CERN openlab

Annual Report
May 2008 – May 2009

LHC Computing Grid
Inaugurated 3rd October 2008

Table of Contents

Cover illustration: LHC Grid Fest ceremonial globe sculpture in the CERN Computer Centre

CERN openlab: the power of partnership

CERN Director General Robert Aymar passed the baton to his successor Rolf Heuer at the close of the CERN Council meeting at the end of 2008. Professor Heuer's mandate began in January 2009.

On 17 May 2006, the second three-year phase of CERN openlab was officially launched at a ceremony at CERN. At that time, I welcomed the industrial partners and contributors to this new phase, and encouraged them to build on experience from the first three-year partnership that had produced many excellent technical results in the field of cluster and Grid computing. Last December 2008, I just completed my five years at the helm of CERN which also coincided with the end of this second phase of openlab. I would like to underline the direct and positive impact of this initiative on the development of the Grid and computing services that support the LHC. The LHC Computing Awards I presented to Intel and Oracle for their outstanding contribution to LHC Computing, as well as the award given by ALICE to HP, are symbols of this successful partnership.

I am convinced that fundamental physics research and excellent infrastructures such as CERN have an essential role to play in modern society, for instance as a source of innovation. In the same spirit, the synergy that CERN openlab creates with leading IT companies is vital and I therefore thank all CERN openlab partners and contributors for their continued support of our joint effort. To conclude, I would like to send my best wishes to CERN openlab-III and to hail the successes that the LHC will no doubt achieve in the near future.

This year is very special to CERN and to me. The first collisions producing data for physics at the LHC are foreseen for this autumn, a very significant milestone for CERN. The year 2009 also coincides with the start of my mandate as Director General, the beginning of the third phase of the CERN openlab and the birthday of an event that changed the world forever: twenty years ago, in March 1989, Tim Berners-Lee handed a document to his supervisor at CERN, Mike Sendall, entitled 'Information Management: a Proposal'. The following year, the World Wide Web was born. With 2009 being declared the European Year of Creativity and Innovation by the European Union, the 20th anniversary of the Web serves as a timely reminder of the powerful role that creativity in basic research plays as a driver of innovation.

Partnering and networking with companies and research institutes is absolutely crucial in fostering innovation and keeping up with the global challenges our society faces. I would therefore like to express my full support to the CERN openlab, which has proven to be a very successful framework for industrial collaboration and an outstanding catalyst for ideas. The openlab gives CERN a means of sharing its vision of the future of scientific computing with leading IT companies, and of gaining deep insights into how industry sees the computing landscape evolving in the future. Expertise and knowledge, valuable to all involved, are created, exchanged and disseminated. I look forward to the future shared benefits as we enter a new age of fundamental physics and understanding of our Universe.

Robert Aymar
Director General
of CERN
2004 – 2008

Rolf Heuer
Director General
of CERN
2009 – 2013

THE CONTEXT

The LHC – In the starting–blocks

At 10:28 on 10 September 2008 at CERN, the world's largest and most complex scientific instrument, the Large Hadron Collider (LHC) circulated its first beam. Twelve years after the project was first approved by CERN council, this impressive start-up marked a key moment in the transition over two decades of preparation to a new era of scientific discovery. The event benefitted from major press and TV coverage from all over the world. One month later, on 21 October 2008, Swiss President Pascal Couchepin and French Prime Minister François Fillon joined by science ministers from CERN's Member States and around the world came to CERN to celebrate the completion of LHC as a project and to inaugurate the accelerator.

The LHC, the world's most powerful particle accelerator, mainly consists of a 27 km ring of superconducting magnets with a number of accelerating structures to boost the energy of the particles along the way. Inside the accelerator, two beams of particles travel at close to the speed of light with very high energies before colliding with one another. The beams travel in opposite directions in separate beam pipes – two tubes kept at ultrahigh vacuum. They are guided around the accelerator ring by a strong magnetic field, achieved using superconducting electromagnets. These are built from coils of special electric cable that operate in a superconducting state, efficiently conducting electricity without resistance or loss of energy.

This requires chilling the magnets to a temperature of just 1.9 degrees above absolute zero (-271.3 °C), colder than outer space. For this reason, much of the accelerator is connected to a distribution system of liquid helium, which cools the magnets, as well as to other supply services. No fewer than 123,000 interconnections were needed for the 27 km ring, including 65,000 electrical connections with superconducting cables.

A replacement magnet for LHC sector 3-4 being lowered in the tunnel

Unfortunately on 19 September, shortly after the successful start-up of the LHC, a faulty electric connection between two of the accelerator magnets caused an incident. This resulted in mechanical damage and release of helium from the magnet cold mass into the tunnel. Detailed studies of the malfunction have allowed the LHC's engineers to identify a means of preventing a similar incident from happening in the future, and to design new protection systems for the machine. A total of 53 magnet units have been removed from the tunnel for cleaning and repair. The current schedule foresees the LHC being cold and ready to produce physics data by next autumn 2009. CERN, the four experiments hosting the detectors and the whole High Energy Physics community are now in the starting-blocks: the countdown for physics data is now on.

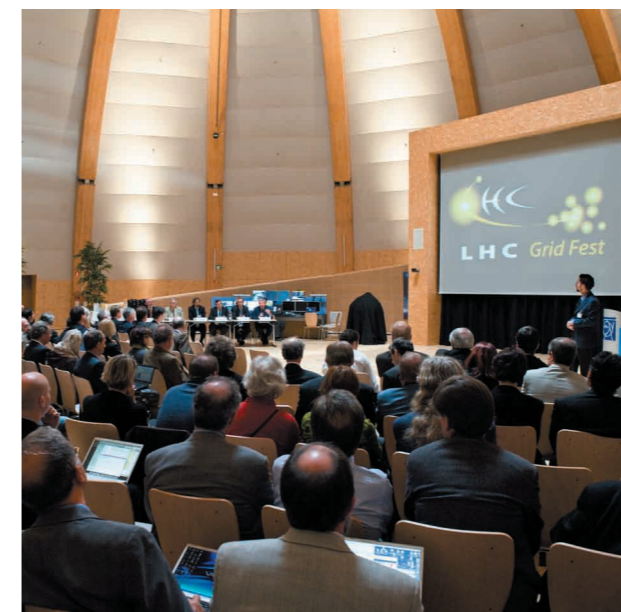Screen monitoring the first beam in the CERN Control Centre
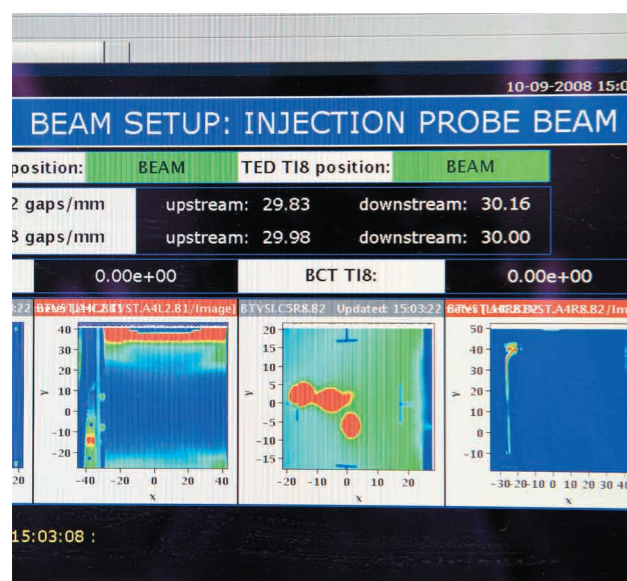
The LHC Grid Fest

At full capacity, the LHC is expected to produce more than 15 million Gigabytes of data each year. Indeed, hundreds of millions of subatomic particles will collide each second, presenting a massive data challenge. The mission of the Worldwide LHC Computing Grid (WLCG) is to build and maintain the data storage and analysis infrastructure for this immense flow of data, thus helping physicists open new frontiers in our understanding of the Universe. This ambitious project supports the offline computing needs of the LHC experiments in a globally distributed manner, connecting and combining the IT power of more than 140 computer centres in 33 countries. The WLCG is already fully running, handling large amounts of cosmic-ray data coming from the LHC experiments. ATLAS, for example, currently stores data at a rate of nearly a Petabyte per month.

The LHC Grid Fest, held on 3 October at CERN and at several sites around the world, commemorated the readiness of the WLCG, marking the end of seven years of development and deployment. More than 250 Grid enthusiasts gathered in the Globe of Science and Innovation at CERN, including representatives from worldwide industrial partners and teams that manage the distributed operations of the WLCG. To illustrate the global nature of the WLCG, the head of the WLCG project, Ian Bird, took an impressive live video tour of many of the major sites, providing a strong reminder of what a challenge it is to run a global 24 hour a day Grid service. At the end of to this trip around the world, Wolfgang von Rüden, Head of CERN's IT Department and Head of CERN openlab, unveiled the ceremonial globe sculpture, a metallic globe with the WLCG data centres illuminated by fibre optics.

Bob Jones, the CERN-based director of the project Enabling Grids for E-Science (EGEE) co-funded by the European Commission, highlighted that computing grids such as the EGEE in Europe and Open Science Grid (OSG) in the U.S. not only contribute their power to the WLCG, but also are key to other scientific projects, covering biology, chemistry, medicine, climate science and more. EGEE currently manages the world's largest multi-science Grid infrastructure contributing to LHC computing. On-site demonstrations were held throughout the day. On this occasion, the invitees were given the opportunity to discover GridMap: this highly successful tool developed by EDS, an HP company and openlab contributor, enables monitoring of Grid infrastructures.

To rise to such an unprecedented computing challenge, new and advanced systems were needed. This required the joint forces of science and industry to expand technological boundaries. As part of the morning ceremony, Intel and Oracle received the prestigious LHC Computing Award in recognition of their outstanding contribution to LHC computing. CERN Director General, Robert Aymar presented the awards to Stephen Pawlowski, Intel's Senior Fellow, and to Stéphane Rousset, Senior Vice President EMEA at Oracle. HP represented by Arnaud Pierson also came on stage with a similar award received from ALICE in 2007. During the afternoon session, CERN openlab partners delivered presentations covering their contributions to WLCG.

