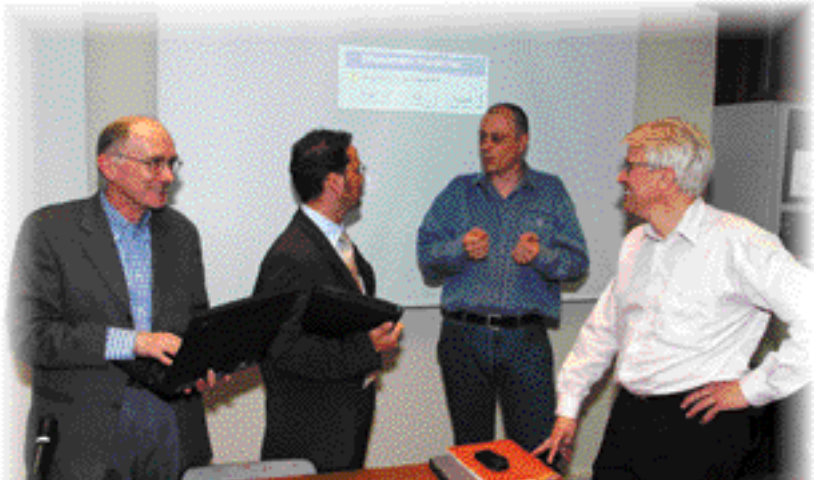

Partner-specific events

OracleWorld Press Panel (San Francisco, Sept 03)
Oracle 10*g* Launch (Switzerland, Sept 03)
EMEA Partner Event (Edinburgh Oct 03)
OracleWorld Press Panel (Paris Oct 03)
Presentations to HP (Cupertino, Nov 03)
Presentation at IBM (Half Moon Bay, Nov 03)
IBM Study tour on storage (CERN, March 04)
Intel Academic Forum (Barcelona, April 04)
Presentation at IBM Almaden Research Labs (April 04)

Public events

CERN openlab hosts regular First Tuesday events at CERN in collaboration with the organization "First Tuesday Suisse Romande". These public events attracted over 150 persons each, and involved speakers from CERN, CERN openlab's industrial partners, as well as other guest speakers. The public is primarily from the regional business and investor communities. The events were also webcast live archived. In the period June 2003 – June 2004, two events were held:

- 3rd FT @ CERN: The Grid gets Real! Included presentations by Philippe Bricard (IBM), Nigel Edwards (HP), 27 September 2002.
- 4th FT @ CERN: Industrial Impact of Information Technology from CERN. Included presentations by John Manley (HP), Monica Marinucci (Oracle) and Patrick Chevaux (Voltaire), 27 May 2003.

The 4th FT@CERN event involved two way webcasts between CERN and a similar industry event run by the PPARC Kite Club at the Queen Elizabeth II Conference Center in Westminster. That event was attended by over sixty senior business representatives from British industry. See www.rezonance.ch for full programmes of events and webcast archives.

VIP and Industry Visits

CERN is regularly visited by prominent delegations from member state ministries and scientific organisations, as well as by companies interested to know more about CERN's technology activities. A presentation of CERN openlab's activities has become a feature of such visits, as it is viewed as an outstanding example of industrial involvement at CERN. Such delegations normally visit the openlab openspace for a briefing, prior to a tour of the computer centre. During the last year, visits included:

- Delegation, Fujitsu Laboratories Ltd (09/03)
- Mr. Demonsang, Director, Banque Populaire des Alpes Innovation (09/03)
- Dr. Fujiwara, Director, Internet Research Institute, Japan (10/03)
- Delegation, Motorola Ventures (02/04)
- Dr. Eli Opper, Research and Development Chief Scientist, Israel (02/04)
- Delegation, Logitech Engineering Wireless Unit (03/04)
- Mrs Letizia Moratti, Minister of Education, Universities and Research, Italy (05/04)
- Dr. Helena Illnerova, President of the Academy of Sciences, Czech Republic (05/04)
- Dr. Fabio Colasanti, Director General for Information Society, European Commission (05/04)

In addition, visiting journalists from European and international newspapers and TV stations are regularly briefed about CERN openlab in the context of CERN's Grid activities.