LHC Grid Fest in the Globe of Science and Innovation

Les Robertson, outgoing head of the WLCG project, received an ovation for his tireless efforts to guide WLCG from inception to its current, mature state

THE CONCEPT

Pushing the limits through innovation

R&D is essential to stay ahead with leading IT solutions. I believe in close collaboration between research and industry as an outstanding catalyst to sustain a steady stream of results. CERN openlab provides an efficient framework for such a collaboration. The motto "you make it, we break it" reflects the spirit of openlab. Throughout the first six years of its existence, openlab has established itself as a reference thanks to the excellent relationship and on-going commitment of all partners and contributors. The combined knowledge and dedication of the engineers from CERN and the companies have produced remarkable results leading to innovation in many areas.

CERN provides access to its complex IT infrastructure and its engineering experience, in some cases even extended to collaborating institutes worldwide. Partners commit to a three-year programme of work and provide three kinds of funding: salaries for young researchers and summer students, products and services, and engineering capacity. Contributors commit to a one-year programme with a lower annual level of funding.

CERN openlab has been organised into successive three-year phases. In openlab-I (2003–2005) we focused on the development of an advanced prototype called opencluster. CERN openlab-II (2006–2008) addressed a range of domains from platforms, databases and Grid to security and networking, with HP, Intel and Oracle as partners and EDS, an HP company, as a contributor. Energy efficient computing, also referred to as "Green IT", became a hot topic as we approached the power limit of the existing data centre. Mastering energy consumption and thermal aspects in large computing centres is one of the major challenges that the information society is facing, with increasing importance in the coming years. Through its collaboration with Intel, CERN openlab has generated white papers on the subject and made them publically available. "Green IT" will remain a central theme of CERN openlab for several years.

Besides technical achievements, a number of joint workshops and press events were organised. Intel, Oracle, HP ProCurve and ETM generously sponsored the various events in the «LHC 2008» programme. Another particularly significant event was the LHC Grid Fest on 3 October 2008, marking the transition from development and deployment to the continuous operation of the Worldwide LHC Computing Grid, to which openlab made significant contributions. Disseminating the expertise and knowledge created in openlab is one of our key objectives. Regular training sessions took place throughout the year, including openlab contributions to the CERN School of Computing. The CERN openlab Summer Student Programme with its specialised lectures is another example matching this objective.

We are now entering the third phase of this programme – openlab-III (2009–2011) – which will not only capitalise on and extend the successful work carried out in openlab-II, but will also tackle new crucial areas. Industrial controls have recently been added to the openlab portfolio with Siemens joining as a new openlab partner. The creation of the CERN openlab Automations and Controls Platform is the direct result of this new partnership. The aim of the project is to improve the security of Programmable Logic Controllers (PLCs), to open Supervision Control And Data Acquisition (SCADA) and PLC software to common software practices, and to develop solutions to deploy SCADA and PLC software in large installations.

I would like to take this opportunity to thank all openlab partners and contributors, past and present, for their ongoing support and engagement in this ambitious partnership. My special thanks go to the team members who make it all happen. I am confident that we will be equally successful in the future and I look forward to another period of exciting developments in CERN openlab-III.

Wolfgang von Rüden
Head of CERN openlab

The openlab team: the source of our creative energy

The openlab team is formed of three complementary groups of people: the young engineers hired by CERN and funded by the partners, technical experts from partner companies involved in the openlab projects, and CERN management and technical experts working partly or fully on the joint activities. At CERN the people involved are not concentrated in a single group but on the contrary, they span multiple units. In the IT Department, the CS (Communication Systems), DES (Databases and Engineering Systems), DI (Departmental Infrastructure), GD (Grid Deployment), GS (Grid Support), IS (Internet Services) and DM (Data Management) groups host openlab activities, as does the ICE (Industrial Controls and Electronics) group in the EN (Engineering) Department, since the beginning of 2009.

This year a Communications Officer dedicated to CERN openlab joined the team. A list of the IT and EN Department staff most closely involved in the CERN openlab activities is given below. The distributed team structure also permits close collaboration with computing experts in the LHC experiments, as well as with engineers and scientists from openlab partners who contribute significant efforts to these activities. Principal liaisons with partners and contributors are listed below. In addition, significant contributions are made by students participating in the CERN openlab student programme, both directly to openlab activities (5 students during summer 2008) and more widely to WLCG, EGEE and other Grid and CERN related activities in the IT Department (13 students). Finally, we would like to thank Jeff Arnold, Senior Software Engineer at Intel, who collaborated with the team for a period of two months during his sabbatical, working on understanding better how memory is used by applications in multicore servers.

CERN openlab Board of Sponsors

Rolf Heuer	CERN (head of Board)
Wolfgang von Rüden	Head of CERN openlab
Frédéric Hemmer	Head of CERN IT Department
Michel Bénard	HP
Richard Dracott	Intel
Vasu Briquez	Oracle
Thomas Hahn	Siemens

CERN openlab Management Unit

François Flückiger	CERN openlab Manager
Sverre Jarp	Chief Technology Officer
Mélissa Le Jeune	Communications Officer
Séverine Pizzera	Administrative Assistant

CERN openlab Fellows and Staff (sponsor indicated)

Gyorgy Balazs	Technical Student (CERN)
Håvard Bjerke	Fellow (Intel)
Daniel Filipe Rodrigues	Fellow (EDS) – Staff (Siemens)
Xavier Gréhant	Fellow (HP)
Omer Khalid	Fellow (Siemens)
Ildefons Magrans de Abril	Fellow (Siemens)
Milosz Marian Hulbój	Fellow (HP ProCurve)
Ryszard Jurga	Staff (HP ProCurve)
Andreas Hirstius	Staff (Intel)
Julien Leduc	Fellow (Intel)
Andrzej Nowak	Fellow(EUMarieCurie)-Staff(Intel)
Eva Dafonte Pérez	Staff (Oracle)
Jose M. Dana	Fellow (HP) – Fellow (EDS)
Svetozar Kapusta	Fellow (Oracle)
Filippo Tilaro	Fellow (Siemens)
Anton Topurov	Fellow (Oracle)
Dawid Wójcik	Fellow (Oracle)

CERN Liaison

Renaud Barillère, Liaison with Siemens
Sverre Jarp, Laison with HP Labs, Intel
Jean-Michel Jouanigot, Liaison with HP ProCurve
Mats Möller, Liaison with Oracle
Jamie Shiers, Liaison with EDS

Other IT Department Staff contribution

Dirk Düllmann	DM Deputy Group Leader
Maria Girone	DM Group
Eric Grancher	DES Deputy Group Leader
Christopher Lambert	DES Group
Jose Carlos Luna Duran	CS Group
Alberto Pace	DM Group Leader
Markus Schulz	GD Group Leader

Industry Partner Liaisons with CERN

Arnaud Pierson	HP
Peter Toft	HP Labs
Dan Ford	HP ProCurve
Alberto Soto	HP ProCurve
Claudio Bellini	Intel
Herbert Cornelius	Intel
Stephan Gillich	Intel
Bjørn Kisbye Engsig	Oracle
Graeme Kerr	Oracle
Monica Marinucci	Oracle
Ulrich Eger	Siemens
Thomas Hahn	Siemens
Juergen Hirte	Siemens

Industry Contributor Liaison with CERN

Max Böhm	EDS
Rolf Kubli	EDS

THE RESULTS

The Platform Competence Centre

The Platform Competence Centre (PCC) – one of the four broad areas of activity in the CERN openlab - was initially developed in the context of CERN openlab-II. In collaboration with Intel and HP labs, it addresses important fields such as thermal optimisation, application tuning and benchmarking. It also has a strong emphasis on teaching. Last year the PCC was able to add to its server farm 64 HP Blade Servers with dual Intel® Xeon® 5400 (Harpertown) processors running at 3.0 GHz. This installation has become the nerve centre of most of the activities mentioned below. The Blade servers, in addition to the existing set of Intel® Itanium® and Xeon® servers, were moved into new racks in the CERN Computer Centre, allowing for optimal cooling and reserving space for future expansion.

To process the enormous amounts of data that will be produced by the LHC, an estimated 5'000 servers will be available in the CERN Computer Centre, and another 6'000 at the detector sites of the four different LHC experiments. Improving the performance of these servers by just 1% is roughly equivalent to saving a sum in the order of a million Swiss francs in hardware purchases. Furthermore, since the CERN Computer Centre facilities are limited both in terms of the electrical input and the cooling output per square metre, the need for intelligent power optimisation is paramount. It is vital that the systems render maximum performance per Watt, so that the data centre power consumption is maintained within the newly increased 2.9 MW limit. For this reason power measurements are conducted on a regular basis and performance monitoring is a crucial activity of the Platform Competence Centre. For instance, within the CERN openlab framework, early versions of Intel® Atom™ and Xeon® 5500 (Nehalem) processors were tested in 2008, the latter well before its commercial launch (details of the Atom tests can be read on page 9). Generally speaking, mastering energy consumption and thermal aspects in large computing centres is one of the major challenges that the information society now faces. Through its collaboration with Intel, CERN openlab has produced white papers made publicly available on the openlab website to the benefits of society at large. Energy efficient computing, also referred to as "Green IT" will remain a central theme of CERN openlab-III.

Another project looks at monitoring both High Energy Physics (HEP) software and our server hardware performance by choosing the best methods and tools for this purpose. As already mentioned, improvements at the percent-level are very rewarding. Therefore, each case is studied with great care, in the hope that the results are relevant to many programmers, both internally and externally to CERN. Thanks to the close cooperation with Intel and HP as well as external collaborators, the research and development results of this project are regularly being disseminated to a broad developer community. On the software side, together with a key developer from HP Labs (now with Google), CERN openlab has been working on improving perfmon2, a universal performance monitoring interface in the Linux kernel, and pfmon, the corresponding client. CERN openlab's contributions included many optimisations related to monitoring and profiling sizeable software frameworks. As a new initiative, openlab spawned another project, gpfmon, a graphical user interface to pfmon. The additional functions on top of the original tool include high level monitoring scenarios, a visual event selection aid, plotting and numerous usability enhancements, which all help to make pfmon attractive to non-expert users. It should also be mentioned that by comparing pfmon output from CERN's production servers with different runs of the new SPEC2006 benchmark, openlab was able to recommend to HEPiX (a forum for HEP scientific computing experts) that the subset consisting of all C++ jobs would be a representative match for use in future WLCG acquisitions. The TOP500 submission reported in last year's Annual Report awarded the CERN cluster with position 96 in the list of the world's most powerful computers in June 2008. This was a gratifying result, but by November the same cluster was downgraded to position 186.

The collection of benchmarks that openlab deploys in the field of optimisation, targeting both multithreading and vectorisation (in order to fully exploit the available resources of modern processors) has been enriched in 2008 and the benchmark repository has now become a project per se. An important extension to the scalable benchmarks came from a joint collaboration with the University of Heidelberg and Intel Brühl in Germany. It started by porting of an existing benchmark that Heidelberg had developed for Track Fitting in the ALICE High Level Trigger to quad-core Xeon processors. This allowed us to demonstrate a throughput increase of almost 32 times. Later on, a second benchmark was developed based on the Cellular Automaton algorithm for Track Finding. This successful exercise is described in an openlab Report which is available on our website.

Virtualisation is currently used at CERN with great benefit in certain areas and was another focus of the Platform Competence Centre in the past. As part of the virtualisation project, virtualisation technologies were tested and solutions to enable virtualisation in Grid environments were developed. Multiple benchmark jobs, both synthetic benchmarks and a set of CERN real-world applications, were used to assess the performance characteristics of many of the virtualisation technologies in the current virtualisation landscape. Measurements showed that using guest domains in Xen (one of the popular virtualisation hypervisors for Linux) had a negligible impact on performance compared to native performance. However, sensitive operations – operations that need the intervention of the hypervisor – can have a moderate to significant impact on performance. Apart from the benchmark activity of the project, a facility called "Content-Based Transfer" (CBT) was an important topic. In particular, the development of a CBT tool enabled to divide the time to transfer a 400 MB Virtual Machine (VM) image to a target machine with 90% identical blocks by a factor of 2.6. The paper, 'Tools and Techniques for Managing Virtual Machine Images', published in the VHPC'08 workshop, proves that the efficiency achieved by this CBT implementation is close to a theoretical maximum. OS Farm, a complementary tool for creating VM images for Xen on demand, can dynamically provide images based on a set of Linux distributions and can thus satisfy the required execution environments for many such deployment scenarios. Both OS Farm and CBT were developed as proofs of concept, but they are freely available for download and can be used under open-source licences. It is believed that virtualisation will gradually penetrate production grids.

Another area where CERN openlab contributes is compiler optimisation, where the aim is to improve performance of a wide range of different jobs by influencing the back-end code generator. The compiler project gets its inspiration from the fact that almost all LHC programs (simulation, reconstruction, data analysis, and so on) are written "in-house" by high energy physicists. One way of improving the performance of such programs on a broad front is to work with the compiler writers to seek improvements that will have an impact on our C++ source codes. The current work in

Gyorgy Balazs with the Atom-based system he built himself from a motherboard obtained from Intel

this area is targeting further improvements in execution time and expanding to more complex benchmarks. Furthermore, in 2008 a major push was made by CERN openlab in the area of multi-threading, which exploits the new multicore processors in a more fine-grained manner with smaller memory requirements. This is an important complementary approach to the traditional multiprocessing for harvesting the full performance of high multicore architectures. In collaboration with the North-Eastern University (NEU) in the USA, multi-threading has been used to make advances in the parallelisation of the Geant4 toolkit for the simulation of the passage of particles through matter. Geant4 is not only widely used in HEP, but also in medical and space science. The 'Multi-threaded Geant4 with Shared Detector' report written by Xin Dong from NEU, a CERN openlab summer student, focuses on this topic. And last but not least, several workshops related to multi-threading and performance optimisation were again organised and are listed in the education section of this report.

To prepare for the future as the LHC gets underway, CERN openlab started several initiatives with Intel to look at future processors such as the forthcoming Xeon generations and possible future languages to support multicore and manycore processors. This work is expected to be one of the foundations for the current openlab period.

The second generation of Intel's 10Gbit Network cards were tested and found to be extremely powerful with peak speeds of 9.49 Gb/s. It is likely that disk servers with a 10Gbit card for direct attachment are of interest to CERN and other Grid partners in a near future. Modern disks require higher speed than what the 1 Gb/s attachment in use today can provide. Research on Solid State Drives (SSD) has also already started (both in the Platform and the Database CC) and a detailed report is expected in the coming months.

Andrzej Nowak and Julien Leduc in the CERN Computer Centre openlab cluster

Intel® Atom™ processor performance tests

Gyorgy Balazs carried out tests as part of the Platform Competence Centre programme. Below are some of the key highlights. The full study is available on the CERN openlab website.

The objective of the tests was to compare an amateur-built Atom N330 single-socket system with an industry-built dualsocket Xeon Core 2 Quad server to gauge the potential of the Atom system for running High Energy Physics applications. Two benchmarks were used; one simulation benchmark from offline and one trigger benchmark from online. The tests of the Atom processor show that it has potential, even though the measurements were conducted on an unoptimised consumer system. Although absolute performance is far behind the Xeon-based server, the priceperformance result obtained with the benchmarks on such a dual-core dual-threaded Atom system were seen as rather encouraging.

Reference: Is the Atom processor ready for High Energy Physics? An initial analysis of the dual core Atom N330 processor. Gyorgy Balazs, Sverre Jarp and Andrzej Nowak, November 2008.

Recent benchmarks comparing three different models of the Intel® Xeon® 5500 generation (released March 2009)
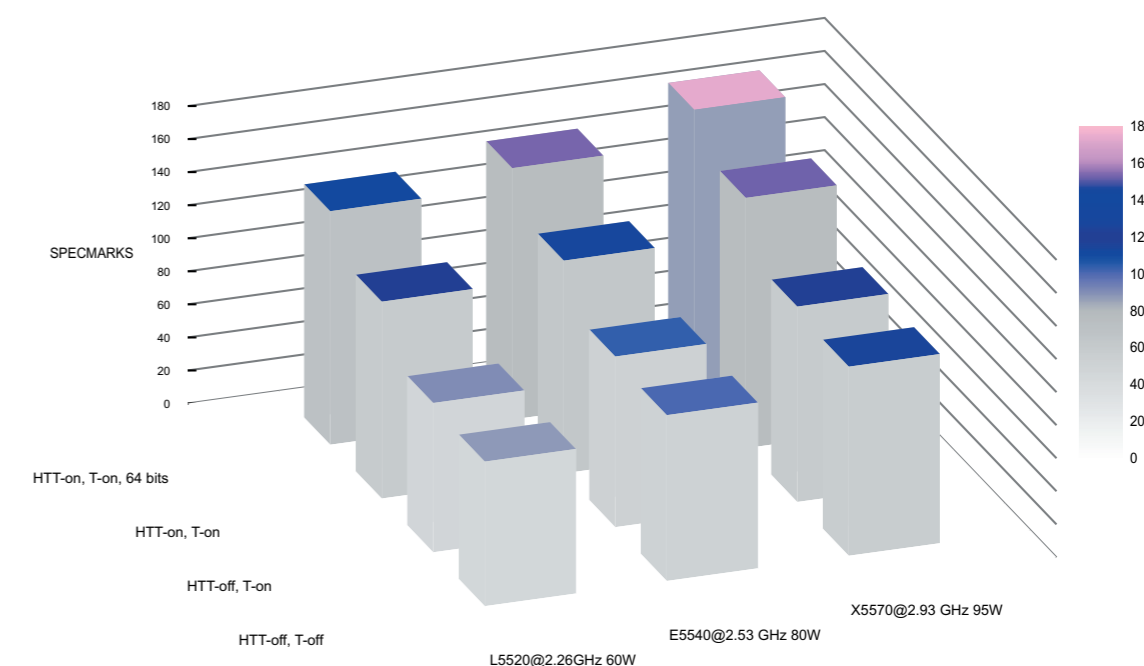

Figure: The Z-axis shows the resulting SPEC2006 rate results when using the C++ subset of the benchmarks. The Y-axis lists processor features that can be switched on for better performance: Intel® Hyper-Threading Technology (Intel® HT Technology) mode, Intel® Turbo Boost Technology (overclocking), and finally 64-bit mode.

The Grid Competence Centre

The Grid Competence Centre focuses on testing and validating partner software solutions in connection with the EGEE project and its middleware stack, gLite. Particle physics collaborations stress grid infrastructures with exceptionally high computing demands for distributed data analysis. Such a demand for controlled, homogeneous environments conflicts with the distribution of resources across different administrative domains. Furthermore, on long-haul networks spanning multiple computer centres, resource allocation is not governed by the same constraints as in local clusters. Latency limits the frequency of communications, and data is not uniformly accessible. We study how this affects resource allocation algorithms.
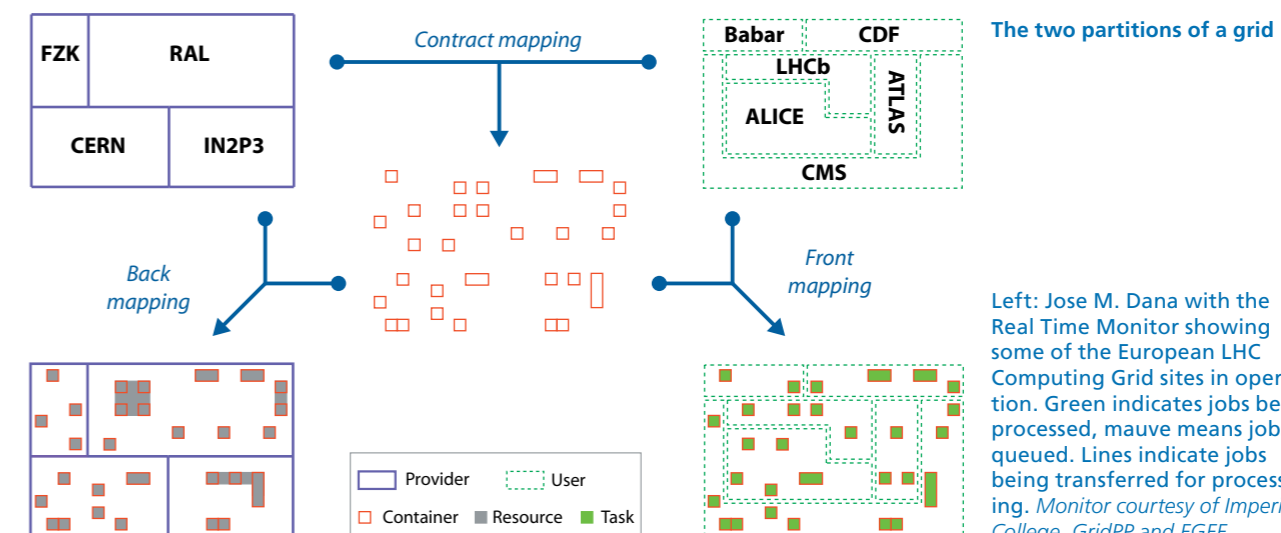
The deployment of virtual-machine based environments separates the concerns of resource usage and back-end servers administration: users get custom nodes decoupled from the physical machines that provide the resources. SmartDomains, based on SmartFrog (HP), provides a means to configure, activate and manage synchronized, distributed environments made up of custom Xen virtual machines. CERN openlab has helped raise awareness of virtualisation techniques in the high-energy physics community, and demonstrating the specific advantages of the Xen virtualisation open source platform. HP is also a long-time contributor to Xen and the first major company to work with this platform.