Italian Minister of Education, Universities and Research, Mrs Letizia Moratti, visiting the openlab openspace at CERN, with Antonino Zichichi of Bologna University, Francois Fluckiger, Manager of CERN openlab, and Fabrizio Gagliardi, Manager of the EGEE Project.

Student activities

Based on a pilot programme in 2002, the CERN openlab student programme was launched in summer 2003 with 11 students from seven European countries. The idea is to have the students work in international teams, on projects related to the applications of Grid technology. One student team worked on issues related to the CERN opencluster, another on solutions for data acquisition from a CERN experiment, and one team worked on developing the Grid café website, described below.

An underlying objective of the openlab student programme is to encourage students and their home institutions to build longer term relationships with CERN openlab, for example by arranging for the students to do their masters project at CERN. As a direct result of the 2003 programme, six Masters students were active in the IT Department during the winter 2003-04, including one working on CERN opencluster. For 2004, 12 students are anticipated to take part in a similar programme, this time with participation of students from the US, as well. HP is co-sponsoring four students.

A related student activity is the CERN School of Computing, which offers a two week programme to some 70 students each year, and which obtained an important grant from the European Commission FP6 Marie Curie programme in 2003, which will provide for grants to assist up to 25% of the students with living and travel allowance over the next four years. Grid computing was a core feature of this school in 2003, and the CERN openlab is given prominent coverage.

In 2003, CERN openlab was part of a successful application by CERN for EU funding from the Marie-Curie programme. As a result, CERN openlab will hire two Ph.D.-level students for two years each, starting in September 2004.

Outreach activities

CERN openlab actively supported the establishment of a Grid Café for the CERN Microcosm exhibition. A related educational website, www.gridcafe.org describes Grid technologies to the general public, including access to Grid demos and pointers to industrial Grid activities. The GridCafe website was launched at Telecom 2003 in Geneva. This website was nominated for a Webby Award in 2004, in the category Technical Achievement.

CERN openlab partners also contributed to the RSIS (Role of Science in the Information Society) conference which took place at CERN the 8th and 9th December 2003, and the SIS-Forum related stand at the ITC4D exhibition (see cern.ch/sis-forum for more). Oracle was a prime sponsor of the RSIS event, which was an official event of the World Summit on the Information Society organized by the UN.

CERN openlab continues to be actively involved in the planning of the CERN Globe of Innovation, a stunning architectural structure that the Swiss government has given to CERN, and that will be used as meeting and networking centre around the theme of science, technology and innovation. The Globe of Innovation is scheduled to open to the public for the 50th anniversary celebrations in October 2004.

Participants in the CERN openlab student programme, summer 2003, visiting the HP-Intel Solutions Centre in Grenoble.

THE FUTURE

Strategy for the second half

The plans for the next 18 months of the three year opencluster project are to increase the number of nodes on the CERN opencluster, to reach a level of 128 nodes. The high-speed switching environment will be correspondingly upgraded, to ensure that 10Gbit data rates can be attained with this size of cluster.

In parallel with this hardware expansion, software activities will continue to focus on participation in the data challenges run by the LHC experiments, and on establishing a meaningful set of benchmarks to evaluate the performance of such an advanced cluster on the LCG. In this context, Oracle's database solutions provide an ideal benchmark for stress testing the various components of the CERN opencluster.

On the storage side, CERN openlab will work together with IBM to extend Storage Tank's capabilities in combination with the CERN opencluster, so it can manage LHC data and provide access to it from any location worldwide. Specifically, planning is underway to integrate Storage Tank in the next ALICE Data Challenge. In order to ensure a throughput of 900 MB/s, Storage Tank will be upgraded to support peak speeds above 1.5 GB/s.

Overall, the CERN opencluster project has now reached the crucial stage where overall integration of the different components of the cluster can take place. Storage Tank has graduated successfully from its initial test phase and is ready to be further integrated with the server cluster. An Infiniband switch large enough to interconnect the whole cluster is due during the summer of 2004, as is the latest version of the Enterasys 10Gb switch, which will allow optimal high-speed networking for the cluster. Oracle 10g has been installed on the cluster and is ready for deployment. This coincides with a period where, thanks to the ported LCG software, opencluster can now be routinely tested on LCG. Thus, the year ahead holds great promise for the testing and validation of these advanced hardware solutions in a global Grid environment.