Physics collaborations need to deploy stable services to allocate their jobs efficiently on their controlled pool of nodes. A mechanism for resilient service re-deployment was prototyped to cope with the heterogeneity and transience of back-end resources. A model is being elaborated to represent the mapping of tasks to resources as a combinatorial optimisation problem, in order to understand and improve the computing throughput of an experiment. This prepares the start of the LHC and the following years of events analysis.

Recent developments conducted by Xavier Gréhant as part of his PhD thesis work at CERN, and to be published jointly with Isabelle Demeure (ENST) in an article called 'Symmetric Mapping: an Architectural Pattern for Resource Supply in Grids and Clouds', during SMTPS'09, May 2009, explores this field in more depth: "In Grid and Cloud computing, multiple sites and organisations are involved in resource allocation. Initially based on voluntary collaborations, grids have not focused on the different interests of the participants. As a result, compromises in performance and flexibility are made on both sides. However, recent developments in Grids ('Towards efficient resource allocation on scientific grids' report, by X. Gréhant, I. Demeure, S. Jarp, available on CERN openlab website) show a trend towards more user control. In commercial Clouds, to the contrary, providers set the rules that determine the access to their resources". The work carried out suggests "that these two trends are compatible and can be used to the benefits of the allocation" thanks to symmetric mapping, a concept detailed on page 13.

"The figure represents a typical grid environment. The rectangles in the background represent providers, CERN, FZK, IN2P3 and RAL. These are research institutions that provide computing resource for the analysis of high energy physics data. The areas on the foreground delimited by dotted lines represent users: Babar, CDF, LHCb, ALICE, ATLAS, CMS. These are Virtual Organisations (VOs), i.e. collaborations of physicists that analyse data. Members of a VO work on behalf of the VO, with a common objective and the same applications, therefore a VO is considered a single user. As represented on the figure, a resource container is mapped on one or an assembly of resources from a unique provider. It supports one or several tasks from a unique user". *Excerpt from 'Symmetric Mapping: an Architectural Pattern for Resource Supply in Grids and Clouds'. Xavier Gréhant and Isabelle Demeure, SMTPS'09.*

The two partitions of a grid

Left: Jose M. Dana with the Real Time Monitor showing some of the European LHC Computing Grid sites in operation. Green indicates jobs being processed, mauve means jobs queued. Lines indicate jobs being transferred for processing. *Monitor courtesy of Imperial College, GridPP and EGEE.*

The end of openlab-II has coincided with the conclusion of the Tycoon two-year project with HP Labs. The integration of Tycoon, a market-based allocation system developed at HP Labs, with the gLite software suite was successfully carried out, resulting in the Tycoon-gLite interface which allows easy deployments of Grid elements in a Tycoon environment. Tycoon-gLite acts as a bridge between Tycoon and the machines (auctioneers) that are ready for deployment of 'special images'. This allows Computing Elements (CEs) and Worker Nodes (WNs) to be deployed on-demand using Tycoon. Briefly, Tycoon-gLite is an interface that offers the possibility of deploying a set of CEs and WNs ready to work for the user: it manages the host certificates, creates the necessary configuration files, deploys the nodes and configures them on-the-fly. Tycoon can then be operated as usual. Some scalability tests were also carried out in order to probe the value of Tycoon within a Grid environment. They permitted to identify improvements and indicated that this design provides Tycoon with good scalability potential. A full report detailing the Tycoon activities at CERN openlab can be downloaded from the openlab website.

The CERN openlab contributor EDS, an HP company, has continued its work in the grid monitoring domain. The GridMap application was developed by EDS in the joint project with CERN openlab during the first year of their collaboration, and is very popular. GridMap provides clear top level management views of the live state of the Grid in a content rich web page, which is interactive and can be used to fully explore the current state of the Grid. A number of upgrades were applied to this application and it has benefitted from various enhancements this year: it is now integrated in

the new Worldwide LHC Computing Grid (WLCG) Topology database and the US Open Science Grid sites. It can show sites with a scheduled downtime in a different colour plus WLCG metrics such as the «Installed Computing Capacity».

The GridMap technology has been reused at CERN to build a number of new LHC experiment specific grid monitoring applications (for instance, the «Critical Services Map», the «Site Status board GridMap» and the «Site Monitoring GridMap»). The development of GridMap has also been influential to Grid communities outside EGEE and WLCG like the EU funded D4Science research project. In addition, the GridMap technology is used within EDS in a number of projects, for example for providing visualisations of SOA services usage and network traffic monitoring. Furthermore CERN hired a Technical Student specifically dedicated to developing applications based on the GridMap technology who will also help to organise the different application branches into a single repository, so that GridMap can be used seamless at CERN beyond the EDS/CERN openlab collaboration.

Enterprise messaging systems are seen as a core enabling technology for the WLCG project as we move into the era of LHC data exploitation. The Messaging System for the Grid (MSG) project evaluated open-source messaging technology for use by WLCG operations. EDS contributed both effort and technical expertise in the testing and evaluation of various messaging implementations. These evaluations led to Apache ActiveMQ being chosen and it is currently being integrated into several projects within the WLCG environment, for example the distributed monitoring system used to check the reliability and availability of the WLCG sites.

GridMap visualising the 'state' of the Grid

Symmetric Mapping

The conclusions of the study on Symmetric Mapping carried out as part of the Grid Competence Centre programme by Xavier Gréhant are listed below. This study presented by Xavier is also part of his PhD thesis work:

"A new architecture for resource allocation developed in the Grid Competence Center programme of work augurs optimal compromises in resource sharing between organisations. This architecture, symmetric mapping, dispatches responsibilities between resource providers and users based on their respective concerns, namely the cost and the value of computing resources. It is in line with late binding, a recent technique for users to gain control in grids, which is gaining momentum in CERN's major virtual organisations. We extrapolated symmetric mapping to be the natural evolution of allocation systems in search for cost-effectiveness. It prepares for enterprise grids and the cloud computing era. Our implementation notably involves SmartFrog (HP Labs) and Condor (UoW)."
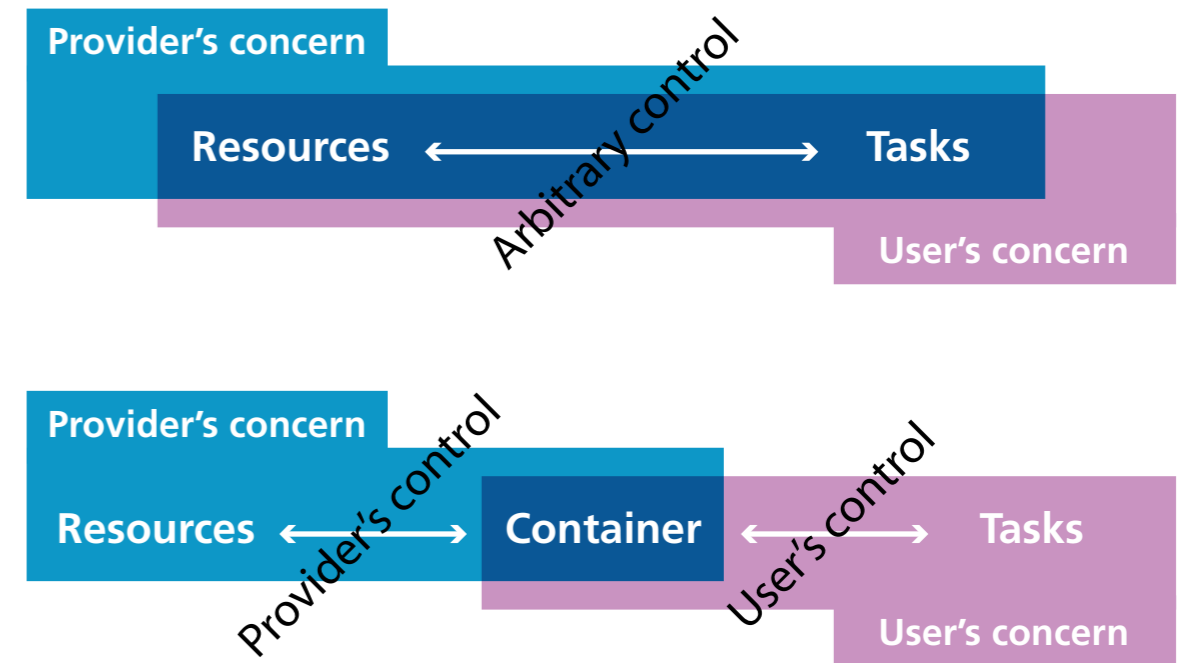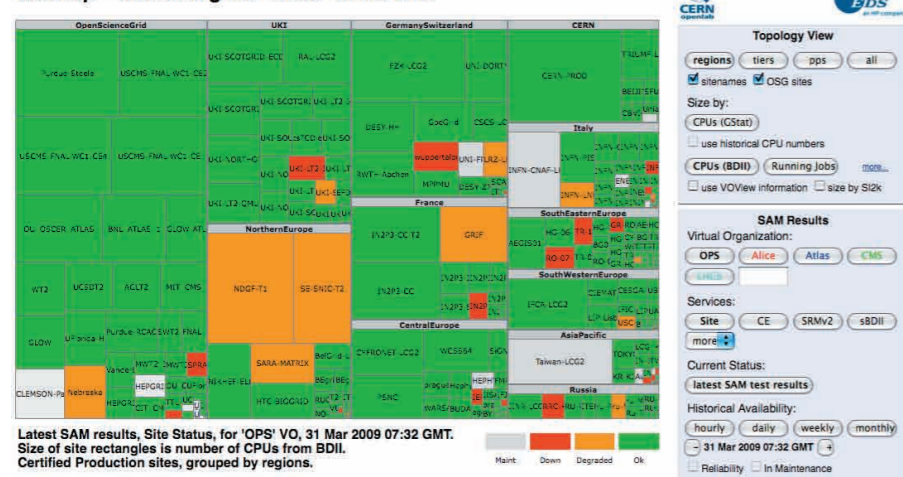

Figure: Traditional grid allocation is partially arbitrary, leading to prejudices for providers and users. Providers are required to support user tasks, while users receive unpredictable service. By inserting each contract in the allocation process in the form of a resource container, the responsibilities are dispatched. Providers choose among their resources to endorse their containers at lowest cost, while users allocate tasks to their subscribed containers for an optimal value.

Reference: Symmetric Mapping: an Architectural Pattern for Resource Supply in Grids and Clouds. Xavier Gréhant and Isabelle Demeure, SMTPS'09.

Relational Database Activity

Physics metadata stored in relational databases play a crucial role in the operation of the LHC experiments and Worldwide LHC Computing Grid (WLCG) services. A large proportion of non-event data such as detector conditions, calibration, geometry and production bookkeeping relies heavily on databases. Also, the core Grid services that catalogue and distribute LHC data cannot operate without a reliable database infrastructure at CERN and the main WLCG sites. The CERN openlab Relational Database Activity focuses on testing and practical implementation of new database techniques for the LHC Grid environment, as well as for CERN's general IT services.

The collaboration between Oracle and CERN is a particularly long-standing and fruitful one. The partnership started 27 years ago and Oracle decided to join CERN openlab from its early beginnings, in 2003. The CERN-Oracle partnership in openlab proved to be an incubator for innovation. Examples of this are the Maximum Availability Architecture technologies (Real Application Clusters – RAC, Streams and Data Guard) which are now used in production on a worldwide scale for key elements of data production and processing for the WLCG. Oracle has helped through several major software upgrades and the migration from Solaris to Linux, as well as giving advice on extending the products' capabilities. In return, the teams at CERN gather a lot of real user experience and share this knowledge at conferences and with the Oracle product and development teams. This has guided product enhancements, and helped Oracle to find new solutions, which they can pass on to other customers.

ATLAS Experiment

One of the latest improvements developed by the Relational Database Activity team streamlines the setup and operations of Streams. Oracle Streams is the main replication technology used for relational data distribution in the framework of WLCG. Different kinds of data (vital for data processing and analysis), such as detector conditions, calibration data, alignment information and accelerator conditions are stored in Oracle RAC databases and are replicated through a distributed database infrastructure between CERN and ten of the eleven tier-1 computer centres all over the world. Thanks to the new set of procedures implemented towards the end of 2008, the work that has to be done to keep the streams replication up and running reliably has been minimised. In particular, these new procedures facilitate and automate the operations required for temporarily suspending the distribution of data to a site – for example during scheduled maintenance – and then re-including it later, together with all pending updates.

The openlab collaboration enabled CERN to standardise and greatly improve the monitoring of the Oracle infrastructure with Oracle Enterprise Manager 10g Grid Control product, resulting in less time spent fire fighting and more time adding value for the organisation. Indeed, regular monitoring of the databases and streams performance is available to the LHC experiments and tier-1 sites, based on Oracle Enterprise Manager and an Oracle Streams monitoring tool developed at CERN. The latest improvements to this monitoring tool developed in-house at CERN, in collaboration with Oracle, have recently been integrated into Oracle Enterprise Manager 10g release 5.

Maria Girone with the Oracle Real Application Cluster (RAC) in the CERN Computer Centre
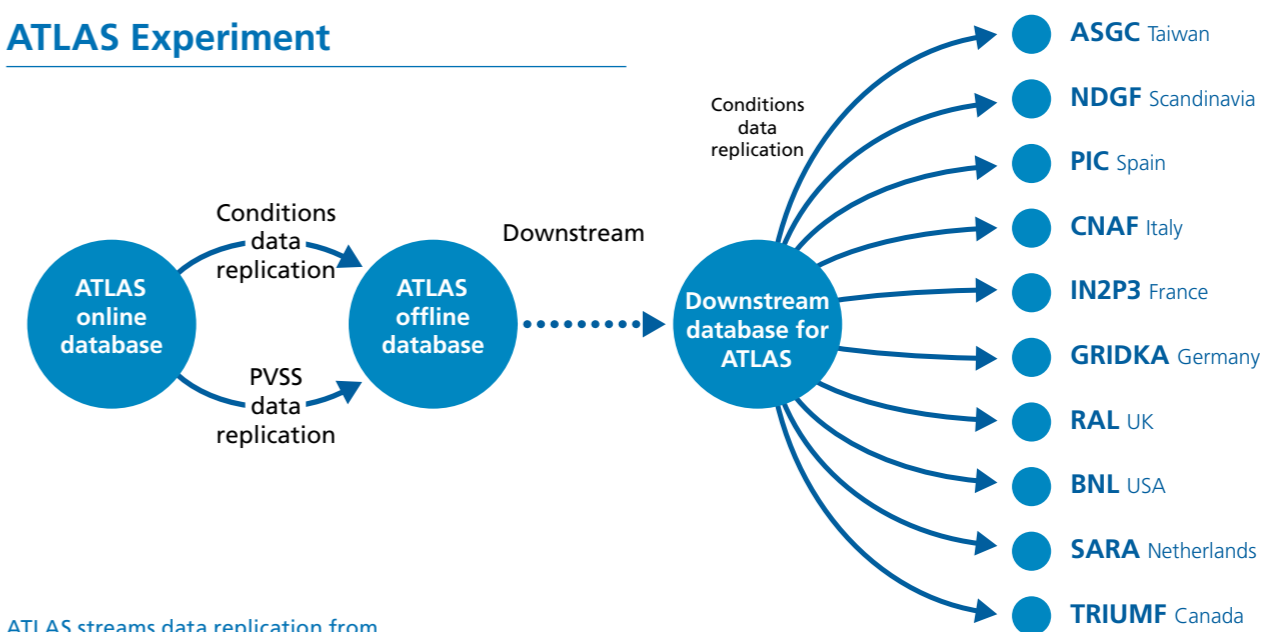
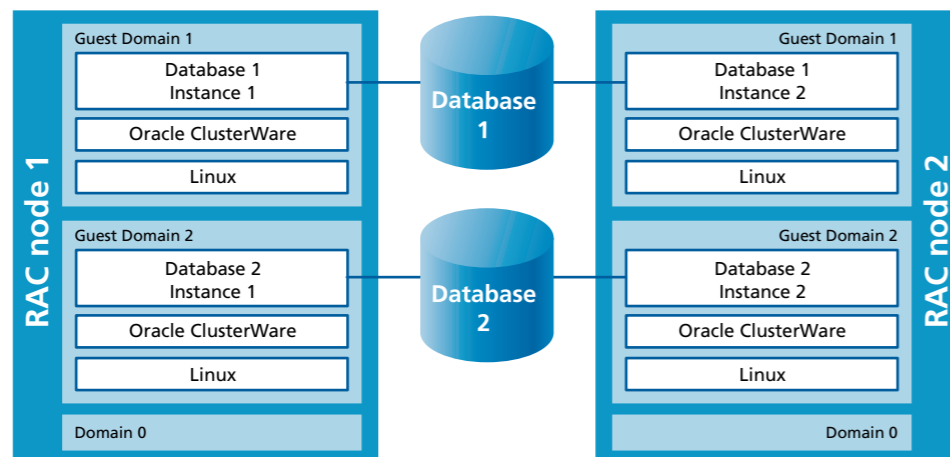ATLAS streams data replication from CERN to the ten tier-1 centres

15

Significant progress was achieved in the area of Oracle Data Guard deployment during 2008. The Data Management (DM) group capitalised on knowledge and experience acquired using Oracle Data Guard during migrations to new hardware. Last summer, real scale tests were performed and the implementation of Oracle Data Guard is now complete. The Data Guard configuration consists of one production (or primary) database and a number of standby databases deployed on different resources. If the production database goes down, the standby database can be declared as primary and be fully functional within a matter of minutes. This results in minimal impact on the service and represents a major improvement with respect to more conventional types of service recovery. The systematic use of Oracle Data Guard for the production databases has enhanced the protection of the experiments' data. It represents a crucial improvement for the service quality, the availability and stability of key services such as the databases for the LHC machine and experiments. Very recently, Active Data Guard – which comes with the latest release of Oracle – has also been tested at CERN. The LHC experiments look forward to using this very promising technology when it is deployed in production.

One of the demanding applications at CERN, requiring a very high and consistent rate for data-insertion, has been tested on the HP Oracle Database Machine, which is able to cope with this demand due to its unique Exadata storage server. The testing showed that this new unique Oracle offering fully provided the required performance. Additionally, Oracle continues to provide early versions of their existing software, which are put through CERN's regression testing including several different and demanding applications. Thanks to this testing, the openlab teams are providing valuable feedback to Oracle and CERN, in return, gains insight into upcoming Oracle technology.

The optimisation of PVSS, a Supervisory Control And Data Acquisition (SCADA) system used widely in the LHC and associated experiments, is also a key objective of the Relational Database Activity. As already reported last year, the out-of-the box performance for running PVSS with Oracle database allowed the archiving of 100 changes per second, but through optimisation of the application code and usage of advanced Oracle features like RAC and partitioning, CERN openlab was able to achieve a practical target of 150,000 changes per second at a stable throughput while queries are being performed on the system. Several of these improvements have recently been integrated by Oracle in the new version 3.8. Oracle and CERN will keep working very closely in this area.

The first project within the openlab collaboration focusing on virtualisation started in early 2006. The emphasis was on the possibility to run a single instance database in a Virtual Machine (VM), as well as testing different Xen schedulers to allocate Central Processing Unit (CPU) slices to different VMs. Virtualisation gives a better usage of hardware, less power consumption, a better isolation of the user applications and offers a greater ease of management. Upon completion of this first project, the tests were expanded by testing Oracle RAC in such an environment. A few months after this study was performed, Oracle announced in November 2007 its own Xen-based virtualisation software – Oracle VM. During summer 2008, together with Andrei Dumitru, a CERN openlab summer student, the team continued the tests and compared RAC performance on Oracle VM versus the native Xen setup tested before. The performance of Oracle VM setup proved to be 5 to 20 % better than on pure Xen, depending on workload, which is an outstanding addition to the high availability features initially provided by RAC. A first implementation for some of CERN's non-critical databases will result in a reduction of hardware and infrastructure costs. From 2009 onwards, with Oracle entering CERN openlab-III, further developments are expected in these domains as well as studies related to the upcoming Oracle 11g release 2.