New project opportunities

While the number of openlab partners in the current opencluster project is considered to be optimal, the creation of a contributor status in CERN openlab provides an opportunity for small and medium-sized enterprises with specific solutions of interest to contribute to the opencluster project. This opportunity will be pursued during 2004.

Over the next year, CERN and its industrial partners will review the possibility to renew the partnership after the initially planned three year period, in order to continue upgrading the opencluster with state-of-the-art technology. Since the LHC Computing Grid will run for over a decade, and require frequent upgrades, there is ample scope to maintain a cutting-edge system like opencluster in the future.

Opportunities to start other projects in the openlab framework have also been the subject of investigation during the last year. Specific Grid-related issues – such as Grid security and mobility – are being considered, as are broader issues of user-centric Grid applications. These topics have implications for long-term commercial applications of Grids. During the coming year, technical proposals for developing a project along these lines will be discussed with current openlab partners as well as other potential partners, including academic institutions.

In conclusion, as long as information technology continues to evolve at breakneck speed, and the demands of CERN's users for advanced IT solutions grows apace, there will surely be a demand for the sort of close collaboration between CERN and leading IT solution providers that CERN openlab has pioneered.

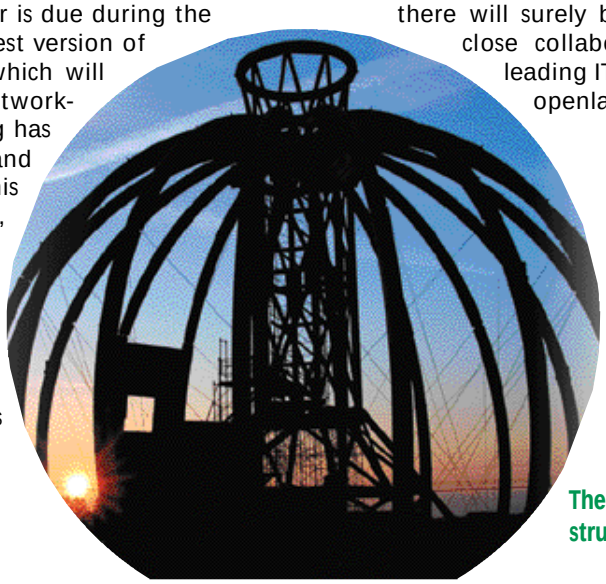

The Globe of Innovation, under construction at CERN, spring 2004.

This is the third annual report of the CERN openlab for DataGrid applications. It was presented to the Board of Sponsors at the Annual Sponsors meeting, June 22nd 2004.

Participants in the Annual Sponsors Meeting, from left to right: Pasquale Di Cesare (IBM), Mike Turnill (Oracle), Les Robertson (CERN), James Purvis (CERN), Dave Pearson (Oracle), Jim Duley (HP), Monica Marinucci Lopez (Oracle), Wolfgang von Rüden (CERN), John Manley (HP), Robert Aymar (CERN), Brian J. Truskowski (IBM), François Flückiger (CERN), Christiane Schlüter (Oracle), David C. Harper (Intel), Jamie Shiers (CERN), Christoph Kälin (Enterasys), Dilip D. Kandlur (IBM), François Grey (CERN), Gerard Vivier (Enterasys), Pierre Mirjolet (Intel), Robert P. Ryan (Enterasys), Sverre Jarp (CERN), Martin Hagger (Intel), Michel Benard (HP).

FOR MORE INFORMATION

Executive Contact

Wolfgang von Rüden
Head of CERN openlab
CERN IT Department
Wolfgang.von.Rueden@cern.ch
Tel: +41 22 767 64 36

Technical Contact

François Fluckiger
Manager of CERN openlab
CERN IT Department
Francois.Fluckiger@cern.ch
Tel: +41 22 767 49 89

Public Relations

François Grey
IT Communications Team
CERN IT Department
Francois.Grey@cern.ch
Tel: +41 22 767 14 83

http://www.cern.ch/openlab

CERN openlab

for DataGrid applications

sponsored by

IBM

intel

hp invent

ORACLE

enterasys

Networks that Know