RAC node 1

Guest Domain 1
Database 1 Instance 1
Oracle ClusterWare
Linux

Guest Domain 2
Database 2 Instance 1
Oracle ClusterWare
Linux

Domain 0

Database 1

Database 2

RAC node 2

Guest Domain 1
Database 1 Instance 2
Oracle ClusterWare
Linux

Guest Domain 2
Database 2 Instance 2
Oracle ClusterWare
Linux

Domain 0

An objective of database virtualisation is to fight underutilisation by combining several clustered databases on two physical servers

16

Networking and Security Activity

The CINBAD (CERN Investigation of Network Behaviour and Anomaly Detection) project was launched in 2007 in collaboration with ProCurve Networking by HP. This is a challenging research activity as it must address large scale issues, requiring the collection and storage of huge quantities of data. The project mission is to understand the behaviour of large computer networks in the context of high performance computing and campus installations such as CERN, whose network counts today roughly 70,000 Gigabit user ports. The goals of the project are to detect traffic anomalies in such systems, perform trend analysis, automatically take counter measures and provide post-mortem analysis facilities. It is divided into three phases: data collection and network management, data analysis and algorithm development, performance and scalability analysis. This research activity is already producing practical results as well as providing additional information to the CERN security team.

The starting point of the project was to define the requirements and ensure a common framework of precise definitions, for example what constitutes an anomaly or a trend. Anomalies are nowadays a fact in computer networks. However anomaly definition is very domain specific and the causes are diverse (network faults, malicious attacks, viruses and worms, misconfiguration, etc.). Numerous brainstorming sessions were organised with HP ProCurve and the team conducted a great number of interviews with experts working in various fields to get different perspectives. The following common denominator emerged: an anomaly is always a deviation of the system from the normal (expected) behaviour (baseline); the normal behaviour (baseline) is never stationary and anomalies are not always easy to define. As a consequence, anomalies are not easy to detect. However some potential anomaly detectors can be identified. Thus, the use of statistical detection methods can be considered. By learning the 'normal behaviour' from network measurements, and continuously updating the 'normal baseline', it is possible to detect new, unknown anomalies. Applying such a method also has some drawbacks as it is still possible to attempt to force a false negative, the selection of suitable input variables is an issue (many anomalies being within 'normal' bounds of the metrics), and finally a false positive can be extremely costly and does not provide a satisfactory anomaly type identification.

HP and CERN had already worked on packet sampling in the 90s. Indeed, sFlow, the industry standard technology for monitoring high speed switched networks, is derived from the collaboration between HP, the University of Geneva and CERN in 1991. Given the size of CERN's networks and considering the fact that HP devices present at CERN are able to perform sFlow sampling, the team decided that the best way to monitor the network on a large scale was to use data statistical analysis by packet sampling. The first results in the area of packet sampling were published at the end of 2007 in a report 'Packet Sampling for Network Monitoring', by Milosz Hulbój and Ryszard Jurga, two openlab team members. The report is available on the CERN openlab website. It emphasises various packet sampling methods and their application to network monitoring and presents the motivations for packet sampling on high-speed network links.

An overview of all known packet sampling techniques is given as well as their respective strengths and weaknesses in terms of reliability and accuracy, with different network traffic parameters. Prior to the publication of the report, it was already known that packet sampling provides reliable estimates in traffic monitoring applications, but not much research had been carried out on whether packet sampling provides a sufficient amount of information for anomaly detection. For this reason, one of the key objectives of this research was to gain some insight into the feasibility of packet sampling in the context of network anomaly detection.

Milosz Hulbój (left) and Ryszard Jurga (right) working on the network switches in the CERN Computer Centre

17

These studies, complemented by an in-depth analysis of the sFlow agents, enabled the CINBAD team to design and implement the sFlow data collector. Given the huge amount of sFlow data (300'000 samples/second) to be collected and analysed, the team decided to benefit from CERN's know-how in data storage and analysis. During the survey on data acquisition, the LHC experiments and Oracle experts were consulted on high performance data storage, data format and representation and analysis principles. This survey led the team to the design of the multi-stage sFlow collector (see figure below) and to implement it. In the summer of 2008, the collection system was successfully tested on a large scale network, using approximately 500 HP switches.

These data collection tests enabled the team to develop tools for analysing the stored data. Various data analysis techniques have been tested, among them a statistical data analysis and time series mining and signature based

approach. In addition to these tools, the team adapted SNORT (open source network intrusion prevention and detection system) to work with sampled sFlow data. This SNORT setup was complemented by open source traffic rules as well as in-house CINBAD rules. Initial data analysis has enabled the team to detect a number of misbehaviours and anomalies in the CERN network. It appears that most of these security anomalies (malicious software, policy violations) originated from end user machines.

During the third year of the project, the CINBAD team will consolidate the activities in order to capture the findings and results of their investigations. Currently the project promises valuable findings, since even incomplete network data provides useful information about the network status and behaviour. Combining statistical analysis with pattern matching has provided encouraging initial results and has motivated to continue this fruitful collaboration.

sFlow data collector
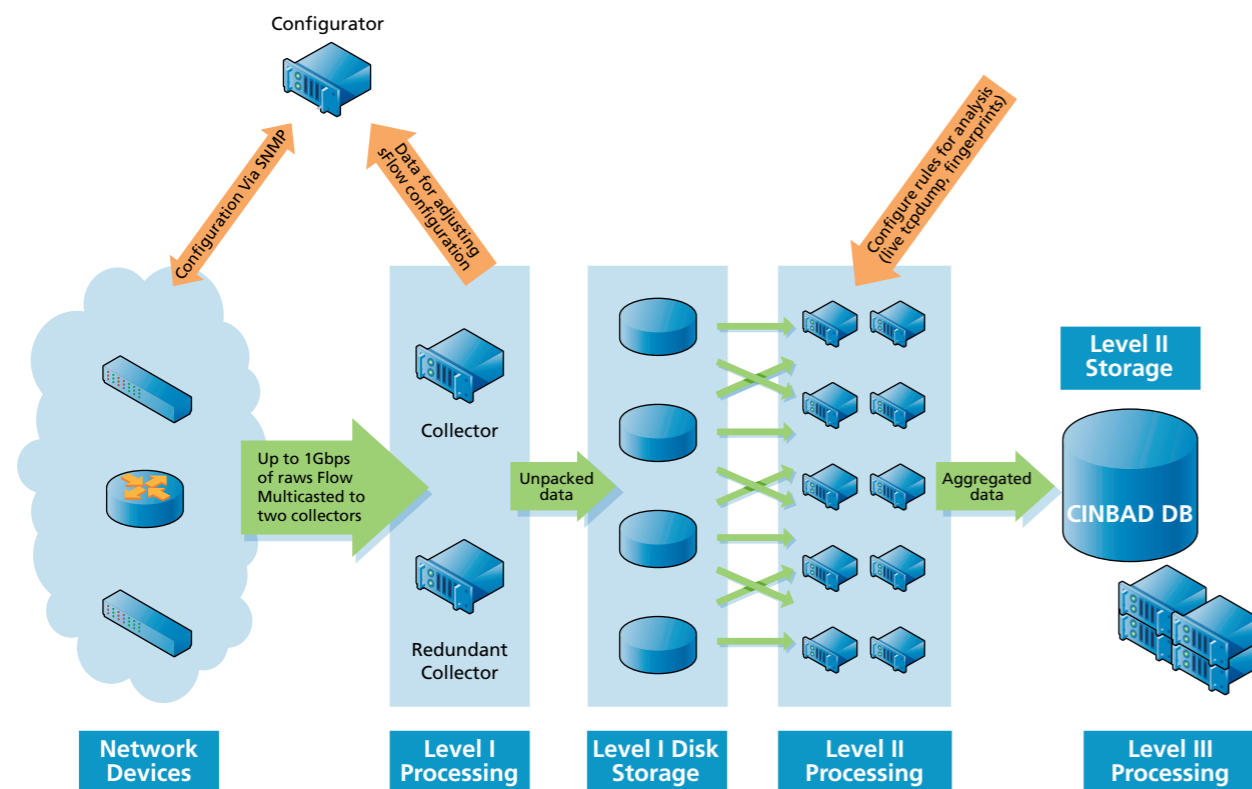
Publications and Presentations

CERN openlab results have been disseminated in a wide range of international conferences, listed below. For a full record of the presentations, consult the CERN openlab website. In addition, key results from CERN openlab have been the subject of a large number of press articles both in the general and IT-specific press and on the Web. These articles are listed on the CERN openlab website.

Presentations

- A. Hirstius/CERN, CPU-Level Performance Monitoring with perfmon/pfmon, HEPiX, CERN, 5 May 2008
- S. Jarp/CERN, A Review of the Current Technical Activities in the CERN openlab, HEPiX, CERN, 7 May 2008
- S. Jarp/CERN, Faire face aux nouvelles architectures de processeurs : la physique des particules est-elle prête ?, LAPP Seminar, Annecy, France, 13 May 2008
- A. Nowak/CERN, High-throughput computing optimization issues at CERN, Bioinformatics in Torun, Torun, Poland, 14 June 2008
- H. Bjerke/CERN, High Throughput Computing for CERN's Large Hadron Collider, ISCA, Beijing, China, 22 June 2008
- S. Jarp/CERN, An Overview of CERN's Approach to Energy Efficient Computing, IDC 'Green IT' Conference, Milan, Italy, 25 June 2008
- X. Gréhant/ENST&CERN and S. Jarp/CERN, Lightweight Task Analysis for Cache-Aware Scheduling on Heterogeneous Clusters, PDPTA, WorldComp, Las Vegas, USA, July 2008
- H. Bjerke/CERN, Tools and Techniques for Managing Virtual Machine Images, VHPC'08, Gran Canaria, Spain, 26 August 2008
- M. Lally/Ingersoll Rand, C. Lambert/CERN, A. Oppenheim/Oracle, One-Stop Asset Tracking, Configuration Analytics, and Policy Compliance: Oracle Enterprise Manager Configuration Management, Oracle Open World Conference, San Francisco, USA, 22 Sept. 2008
- D. Rodrigues/CERN, Messaging System for the Grid, EGEE'08, Istanbul, Turkey, 24 September 2008
- S. Jarp/CERN, Faire face aux nouvelles architectures de processeurs : la physique des particules est-elle prête ?, JI'08, Obernai, France, 30 September 2008
- A. Topurov/CERN, CERN Experience with Virtualization of Oracle RAC with Native Xen and Oracle VM, TrivadisOpen, Zurich, Switzerland, 22 October 2008
- S. Jarp/CERN, Forget multicore! The future is manycore: An outlook to the explosion of parallelism likely to occur in the LHC era, ACAT'08, Erice, Italy, 6 November 2008
- E. Grancher/CERN, Oracle and storage IOs, explanations, experience at CERN and SSD tests, UKOUG conference, Birmingham, UK, 2 December 2008
- A. Topurov/CERN, CERN Experience with Virtualization of Oracle RAC with Native Xen and Oracle VM, UKOUG Conference, Birmingham, UK, 2 December 2008
- E. Grancher/CERN, Learning from failures, design errors, problematic recoveries and downtimes of Oracle databases, experience at CERN, UKOUG conference, Birmingham, UK, 3 December 2008
- L. Canali/CERN and D. Wojcik, Implementing ASM without HW RAID, a user's experience, UKOUG Conference, Birmingham, UK, Dec. 2008
- M. Hulbój/CERN and R. Jurga/CERN, Statistical analysis, traffic signatures and spoofed mac/ip detection, HP ProCurve Architect Forum, Roseville, USA, 8 January 2009
- J. M. Dana/CERN and W. A. Romero/Summer Student, Performance Monitoring of the Software Frameworks for LHC Experiments, EELA-2 Conference, Bogotá, Colombia, 25-26 February 2009
- X. Gréhant/ENST&CERN, Efficient and non Obtrusive Resource Allocation in Grids, Winter School on Hot Topics in Distributed Computing, La Plagne, France, 15-20 March 2009
- E. Grancher/CERN, Oracle and Storage IOs, Explanations and Experience at CERN, CHEP'09, Prague, Czech Rep. , 21-27 March 2009
- A. Nowak/CERN, An Update on Perfmon and the Struggle to

get into the Linux Kernel, CHEP'09, Prague, Czech republic, 21-27 March 2009
- S. Jarp/CERN, Experience with Low-Power x86 Processors (ATOM) for HEP Usage, CHEP'09, Prague, Czech rep. , 21-27 March 2009
- M. Girone/CERN, Distributed Database Services – a Fundamental Component of the WLCG Service for the LHC experiments – Experience and Outlook, CHEP'09, Prague, Czech republic, 21-27 March 2009
- I. Demeure/ENST and X. Gréhant/ENST&CERN, Symmetric Mapping: an Architectural Pattern for Resource Supply in Grids and Clouds, SMTPS, IPDPS, Rome, Italy, May 2009

Publications

- X. Gréhant/ENST&CERN and S. Jarp/CERN, Lightweight Task Analysis for Cache-Aware Scheduling on Heterogeneous Clusters, PDPTA, WorldComp, July 2008
- H. Bjerke/CERN, Tools and Techniques for Managing Virtual Machine Images, VHPC'08, August 2008
- A. Hirstius/CERN, The Large Hadron Collider, Physics World, Nov. 2008
- J. M. Dana/CERN and W. A. Romero/Summer Student, Performance Monitoring of the Software Frameworks for LHC Experiments, EELA-2 Conference, Bogotá, Colombia, February 2009
- I. Demeure/ENST and X. Gréhant/ENST&CERN, Symmetric Mapping: an Architectural Pattern for Resource Supply in Grids and Clouds, SMTPS, IPDPS, May 2009

Posters

- J. Casey/CERN, D. Rodrigues/CERN, Messaging Systems for Grid, EGEE'08, Istanbul, Turkey, September 2008
- M. Girone/CERN , Distributed Databases Services in the EGEE grid and beyond, EGEE'08, Istanbul, Turkey, September 2008

CERN openlab Reports

- N. Basha/Summer Student, CINBAD Investigation of Different Packet Filters, August 2008
- X. Dong/Summer Student, Multi-Threaded Geant4 with Shared Detector, August 2008
- P-L. Hémery/Summer Student, Improving Display and Customization of Timetable in Indico, August 2008
- W. A. Romero/Summer Student, Performance Monitoring of the Software Frameworks for LHC Experiments, August 2008
- K. Sarnowska/Summer Student, The SNARL Service: Standards-based Naming for Accessing Resources in an LFC, August 2008
- A. D. Dumitru/Summer Student, Oracle RAC Virtualization, September 2008
- A. D. Dumitru/Summer Student, A. Topurov/CERN, Oracle RAC Virtualization – Installation Guide, September 2008
- E. Grancher/CERN, A. Topurov/CERN, CERN PVSS Tests on SAGE/Exadata, November 2008
- G. Balazs/CERN, S. Jarp/CERN, A. Nowak/CERN, Is the Atom Processor Ready for High Energy Physics? An Initial Analysis of the Dual Core Atom N330 Processor, November 2008
- G. Balazs/CERN, Performance Monitoring of Production Batch Servers, January 2009
- J. M. Dana/CERN, Tycoon Activities at CERN openlab, Jan. 2009
- H. Bjerke/CERN, Efficient Parallel Tracking, January 2009
- G. Balazs/CERN, Increasing Power Efficiency of Computer Centres, February 2009

Education

The education programme is an essential pillar of the openlab activities. Indeed, CERN openlab is a framework for industrial collaboration where expertise and knowledge is created, exchanged and disseminated. Regular training sessions take place, involving CERN and partner experts in advanced domains such as optimized programming of modern processors. CERN openlab also contributes to the activities of the CERN School of Computing.

This year, CERN openlab expanded and enriched its portfolio of workshops. A new type of workshop related to performance tuning and hardware architecture was offered to CERN programmers. The tried-and-tested formula of the multithreading workshop involving a mix of experienced instructors from Intel and CERN proved to be very successful again this year. With the participation of Oracle, a Database Performance Seminar was also organised for the first time. A tutorial on SIMD (Single Instruction Multiple Data) Programming and an introduction to Many-Core Programming were delivered as part of the ALICE workshop on Advanced Programming. These workshops are proposed to users affiliated with CERN, free of charge. These lectures and workshops are listed below.

- **Multi-threading and Parallelism Workshop**, 29-30 May 2008, CERN, S. Jarp/CERN, A. Nowak/CERN
- **Database Performance Seminar**, 6 June 2008, CERN, RAC scalability for logging application, E. Grancher/CERN
- **Database Performance Seminar**, 6 June 2008, CERN, Scalable Application Design, B. K. Engsig/Oracle
- **Database Performance Seminar**, 6 June 2008, CERN, Oracle Streams performance, E. Dafonte Pérez/CERN
- **Database Developers' Workshop**, 8 July 2008, CERN, many openlab team members from CERN IT Data Management Group delivered lectures.
- **CERN School of Computing**, August September 2008. Gjøvik, Norway. Three lectures: Computer Architecture and Performance Tuning, S. Jarp/CERN, A. Nowak/CERN.
- **CERN School of Computing**, August September 2008. Gjøvik, Norway. Three lectures: Networking QoS and Performance, F. Flückiger/CERN.
- **Machine Learning Summer School**, 8 September 2008, Ile de Ré, France, The CINBAD Project, M. Hulbój/CERN, R. Jurga/CERN
- **Computer Architecture and Performance Tuning Workshop**, 14 October 2008, S. Jarp/CERN, A. Nowak/CERN
- **Distributed Database Operations Workshop**, 11-12 November 2008, CERN, many openlab team members from CERN IT Data Management Group delivered lectures.
- **Multi-threading and Parallelism Workshop**, 11-12 November 2008, J. Arnold/Intel, S. Jarp/CERN, A. Nowak/CERN.
- **ALICE Workshop on Advanced Programming**, 4-6 December 2008, Introduction to Many-Core Programming, S. Jarp/CERN.
- **ALICE Workshop on Advanced Programming**, 4-6 December 2008, SIMD Programming tutorial, H. Bjerke/CERN.
- **Computer Architecture and Performance Tuning Workshop**, 12-13 March 2009, J. Arnold/Intel, S. Jarp/CERN, A. Nowak/CERN.
- **Lectures for grad students of Télécom-ParisTech, Grilles de calcul**, April 2009, X. Gréhant/CERN&ENST
- **Distributed Database Workshop**, 20-21 April 2009, CERN, many openlab team members from CERN IT Data Management Group delivered lectures.

Jeff Arnold, Senior Software Engineer at Intel, delivering a lecture during latest CERN openlab Computer Architecture and Performance Tuning Workshop

These direct training activities are complemented by the CERN openlab Student Programme. This programme was launched in 2002 to enable undergraduate, Masters and Ph.D. students to get hands-on experience with Grid technology and other advanced openlab-related topics. In 2008, the programme accepted 13 computer science and physics students from 12 countries for two months, during the period June to September. The students worked on cutting-edge computing technologies supervised by openlab staff, other groups in the IT Department as well as staff from the WLCG and EGEE Grid projects.

Visits were organised to the CERN Control Centre, ATLAS, LINAC 2, the Anti-matter Factory as well as a study tour to the Ecole Polytechnique Fédérale de Lausanne (EPFL). A dedicated lecture series for the students was given by CERN and external experts. Several of this year's students were co-funded by Intel and Oracle.

CERN openlab Summer Student Programme Teaching Series, July-August 2008:

- Grid overview and gLite details, A. Unterkircher/CERN
- Server Hardware, A. Hirstius/CERN
- Virtualisation, H. Bjerke/CERN
- Compilers, J. M. Dana/CERN
- Creating Secure Software, S. Lopienski/CERN
- Computer Architecture and Performance Optimisation, S. Jarp & A. Nowak/CERN
- LCG, L. Poncet/CERN
- Networking, R. Jurga/CERN,
- Oracle Database Architecture, B. K. Engsig/ORACLE

The CERN openlab summer students 2008 are listed below, with home institute and project topic:

- Alejandro Alvarez Ayllón, Spain, University of Cádiz, (CERN School of Computing)
- Natalya Basha, Russia, International University of Dubna, (CINBAD investigation of different packet filters)
- Paola Cecuk, Croatia, University of Electrical Engineering and Computing – Zagreb, (Study of SPEC2006 sub-benchmarks and analysis of 'perfmon' performance monitoring data collected on batch nodes)
- Xin Dong, China, Northeastern University - Boston, (Geant4 multithreaded implementation)
- Andrei-Daniel Dumitru, Romania, University of Bucharest, (Database virtualization)
- Ilya Gorbunov, Russia, Moscow State University, (Virtual software system for physics simulation)
- Pierre-Luc Hémery, France, ENST, (Improving display and customization of timetables in Indico)
- David Horat, Spain, University of Las Palmas de Gran Canaria, (SVN pilot service consolidation)
- Shuaib Khan, Pakistan, National University of Sciences and Technology - Rawalpindi, (Performance monitoring system development)
- Spiros Koulouzis, Greece, University of Amsterdam, (System to automatically replicate grid data between sites)
- Christopher Parker, South Africa, University of Cape Town, (Deduplication and publication workflow for CDS collections)
- William Alberto Romero Ramirez, Colombia, University de los Andes – Bogotá, (Benchmarks and performance monitoring)
- Karolina Sarnowska, Poland, USA, University of Virginia – Charlottesville, (SNARL service: standards-based naming for accessing resources in an LFC)

CERN openlab students at CERN, August 2008

Events and Outreach

As well as the excellent technical results that CERN openlab provides, the partnership gives CERN a means to share a vision of the future of scientific computing with its partners, through joint workshops and events, as well as to disseminate this to a wider audience, including partner clients, the press and the general public.

CERN is regularly toured 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.

TIC Région Rhône Alpes Visit, CERN, 2 July 2008
Introduction to CERN openlab as part of the Computer Centre visit and a meeting focusing on LHC Computing.

Intel Press Event for Dunnington Launch, Zurich, 16 September 2008
Presentation given by Sverre Jarp, 'CERN openlab and Intel: Today and Tomorrow', and participation in the roundtable discussion.

HP ProCurve Customer Visit, CERN, 2 October 2008
Tour of CERN and review of CERN openlab results for HP ProCurve corporate customers.

LHC Grid Fest, CERN, 3 October 2008
Tour of CERN given to HP ProCurve and Intel Press with interview access to senior people from CERN and partners. LHC Computing Awards presented to Intel and Oracle by Dr Robert Aymar, Director General of CERN, presentation of the ALICE award given to HP in 2007.
Partner and CERN openlab presentations:
- B. Johnson, HP, 'HP and CERN openlab: an R&D Partnership'.
- S. Pawlowski, Intel, 'The Drive to Improved Performance/ watt and Increasing Compute Density'.
- S. Rousset, Oracle, 'When Research Meets Industry'.
- W. von Rüden, CERN. 'CERN's Partnership with IT Industry'.

Industry Award Day, CERN, 20 October 2008
Focus on the contribution from CERN openlab partners to the LHC Computing and speeches from Patrick P. Gelsinger from Intel and Sergio Giacoletto from Oracle.

HP ProCurve Collaborators' Visit, CERN, 19 Nov. 2008
Tour of CERN and review of CERN openlab results for HP ProCurve collaborators.

Visit of Craig Barrett, Intel Chairman of the Board, plus IT colloquium, CERN, 27 January 2009
Tour of CERN, CERN IT colloquium presentation 'Inspiring Innovation', and interview with the press.

Stéphane Rousset from Oracle receiving the LHC Computing Award during the LHC Grid Fest

Patrick P. Gelsinger from Intel delivering its presentation during the Industry Award Day

The Board of Sponsors

This is the third annual report of the second phase of CERN openlab, which was launched in 2006, and the first one of its third phase, which was launched in 2009. It was presented to the Board of Sponsors at the Annual Sponsors meeting, 02-03 April 2009. Present at the meeting were:

Front row, left to right: Sverre Jarp (CERN), Alberto Pace (CERN), Mats Möller (CERN), François Flückiger (CERN), Séverine Pizzera (CERN), Eva Dafonte Pérez (CERN), Monica Marinucci (Oracle), Wolfgang von Rüden (CERN), Juergen Hirte (Siemens), Arnaud Pierson (HP), Mélissa Le Jeune (CERN), Renaud Barillère (CERN), Andrzej Nowak (CERN).

Second row, left to right: Jamie Shiers (CERN), Ian Bird (CERN), Sergio Cittolin (CERN), Claudio Bellini (Intel), Stephan Gillich (Intel), Bjørn Kisbye Engsig (Oracle), Rolf Heuer (CERN), Thomas Hahn (Siemens), Bernard Reichl (ETM, a Siemens company), Rolf Sauter (HP), Jean-Michel Jouanigot (CERN), Dan Ford (HP ProCurve), Max Böhm (EDS, an HP company), Xavier Gréhant (CERN).

THE FUTURE

New projects, new partner joining

CERN openlab continues to develop on several fronts with the start of the third phase. New challenging projects have already been initiated with the partners. Additional team members have recently joined and the structure is now in place to collaborate and work on bringing these projects to fruition. The projects are now structured in four Competence Centres (CC): Automation and Controls CC, Database CC, Networking CC, Platform CC.

Siemens joined the third phase of CERN openlab as a partner, enriching the framework activities portfolio with a new dimension and giving birth to the Automation and Controls Competence Centre. The collaboration focuses on security, opening automation tools towards software engineering, and handling large environments.

Security

Due to the growing usage of Ethernet and TCP/IP in automation devices, and because of the move away from proprietary or dedicated networks, the automation devices and control applications have to become resistant to the common threats on Ethernet cable. These threats can be deliberate (attackers), collateral (viruses and worms), or accidental (misconfiguration, such as an error in the IP address or broken devices flooding the network). Siemens and CERN have an inherent interest that automation devices (e.g. PLCs) survive these kinds of attacks and they will be investigating the resistance of the devices. More specifically, resistance to malicious network traffic will be analysed through robustness tests and vulnerability tests. There is a similar interest for PVSS-based control applications.

Software engineering

PVSS is used at CERN for large distributed control systems, some with more than 150 computers. PVSS and the controls applications have to be deployed on these computers which requires an initial installation as well as regular upgrades and patches for both PVSS software and control applications. However, PVSS does not offer any native facility to deal with this type of environment. Furthermore, the tools developed at CERN for software deployment – CERN Management Framework (CMF) and Linux for Controls – are only appropriate for the initial PVSS installation. Thus, the installation of the applications (software + configuration) has to be handled manually, which is also true for the PLC layer. CERN and Siemens will collaborate to build adequate deployment solutions for the supervision and process control layers of large control applications.

Large environment

CERN control systems will be operated and maintained for a very long period of time (about 15 years). Thus, the PVSS-based control applications and the user programs running in the corresponding PLCs have to follow the same cycle. During this time they have to be managed, maintained, adapted and extended. In addition, the increasing complexity of the control systems sometimes leads to the production of the PVSS and PLC code with the help of external tools (e.g. model-driven ones). Siemens and CERN are interested in following the convergence trends between the automation and the information technology worlds.

ETM Professional Control is a subsidiary of Siemens AG, organisationally and functionally assigned to the Industry Automation Division. The company develops the process visualisation and control system PVSS, used for the automation of complex, geographically widely distributed Systems, with high security and availability requirements.

Ildefons Magrans de Abril and Daniel Rodrigues in front of PVSS screens in the CERN Control Centre

About the partners

CERN openlab partners

HP, the world's largest technology company, simplifies the technology experience for consumers and businesses with a portfolio that spans printing, personal computing, software, services and IT infrastructure. More information about HP is available at www.hp.com.

Intel, the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live. Additional information about Intel is available at www.intel.com/pressroom.

For 30 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.

Siemens AG (Berlin and Munich) is a global powerhouse in electronics and electrical engineering, operating in the industry, energy and healthcare sectors. The company has around 430,000 employees (in continuing operations) working to develop and manufacture products, design and install complex systems and projects, and tailor a wide range of solutions for individual requirements. For over 160 years, Siemens has stood for technological excellence, innovation, quality, reliability and internationality. www.siemens.com.

CERN openlab contributor

EDS, an HP company, is a leading global technology services provider, delivering business solutions to its clients. EDS founded the information technology outsourcing industry nearly 50 years ago. Today, EDS delivers a broad portfolio of information technology, applications and business process outsourcing services to clients in the manufacturing, financial services, healthcare, communications, energy, transportation, and consumer and retail industries, and to governments around the world. www.eds.com

For more information:

Executive Contact:
Wolfgang von Rüden, Head of CERN openlab,
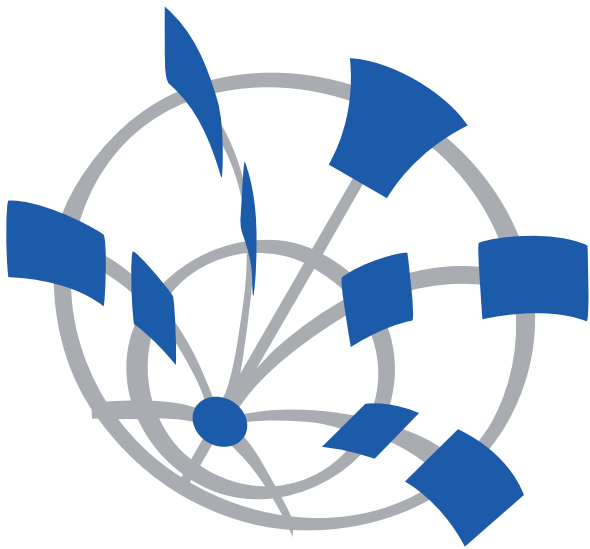Wolfgang.von.Rueden@cern.ch

Technical Contact:
François Flückiger, Manager of CERN openlab,
Francois.Fluckiger@cern.ch

http://www.cern.ch/openlab

CERN, the European Organization for Nuclear Research, is the world's leading laboratory for particle physics. It 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.
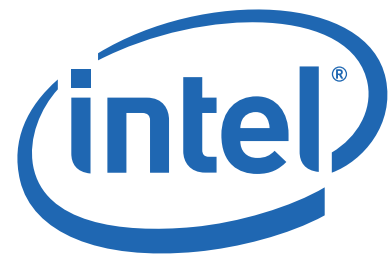
CERN openlab

www.cern.ch/openlab

hp invent

intel®

ORACLE®

SIEMENS

EDS
an HP company